Heisenberg-picture TQFT

Theo Johnson-Freyd, Northwestern University

Representation Theory, Geometry, and Combinatorics
UC Berkeley, 2 May 2014

The two goals for this talk are:

1. Motivate and then define the "Heisenberg picture" of (functorial, bordism-full) QFT.
2. Describe some examples in three dimensions, namely a version of quantum Chern–Simons theory which is fully extended and "at generic level."

A word of warning: this work is still in progress, although I hope the papers will appear soon.

1 Pictures of quantum mechanics

To motivate the "Heisenberg picture" of qft, let me begin by recalling some flavor of the Schrödinger and Heisenberg pictures of quantum mechanics. As an aside, what is a "picture"? In physics, systems just *are*, and our job is to construct mathematical models that capture some of their qualities. So names are supposed to suggest meaning. In mathematics, instead we give definitions and axioms. The point is that two different pictures of the same system should be related, but don't necessarily capture exactly the same data. Anyway, I'll word things like a mathematician.

The most familiar picture of quantum mechanics is probably the Schrödinger picture.

Definition: A *Schrödinger-picture quantum mechanical system* consists of the following data:

1. A Hilbert space $\mathcal{H}$, called the "space of states."

This describes the kinematics. For our purposes, the word "Hilbert" really doesn't matter. We might as well say:

1'. An abelian group $\mathcal{H}$, called the "space of states."

What matters is that $\mathcal{H}$ is an object of linear algebra.

The dynamics is encoded by:

2. For each interval $[t_1, t_2] \subseteq \mathbb{R}$, a unitary operator $U_{t_2,t_1} : \mathcal{H} \to \mathcal{H}$. These should satisfy a "group law" $U_{t_3,t_2} U_{t_2,t_1} = U_{t_3,t_1}$.

Again, it's the group law that really matters — you can replace "unitary" by "linear."

Finally, we should be able to set up experiments and ask questions, so there should be:

3. Some distinguished states $|v_i\rangle \in \mathcal{H}$ that we know how to prepare, and some distinguished costates $\langle w_j| \in \mathcal{H}^\vee$ that we know how to postpare (post-pair?). Here $\mathcal{H}^\vee$ is the dual abelian group or module or whatever over whatever ground ring we're working. In applications, maybe there's extra structure, like some subset of the distinguished states is an orthonormal basis, or $\langle w_j| = \langle w_j, -\rangle$ for $|w_j\rangle$ a distinguished state, or whatever.

   Moreover, there might be some distinguished manipulations we can do, transforming states in ways that aren't just "flow for some time." These would be encoded by elements of $\mathrm{End}(\mathcal{H})$.

This is set up already to suggest the Atiyah–Segal–others picture of QFT. There is a *spacetime category* SPACETIME whose objects $*_t$ are points labeled by times $t \in \mathbb{R}$, and whose morphisms are closed intervals in $\mathbb{R}$. Data 1 and 2 above together package into a functor SPACETIME $\to$ VECT, where VECT is the category of "Hilbert spaces" (abelian groups or whatever). To incorporate data 3, we can adjoin to SPACETIME a little more stuff, namely: another object, called $\emptyset$; some labeled morphisms $v_i : \emptyset \to *_t$ and $w_j : *_t \to \emptyset$; and maybe some other labeled morphisms between different points. Then you can draw diagrams of processes.

There is one defect to the Schrödinger picture that I want to highlight. Namely, it's not literally true that the space of physical states *is* $\mathcal{H}$. Rather, the physical states are elements of the projectivization $\mathbb{P}\mathcal{H}$. So perhaps one should say that $\mathbb{P}\mathcal{H}$ is assigned to the point, and projective operators rather than unitary ....

The Heisenberg picture is perhaps most familiar from deformation quantization, which I'm not going to address. It says that "space of states" is secondary, and most important is "algebra of observables."

Definition: A *Heisenberg-picture QM system* consists of:

1. A (C^* or von Neumann or…) associative algebra $\mathcal{A}$, called the "algebra of observables."
2. For each interval $[t_1, t_2] \subseteq \mathbb{R}$, an isomorphism $F_{t_2,t_1} : \mathcal{A} \to \mathcal{A}$. These should satisfy a group law.
3. [Something I will come back to.]

The Heisenberg picture is more general than the Schrödinger picture: it captures classical mechanics as well (when $\mathcal{A}$ is commutative). Can we compare the two pictures directly? Yes. Indeed, there is a functor from Schrödinger-picture QM systems to Heisenberg-picture QM systems:

1. $\mathcal{H} \quad \mapsto \quad \mathcal{A} = \mathrm{End}(\mathcal{H})$.
2. $U \quad \mapsto \quad F =$ conjugate by U.
3. How can you encode a vector $|v\rangle \in \mathcal{H}$ in terms of $\mathrm{End}(\mathcal{H})$? You can't, really. But you can encode the line it spans, which really is the only physical data. Namely, you encode it as the annihilation ideal $I = \{a \in \mathcal{A} \text{ s.t. } a|v\rangle = 0\}$.

So this suggests that 3 above should be:

3'. Some (left or right) ideals in $\mathcal{A}$,

since at least that works for the $|v_i\rangle$s and the $\langle w_j|$s.

This is a very successful encoding, since the Stone–von Neumann theorem says that certain algebras that show up a lot in QM have "unique" irreps. How unique? Up to a scalar, of course, and this exactly accommodates the $\mathbb{P}$rojectivization. (Compare with the Brauer group from algebraic geometry: central simple algebras = projective bundles.)

Actually, there's a bit of a problem with 2. The functorial picture makes clear that there's no reason for U to be invertible. So what to do with it? Indeed, consider 2-d qft, and the pair of pants. In the Atiyah–Segal picture, the pair of pants gives a "multiplication" on $\mathcal{H}$. This certainly will not translate into a homomorphism of algebras. Or how about classical field theory (= PDEs)? Then the pair of pants gives: a span of spaces (namely, space of classically-allowed fields), which is a cospan of algebras of functions.

So how can we understand these examples in terms of algebras of functions? I hope that cospans already look to you like bimodules. Also, suppose we have a homomorphism of algebras $F : \mathcal{A} \to \mathcal{B}$. Then there is a *modulation* of F, which is the vector space $\mathcal{B}$ as an $\mathcal{B}$–$\mathcal{A}$ bimodule, where the right action is the canonical one and the left action is via F. So let's try to modify:

2'. In a Heisenberg QM system, to an interval $[t_1, t_2]$ assign an $\mathcal{A}$–$\mathcal{A}$ bimodule M_{t_2,t_1}. When turning a Schrödinger-picture QM system into a Heisenberg one, send $U \mapsto$ modulate(conjugate by U).

The second part's still not quite right, since it still needs "conjugate by U" to be a well-defined homomorphism. Also, you should be suspicious of the modulation. Indeed, lots of different homomorphisms can have the same (i.e. isomorphic) modulation. The statement is: a $\mathcal{B}$–$\mathcal{A}$ bimodule M is a modulation if after forgetting the $\mathcal{A}$ action the resulting $\mathcal{B}$-module is isomorphic to the rank-one free $\mathcal{B}$-module ${}_{\mathcal{B}}\mathcal{B}$. Two homorphisms have isomorphic modulations iff they differ by an inner automorphism of $\mathcal{B}$, i.e. conjugation by $b \in \mathcal{B}$. In this case the isomoprhism is multiplication by b.

In particular, the second sentence of 2' loses *all* data about U. This is a disaster.

How can we keep U in the game? At minimum, let's try to kill the "multiplication by b" isomorphism between different modulations. To do this, track the behavior of $1 \in \mathcal{B}$. If we demand that isomorphisms fix 1, then they can't multiply by b. And sure enough, the pointed module $(\mathcal{B}_F, 1 \in \mathcal{B})$ recovers F, by $F(a) = 1 \triangleleft a$. Thus:

2". In a Heisenberg QM system, to an interval $[t_1, t_2]$ assign a *pointed* $\mathcal{A}$–$\mathcal{A}$ bimodule $(M, m \in M)$.

Ok, so what about when U isn't invertible? Well, when U is invertible, then you can check that the modulation of "conjugate by U" is isomorphic to: $\mathcal{A} = \mathrm{End}(\mathcal{H})$ as an $\mathcal{A}$–$\mathcal{A}$ bimodule in the usual way, but pointed by U rather than by 1. This tells you what to do for general linear maps.

Finally, we have:

3”. In place of the labeled manipulations $\mathcal{H} \to \mathcal{H}$, use the corresponding pointed bimodules. In place of annihilation ideals I, use the cyclic module $\mathcal{A}/I$, pointed by $1 \in \mathcal{A}$.

2 Heisenberg picture QFT

To summarize:

Definition redux: A *Schrödinger QM system* is a functor $\text{Spacetime} \to \text{Vect}$, where the domain is the category of intervals in $\mathbb{R}$ (and perhaps some extra stuff like $\emptyset \to *$) and the target is the category of “Hilbert spaces.” A *Heisenberg QM system* is a functor $\text{Spacetime} \to \text{PointedBimod}$, where the target is the bicategory with:

Objects Associative algebras.

1-morphisms Pointed bimodules.

2-morphisms Homomorphisms of pointed bimodules. I.e. a 2-morphism $(M, m \in M) \to (N, n \in N)$ is a bimodule homomorphism $f : M \to N$ such that $f(m) = n$.

It’s worth emphasizing that PointedBimod *is not* the Morita bicategory you’re most familiar with. The latter allows arbitrary bimodule homomorphisms, which may ignore the pointing. There is a forgetful functor $\text{PointedBimod} \to \text{Morita}$, but it is not an equivalence.

There’s one final generalization that I want to mention. It’s sometimes very important to embed the Morita category into the bicategory of categories. More specifically, consider the bicategory with:

Objects Linear categories (perhaps plus a “locally presentable” condition).

1-morphisms Linear cocontinuous functors (or just the left adjoints; these are equivalent under the “locally presentable” condition).

2-morphisms Linear natural transformations.

I will call (versions of) this category 2Vect, and set $1\text{Vect} = \text{Vect}$.

Eilenberg–Watts Theorem: There is a full and faithful embedding $\text{Morita} \hookrightarrow 2\text{Vect}$ sending an algebra $\mathcal{A}$ to its category of left modules and a bimodule ${}_{\mathcal{B}}M_{\mathcal{A}}$ to the functor $M \otimes_{\mathcal{A}}$.

There is a corresponding version for pointed bimodules:

Corollary: There is a full and faithful embedding $\text{PointedBimod} \hookrightarrow \text{E}_0(2\text{Vect})$.

I should explain the target of this. In general, an E_0-*algebra* in a symmetric monoidal category $\mathcal{C}$ is a “pointed object” i.e. an object $X \in \mathcal{C}$ along with a morphism $x : \mathbf{1} \to X$, where $\mathbf{1}$ is the unit object. If you unpack this, you find that an E_0 object in 2Vect is a linear (locally presentable) category with a distinguished object:

Objects Category X with object $x \in X$.

But you have to be careful with the 1-morphisms. The correct answer is:

1-morphisms A morphism $(X, x \in X) \to (Y, y \in Y)$ is a pair consisting of a (linear, cocontinuous) functor $F : X \to Y$ together with a moprhism $f : y \to F(x)$.

Depending on your taste, you might call this either a "lax" or "oplax" functor of E_0 categories. Finally:

2-morphisms A morphism $(F, f) \to (G, g)$ is a pair consisting of a (linear) natural transformation $\eta : F \to G$ along with an equality $\eta(x) \circ f = g : y \to G(x)$.

If 2VECT weren't just a bicategory, but instead some higher n-category, then you wouldn't have an equality, but instead a 3-morphism $g \Rightarrow \eta(x) \circ f$.

At this point, I should mention that getting all these (op)lax things right is subtle, both because you have to find the right directions for everything, but also because:

Exercise: There does not exist a tricategory whose objects are bicategories, 1-morphisms are strong functors, 2-morphisms are *lax* natural transformations, and 3-morphisms are modifications.

So this is a warning that you can't just wave your hands and shout "abstract nonsense!" — you have to actually work. The details follow from joint work in progress with Claudia Scheimbauer.

Anyway, the embedding in the corollary is $\mathcal{A} \mapsto \mathcal{A}$-Mod, pointed by ${}_{\mathcal{A}}\mathcal{A}$.

Why did I bring this up? Because there might be "non-affine quantum spaces". Recall that a scheme is fully determined by its category of quasicoherent sheaves; this category in turn is pointed by $\mathcal{O}$. Similarly, a "non-affine quantum space" is an E_0 category.

To define n-dimensional fully-extended TQFT, we need to work with higher categories. Let $\mathcal{C}$ be a symmetric monoidal $(\infty, 1)$-category. The following is from Claudia's thesis:

Theorem (Scheimbauer): For each n, there is an (∞, n)-category $\mathrm{ALG}_n(\mathcal{C})$, with:

Objects E_n algebras in $\mathcal{C}$. An E_n algebra is to $\mathbb{R}^n$ as an associative algebra is to $\mathbb{R}^1$. In particular, E_2 categories are precisely braided monoidal categories.

1-morphisms E_{n-1} bimodules between E_n algebras.

...

$(n-1)$**-morphisms** (Some type of) associative algebras.

n**-morphisms** Pointed bimodules.

This n-category was predicted in many papers, but Claudia's is the first rigorous construction.

It follows from our joint work that:

Theorem (JF–Scheimbauer): Suppose that $\mathcal{C}$ is an (∞, k)-category. Then $\mathrm{ALG}_n(\mathcal{C})$ is the bottom n-dimensions of an $(\infty, n+k)$-category.

E.g. when $\mathcal{C} = 1\text{VECT}$ and $n = 1$, the 2-category $\text{ALG}_1(1\text{VECT})$ is precisely POINTEDBIMOD. When $\mathcal{C} = 2\text{VECT}$ and $n = 0$, you get $\text{E}_0(2\text{VECT})$.

So, with this, I can define Heisenberg-picture QFT. First I'll recall the Schrödinger picture.

Definition: Fix a symmetric monoidal (∞, n)-category SPACETIME of n-dimensional cobordisms (possibly with geometric structure, depending on the type of QFT). A *n-dimensional Schrödinger-picture QFT* is a symmetric monoidal functor SPACETIME $\to n$VECT, where nVECT is an appropriate (∞, n)-category that in the top spots looks like VECT. A *n-dimensional Heisenberg-picture QFT* is a symmetric monoidal functor

$$\text{SPACETIME} \to \text{ALG}_m(k\text{VECT})$$

where $m + k = n + 1$. (Raising k and lowering m allows more and more "non-affine-ness". Note that the RHS is an $(\infty, n + 1)$-category.)

In both cases, a *TQFT* is when SPACETIME $= \text{BORD}_n$ is the fully-extended framed Bordism category.

There is a famous result of Jacob Lurie's, which is certainly true, but not yet completely rigorously proven:

Theorem (Lurie): A symmetric monoidal functor $\text{BORD}_n \to \mathcal{C}$ is determined by its value on a point. Said value must by *n-dualizable* (and may be any n-dualizable object in $\mathcal{C}$). Dualizability is an inductive algebraic condition, with finitely much to check; it can be described entirely in terms of bicategories, and depends only on the first n dimensions of $\mathcal{C}$, even if $\mathcal{C}$ is a $(> n)$-category.

By directly (and rigorously!) constructing the TQFT, Claudia has proved:

Theorem (Scheimbauer): Let $\mathcal{C}$ be an (∞, k)-category for $k \geq 1$. Every object of $\text{ALG}_n(\mathcal{C})$ is n-dualizable.

On the other hand,

Theorem: Let $\mathcal{C}$ be an (∞, k)-category. The only $(n + k)$-dualizable object in the $(\infty, n + k)$-category $\text{ALG}_n(\mathcal{C})$ is the unit object.

This is why I have $m + k = n + 1$, and not n, above.

3 Examples when $n = 3$

Claudia's result implies that you get 3-dimensional Heisenberg-picture TQFTs any time you have an "E_3-algebra object" in some setting. But I don't know any examples. For comparison, a E_2-algebra in CAT (or 2VECT) is a braided monoidal category, whereas E_k categories for all $k \geq 3$ are symmetric monoidal (and hence not truly quantum). There are probably E_3 *bicategories*, but I've never met any. Actually, this is something someone should be on the lookout for.

On the other hand, 2VECT is a 2-category, so from what I've said there might be some 3-dualizable objects in $\text{ALG}_2(\text{2VECT})$, but probably not all are. I have proven a sufficiency condition that I'll describe.

First, I should give the precise set-up, in case you care.

Let $\mathcal{S}$ be a symmetric monoidal locally presentable category. Local presentability is a technical condition that you can ignore — most categories you've met satisfy it. It puts you in a context where monoidal structures are by default closed ($\otimes x$ has an adjoint). A perfect example is $\mathcal{S} = \text{ABGP}$ or MOD_R for any commutative ring R (or quasicoherent sheaves, or ...).

Definition: The symmetric monoidal bicategory $\text{2MOD}_{\mathcal{S}}$ has

Objects Locally presentable categories with $\mathcal{S}$ action (which is "closed" in both ways it can be).

1-morphisms $\mathcal{S}$-linear cocontinuous functors (automatically left adjoints).

2-morphisms $\mathcal{S}$-linear natural transformations.

Tensor structure is given by the usual universal property, if you do it correctly.

Then the objects of $\text{ALG}_2(\text{2MOD}_{\mathcal{S}})$ are $\mathcal{S}$-linear braided (closed) monoidal categories. Just to emphasize, this is a 4-category. In dimensions 0,1,2, it looks like a categorified version POINTEDBIMOD. In dimensions 3 and 4 it is a version of "pointed functor" and "pointed natural transformation."

My theorem is:

Theorem: Suppose that $\mathcal{C} \in \text{ALG}_2(\text{2MOD}_{\mathcal{S}})$ satisfies:

(1) Forgetting the braided monoidal structure, $\mathcal{C}$ is 1-dualizable as an object of $\text{2MOD}_{\mathcal{S}}$.

(2) $\mathcal{C}$ has a generating set consisting entirely of 1-dualizable objects — dualizable "inside $\mathcal{C}$" with respect to $\mathcal{C}$'s monoidal structure. (In all the examples I know, there are not more than a set worth of dualizable objects total, so this condition is equivalent to saying that every object of $\mathcal{C}$ is a colimit of dualizable objects.)

Then $\mathcal{C}$ is 3-dualizable in $\text{ALG}_2(\text{2MOD}_{\mathcal{S}})$.

For example, the category of comodules for any cosemisimple Hopf algebra works. "Comodules" guarantees condition (2), and "cosemisimple" gives condition (1). In particular, the (co)module theories for "quantized function algebras," which are a version of quantum groups, works.

If you will grant Lurie's cobordism theorem as true, then it follows that any such $\mathcal{C}$ defines a framed 3-dimensional TQFT valued in $\text{ALG}_2(\text{2MOD}_{\mathcal{S}})$. Just to emphasize the structure of $\text{ALG}_2(\text{2MOD}_{\mathcal{S}})$, such a TQFT necessarily assigns:

given as input a ...	the TQFT assigns ...
point	$\mathcal{C}$ as a braided monoidal category, up to equivalence of braided monoidal categories
closed curve	a (locally presentable, $\mathcal{S}$-enriched) monoidal category
closed surface	a (locally presentable, $\mathcal{S}$-enriched) category with a distinguished object
closed 3-manifold	an object in $\mathcal{S}$ along with a distinguished "element," i.e. a map to it from the unit object

The answers in dimensions 1 and 2 can be computed immediately from Scheimbauer's theorem (which moreover doesn't need Lurie's theorem): they are given by what's called "topological chiral homology with coefficients in $\mathcal{C}$," and there are straightforward explicit presentations of the results. Dimension 3 is a bit more complicated, but not too much.

The most important example of $\mathcal{C}$ satisfying the conditions of my theorem is the following version of Temperley–Lieb.

Definition: Let $\mathcal{S} = \mathrm{MOD}_{\mathbb{Z}[q,q^{-1}]}$ denote the symmetric monoidal category of $\mathbb{Z}[q,q^{-1}]$-modules. The category TL is the (locally presentable) $\mathcal{S}$-enriched monoidal category freely generated by a self-dual object X of dimension $-q^2-q^{-2}$. Here's how to build TL in detail. First, you take the free monoidal category generated by one object X, and adjoin morphisms $\mathrm{ev} : X\otimes X \to \mathbf{1}$ and $\mathrm{coev} : \mathbf{1} \to X \otimes X$ subject to the zig-zag relations $(\mathrm{id}_X \otimes \mathrm{ev}) \circ (\mathrm{coev} \otimes \mathrm{id}_X) = \mathrm{id}_X = (\mathrm{ev} \otimes \mathrm{id}_X) \circ (\mathrm{id}_X \otimes \mathrm{coev})$. Then you formally "linearize," replacing every hom-set by the free $\mathbb{Z}[q,q^{-1}]$-module with that hom-set as basis. Then you quotient by the relation $\mathrm{ev} \circ \mathrm{coev} = -q^2 - q^{-2}$. If you stop at this point, you get the usual category called "Temperley-Lieb," but I don't want to stop here. Instead, I want you to construct the free $\mathbb{Z}[q,q^{-1}]$-linear cocompletion of this category: include all colimits (coproducts and cokernels).

Remark: The Yoneda theorem says that my version of Temperley–Lieb is the category of "modules" of the usual Temperley–Lieb, just like MOD_R for any ring R is the free $\mathbb{Z}$-linear cocompletion of the category with one object $*$ and $\hom(*,*) = R$.

Anyway, TL is braided by declaring that the braiding of X past itself is $q\,\mathrm{id}_{X\otimes X} + q^{-1}\,\mathrm{coev} \circ \mathrm{ev} : X \otimes X \to X \otimes X$.

Being a free ($\mathcal{S}$-linear) cocompletion of a ($\mathcal{S}$-linear) category, this TL is 1-dualizable over $\mathcal{S}$, with dual the free ($\mathcal{S}$-linear) cocompletion of the opposite category. Free cocompletions are always generated by the original category, which consists entirely of dualizable objects, namely powers of the generator X. Thus TL satisfies the conditions of my theorem.

What assignments does it make? I said already that in dimensions ≤ 2, the assignments can be computed via factorization homology. For a more explicit description, let's look first at dimension 2, where the TQFT assigns to a surface Σ a pointed category $\int_\Sigma \mathrm{TL}$. It is the appropriate cocompletion of the category whose objects are marked points in Σ (labeled by objects in TL, but that's just one object up to multiplicities) and morphisms are graphs embedded in $\Sigma \times [0,1]$ with vertices labeled by morphisms in TL, modulo isotopy. (There's also some framing data to keep track of.) You can project such a graph to a "knotted graph diagram" on Σ, and then use the braiding rule to resolve all crossings.

By the way, you can say very explicitly what "appropriate" means: it means "the $\mathbb{Z}[q,q^{-1}]$-linear presheaves on the non-cocompleted category."

The category is pointed by the "empty" object, which I will call $|\text{vac}\rangle_\Sigma \in \int_\Sigma \text{TL}$. You might have heard of the *skein algebra* for the surface Σ — it is spanned by crossingless diagrams on Σ, with multiplication given by stacking diagrams and resolving crossings. But it is also precisely the ring $R_\Sigma = \text{End}(|\text{vac}\rangle_\Sigma)$. It is an important piece of quantum topology. Abstract category theory says that the part of $\int_\Sigma \text{TL}$ generated by $|\text{vac}\rangle_\Sigma$ is on the nose a copy of R_Σ-MOD. Indeed, R_Σ is the ring whose module theory "best approximates" the pointed category $\int_\Sigma \text{TL}$. So the TQFT includes at least the data of the skein algebra R_Σ, but actually knows a bit more, since $\int_\Sigma \text{TL}$ is not generated by the object $|\text{vac}\rangle_\Sigma$.

What about a 3-manifold M with boundary ∂M? The TQFT will assign to it a pointed object $\int_M \text{TL} \in \int_{\partial M} \text{TL}$. By "pointed object" I mean that it comes equipped with a morphism $|\text{vac}\rangle_M : |\text{vac}\rangle_{\partial M} \to \int_M \text{TL}$. You can say very explicitly what this object is as a presheaf on the non-cocompleted category above: it assigns to any configuration of points in ∂M the $\mathbb{Z}[q, q^{-1}]$-module of tangles in M ending at those points, modulo the skein relations.

Look now at $\hom(|\text{vac}\rangle_{\partial M}, \int_M \text{TL})$. It automatically picks up an action by $R_{\partial M} = \text{End}(|\text{vac}\rangle_{\partial M})$, and so it is an object of $R_{\partial M}$-MOD. And it's automatically pointed by $|\text{vac}\rangle_M$. What $R_{\partial M}$-module is it? It's precisely the *skein module* for M, generated by links in M modulo the skein relation. So this TQFT knows about skein modules, but actually knows more: for most Ms, the object $\int_M \text{TL}$ is not in the part of $\int_{\partial M} \text{TL}$ generated by $|\text{vac}\rangle_{\partial M}$.

Finally, you can see directly what happens when you specialize q to values. All of the constructions are very well-behaved under specialization, because that's a form of "colimit." So you can, for example, specialize $q \mapsto 1$. We can also base-extend, if we want — everything is well-behaved under switching from $\mathbb{Z}$-coefficients to, say, $\mathbb{C}$-coefficients. Then $\mathcal{S} \mapsto \text{VECT}$ and $\text{TL} \mapsto \text{MOD}_{\text{SL}(2,\mathbb{C})}$. The category $\int_\Sigma \text{TL}$ specializes to the category of quasicoherent sheaves on the *stack* of $\text{SL}(2, \mathbb{C})$-local systems on Σ. Note that this stack has affinization the usual SL(2)-character variety, which is the spectrum of the specialization of R_Σ. The point is that this stack isn't affine, and that's why the TQFT is more data than just the skein algebras/modules.

Let me end on some optimism about future applications. This TQFT is powerful enough to reproduce a bunch of known results about colored Jones polynomials, quantum A-polynomial, and so on. I haven't been able to pull new results out of it, but it's possible that they're there, using the robustness of my construction under specializations and base changes. The missing piece right now is a really solid description of $\int_{T^2} \text{TL}$ — I can give a "generators and relations" description of the category, but what I'd want is, say, a list of all simple objects, and understanding extensions, and that type of thing.