Introduction

The goal of this book is to communicate a certain *Weltanschauung* uncovered in pieces by many different people working in bordism theory, and the goal just for this introduction is to tell a story about one theorem where it is especially apparent.

To begin, we will define a homology theory called *bordism homology*. Recall that the singular homology of a space X comes about by probing X with simplices: Beginning with the collection of continuous maps $\sigma\colon \Delta^n \to X$, we take the free $\mathbb{Z}$-module on each of these sets and construct a chain complex

$$\cdots \xrightarrow{\partial} \mathbb{Z}\{\Delta^n \to X\} \xrightarrow{\partial} \mathbb{Z}\{\Delta^{n-1} \to X\} \xrightarrow{\partial} \cdots .$$

Bordism homology is constructed analogously, but using manifolds Z as the probes instead of simplices:[1]

$$\begin{aligned} \cdots &\xrightarrow{\partial} \{Z^n \to X \mid Z^n \text{ a compact } n\text{-manifold}\} \\ &\xrightarrow{\partial} \{Z^{n-1} \to X \mid Z^{n-1} \text{ a compact } (n-1)\text{-manifold}\} \\ &\xrightarrow{\partial} \cdots . \end{aligned}$$

Lemma 1 ([Koc78, section 4]) *This forms a chain complex of monoids under disjoint union of manifolds, and its homology is written $MO_*(X)$. These are naturally abelian groups,*[2] *and moreover they satisfy the axioms of a generalized homology theory.* □

In fact, we can define a bordism theory MG for any suitable family of structure

[1] One does not need to take the free abelian group on anything, since the disjoint union of two manifolds is already a (disconnected) manifold, whereas the disjoint union of two simplices is not a simplex.

[2] For instance, the inverse map comes from the cylinder construction: For a manifold M, the two components of $\partial(I \times M)$ witness the existence of an inverse to M in the bordism groups.

groups $G(n) \to O(n)$. The coefficient ring of MG, or its value $MG_*(*)$ on a point, gives the ring of G-bordism classes, and generally $MG_*(Y)$ gives a kind of "bordism in families over the space Y." There are comparison morphisms for the most ordinary kinds of bordism, given by replacing a chain of manifolds with an equivalent simplicial chain:

$$MO \to H\mathbb{Z}/2, \qquad MSO \to H\mathbb{Z}.$$

In both cases, we can evaluate on a point to get ring maps, called *genera*:

$$MO_*(*) \to \mathbb{Z}/2, \qquad MSO_*(*) \to \mathbb{Z},$$

neither of which is very interesting, since they are both zero in positive degrees.

However, having maps of homology theories (or, really, of spectra) is considerably more data than just the genus. For instance, we can use such a map to extract a theory of integration by considering the following special case of oriented bordism, where we evaluate MSO_* on an infinite loopspace:

$$\begin{aligned} MSO_n K(\mathbb{Z}, n) &= \{\text{oriented } n\text{-manifolds mapping to } K(\mathbb{Z}, n)\} \,/\, \sim \\ &= \left\{ \begin{array}{c} \text{oriented } n\text{-manifolds } Z \\ \text{with a specified class } \omega \in H^n(Z;\mathbb{Z}) \end{array} \right\} \Big/ \sim . \end{aligned}$$

Associated to such a representative (Z, ω), the yoga of stable homotopy theory then allows us to build a composite

$$\begin{aligned} \mathbb{S} &\xrightarrow{(Z,\omega)} MSO \wedge (\mathbb{S}^{-n} \wedge \Sigma^\infty_+ K(\mathbb{Z}, n)) \\ &\xrightarrow{\text{colim}} MSO \wedge H\mathbb{Z} \\ &\xrightarrow{\varphi \wedge 1} H\mathbb{Z} \wedge H\mathbb{Z} \\ &\xrightarrow{\mu} H\mathbb{Z}, \end{aligned}$$

where φ is the orientation map. Altogether, this composite gives us an element of $\pi_0 H\mathbb{Z}$, i.e., an integer.

Lemma 2 *The integer obtained by the above process is* $\int_Z \omega$. □

Many theorems accompany this definition of $\int_Z \omega$ for free, entailed by the general machinery of stable homotopy theory. The definition is also very general: Given a ring map off of any bordism spectrum, a similar sequence of steps will furnish us with an integral tailored to that situation.

In the case of the trivial structure group $G = e$, this construction gives the bordism theory of stably framed manifolds, and the Pontryagin–Thom theorem amounts to an equivalence $\mathbb{S} \xrightarrow{\simeq} Me$. Through this observation, these techniques gain widespread application in stable homotopy theory. For a ring spectrum E, we can reconsider the unit map as a ring map

$$Me \xrightarrow{\simeq} \mathbb{S} \to E,$$

and by following the same path of ideas we learn that E is therefore equipped with a theory of integration for framed manifolds:

$$\int : \left\{ \begin{array}{c} \text{stably framed } n\text{-manifolds } Z \\ \text{with a specified class } \omega \in E^{n-m}(Z) \end{array} \right\} \Big/ \sim \;\to \pi_m E.$$

Sometimes, as in the examples above, this unit map factors:

$$\mathbb{S} \simeq Me \to MO \to H\mathbb{Z}/2.$$

This is a witness to the overdeterminacy of $H\mathbb{Z}/2$'s integral for framed bordism: If the framed manifold is pushed all the way down to an unoriented manifold, there is still enough residual data to define the integral.[3] Given any ring spectrum E, we can ask the analogous question: If we filter O by a decreasing system of structure groups, through what stage does the unit map $Me \to E$ factor? For instance, the map

$$\mathbb{S} = Me \to MSO \to H\mathbb{Z}$$

considered above does *not* factor further through MO – an orientation is *required* to define the integral of an integer-valued cohomology class. Recognizing $SO \to O$ as the zeroth Postnikov–Whitehead truncation of O, we are inspired to use the rest of the Postnikov filtration as our filtration of structure groups. Here is a diagram of this filtration and some interesting minimally factored integration theories related to it, circa 1970:

$$\begin{tikzcd}
Me \arrow[r] \arrow[drrr] \arrow[drrrr] \arrow[drrrrr] & \cdots \arrow[r] & MString \arrow[r] & MSpin \arrow[r] \arrow[d] & MSO \arrow[r] \arrow[d] & MO \arrow[d] \\
& & & kO & H\mathbb{Z} & H\mathbb{Z}/2.
\end{tikzcd}$$

This is the situation homotopy theorists found themselves in some decades ago, when Ochanine and Witten proved the following mysterious theorem using analytical and physical methods:

Theorem 3 (Ochanine [Och87b, Och91], Witten [Wit87, Wit88]) *There is a map of rings*

$$\sigma : MSpin_* \to \mathbb{C}((q)).$$

[3] It is literally more information than this – even unframeable unoriented manifolds acquire a compatible integral.

Moreover, if Z is a Spin manifold such that half[4] *its first Pontryagin class vanishes – that is, if Z lifts to a String manifold – then $\sigma(Z)$ lands in the subring $MF \subseteq \mathbb{Z}[\![q]\!]$ of q-expansions of modular forms with integral coefficients.* □

However, neither party gave indication that their result should be valid "in families," and no attendant theory of integration was formally produced. From the perspective of a homotopy theorist, it was not clear what such a claim would mean: To give a topological enrichment of these theorems would mean finding a ring spectrum E such that $E_*(*)$ had something to do with modular forms.

Around the same time, Landweber, Ravenel, and Stong began studying *elliptic cohomology* for independent reasons [LRS95]; some time much earlier, Morava had constructed an object "K^{Tate}" associated to the Tate elliptic curve [Mor89, section 5]; and a decade later Ando, Hopkins, and Strickland [AHS01] put all these together in the following theorem:

Theorem 4 ([AHS01, Theorem 2.59]) *If E is an "elliptic cohomology theory," then there is a canonical map of homotopy ring spectra $MString \to E$ called the σ-orientation (for E). Additionally, there is an elliptic spectrum K^{Tate} whose σ-orientation gives Witten's genus $MString_* \to K^{\mathrm{Tate}}_*$.* □

We now come to the motivation for this text: The homotopical σ-orientation was actually first constructed using formal geometry. The original proof of Ando, Hopkins, and Strickland begins with a reduction to maps of the form

$$MU[6, \infty) \to E.$$

They then work to show that in especially good cases they can complete the missing arrow in the diagram

$$\begin{tikzcd} MU[6,\infty) \arrow[r] \arrow[dr] & MString \arrow[d, dotted] \\ & E. \end{tikzcd}$$

Leaving aside the extension problem for the moment, their main theorem is the following description of the cohomology ring $E^* MU[6, \infty)$:

Theorem 5 (Ando–Hopkins–Strickland [AHS01], see Singer [Sin68] and Stong [Sto63]) *For E an even-periodic cohomology theory, there is an isomorphism*

$$\operatorname{Spec} E_* MU[6, \infty) \cong C^3(\widehat{\mathbb{G}}_E; \mathcal{I}(0)),$$

where "$C^3(\widehat{\mathbb{G}}_E; \mathcal{I}(0))$" is the affine scheme parametrizing cubical structures on

[4] It is a special property of *Spin* manifolds that this class is always divisible by 2.

the line bundle $\mathcal{I}(0)$ over $\widehat{\mathbb{G}}_E$. When E is taken to be elliptic, so that there is a specified elliptic curve C and a specified isomorphism $\widehat{\mathbb{G}}_E \cong C_0^\wedge$, the theory of elliptic curves gives a canonical such cubical structure and hence a preferred class $MU[6, \infty) \to E$. This assignment is natural in the choice of elliptic E. □

Our real goal is to understand theorems like this last one, where algebraic geometry asserts some real control over something squarely in the domain of homotopy theory. In the course of this text, we will work through a sequence of case studies where this perspective shines through most brightly. In particular, rather than taking an optimal route to the Ando–Hopkins–Strickland result, we will use it as a gravitational slingshot to cover many ancillary topics which are also governed by the technology of formal geometry. We will begin by working through Thom's calculation of the homotopy of MO, which holds the simultaneous attractive features of being approachable while revealing essentially all of the structural complexity of the general situation, so that we know what to expect later on. Having seen that through, we will then venture on to other examples: the complex bordism ring, structure theorems for finite spectra, unstable cooperations, and, finally, the σ-orientation and its extensions. Again, the overriding theme of the text will be that algebraic geometry is a good organizing principle that gives us one avenue of insight into how homotopy theory functions: It allows us to organize "operations" of various sorts between spectra derived from bordism theories.

We should also mention that we will specifically *not* discuss the following aspects of this story:

- Analytic techniques will be completely omitted. Much of modern research stemming from the above problem is an attempt to extend index theory across Witten's genus, or to find a "geometric cochains" model of certain elliptic cohomology theories. These often mean heavy analytic work, and we will strictly confine ourselves to the domain of homotopy theory.
- As sort of a subpoint (and despite the motivation provided in this introduction), we will also mostly avoid manifold geometry. Again, much of the contemporary research about *tmf* is an attempt to find a geometric model, so that geometric techniques can be imported – including equivariance and the geometry of quantum field theories, to name two.
- In a different direction, our focus will not linger on actually computing bordism rings MG_*, nor will we consider geometric constructions on manifolds and their behavior after imaging into the bordism ring. This is also the source of active research: the structure of the symplectic bordism ring remains, to a large extent, mysterious, and what we do understand of it comes through a mix of formal geometry and raw manifold geometry. This could be a topic

that fits logically into this document, were it not for time limitations and the author's inexpertise.

- The geometry of E_∞-rings will also be avoided, at least to the extent possible. Such objects become inescapable by the conclusion of our story, but there are better resources from which to learn about E_∞-rings, and the pre-E_∞ story is not told so often these days. So, we will focus on the unstructured part, relegate the E_∞-rings to Appendix A, and leave their details to other authors.
- As a related note, much of the contents of this book could be thought of as computational foundations for the derived algebraic geometry of even-periodic ring spectra. We will make absolutely no attempt to set up such a theory here, but we will endeavor to phrase our results in a way that will, hopefully, be forward-compatible with any such theory arising in the future.
- There will be plenty of places where we will avoid making statements in maximum generality or with maximum thoroughness. The story we are interested in telling draws from a blend of many others from different subfields of mathematics, many of which have their own dedicated textbooks. Sometimes this will mean avoiding stating the most beautiful theorem in a subfield in favor of a theorem we will find more useful. Other times this will mean abbreviating someone else's general definition to one more specialized to the task at hand. In any case, we will give references to other sources where you can find these cast in starring roles.

Finally, we must mention that there are several good companions to these notes. Essentially none of the material here is original – it is almost all cribbed either from published or unpublished sources – but the source documents are quite scattered and individually dense. We will make a point to cite useful references as we go. One document stands out above all others, though: Neil Strickland's *Functorial Philosophy for Formal Phenomena* [Strb]. These lecture notes can basically be viewed as an attempt to make it through this paper in the span of a semester.

Conventions

Throughout this book, we use the following conventions:

- Categories will be consistently typeset as in the examples

$$\text{Spaces}, \quad \text{FormalGroups}, \quad \text{GradedHopfAlgebras}.$$

- $\mathsf{C}(X, Y)$ will denote the mapping set of arrows $X \to Y$ in a category C. If C is an ∞-category, this will be interpreted instead as a mapping *space*. If C has a self-enrichment, we will often write $\underline{\mathsf{C}}(X, Y)$ (or, e.g., $\underline{\mathrm{Aut}}(X)$) to distinguish the internal mapping object from the classical mapping set $\mathsf{C}(X, Y)$. As a first exception to this uniform notation, we will sometimes abbreviate $\underline{\text{Spaces}}(X, Y)$ to $F(X, Y)$, and similarly we will sometimes abbreviate $\underline{\text{Spectra}}(X, Y)$ to $F(X, Y)$, with "F" short for "function." As a second exception, for two formal groups $\widehat{\mathbb{G}}$ and $\widehat{\mathbb{H}}$, we denote the function scheme by

$$\underline{\text{FormalGroups}}(\widehat{\mathbb{G}}, \widehat{\mathbb{H}}),$$

even though this is a *scheme* rather than a *formal scheme*.
- Following Lurie, for an object $X \in \mathsf{C}$ we will write $\mathsf{C}_{/X}$ for the slice category of objects *over* X and $\mathsf{C}_{X/}$ for the slice category of objects *under* X.
- For a spectrum E, we will write $E^*(X)$ for the unreduced E-cohomology of a space X and $E_*(X)$ for the unreduced E-homology of X. We denote the reduced E-cohomology of a pointed space X by $\widetilde{E}^*(X)$ and the reduced E-homology by $\widetilde{E}_*(X)$. Finally, for F another spectrum, we write $E^*(F)$ and $E_*(F)$ for the E-cohomology and E-homology respectively of F. Altogether, these satisfy the relations

$$E_*(X) = E_*(\Sigma^\infty_+ X) = E_*(\Sigma^\infty X) \oplus E_* = \widetilde{E}_*(X) \oplus E_*,$$

and similarly for cohomology.
- For a spectrum E, we will write $\underline{E}_n$ for the nth space in the Ω-spectrum representing E. The homotopy type of this space is determined by the formula

$$h\text{Spaces}(X, \underline{E}_n) = h\text{Spaces}(X, \Omega^\infty \Sigma^n E) = h\text{Spaces}(\Sigma^\infty X, \Sigma^n E) = \widetilde{E}^n(X).$$

- For a ring spectrum E, we will write $E_* = \pi_* E$ for its coefficient ring, $E^* = \pi_{-*} E$ for its coefficient ring with the opposite grading, and $E_0 = E^0 = \pi_0 E$ for the zeroth degree component of its coefficient ring. This allows us to make sense of expressions like $E^*[\![x]\!]$, which we interpret as

$$E^*[\![x]\!] = (E^*)[\![x]\!] = (\pi_{-*}E)[\![x]\!] = \left\{ \sum_{j=0}^{\infty} a_j x^j \,\middle|\, \begin{array}{l} a_j \text{ is of degree } * - j|x| \\ \text{for some fixed degree } * \end{array} \right\}.$$

- For a space or spectrum, we will write $X[n, \infty) \to X$ for the upward nth Postnikov truncation over X and $X \to X(-\infty, n)$ for the downward nth Postnikov truncation under X. There is thus a natural fiber sequence

$$X[n, \infty) \to X \to X(-\infty, n).$$

  This notation extends naturally to objects like $X(a, b)$ or $X[a, b]$, where $-\infty \le a \le b \le \infty$ denote the (closed or open) endpoints of any interval.
- We will write S^n for the nth sphere when considered as a space and $\mathbb{S}^n$ for its suspension spectrum. We will often abbreviate $\mathbb{S}^0$ to simply $\mathbb{S}$.
- We prefer the notation $\mathcal{O}_X$ for the ring of functions on a scheme X and $\mathcal{I}_D$ for ideal sheaf determined by a subscheme D, but we will also denote these by the synonyms $\mathcal{O}(X)$ and $\mathcal{I}(D)$ when the subscripts reach sufficient complexity.
- We write KO and KU for periodic real and complex K-theory, and we write kO and kU for their respective connective variants. (Other authors write ko and ku, or bo and bu, or even the ill-advised BO and BU for these spectra.)
- We primarily treat 2-periodic spectra, though "in the wild" many of the spectra we consider here are taken to have lower periodicity (e.g., $E(d)$ is typically taken to have periodicity $2(p^d - 1)$) or no periodicity at all (e.g., the ordinary homology spectrum $H\mathbb{F}_2$). Where confusion might otherwise arise, we have done our best to insert a "P" into the names of our standard spectra as a clear indication that we are speaking about the 2-periodic version.

In all these cases, I have done my best to be absolutely consistent in these regards, and I apologize profusely for any erratic typesetting that might have slipped through.

Case Study 1

Unoriented Bordism

A simple observation about the bordism ring $MO_*(*)$ (or $MO_*(X)$ more generally, for any space X) is that it consists entirely of 2-torsion: Any chain $Z \to X$ can be bulked out to a constant cylinder $Z \times I \to X$, which has as its boundary the chain $2 \cdot (Z \to X)$. Accordingly, $MO_*(X)$ is always an $\mathbb{F}_2$-vector space. Our goal in this case study is to arrive at two remarkable calculations: First, in Corollary 1.5.7 we will make an explicit calculation of this $\mathbb{F}_2$-vector space in the case of the bordism homology of a point; and second, in Lemma 1.5.8 we will show that there is a natural isomorphism

$$MO_*(X) = H\mathbb{F}_{2*}(X) \otimes_{\mathbb{F}_2} MO_*(*).$$

Our goal in discussing these results in the first case study of the book is to take the opportunity to introduce several key concepts that will serve us throughout. First and foremost, we will require a definition of bordism spectrum that we can manipulate computationally, using just the tools of abstract homotopy theory. Once that is established, we immediately begin to bring algebraic geometry into the mix: The main idea is that the cohomology ring of a space is better viewed as a scheme (with plenty of extra structure), and the homology groups of a spectrum are better viewed as a representation for a certain elaborate algebraic group. This data actually finds familiar expression in homotopy theory: We show that a form of group cohomology for this representation forms the input to the classical Adams spectral sequence, which classically takes the form

$$\operatorname{Cotor}^{*,*}_{\mathcal{A}_*}(\mathbb{F}_2, H\mathbb{F}_{2*}(Y)) \Rightarrow \pi_*(Y),$$

converging for certain very nice spectra Y – including, for instance, $Y = MO$. In particular, we can bring the tools from the preceding discussion to bear on the homology and cohomology of MO, where we make an explicit calculation of its representation structure. Finding that it is suitably free, we thereby gain control of the Adams spectral sequence, finish the computation, and prove the desired result.

Our *real* goal in this case study, however, is to introduce one of the main phenomena guiding this text: There is some governing algebro-geometric object, the formal group $\mathbb{RP}^{\infty}_{H\mathbb{F}_2}$, which exerts an extraordinary amount of control over everything in sight. We will endeavor to rephrase as much of this classical computation as possible so as to highlight its connection to this central object, and we will use this as motivation in future case studies to pursue similar objects, which will lead us down much deeper and more rewarding rabbit holes. The counterbalance to this is that, at least for now, we will not introduce concepts or theorems in their maximum generality.[1] Essentially everything mentioned in this case study will be re-examined more thoroughly in future case studies, so the reader is advised to look to those for the more expansive set of results.

1.1 Thom Spectra and the Thom Isomorphism

Our goal is a sequence of theorems about the unoriented bordism spectrum MO. We will begin by recalling a definition of the spectrum MO using just abstract homotopy theory, because it involves ideas that will be useful to us throughout the text and because we cannot compute effectively with the chain-level definition given in the Introduction.

Definition 1.1.1 For a spherical bundle $S^{n-1} \to \xi \to X$, its *Thom space* is given by the cofiber

$$\xi \to X \xrightarrow{\text{cofiber}} T_n(\xi).$$

"Proof" of definition There is a more classical construction of the Thom space: Take the associated disk bundle by gluing an n-disk fiberwise, and add a point at infinity by collapsing ξ:

$$T_n(\xi) = (\xi \cup_{X \times S^{n-1}} (X \times D^n))^+.$$

To compare this with the cofiber definition, recall that the thickening of ξ to an n-disk bundle is the same as taking the mapping cylinder on $\xi \to X$. Since the inclusion into the mapping cylinder is now a cofibration, the quotient by this subspace agrees with both the cofiber of the map and the introduction of a point at infinity. □

Before proceeding, here are two important examples:

[1] For an obvious example, everything in this case study will be done relative to Spec $\mathbb{F}_2$.