Introduction

The motivating problem for these lectures comes from condensed matter physics:

($\star$) Classify invertible gapped phases of matter

Phases of matter are familiar in everyday life. Ice, water, and vapor are different forms of H_2O, more different from each other than, say, water at 1°C and water at 99°C. If we consider "forms of H_2O" as a function of temperature, then there are two special temperatures—0°C and 100°C—at which there is a transition, from solid to liquid and from liquid to vapor. The real line of temperatures minus those two transition points has three path components: the three phases of solid, liquid, and vapor. In other words, forms of H_2O connected continuously by a path are considered to be in the same phase. So far we have only introduced one parameter—temperature—with pressure fixed to be the one you are experiencing right now. If we also allow pressure to vary, then there are only two phases of H_2O.

Our task is to (1) build a mathematical model of the physics problem, (2) solve the mathematics questions which arise in the model, and (3) apply the solution back to physics. This is the classic Three-Step Procedure of external applications of mathematics. From the example of H_2O we see the outlines of a mathematical framework: there is a space $\mathcal{M}$ of "systems" with a "singular" locus $\Delta \subset \mathcal{M}$, and we are interested in $\pi_0(\mathcal{M}\backslash\Delta)$. But what are the "systems"? How do we construct the space $\mathcal{M}$ which parametrizes them? And which systems are "singular"? What is clear is that problem ($\star$) is topological in nature. For example, the answer only depends on the homotopy type of $\mathcal{M}\backslash\Delta$, so our solution need only construct a homotopy type, not a more precise algebraic or smooth space.

There are many physical models of a given "system" in nature. They roughly fall into two boxes: discrete and continuous. Discrete models include various types of lattice models (discrete space, continuous time) and statistical mechanics models (discrete space, discrete time); they are prevalent in condensed matter physics. There is an extensive mathematical literature on these systems, but not as far as we know general definitions which apply directly to the problem at hand. (Nonetheless, Alexei Kitaev has made great strides in this direction.) Field theories, such as those of Maxwell and Einstein, are continuous models of nature. There are general physical principles which guide passage between discrete and continuous models. Our mathematical model of 'phases of matter' is grounded in field theory, and more specifically in an Axiom System introduced by Graeme Segal. This structural vision of field theory is the starting point for constructing a space $\mathcal{M}$ which parametrizes mathematical objects that in turn model some physical reality. The model retains minimal information—the long range physics—which on the one hand is robust enough to determine the phase and on the other hand is flabby enough to be amenable to topological techniques.

These lectures are based on a joint paper [**FH1**] with Mike Hopkins. In these lectures we offer complementary and supplementary background, motivation, and results; we leave out several detailed proofs which are in [**FH1**]. In this extended introduction we outline the lectures which make up this volume. The reader may wish to flip back and forth between this broad idea sketch and the more detailed lectures which follow.

Moduli spaces and deformation classes

Problems analogous to ($\star$) are familiar in mathematics. As a simple example, fix a positive integer n and let $\mathcal{M}_n$ be the space of configurations of n points on the real line $\mathbb{R}$. The position of the i^{th} point is a function $x^i\colon \mathcal{M}_n \to \mathbb{R}$. Together the positions define an isomorphism $(x^1,\dots,x^n)\colon \mathcal{M}_n \xrightarrow{\cong} \mathbb{R}^n$. We say $\mathcal{M}_n$ is the *moduli space* for this problem. There are natural "questions" which take the form of functions on $\mathcal{M}_n$. 'What is the distance between the 1^{st} and 3^{rd} points?' is the function $|x^1 - x^3|$. So far there is no interesting topology: $\mathcal{M}_n$ is contractible. Let $\Delta \subset \mathcal{M}_n$ be the locus of n-tuples $x = (x^1,\dots,x^n)$ in which not all x^i are distinct—the union of all diagonals. Configurations in $\mathcal{M}_n\backslash\Delta$ satisfy a "gap condition", and now there is nontrivial topology: $\mathcal{M}_n\backslash\Delta$ has $n!$ contractible components. A gapped configuration $x \in \mathcal{M}_n\backslash\Delta$ determines a permutation $\sigma(x) \in \mathrm{Sym}_n$, and the permutation is a complete invariant of the path component, or *deformation class*. In other words, $\sigma\colon \pi_0(\mathcal{M}_n\backslash\Delta) \to \mathrm{Sym}_n$ is an isomorphism.

It is worth pausing to contemplate the sophisticated mathematical theory underlying this example. Nowadays we confidently write '$\mathbb{R}$' because we have in hand a rich theory of real numbers. Historically it was not always so. Only after many hard-fought struggles and contradictions could victory be declared in the form of the three-word characterization of the real numbers: *complete ordered field*. That profound hard-won phrase is the starting point of every undergraduate real analysis class. The order on the real numbers underlies the isomorphism σ. But do you see any infinite decimals, Cauchy sequences, or Dedekind cuts? The topological problem 'Compute $\pi_0(\mathcal{M}_n\backslash\Delta)$' does not require such precision, whereas geometric questions about distance between points do. The characterization of the nature of real numbers is more useful for qualitative questions than any detailed construction would be.

This toy problem exhibits several features common to moduli problems. First, there are discrete parameters, here the positive integer n. Second, there is a singular locus Δ; off of Δ the parametrized objects satisfy a nonsingularity condition. Third, there are interesting functions on the moduli space which encode geometric information about the parametrized objects. Finally, there is a complete invariant of the deformation class, which is an isomorphism to a known or computable set. In this toy problem, as well as in problem ($\star$), there is a natural group structure: π_0 is a group, not merely a set. (It should be said the group structure is not obvious if we view points of $\mathcal{M}_n\backslash\Delta$ as configurations.) A known complete invariant of π_0 is not present in all situations.

As another example we might ask to parametrize "1-dimensional metric shapes". This is ill-defined as stated. Better said, it is *not* defined as stated: to make a mathematical theory we must provide a definition. Whereas the real numbers are characterized uniquely (complete ordered field), here there is no uniqueness. Still, a definition should capture general features, even though it does not characterize. We might decide that, intuitively, we do not want to allow the 1-dimensional shape to cross itself. That is a "gap condition". Once more there is a hard-won mathematical definition at hand: 1-dimensional smooth Riemannian manifold. Furthermore, we know what a family of such objects is: a smooth fiber bundle $\pi\colon X \to S$ with a Riemannian structure. (We leave the reader to ponder what the Riemannian structure is; the definition should lead to a unique Levi-Civita connection on the relative tangent bundle.) But now there is a second kind of "singularity" which is still in the game. Namely, there exists a smooth fiber bundle over $S = \mathbb{R}$ such that the fiber at $s < 0$ is diffeomorphic to a circle and the fiber at $s > 0$ is diffeomorphic to the disjoint union of two circles. The "singularity" at $s = 0$ is a noncompact fiber: two lines. To rule out the transition from one circle to two circles, restrict to proper fiber bundles. Demanding compactness is again a kind of gap condition—think spectrum of the Laplace operator. The moduli space $\mathcal{M}$ of closed Riemannian 1-manifolds has all the features enumerated in the previous paragraph. The discrete parameter is the dimension, the singular locus was already eliminated, functions such as total length are interesting geometric invariants, and the number of components is a complete invariant. There is a natural commutative monoid structure on $\pi_0\mathcal{M}$ given by disjoint union, but there are no inverses and so $\pi_0\mathcal{M}$ is not a group.

Problem ($\star$) is about families of quantum mechanical systems. There is a basic dichotomy determined by the energy spectrum: a system is *gapped* if there is a gap in the spectrum of the Hamiltonian above the minimal energy and otherwise is *gapless*. This problem only considers gapped systems. In this context the singular locus parametrizes *phase transitions*, which are bifurcated into first-order and higher-order. A phase transition occurs along a path of gapped quantum systems when the energy gap is closed—the energy spectrum comes down to the minimum. If discrete spectrum goes down the transition is first-order; if continuous spectrum, then it is higher-order. In any case we throw them all out and define two systems to be in the same phase if they can be joined by a continuous path of gapped systems with no phase transition. Problem ($\star$) includes another adjective—invertible—which we discuss below; the classification question makes sense in the absence of invertibility. Of course, there are interesting analytic questions which are expressed in terms of correlation functions, which are functions on $\mathcal{M}$, but problem ($\star$) is topological so we do not need anything so precise. The broad rough outline, then, is: construct a moduli space $\mathcal{M}$ for invertible gapped systems; throw out a locus Δ of phase transitions; compute $\pi_0(\mathcal{M}\backslash\Delta)$. In fact, our transformation to a problem in field theory obviates the need to consider Δ.

Axiom System for field theory

As stated above, we attack ($\star$) by shifting to a problem in field theory. Previously we used 'characterization' and 'definition' as monikers for the mathematical starting point, but for field theory we use 'axiom system'. Why? Certainly there is not a unique field theory, even with discrete parameters fixed, so 'characterization' is inappropriate. We shy away from 'definition', which to our ears suggests *stare decisis*—settled law—and the situation for quantum field theory is hardly that! Also, the Axiom System for field theory is perhaps more analogous to the Eilenberg-Steenrod axioms for (generalized) homology theories than, say, to the definition of a smooth Riemannian manifold. The Eilenberg-Steenrod axioms, which *do* play the formal role of a 'definition', tell what a homology theory is without constructing one. Their power lies in their simplicity. One can check whether a construction satisfies the axioms, and in this way know that highly disparate constructions yield the same mathematical object, something that was not at all apparent before the Eilenberg-Steenrod 1945 paper.

The analogy between pre-1945 algebraic topology and present-day quantum field theory is not perfect, but consider the strong commonalities: multiple starting points, constructions, and approaches. The Axiom System for field theory, introduced by Segal in the 1980s for conformal theories in two spacetime dimensions and later adapted by Atiyah for topological theories in all dimensions, is flexible. It applies to both classical and quantum theories in all dimensions. In the Axiom System a field theory is a map, and as such has a domain and codomain; part of the flexibility is the freedom to vary them. The domain and codomain are each a symmetric monoidal category, and the map is a symmetric monoidal functor. The domain is a bordism category of smooth manifolds equipped with fields. 'Field' has a precise definition[1] and includes traditional scalar fields of physics as well as topological structures (orientation, spin structure) and more exotic possibilities. The codomain for a physically relevant field theory is an appropriate category of complex topological vector spaces. This choice goes back to the early days of quantum mechanics: linearity encodes superposition and complex numbers encode interference. Ergo the Axiom System in a nutshell: *a field theory is a linear representation of a geometric bordism category.* Other choices for the codomain, such as the category of abelian groups, have proved useful in mathematical contexts. Indeed, the viewpoint of the Axiom System has proved its value many times over in *mathematics*: in low dimensional topology, symplectic geometry, geometric representation theory, category theory, etc. The story of these lectures is one small application to *physics*, and there are many more indications there of its pertinence and utility. However, it has not unified the disparate points of view on quantum field theory, and there are few rigorous examples; in non-topological contexts the Axiom System is not as established as, say, the Eilenberg-Steenrod axioms.

We explain two routes to the Axiom System in these lectures. The first route, explained in Lecture 1, is through classical bordism theory: a topological field theory is a categorification of a classical bordism invariant. The second is through quantum field theory. An essential point is that what is being axiomatized is *Wick-rotated* quantum theory—physics with purely imaginary time. In Lecture 2 we introduce the basic characters in quantum theory in the context of quantum mechanics: states, observables, and correlation functions. In fact, we give a unified picture of mechanics—classical and quantum—which goes back to early mathematical work on quantum theory. Wick rotation is straightforward in quantum mechanics, and one can already see the Axiom System emerging in this 1-dimensional case of field theory. Lecture 3 takes up relativistic quantum field theory, but only from a very structural perspective in order to explain where the Axiom Systems sits. Our starting point is definitely non-topological, so it is not surprising that the Axiom System applies to non-topological field theories. Although we invoke non-topological field

[1] which we do not include in these lectures; see [**FT1**, Appendix].

theories at many points in the subsequent lectures, the technical work is for topological field theory. The reader might mistakenly infer from our words that the Wick-rotated non-topological theories we refer to, such as Yang-Mills + Chern-Simons in three dimensions, are completely well-defined objects which have been mathematically worked out. Not so! Even the topological parts of field theory are under rapid development.

The Axiom System is, in a sense, a Wick-rotated version of the Schrödinger approach to quantum mechanics and to quantum field theory (Wightman *et al.*). That is, it emphasizes states and time-evolution of states, albeit imaginary time-evolution. By contrast, the Heisenberg approach to quantum mechanics emphasizes algebras of observables, as does that approach to quantum field theory (Haag *et al.*). There are also modern mathematical axiom systems based on the Heisenberg approach, most prominently in work of Costello and collaborators. It should also be said that the Axiom System does not distinguish classical and quantum; a classical field theory fits the formal axioms as a very special case—classical field theories are invertible. Many quantum theories have semiclassical limits in which they are described via quantization, say in terms of path integrals. Detailed descriptions in terms of fluctuating fields furnish important information about a quantum field theory, but they do not enter our approach to the classification problem ($\star$). Nevertheless, one of the original goals of the Axiom System was precisely to capture the formal properties of canonical quantization and the path integral.

Symmetries

The sample moduli problems introduced above all have discrete parameters—number of points on a line, dimension of a Riemannian manifold—and a moduli space is constructed for fixed values of these parameters. The discrete parameters in the phases of matter problem ($\star$) are the dimension of *space* and the symmetry group. Analogous parameters are present in effective long-range field theories, so it is important that we understand how these parameters manifest in the Axiom System. Dimension is a fundamental parameter evident in the domain bordism category. The dimension n, which is the dimension of *spacetime*, strongly affects not only topological theories, but also analytic aspects of usual quantum field theories. Symmetry is more complicated, and we devote much effort in these lectures and in [**FH1**] to this topic.

The initial arena for a relativistic quantum field theory is Minkowski spacetime $\mathbb{M}^n$, an n-dimensional affine space equipped with a translation-invariant Lorentz metric and a time-orientation. Its automorphism group $\mathcal{I}_n$ is the subgroup of the affine group which preserves the metric and time-orientation. A quantum field theory is a structure over $\mathbb{M}^n$, so its automorphism group $\mathcal{G}_n$ comes equipped with a homomorphism $\rho\colon \mathcal{G}_n \to \mathcal{I}_n$. In other words, a symmetry of a relativistic quantum field theory induces a symmetry of the underlying spacetime. Relativistic invariance is the requirement that the image of ρ contain the identity component of the Lie group $\mathcal{I}_n$. This geometrically natural setup places $\mathcal{I}_n$ or its identity component as a *quotient* of $\mathcal{G}_n$, whereas in traditional approaches to quantum field theory one assumes that the Poincaré group is a *subgroup* of $\mathcal{G}_n$. (The Poincaré group is a double cover of the identity component of $\mathcal{I}_n$.) In Lecture 3 we track symmetry through Wick rotation to Euclidean space and then to curved Riemannian manifolds. As we know from differential geometry, at this last stage it is natural to divide by translations and use the quotient as the structure group of a smooth manifold. Doing so we obtain from $\mathcal{G}_n$ a *compact* Lie group H_n and a homomorphism $\rho\colon H_n \to \mathrm{O}_n$ whose image is either SO_n or O_n. We call the pair (H_n, ρ_n) the *symmetry type* of the theory. It is an important discrete invariant in any field theory; we advocate articulating it explicitly in every example.

For the theories in these lectures $\mathcal{G}_n$ is a Lie group, but in general $\mathcal{G}_n$ may be a *super* Lie group and may include "higher" symmetries. In the Axiom System the symmetry type is manifest in the domain bordism category, which consists of Riemannian manifolds equipped with an H_n-structure in the sense of Cartan. In physics the H_n-structure is a "background field". For example, if $H_n = \mathrm{O}_n \times K$ for a compact Lie group K, then the background fields are a Riemannian metric and a principal K-bundle with connection. More general background fields can also be considered part of the "symmetry type" of a theory. But, as already stated, in these lectures we stick to compact Lie groups H_n and their associated background gauge fields and Riemannian metrics.

The rigidity of compact Lie groups leads to general structure theorems (Proposition 3.16, Theorem 3.24) and to classification theorems (Example 3.22, Theorem 10.2). The discrete parameters of dimension and symmetry type appear in our main theorems about invertible reflection positive field theories (§8.6).

Extended Locality

In Minkowski spacetime $\mathbb{M}^n$ one expression of locality is that observables supported in spacelike separated regions commute. After Wick rotating to Euclidean space $\mathbb{E}^n$, every pair of regions is spacelike separated and so operators with disjoint supports commute. Another expression of locality in $\mathbb{M}^n$ is cluster decomposition, a factorization property of correlation functions. The Axiom System for an n-dimensional Wick-rotated theory on Riemannian manifolds captures locality in codimension one. That is, if we cut a closed n-manifold X along a codimension one separating submanifold $Y \subset X$, so that we obtain two manifolds X_1, X_2 with common boundary Y, then a correlation function on X factors as a product of a correlation function on X_1 and a correlation function on X_2, assuming the supports of all observables are disjoint from Y. However, in all but exceptionally simple cases the correlation functions on X_1, X_2 are not complex numbers but rather lie in a complex vector space, the state space obtained by quantization on Y. Composition in a bordism category encapsulates this factorization of manifolds. Since a field theory is a functor out of a bordism category, the factorization of correlation functions in codimension one follows.

The strong locality of quantum field theory is perhaps more evident in the Haag approach, in which one essentially considers the theory built up from information on arbitrary open subsets of $\mathbb{M}^n$. This suggests that in Wick-rotated field theory one should be able to reconstruct everything from invariants attached to small balls. Starting with an n-manifold X, we must make cuts in n "directions" to express X as a union of balls. An algebraic structure which encapsulates these cuts has n composition laws. In this way n-categories enter the picture: the domain of an *extended* field theory is a bordism n-category. What is not known is a natural choice of codomain n-category, so that choice is left flexible, though in physical examples it is still constrained by superposition and interference. An extended field theory is then a symmetric monoidal functor between symmetric monoidal n-categories.

Lecture 5 is an exposition of extended locality and the extended Axiom System. The ideas are most developed in topological field theory. A basic theorem, the *cobordism hypothesis*, is a precise version of the statement that a field theory can be reconstructed from its restriction to a small ball. (In a topological theory one usually shrinks the ball to a point.)

Extended locality in this form was introduced in the mathematical literature in the early 1990s in connection with 3-dimensional quantum Chern-Simons theory. It appeared earlier in the physics literature in the form of extended observables, such as line operators.

Invertibility and homotopy theory

There is a composition law for quantum systems: conjunction without interaction. The state space of the composite is the tensor product of constituent state spaces and the Hamiltonian is the sum of constituent Hamiltonians. In the (extended) Axiom System the symmetric monoidal structures on the domain and codomain categories as well as on the functor between them combine to give the Wick-rotated version of this composition law. The trivial theory is a unit for the composition law; on each closed space there is a single quantum state and zero Hamiltonian, and in terms of the Axiom System it is the theory whose values are tensor units. Thus invertibility is defined. It is immediate that a theory is invertible if and only if it factors through the maximal subgroupoid of the codomain: the state space of every closed $(n-1)$-manifold is 1-dimensional, for example. But then one can "localize" and factor through the groupoid quotient of the domain bordism category. In this way an invertible field theory becomes a symmetric monoidal functor between (higher) Picard groupoids. ('Picard' is short for 'symmetric monoidal with invertible objects'.)

Enter stable homotopy theory. One passes freely between higher groupoids and topological spaces, or rather homotopy types, via the homotopy hypothesis. A Picard structure on a higher groupoid goes over to an infinite loop structure on a topological space. After that transmogrification, an invertible field theory is an infinite loop map of infinite loop spaces. This is a far cry from our starting point in physics(!), yet is the result of a step-by-step progression. So as not to muddy the mathematical waters, in Lecture 6 we take this homotopical incarnation of an invertible field theory as an *ansatz*. But in reality it is a theorem derived from the Axiom System, at least for topological field theories; we give appropriate references in that lecture.

We begin in Lecture 6 with non-extended topological field theories. In Theorem 6.27 we use an elementary Morse theory argument to prove that the partition function of an invertible field theory is a bordism invariant. However, it is not in general a *Thom* bordism invariant, but rather a *Reinhart* bordism invariant. That same conclusion follows from a deeper theorem which identifies the result of inverting all morphisms in a bordism n-category. It is the infinite loop space associated to a *Madsen-Tillmann spectrum*, which is then the domain spectrum of a field theory (Ansatz 6.89). Since we do not have in hand a canonical codomain n-category for a field theory, we cannot take its maximal subgroupoid to determine a canonical codomain spectrum for an invertible field theory. However, there *is* a natural choice for the codomain spectrum: the *character dual* to the sphere spectrum. It is characterized in field-theoretic terms by the property that the partition function determines the theory, something not true in general but a desirable property. Magically, the boson/fermion dichotomy of states falls out automatically (Remark 6.91).

At this point we have a homotopy type for a space of invertible topological field theories, but for three reasons it is not the correct homotopy type to apply to problem ($\star$). First, the topology is wrong: π_0 is the abelian group of *isomorphism* classes of topological theories, not the group of *deformation* classes. It is as if we ask about deformation classes of nonzero complex numbers but use the discrete topology rather than the usual continuous topology. We introduce the Anderson dual to the sphere spectrum as a substitute for the continuous topology; see §6.8. This leads us to introduce *continuous* invertible topological field theories (§6.10), which we argue capture the deformation class of an invertible theory. The second consideration which tells we have the wrong homotopy type is that the low-energy description of an invertible gapped quantum system is not necessarily a discrete *topological* theory, as we discuss below. And, importantly, the Axiom System does not incorporate *unitarity*, an important property of quantum systems to which we now turn.

Extended unitarity

The two pillars of quantum field theory are locality and unitarity. We explained above that a strong form of locality is implemented in Wick-rotated field theory by an extended Axiom System. Our mission now is to implement unitarity in the Axiom System, both in non-extended and extended forms, a topic we take up in Lecture 7 and Lecture 8, respectively. We only succeed in defining extended unitarity for *invertible topological* field theories; it is an interesting open question to define extended unitarity in general.

It is well-known that *reflection positivity* in Euclidean field theory is the Wick-rotated manifestation of unitarity in relativistic field theory. In fact, reflection and positivity are separate concepts. In the non-extended version discussed in Lecture 7, reflection is a structure and positivity is a condition. Reflection is implemented in the Axiom System as an involution on both the domain bordism category and the codomain category of complex vector spaces. If the domain consists of oriented manifolds, then the reflection involution is orientation-reversal. We define an analog for any symmetry type (H_n, ρ_n) in terms of a co-extension of the Lie group H_n; see Theorem 7.13. On the codomain the reflection involution is complex conjugation. A reflection structure is equivariance data for the functor which defines the field theory. This realizes the slogan "orientation-reversal maps to complex conjugation". A *reflection structure* induces a nondegenerate hermitian inner product on the state space attached to every closed $(n-1)$-manifold. In a non-extended theory positivity is the requirement that all of these hermitian inner products be positive definite. The reflection structure/positivity condition in the Axiom System on curved manifolds is a direct generalization of standard reflection positivity in Euclidean field theory.

In Lecture 8 we narrow the focus to invertible topological field theories. Recall that we model an extended invertible theory in stable homotopy theory as a map between appropriate spectra. An extended reflection structure in this invertible case is a lift of this map to an equivariant map between spectra with involutions. We give arguments to justify specific involutions in the domain and codomain, as appear in (8.45), (8.46). Not surprisingly, extended positivity is no longer just a condition in codimension one; it is also data in lower codimensions. To formulate it in our stable homotopy world, we first recast naive positivity of hermitian inner products in categorical terms (§7.1). This motivates the involution which models complex conjugation (§8.3), Definition 8.33 of a spectrum of positive definite Hermitian lines, and finally Definitions 8.53 and 8.55 of a space, or homotopy type, of reflection positive invertible topological field theories of fixed dimension and symmetry type. It turns out that for an extended invertible n-dimensional field theory with reflection structure, an extended positivity structure is a trivialization of an associated "real" $(n-1)$-dimensional field theory (Definition 8.62).

Non-topological invertible theories

Throughout the lectures we use non-topological field theories to guide our modeling, although the mathematical theorems pertain only to topological theories. Non-topological *invertible* field theories are relevant for the solution to ($\star$), as we explain in Lecture 9. A standard hypothesis is that some long range behavior of a physical system, including its phase, is captured by a scale-invariant field theory. If the physical system is gapped, then this effective field theory is almost topological, but it may be off by a non-topological invertible theory; we coin the term "topological*" for this class of theories. If the entire effective field theory is invertible, then there is no reason for it to be topological. So to solve ($\star$) we should produce a homotopy type of not-necessarily-topological invertible reflection positive theories. What we argue in Lecture 9 is that this is the homotopy type of *continuous* invertible reflection positive theories. Our proposal (Conjecture 9.34) has a more specific incarnation: invertible not-necessarily-topological field theories correspond to appropriate cocycles for *generalized differential cohomology*. The partition function of a theory is a secondary geometric invariant, the "integral" of a generalized differential cohomology class, and the theory itself provides a fully local description of the secondary invariant. A typical example is the exponentiated η-invariant of Atiyah-Patodi-Singer, which from this point of view is fully local.

Theorems

At this point we have carried out the First Step in the Three-Step Procedure of applications of mathematics: we have built a mathematical model of the physics problem ($\star$). That model consists of a well-defined homotopy type $\mathcal{M}$ of appropriate field theories. The Second Step is to prove mathematical theorems which determine $\pi_0\mathcal{M}$. In §8.6 we state these results but do not include the proofs, which may be found in [**FH1**, §8]. We not only determine π_0 but identify the entire homotopy type in familiar terms, namely as maps from a Thom spectrum to the Anderson dual to the sphere spectrum. (This is the answer for continuous theories; for discrete theories there is a slightly more complicated answer.) Therefore, the entire effect of imposing unitarity in its Wick-rotated manifestation is to replace a Madsen-Tillmann spectrum by a Thom spectrum. We remark that in our general study of symmetry we prove (Theorem 3.24) that a symmetry type (H_n, ρ_n) has a stabilization (H, ρ) as $n \to \infty$. This produces a sequence (8.7) of Madsen-Tillmann spectra which limit to a Thom spectrum, and leads to the notion of a *stable* invertible topological field theory, i.e., one which factors through the Thom spectrum. Thus a reflection positive continuous invertible theory is stable.

Part B of the Second Step in the Three-Step Procedure is to compute the abelian group $\pi_0\mathcal{M}$ of deformation classes for physically relevant values of the discrete parameters (dimension and symmetry type). Fortunately, the mathematical techniques to make the computations are already well-developed. A major tool is the Adams spectral sequence, but it is not the only one. In Lecture 10 we present some special computations; many more computations are contained in the references, which also include pedagogical introductions to the Adams spectral sequence.

While these are the mathematical theorems and techniques directly applicable to ($\star$), we also take the opportunity to present other classification theorems in Lecture 4. The main techniques there are Morse and Cerf theory, not homotopy theory. The first theorem of this type is the classification of oriented 2-dimensional topological theories in terms of Frobenius algebras. We quote without proof an analogous result in a non-topological case, but where the theory only depends on an area form, not on a full Riemannian metric. We also sketch the proof of a 1-dimensional case of the cobordism hypothesis.

Free spinor fields

We are ready for Step Three of the Three-Step Procedure: application of the mathematical theorems to the physics problem ($\star$). We test our theorems and homotopy theoretical computations against known results in the physics literature, and we also derive new results. What is known and what is new is entirely a function of the discrete parameters: dimension and symmetry type. To apply our field theoretic theorems to discrete systems in condensed matter physics, we must understand how these parameters match up in the two descriptions of the same physical system; this is one of the topics treated in §9.1. The condensed matter literature primarily deals with spatial dimensions $d = 0, 1, 2, 3$, which of course corresponds to spacetime dimensions $n = 1, 2, 3, 4$. There are many computations in the physics literature for various symmetry types in varying dimension. By now many of the corresponding homotopy theoretic computations have been made, and there is complete agreement; see the references to Lecture 10. New results have also been obtained, and essentially any case in low dimensions can be computed on demand by an appropriately skilled young homotopy theorist.

In Lecture 10 we focus on symmetry types for which there is a notion of a "free fermion" system. There is a map from free systems to interacting systems, and thus three pieces of data for each symmetry type: the map from free to interacting together with its domain and codomain. Hence for these special symmetry types the test against physics literature is richer; for other symmetry types there is only one piece of data to check. Our first task is to classify appropriate symmetry types. It turns out there are ten of them (Theorem 10.2), another instance of the famous 10-fold way, which goes back to Dyson. We remark that the two flavors of pin group occur among the ten fermionic symmetry types, and our ideas about free fermions provide insight into the theory of Dirac operators on pin manifolds. A classical free fermion field theory, massless or massive, is specified by an appropriate Clifford module, as we recount in §10.2. The crucial Lemma 10.21 tells how a nondegenerate mass is equivalent to an extra Clifford generator. To pass from the classical free fermion to its free field quantization requires a choice beyond the Clifford data, as we explain shortly in the discussion about anomalies. In the massive case, which is the effective field theory of a gapped free fermion lattice system, Conjecture 10.25 tells the deformation class of the invertible low energy approximation in terms of the Clifford data. The conjectural formula employs (1) the connection between Clifford modules and KO-theory, and (2) the map from spin bordism to KO-theory, both of which were elucidated by Atiyah-Bott-Shapiro. In §10.3 we present the results of computations for two of the ten symmetry types.

Anomalies

Bonus Lecture 11 is a piece on the general topic of anomalies in quantum field theory; it was not part of the CBMS conference. There is a rich theory of anomalies which exhibits a multitude of approaches. We focus on geometric aspects, a point of view which fits in best with these lectures. Also, we use free spinor fields as motivation for the discussion of anomalies in general. The modern view on anomalies is: to an n-dimensional field theory F is canonically attached a truncated invertible $(n+1)$-dimensional field theory α_F, its *anomaly*. (The truncation means we only evaluate on manifolds of dimension $\leqslant n$; in many cases α_F extends to a full non-truncated theory.) This is half of Thesis 11.29. The other half asserts that to produce a well-defined quantum field theory one must specify a trivialization τ of the anomaly α_F. The ratio of two trivializations is an invertible n-dimensional field theory δ. Often what appears to be natural to define is the *relative* theory F with anomaly α_F; an absolute theory $\tau \circ F$ is only determined up to tensoring by an invertible theory δ. Quantum mechanics illustrates this indeterminacy (Example 11.34); shifting by δ tensors the state space by a line and shifts the energy by a constant. This indeterminacy also helps explain why a gapped physical system may have an effective long range field theory which is topological only up to tensoring by a possibly-non-topological invertible theory.

We return to free massive spinor fields in Example 11.36. We explain how to an n-dimensional massive free spinor field theory is canonically associated an $(n-1)$-dimensional massless free spinor field theory. The n-dimensional anomaly theory of the latter is a canonical choice for the long range effective field theory of the former. It is this choice which is used in Conjecture 10.25. This maneuver implicitly defines a canonical trivialization τ of the anomaly of a massive free spinor field theory, but in families of theories—for example with variable mass—one does not necessarily use this canonical choice.

Final remarks

An appendix to the lectures summarizes some relevant facts about 1-categories. In these lectures we treat higher categories heuristically and leave the full theory, including construction of models, to the literature.

Problem ($\star$) concerns physical systems which exist on any manifold which represents space. One can also ask for the classification of phases on a particular space:

($\star\star$) Classify invertible gapped phases of matter on a space Y, possibly with group action

The theory developed in [**FH1**], recounted in these lectures, leads to a solution to problem ($\star\star$), again based on field theoretic ideas. The answer is a Borel-Moore (Borel equivariant) homology group of Y, as we explain in [**FH3**].

As should be crystal clear by now, these lectures are based on joint work with Mike Hopkins, whom I thank for many enjoyable collaborations. He and several other confederates and correspondents—including the Aspen 7 Gang—Mike Freedman, Anton Kapustin, Alexei Kitaev, Greg Moore, and Constantin Teleman, in addition to Mike Hopkins and me; David Ben-Zvi; Jonathan Campbell; Jacques Distler; Davide Gaiotto; Zheng-Cheng Gu; Matt Hastings; Jacob Lurie; Max Metlitski; Andy Neitzke; Graeme Segal; Nati Seiberg; Stephan Stolz; Peter Teichner; Ulrike Tillmann; Senthil Todadri; Kevin Walker; Xiao-Gang Wen; and Ed Witten—have influenced the ideas and work recounted in these lectures, though none of them should be held responsible for any errors of fact or judgment. The reader will join me in thanking Arun Debray for his expert proofreading, which greatly improved the manuscript. I am grateful to all.