SHEAVES IN HOMOTOPY THEORY

CHRISTOPHER L. DOUGLAS

Contents

Background and Motivation

In the chapter on the Landweber Exact Functor Theorem, we constructed a presheaf $\mathcal{O}^{\text{hom}}$ of homology theories on the moduli stack of elliptic curves, as follows. A map $f : \text{Spec}(R) \to \mathcal{M}_{ell}$ from an affine scheme to the moduli stack of elliptic curves provided an elliptic curve C over the ring R, and this elliptic curve had an associated formal group $\hat{C} : MP_0 = MU_* \to R$. Provided the map $f : \text{Spec}(R) \to \mathcal{M}_{ell}$ was flat, the functor $Ell_{C/R}(X) = MP_*(X) \otimes_{MP_0} R$ was a homology theory. The value of the presheaf $\mathcal{O}^{\text{hom}}$ on an elliptic curve was defined to be the homology theory associated to the formal group of that elliptic curve: $\mathcal{O}^{\text{hom}}(f) = Ell_{C/R}$. Recall that such a presheaf is by definition simply a contravariant functor:

$$\begin{array}{ccc} (\text{Affine Schemes}/(\mathcal{M}_{ell}))^{\text{op}} & \xrightarrow{\mathcal{O}^{\text{hom}}} & \text{Homology Theories} \\ C/R & \longmapsto & Ell_{C/R} \end{array}$$

The presheaf $\mathcal{O}^{\text{hom}}$ nicely encodes all the homology theories built from the formal groups of elliptic curves. The only problem is that there are many such theories, and they are related to one another in complicated ways. We would like instead a "global" or "universal" elliptic homology theory. The standard way of building a global object from a presheaf is of course to take global

Date: 23 July 2007.

sections. Unfortunately, the site of affine schemes over the moduli stack of elliptic curves has no initial object, and therefore no notion of global sections. We would like to find a homology theory $\mathcal{O}^{\mathrm{hom}}(\mathcal{M}_{ell})$ associated to the whole moduli stack. One might guess that this homology theory should be the limit $\lim_{U\in\mathfrak{U}}\mathcal{O}^{\mathrm{hom}}(U)$ of the theories $\mathcal{O}^{\mathrm{hom}}(U)$, where $\mathfrak{U}$ is an affine cover of the moduli stack. The category of homology theories is not complete, though, and this limit does not exist.

Thanks to Brown representability, we know that homology theories can be represented by spectra. The category of spectra is rather better behaved than the category of homology theories—for instance, it has limits and their homotopically meaningful cousins, homotopy limits. If we can show that the presheaf $\mathcal{O}^{\mathrm{hom}}$ is the presheaf of homology theories associated to a presheaf $\mathcal{O}^{\mathrm{top}}$ of spectra, then we can build a global spectrum, thus have a global homology theory, using a homotopy limit construction. The main theorem is that there is indeed an appropriate presheaf of spectra:

Theorem 0.1 (Hopkins-Miller)**.** *There exists a sheaf $\mathcal{O}^{\mathrm{top}}$ of E_∞ ring spectra on $(\mathcal{M}_{ell})_{\acute{e}t}$, the moduli stack of elliptic curves in the étale topology, whose associated presheaf of homology theories is the presheaf $\mathcal{O}^{\mathrm{hom}}$ built using the Landweber Exact Functor Theorem.*

That $\mathcal{O}^{\mathrm{top}}$ is a sheaf and not merely a presheaf entails, for example, that its value $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ on the whole moduli stack is determined as a homotopy limit $\mathrm{holim}_{U\in\mathfrak{U}}\mathcal{O}^{\mathrm{top}}(U)$ of its value on the open sets in a cover $\mathfrak{U}$ of the moduli stack. The spectrum $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ represents the homology theory we were hunting for, and warrants a special name:

Definition 0.2. *$TMF := \mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$.*

The first goal of this chapter is to explain what it means to have a sheaf of E_∞ ring spectra on the moduli stack of elliptic curves. Note that we would have been happy with a sheaf of (not necessarily E_∞ ring) spectra. That the theorem produces a sheaf of E_∞ ring spectra is an artifact of the ingenious proof: it turns out to be easier to handle the obstruction theory for sheaves of E_∞ ring spectra than the obstruction theory for sheaves of ordinary spectra.

Once we have the sheaf $\mathcal{O}^{\mathrm{top}}$, we would like to understand the global homology theory TMF. In particular, we would like to compute the coefficient ring $TMF_* = \pi_*(\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell}))$. The spectrum $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ is, as described above, built as a homotopy limit out of smaller pieces $\{\mathcal{O}^{\mathrm{top}}(U)\}_{U\in\mathfrak{U}}$. There is a spectral sequence that computes the homotopy groups of the homotopy limit $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ in terms of the homotopy groups of the pieces $\mathcal{O}^{\mathrm{top}}(U)$. The E_2 term of this spectral sequence is conveniently expressed in terms of the sheaf cohomology of the sheafification of the presheaf on $\mathcal{M}_{ell}$ given by $U \mapsto \pi_*(\mathcal{O}^{\mathrm{top}}(U))$.

Proposition 0.3. *There is a strongly convergent spectral sequence*

$$E_2 = H^q(\mathcal{M}_{ell}, \pi_p^\dagger \mathcal{O}^{\mathrm{top}}) \Longrightarrow \pi_{p-q} TMF.$$

Here $\pi_p^\dagger\mathcal{O}^{\mathrm{top}}$ is the sheafification of the presheaf $\pi_p\mathcal{O}^{\mathrm{top}}$.

The second goal of this chapter is to construct this “descent spectral sequence”. In the chapter on the homotopy of TMF we will evaluate the E_2 term of this and related spectral sequences, and illustrate how one computes the numerous differentials.

In the following section 1, we review the classical notion of sheaves, discuss a homotopy-theoretic version of sheaves, and describe stacks as sheaves, in this homotopy sense, of groupoids. Then in section 2 we describe what it means to have a sheaf on a stack and recall the notion of homotopy limit needed to make sense, in particular, of sheaves of spectra. In the final section 3, we discuss sheaf cohomology and Cech cohomology on a stack and construct the descent spectral sequence for a sheaf of spectra on a stack.

In writing this chapter, we have benefited enormously from discussions with Andrew Blumberg and Andre Henriques, and from reading work of Dan Dugger, Phil Hirschhorn, Paul Goerss and Rick Jardine, and, of course, Mike Hopkins.

1. Sheaves and Stacks

A sheaf S of sets on a space X is a way of functorially associating a set $S(U)$ to each open subset U of X. It is natural to generalize this notion in two directions, by considering sheaves of objects besides sets and by considering sheaves on objects besides spaces. Stacks are, at root, representing objects for moduli problems in algebraic geometry, and as such might seem to have little to do with sheaves. However, stacks can naturally be viewed as sheaves of groupoids on the category of schemes, and this perspective is useful when discussing, as we will in section 2, sheaves on a stack.

1.1. Sheaves.

We begin with the classical notions of presheaves and sheaves of sets on a space.

Definition 1.1. Given a space X, let $\mathcal{X}$ denote the category whose objects are open subsets U of X and whose morphisms are inclusions $U \hookrightarrow V$ of open subsets. A presheaf of sets S on the space X is a contravariant functor from the category $\mathcal{X}$ to the category Set of sets.

Explicitly, the presheaf provides a set $S(U)$ for each open U and a restriction map $S(V) \xrightarrow{r_{VU}} S(U)$ for each inclusion $U \hookrightarrow V$ such that the composite $S(W) \xrightarrow{r_{WV}} S(V) \xrightarrow{r_{VU}} S(U)$ is the restriction map r_{WU}. The prototypical example is the presheaf of real valued functions on the space: $S(U) = \mathrm{Map}(U, \mathbb{R})$; here the restriction maps are restriction of functions. This presheaf has the special property that an element of $S(U)$, that is a function, is uniquely determined by its restriction to any open cover of U by smaller open sets $\{U_i \hookrightarrow U\}_{i \in I}$—such a presheaf is called a sheaf.

Definition 1.2. A sheaf of sets on a space is a presheaf of sets S on a space X such that for all open sets $U \subset X$ and all open covers $\{U_i \hookrightarrow U\}_{i\in I}$ of U, the set $S(U)$ is given by the following limit:

$$S(U) = \lim \left(\prod_i S(U_i) \rightrightarrows \prod_{i,j} S(U_{ij}) \substack{\rightarrow\\\rightarrow\\\rightarrow} \prod_{i,j,k} S(U_{ijk}) \substack{\rightarrow\\\rightarrow\\\rightarrow\\\rightarrow} \cdots \right)$$

Here the intersection $U_i \cap U_j$ is denoted U_{ij}, the triple intersection $U_i \cap U_j \cap U_k$ is denoted U_{ijk}, and so forth. To be clear, the products above occur over unordered tuples of not-necessarily-distinct elements of the indexing set I, and the diagram indexing the limit is the full standard cosimplicial diagram.

Remark 1.3. We emphasize that the limit diagram in this definition does contain codegeneracy maps, despite their frequent omission from the notation. For example, if the index set has order two, then the limit, written out, is $\lim \left(S(U_1) \times S(U_2) \substack{\rightarrow\\\leftarrow\\\rightarrow} S(U_{12}) \times S(U_{11}) \times S(U_{22}) \substack{\rightarrow\\\leftarrow\\\rightarrow\\\leftarrow\\\rightarrow} \cdots \right)$, not $\lim \left(S(U_1) \times S(U_2) \rightrightarrows S(U_{12}) \substack{\rightarrow\\\rightarrow\\\rightarrow} * \substack{\rightarrow\\\rightarrow\\\rightarrow\\\rightarrow} \cdots \right)$.

Remark 1.4. The classical definition of a sheaf demands that $S(U)$ be the limit $\lim \left(\prod_i S(U_i) \rightrightarrows \prod_{i,j} S(U_{ij}) \right)$. This truncated limit is equal to the limit in definition 1.2. However, only the full limit generalizes well when we consider sheaves of objects other than sets.

We can define a presheaf of sets on a category $\mathcal{C}$, not necessarily the category of open subsets of a space, simply as a contravariant functor from $\mathcal{C}$ to Set. Moreover we can give a definition of sheaves of sets on $\mathcal{C}$ provided we have a notion of covers in the category. A Grothendieck topology on a category $\mathcal{C}$ provides such a notion:

Definition 1.5. A Grothendieck topology on a category $\mathcal{C}$ is a collection of sets of morphisms $\{\{U_i \to U\}_{i\in I}\}$; these sets of morphisms are called covering families. The collection of covering families is required to satisfy the following axioms: 1) $\{f : V \to U\}$ is a covering family if f is an isomorphism; 2) if $\{U_i \to U\}_{i \in I}$ is a covering family, and $g : V \to U$ is a morphism, then $\{g^* U_i \to V\}$ is a covering family; 3) if $\{U_i \to U\}$ is a covering family and $\{V_{ij} \to U_i\}$ is a covering family for each i, then $\{V_{ij} \to U\}$ is a covering family. A pair of a category $\mathcal{C}$ and a Grothendieck topology on $\mathcal{C}$ is called a Grothendieck site.

The basic example of a Grothendieck site is of course the category of open subsets of a space, with morphisms inclusions, together with covering families the sets $\{U_i \to U\}$ where $\{U_i\}$ is an open cover of U. More interesting are the various Grothendieck topologies on the category Sch of schemes. For example, in the étale (respectively flat) topology, the covering families are the sets $\{U_i \to U\}$ such that $\coprod_i U_i \to U$ is an étale (respectively flat) covering map. A sheaf on a Grothendieck site $\mathcal{C}$ is of course a presheaf $S : \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ such that for all covering families $\{U_i \to U\}$ of the site, the set $S(U)$ is the limit $\lim \left(\prod_i S(U_i) \rightrightarrows \prod_{i,j} S(U_{ij}) \substack{\to\\\to\\\to} \cdots \right)$.

Next we consider sheaves on a Grothendieck site $\mathcal{C}$ taking values in a category $\mathcal{D}$ other than sets. We are interested in categories $\mathcal{D}$ that have some notion of homotopy theory—these include the categories of groupoids, spaces, spectra, and E_∞ ring spectra. More specifically, we need the category $\mathcal{D}$ to come equipped with a notion of homotopy limits and a notion of weak equivalences. We will discuss homotopy limits in detail in section 2.3. For now we content ourselves with a brief example illustrating the idea that homotopy limits in, for example, spaces behave like limits with a bit of homotopical wiggle room:

Example 1.6. Suppose we are interested in the diagram of spaces $X \rightrightarrows Y$, where the two maps are f and g. The limit of this diagram is the space of points of X whose image in Y is the same under the two maps: $\lim(X \rightrightarrows Y) = \{x \in X \text{ s.t. } f(x) = g(x)\}$. The homotopy limit, by contrast, only expects the two images to be the same up to chosen homotopy:

$$\operatorname{holim}(X \rightrightarrows Y) = \{(x \in X, h_x : [0,1] \to Y) \text{ s.t. } h_x(0) = f(x), h_x(1) = g(x)\}.$$

A presheaf on the site $\mathcal{C}$ with values in the category $\mathcal{D}$ is a contravariant functor $F : \mathcal{C}^{\mathrm{op}} \to \mathcal{D}$. These presheaves are also referred to as presheaves of objects of $\mathcal{D}$ on the site $\mathcal{C}$: for example, "presheaves of sets", "presheaves of spaces", "presheaves of spectra".

Definition 1.7. A sheaf on the site $\mathcal{C}$ with values in the category $\mathcal{D}$ is a presheaf F such that for all objects U of $\mathcal{C}$ and all covers $\{U_i \to U\}_{i \in I}$, the map

$$F(U) \xrightarrow{\simeq} \operatorname{holim} \left(\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_{ij}) \substack{\to\\\to\\\to} \prod_{i,j,k} F(U_{ijk}) \substack{\to\\\to\\\to\\\to} \cdots \right)$$

is a weak equivalence. The products here occur over unordered tuples of not-necessarily-distinct elements of the indexing set I; in particular, the indexing diagram does contain codegeneracy maps.

We will be particularly interested in the case where $\mathcal{C}$ is the étale site $(\mathcal{M}_{ell})_{\text{ét}}$ on the moduli stack of elliptic curves and $\mathcal{D}$ is the category of E_∞ ring spectra. We describe this particular site in section 2.2 and discuss homotopy limits of (E_∞ ring) spectra in section 2.3.

1.2. **Stacks.** A scheme X represents a functor $\mathsf{Sch}^{\mathrm{op}} \to \mathsf{Set}$ by $Y \mapsto \operatorname{Hom}(Y, X)$. A moduli problem, such as "What are the elliptic curves over a scheme?", also associates a set to each scheme, for example by $Y \mapsto \{\text{ell curves}/Y\}/\text{iso}$. Unfortunately, many such moduli problems are not representable by schemes. To manage this situation, we keep track of not just the moduli set but the moduli groupoid. We therefore consider, for example, the association taking $Y \in \mathsf{Sch}$ to $\{\text{ell curves}/Y, \text{with isoms}\} \in \mathsf{Gpd}$. This association is very nearly a presheaf of groupoids on the category of schemes; (it is not a presheaf because pullback is only functorial up to isomorphism.) It moreover has the sheaf-like property that elliptic curves over a scheme Y can be built by gluing together elliptic curves on a cover of Y. Altogether, this suggests that sheaves of groupoids are a reasonable model for studying moduli problems.

Definition 1.8. A stack on the site $\mathcal{C}$ is a sheaf of groupoids on $\mathcal{C}$.

Recall that this definition means that for a presheaf $F : \mathcal{C}^{\mathrm{op}} \to \mathsf{Gpd}$ to be a stack, the map $F(U) \to \operatorname{holim} \left(\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_{ij}) \substack{\to\\\to\\\to} \cdots \right)$ must be a weak equivalence for all covers $\{U_i \to U\}$. In order to unpack this condition, we need to know what the weak equivalences and the homotopy limits are in the category of groupoids. A weak equivalence of groupoids is simply an equivalence of categories. The following proposition identifies the needed homotopy limit.

Proposition 1.9 ([6])**.** *The homotopy limit of groupoids* holim$\left(\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_{ij}) \overset{\rightarrow}{\rightrightarrows} \cdots\right)$ *associated to a cover* $\mathfrak{U} = \{U_i \to U\}$ *is the groupoid* $\mathrm{Desc}(F, \mathfrak{U})$ *defined as follows. The objects of* $\mathrm{Desc}(F, \mathfrak{U})$ *are collections of objects* a_i *in* $\mathrm{ob}(F(U_i))$ *and morphisms* $\alpha_{ij} : a_i|_{U_{ij}} \to a_j|_{U_{ij}}$ *in* $\mathrm{mor}(F(U_{ij}))$ *such that* $\alpha_{jk}\alpha_{ij} = \alpha_{ik}$*. The morphisms in* $\mathrm{Desc}(F, \mathfrak{U})$ *from* $\{a_i, \alpha_{ij}\}$ *to* $\{b_i, \beta_{ij}\}$ *are collections of morphisms* $m_i : a_i \to b_i$ *in* $\mathrm{mor}(F(U_i))$ *such that* $\beta_{ij}m_i = m_j\alpha_{ij}$*.*

Stacks are often defined to be presheaves of groupoids such that the natural map $F(U) \to \mathrm{Desc}(F, \mathfrak{U})$ is an equivalence of categories for all covers $\mathfrak{U}$ of U. The above proposition establishes that the more conceptual homotopy limit definition agrees with the descent definition.

Example 1.10. We will be concerned primarily with the moduli stack of elliptic curves $\mathcal{M}_{ell}$. This is a stack on the category of schemes in any of the flat, étale, or Zariski topologies. Roughly speaking, the stack associates to a scheme Y the groupoid of elliptic curves over Y. (Precisely, this association must be slightly rigidified, a la [7, p.26].) Alternately, $\mathcal{M}_{ell}$ is the stack associated to the Hopf algebroid $(A, \Gamma) := (\mathbb{Z}[a_1, a_2, a_3, a_4, a_6][\Delta^{-1}], A[u^{\pm 1}, r, s, t][\Delta^{-1}])$.

2. Sheaves on Stacks

We would like to understand sheaves of spectra on the moduli stack of elliptic curves in the étale topology. First we consider the general notion of a sheaf on a stack, then describe the étale site of the moduli stack of elliptic curves, and finally discuss sheaves of spectra in particular.

2.1. The Site of a Stack.

Suppose $\mathcal{C}$ is a Grothendieck site and X is an object of $\mathcal{C}$. We have a notion of sheaves on X, which are by definition sheaves on the site $\mathcal{C}/X$ whose objects are maps $Y \to X$ in $\mathcal{C}$ and whose covers are inherited from $\mathcal{C}$. We would like a notion of sheaves on a stack $\mathcal{M}$ on the site $\mathcal{C}$. In order to consider objects of $\mathcal{C}$ over $\mathcal{M}$, we need objects of $\mathcal{C}$ to live in the same place as stacks on $\mathcal{C}$. This is accomplished by the following functors:

$$\begin{gathered} \mathcal{C} \hookrightarrow \mathrm{Pre}\,\mathcal{C} \hookrightarrow \mathrm{Pre}_{\mathsf{Gpd}}\,\mathcal{C} \\ U \mapsto \mathrm{Hom}(-, U) \mapsto \mathrm{Hom}(-, U) \text{ with id} \end{gathered}$$

Thinking of object of $\mathcal{C}$ as the presheaves of groupoids they represent, we can consider the site of objects over $\mathcal{M}$.

Definition 2.1. Let $\mathcal{M}$ be a stack on the site $\mathcal{C}$. The site $\mathcal{C}/\mathcal{M}$ has objects the morphisms $U \to \mathcal{M}$ in presheaves of groupoids on $\mathcal{C}$. The morphisms in $\mathcal{C}/\mathcal{M}$ from $U \xrightarrow{a} \mathcal{M}$ to $V \xrightarrow{b} \mathcal{M}$ are the pairs (c, ϕ) where $c : U \to V$ is a morphism of $\mathcal{C}$ and ϕ is a natural isomorphisms between the functors a and bc. The covering families of $U \to \mathcal{M}$ in $\mathcal{C}/\mathcal{M}$ are the sets of morphisms $\{c_i : U_i \to U, \phi_i\}$ such that $\{c_i\}$ is a covering family in $\mathcal{C}$. Schematically, the definition is

$$\mathcal{C}/\mathcal{M} = \begin{cases} \mathrm{obj} = \{U \to \mathcal{M} \text{ in } \mathrm{Pre}_{\mathsf{Gpd}}\,\mathcal{C}\} \\ \mathrm{mor} = \left\{ \begin{tikzcd} U \arrow[dr, "a"] \arrow[d, "c"'] & \\ V \arrow[r, "b"'] \arrow[ur, phantom, "\Downarrow \phi"] & \mathcal{M} \end{tikzcd} \right\} \\ \mathrm{cov} = \left\{ \left\{ \begin{tikzcd} U_i \arrow[dr] \arrow[d] & \\ U \arrow[r] \arrow[ur, phantom, "\Downarrow"] & \mathcal{M} \end{tikzcd} \right\} \text{ s.t. } \left\{ \begin{tikzcd} U_i \arrow[d] \\ U \end{tikzcd} \right\} \text{ cov in } \mathcal{C} \right\} \end{cases}$$

Definition 2.2. For a stack $\mathcal{M}$ on the site $\mathcal{C}$, a sheaf on $\mathcal{M}$ with values in $\mathcal{D}$ is a sheaf on $\mathcal{C}/\mathcal{M}$ with values in $\mathcal{D}$.

For example, we might consider $\mathcal{M}_{ell}$ as a stack on schemes in the étale topology $\mathsf{Sch}_{\text{ét}}$ and then consider sheaves on $\mathcal{M}_{ell}$ with values in the category of spectra. However, in the end this is not the notion of sheaves on the moduli stack that we want, so we need to modify the site $\mathsf{Sch}_{\text{ét}}/\mathcal{M}_{ell}$.

2.2. **The site $(\mathcal{M}_{ell})_{\text{ét}}$.** We adjust the site $\mathsf{Sch}_{\text{ét}}/\mathcal{M}_{ell}$ in two ways. First, to enable later obstruction theory arguments we need to restrict the objects of our site to be étale, not arbitrary, maps to the moduli stack. Second, it will be convenient if our sheaves take values not only on schemes over the moduli stack but also on stacks over the moduli stack; in particular we will then be able to evaluate a sheaf on the moduli stack itself, producing a spectrum of global sections.

We will be interested in étale maps between stacks and étale covers of stacks. These notions are derived from the corresponding notions for schemes. Recall that a map $X \to Y$ between schemes is étale if it is flat and unramified or equivalently smooth of relative dimension zero. The topologist can think of étale maps as being the algebro-geometric analog of local homeomorphisms. A collection of étale maps $\{U_i \to U\}$ is an étale cover if for all algebraically closed fields k and all maps $f : \operatorname{Spec} k \to U$ there exists an i with a lift of f to a map $\tilde{f} : \operatorname{Spec} k \to U_i$:

$$\left\{\begin{tikzcd} U_i \arrow[d] \\ U \end{tikzcd}\right\} \text{ étale cover if } \begin{tikzcd} & U_i \arrow[d] \\ \operatorname{Spec} k \arrow[r, "\forall f"'] \arrow[ur, dashed, "\exists \tilde{f}"] & U \end{tikzcd}$$

Note that these collections are the covers in the site $\mathsf{Sch}_{\text{ét}}$.

Definition 2.3. A map of stacks $f : \mathcal{N} \to \mathcal{M}$ is étale if for all maps $V \to \mathcal{M}$ from a scheme to $\mathcal{M}$, the pullback f^*V is a scheme and the induced map $f^*V \to V$ is étale:

$$\mathcal{N} \xrightarrow{f} \mathcal{M} \text{ étale if } \begin{tikzcd} f^*V \arrow[r, "\boxed{\text{ét}}"] \arrow[d] & V \arrow[d, "\forall"] \\ \mathcal{N} \arrow[r, "f"'] & \mathcal{M} \end{tikzcd}$$

Étale maps $\mathcal{N} \to \mathcal{M}_{ell}$ to the moduli stack of elliptic curves will be the objects of the étale site of the moduli stack. Roughly, the morphisms are maps of stacks $\mathcal{N}' \to \mathcal{N}$ over the moduli stack, and covers are collections of maps $\{\mathcal{N}_i \to \mathcal{N}\}$ over the moduli stack that are étale covers in their own right. (A collection of maps of stacks is an étale cover if it satisfies a lifting property precisely analogous to the one for étale covers of schemes.) In more detail, the étale site is defined as follows. Note that we are considering all stacks, the moduli stack $\mathcal{M}_{ell}$ included, as stacks on the étale site of schemes $\mathsf{Sch}_{\text{ét}}$.

Definition 2.4. The objects of the étale site of the moduli stack of elliptic curves $(\mathcal{M}_{ell})_{\text{ét}}$ are the étale morphisms $\mathcal{N} \to \mathcal{M}_{ell}$ from a stack $\mathcal{N}$ to the moduli stack. The morphisms from $\mathcal{N} \xrightarrow{a} \mathcal{M}_{ell}$ to $\mathcal{N}' \xrightarrow{b} \mathcal{M}_{ell}$ are equivalence classes of pairs (c, ϕ), where $c : \mathcal{N} \to \mathcal{N}'$ is a map of stacks and ϕ is a natural isomorphism between a and bc. A natural isomorphism ψ from $c : \mathcal{N} \to \mathcal{N}'$ to $d : \mathcal{N} \to \mathcal{N}'$ can be viewed as an isomorphism from the pair (c, ϕ) to the pair $(d, \psi\phi)$—these two pairs are therefore considered equivalent as morphisms between $\mathcal{N} \to \mathcal{M}_{ell}$ and $\mathcal{N}' \to \mathcal{M}_{ell}$. A collection of morphisms $\{[(c_i : \mathcal{N}_i \to \mathcal{N}, \phi_i)]\}$ of stacks over $\mathcal{M}_{ell}$ is a cover of $(\mathcal{M}_{ell})_{\text{ét}}$ if for all algebraically closed fields k and all maps $f : \operatorname{Spec} k \to \mathcal{N}$, there exists an i, a representative (c, ϕ) of the equivalence class $[(c_i, \phi_i)]$, and a lift of f to a map $\tilde{f} : \operatorname{Spec} k \to \mathcal{N}_i$ such that $f = c\tilde{f}$. Schematically, we have

$$(\mathcal{M}_{ell})_{\text{ét}} = \begin{cases} \text{obj} = \{\mathcal{N} \xrightarrow{\text{ét}} \mathcal{M}_{ell}\} \\ \text{mor} = \left\{\begin{tikzcd} \mathcal{N} \arrow[d, "c"'] \arrow[dr, "a"] \arrow[dr, phantom, shift right=3, "\Swarrow \phi"] & \\ \mathcal{N}' \arrow[r, "b"'] & \mathcal{M}_{ell} \end{tikzcd}\right\} \Big/ \left\{\begin{tikzcd} \mathcal{N} \arrow[d, bend right, "d"'] \arrow[d, bend left, "c"] \arrow[d, phantom, "\overset{\psi}{\Leftarrow}"] \\ \mathcal{N}' \end{tikzcd}\right\} \\ \text{cov} = \left\{\left\{\begin{tikzcd} \mathcal{N}_i \arrow[d, "c_i"'] \arrow[dr] \arrow[dr, phantom, shift right=3, "\Swarrow \phi_i"] & \\ \mathcal{N} \arrow[r] & \mathcal{M}_{ell} \end{tikzcd}\right\} \text{ s.t. } \begin{tikzcd} & \mathcal{N}_i \arrow[d] \\ \operatorname{Spec} k \arrow[r, "\forall f"'] \arrow[ur, dashed, "\exists \tilde{f}"] & \mathcal{N} \end{tikzcd}\right\} \end{cases}$$

Though the definition of this étale site $(\mathcal{M}_{ell})_{\text{ét}}$ is complicated by the introduction of stacks over the moduli stack, the site $(\mathcal{M}_{ell})_{\text{ét}}$ still contains the fundamental objects of study, namely elliptic curves over schemes; sheaves on $(\mathcal{M}_{ell})_{\text{ét}}$ should be thought of primarily as assignments of sets (or, in a moment, spectra) to these elliptic curves.

We are hunting for a sheaf of spectra $\mathcal{O}^{\text{top}}$ on the étale site $(\mathcal{M}_{ell})_{\text{ét}}$ of the moduli stack of elliptic curves. Such a sheaf was defined in section 1.1 as a presheaf $F : (\mathcal{M}_{ell})_{\text{ét}}^{\text{op}} \to \mathsf{Spec}$ such that, for all objects U of $(\mathcal{M}_{ell})_{\text{ét}}$, the natural map

$$F(U) \to \operatorname{holim}\left(\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_{ij}) \mathrel{\substack{\textstyle\rightarrow\\[-0.6ex]\textstyle\rightarrow\\[-0.6ex]\textstyle\rightarrow}} \prod_{i,j,k} F(U_{ijk}) \mathrel{\substack{\textstyle\rightarrow\\[-0.6ex]\textstyle\rightarrow\\[-0.6ex]\textstyle\rightarrow\\[-0.6ex]\textstyle\rightarrow}} \cdots\right)$$

is a weak equivalence. We need therefore to understand in detail the notion of homotopy limits in the category of spectra.

2.3. **Homotopy Limits and Sheaves of Spectra.** A sheaf of sets on a space X is a functor $F : \{U \subset X\}^{\text{op}} \to \mathsf{Set}$ whose value on large open sets is determined as a limit of the value on smaller open sets. We could take a similar definition for sheaves of spectra $F : \{U \subset X\}^{\text{op}} \to \mathsf{Spec}$, but the limit condition ignores the topological structure of spectra, and the value of the sheaf on large open sets would not capture any information about the topological behavior of the sheaf on small open sets. Instead of "gluing" the values of F together with a limit, we glue them together with a homotopy limit. The homotopy limit takes the various values of F and thickens them up with a bit of padding, so that they aren't too badly damaged, homotopically speaking, by the gluing process.

In section 2.3.1, we describe colimits and limits in terms of tensors and cotensors, and use this framework to give a concise description of homotopy colimits and limits. In section 2.3.2, we fess up to the fact that even the homotopy limit is not always appropriately homotopy invariant, and this leads us into a discussion of derived limits and corrected homotopy limits. In section 2.3.3, we specialize to the case of spectra, describing the categories of orthogonal spectra, symmetric spectra, and S-modules and specifying the tensors and cotensors needed for homotopy colimits and limits in these categories.

2.3.1. *Limits and Homotopy Limits.* Limit and colimit are brutal operations in the category of spaces: they tend to destroy homotopical information, and they are not invariant under homotopies of maps. For example, the colimit of the diagram $* \leftarrow S^2 \to *$ is a point and has no recollection of the homotopy type of the middle space S^2; the limit of the diagram $* \rightrightarrows [0,1]$, where both maps send the point to 0, is also a point, but becomes empty if we deform one of the two maps away from 0.

We would like homotopy versions of limit and colimit that have more respect for the homotopical structure of spaces. We take our cue from two fundamental examples.

Example 2.5. The colimit of the diagram of spaces $A \xleftarrow{f} C \xrightarrow{g} B$ is $(A \sqcup B)/(f(c) \sim g(c), c \in C)$. We can homotopify this construction by, instead of directly identifying $f(c)$ and $g(c)$, putting a path between them. This is the double mapping cylinder construction and is an example of a homotopy colimit:

$$\begin{aligned}\operatorname{hocolim}(A \xleftarrow{f} C \xrightarrow{g} B) &= \operatorname{colim}(A \leftarrow C \to C \times [0,1] \leftarrow C \to B)\\ &= (A \sqcup C \times [0,1] \sqcup B)/\{(c,0) \sim f(c), (c,1) \sim g(c)\}.\end{aligned}$$

The suspension functor is a special case of this homotopy colimit, when $A = B = *$.

Example 2.6. The limit of the diagram of spaces $X \xrightarrow{f} Z \xleftarrow{g} Y$ is $\{(x,y) \in X \times Y | f(x) = g(y)\}$. Instead of expecting $f(x)$ and $g(y)$ to be equal in this limit, we can merely demand that they be

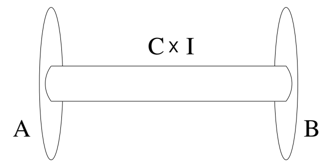

connected by a chosen path. This is the double path space construction and is a homotopy limit:

$$\begin{aligned}\operatorname{holim}(X \xrightarrow{f} Z \xleftarrow{g} Y) &= \lim\left(X \to Z \leftarrow Z^{[0,1]} \to Z \leftarrow Y\right)\\ &= \{(x, m : [0,1] \to Z, y) | f(x) = m(0), g(y) = m(1)\}.\end{aligned}$$

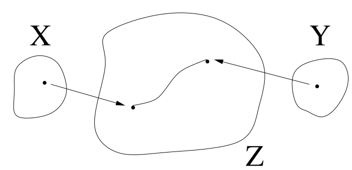

The loop functor is the special case of $X = Y = *$ and $f = g$.

In order to generalize these homotopical constructions to other colimits and limits, it is convenient to have a concise description of the colimit and limit functors. Let $\mathcal{C}$ be the category we are working in, typically spaces or spectra or more generally a simplicial model category, let I be a small category, and let $X : I \to \mathcal{C}$ be a diagram in $\mathcal{C}$ indexed by I. The colimit and limit can be explicitly constructed as follows:

$$\operatorname{colim}_I X \cong X \otimes_I *_I$$
$$\lim_I X \cong \hom^I(*_I, X)$$

Here $*_I$ is the I-diagram of simplicial sets with $*_I(i) = *$ for all $i \in I$. (Full disclosure: here we are using $*_I$ to refer both to this trivial I-diagram and to the trivial I^{op}-diagram.) The constructions $\otimes_I$ and $\hom^I$ are the tensor and cotensor on the diagram category $\mathcal{C}^I$; these are special cases of, respectively, coends and ends, and are discussed in the following remark and example.

Remark 2.7. We recall the tensor and cotensor on the diagram category $\mathcal{C}^I$, following Hirschhorn [5, §18.3.1]. For X an I-diagram in $\mathcal{C}$ and A an I^{op}-diagram in simplicial sets, the tensor of X and A is as follows:

$$X \otimes_I A := \operatorname{colim}\left(\coprod_{i \xrightarrow{\alpha} j} X(i) \otimes A(j) \rightrightarrows \coprod_i X(i) \otimes A(i)\right)$$

This is the coend $\int^i X(i) \otimes A(i)$. Here $\otimes$ is the tensor action of simplicial sets on the category $\mathcal{C}$.

For X again an I-diagram in $\mathcal{C}$ and A an I-diagram in simplicial sets, the cotensor of A and X is as follows:

$$\hom^I(A, X) := \lim\left(\prod_i X(i)^{A(i)} \rightrightarrows \prod_{j \xrightarrow{\alpha} i} X(i)^{A(j)}\right)$$

This is the end $\int_i X(i)^{A(i)}$. Here the superscript refers to the cotensor coaction of simplicial sets on the category $\mathcal{C}$.

Example 2.8. When $\mathcal{C}$ is the category of spaces, the colimit tensor expression $X \otimes_I *_I$ boils down to the space $\left(\coprod_{i \in I} X(i)\right) / \{x \sim \alpha(x) \, \forall \, (i \xrightarrow{\alpha} j) \in I, x \in X(i)\}$—this is the disjoint union of all the objects in the diagram, mod equivalences introduced by the arrows of the diagram. Similarly, the limit cotensor $\hom^I(*_I, X)$ is simply $\{(x_i) \in \prod_{i\in I} X(i) | \alpha(x_j) = x_i \,\forall\, (j \xrightarrow{\alpha} i) \in I\}$. This last space can conveniently be thought of as the space of maps of diagrams from the trivial I-diagram to X, and this justifies the "hom" notation for the cotensor.

The point of all this abstract hoopla is that we can replace $*_I$ in the constructions $X \otimes_I *_I$ and $\hom^I(*_I, X)$ by a diagram of larger contractible spaces—this replacement gives us the homotopical wiggle room we were looking for and produces the homotopy colimit and homotopy limit. The minimal natural choices for these contractible spaces come from the nerves of over and under categories in the diagram. Specifically we have the following definitions:

$$\operatorname{hocolim}_I X := X \otimes_I N(-/I)^{\mathrm{op}}$$
$$\operatorname{holim}_I X := \hom^I(N(I/-), X)$$

Here $N(I/-)$ and $N(-/I)^{\mathrm{op}}$ are respectively the functors taking i to the nerve of the over respectively opposite under categories I/i and $(i/I)^{\mathrm{op}}$. See section 2.3.2 for a discussion of why these are sensible replacements for the trivial diagram $*_I$. It is worth writing out this tensor and cotensor:

Definition 2.9. For X an I-diagram in the simplicial model category $\mathcal{C}$, the homotopy colimit and limit are defined as follows:

$$\operatorname{hocolim}_I X = \operatorname{colim}\left(\coprod_{i \xrightarrow{\alpha} j} X(i) \otimes N(j/I)^{\mathrm{op}} \rightrightarrows \coprod_i X(i) \otimes N(i/I)^{\mathrm{op}} \right)$$
$$\operatorname{holim}_I X = \lim\left(\prod_i X(i)^{N(I/i)} \rightrightarrows \prod_{j \xrightarrow{\alpha} i} X(i)^{N(I/j)} \right).$$

The casual reader can safely ignore the "op" here, referring if desired to Hirschhorn [5, Remark 18.1.11] for a description of how and why it arises and also for a comparison of these definitions to the original treatment of homotopy colimits and limits by Bousfield and Kan.

The reader is invited to check that this definition specializes to the description of the homotopy colimit and limit in example 2.5 and example 2.6. Such a specialization requires, of course, knowing the tensor and cotensor on spaces, namely, for a space Y and simplicial set B, that $Y \otimes B = Y \times |B|$ and $Y^B = \mathrm{Map}(|B|, Y)$.

To pin down the homotopy limit and colimit of spectra, it remains only to specify the tensor and cotensor on some particular category of spectra—see section 2.3.3 for these constructions in orthogonal spectra, symmetric spectra, and S-modules. Note that once we have a complete picture of homotopy limits of spectra, we have, combining definitions 1.7 and 2.4, our desired notion of sheaves of spectra on the moduli stack of elliptic curves.

2.3.2. *Derived Limits and Corrected Homotopy Limits.* Unfortunately, the above definitions of homotopy limit and colimit do not always behave as well as we would like, particularly when we are working in categories other than spaces or simplicial sets. In particular, they are not always homotopy invariant and so do not induce functors on the level of homotopy categories. Problems tend to arise when the objects of our diagram are not fibrant or not cofibrant. In this section we discuss these technicalities and describe and differentiate the four relevant notions: limits, derived limits, homotopy limits, and corrected homotopy limits (and, of course, their co- analogs).

As before, let $\mathcal{C}$ be a simplicial model category, I a small category, and $X : I \to \mathcal{C}$ an I-diagram in $\mathcal{C}$. As discussed above, the limit is the functor

$$\lim : \mathcal{C}^I \to \mathcal{C}$$
$$X \mapsto \hom^I(*_I, X)$$

This functor is not homotopy invariant, in that an objectwise weak equivalence $X \xrightarrow{\sim} Y$ of I-diagrams need not induce a weak equivalence of their limits. The most straightforward way to attempt to fix this problem is to derive the limit functor.

The diagram category $\mathcal{C}^I$ can itself have a model structure, and it may even have several depending on particular properties of $\mathcal{C}$ and I. If $\mathcal{C}$ is a combinatorial model category, then $\mathcal{C}^I$ has an injective model structure, where the weak equivalences and cofibrations are detected objectwise [1, 8, 11]. If $\mathcal{C}$ is merely a cofibrantly generated model category, then $\mathcal{C}^I$ has a projective model structure, where the weak equivalences and fibrations are detected objectwise [5, p.224]. If the diagram I is Reedy, then for any model category $\mathcal{C}$, the category of diagrams $\mathcal{C}^I$ has a Reedy model structure, where only the weak equivalences are detected objectwise. In the following, we assume without comment that $\mathcal{C}$ and I have appropriate structure to ensure the existence of injective, projective, or Reedy model structures, as needed.

Limit is a right Quillen functor from the injective model structure on $\mathcal{C}^I$ to $\mathcal{C}$ [11, p.16]. It therefore makes sense to take the total right derived functor of the limit:

$$R\lim : \mathrm{Ho}(\mathcal{C}^I) \to \mathrm{Ho}(\mathcal{C})$$
$$[X] \mapsto [\hom^I(*_I, FX)]$$

Here F is fibrant replacement in the injective model structure, and brackets refer to the object in the homotopy category. By construction this derived limit is homotopy invariant and so is in a sense the "right" replacement for the limit. The derived limit and the homotopy limit are occasionally conflated in the literature, but they are distinct functors and are easy to tell apart because the derived limit is a functor from $\mathrm{Ho}(\mathcal{C}^I)$ to $\mathrm{Ho}(\mathcal{C})$, while the homotopy limit is a functor from $\mathcal{C}^I$ to $\mathcal{C}$. In retrospect, it might have made more sense to call the derived limit the "homotopy limit" and to have a different name for the particular not-always-homotopy-invariant functor now called the homotopy limit—but it is much too late for a terminological switcheroo.

The biggest disadvantage of the derived limit is that it can be quite difficult to calculate the fibrant replacement FX. In general, such a calculation is hopeless, but if the diagram is particularly simple, we can proceed as follows. Suppose the diagram I is Reedy and has cofibrant constants. (A diagram I is said to have cofibrant constants if the constant I-diagram at any cofibrant object of any model category is Reedy cofibrant.) In this case, the limit is right Quillen not only with respect to the injective model structure, but also with respect to the Reedy model structure [5, Thm 15.10.8]. The derived limit (which up to homotopy does not depend on the model structure we use) can therefore be described as $R\lim X = [\hom^I(*_I, F_R X)]$, where F_R is fibrant replacement in the Reedy model structure. We can then go about explicitly calculating the Reedy fibrant replacement of our diagram.

Example 2.10. If the diagram X is $Y \xrightarrow{f} Z \xleftarrow{g} W$, then a Reedy fibrant replacement $F_R X$ is a diagram $Y' \xrightarrow{f'} Z' \xleftarrow{g'} W'$ with an objectwise weak equivalence $X \to F_R X$ such that Y', Z', and W' are fibrant, and f' and g' are fibrations. The derived limit $R\lim X$ of the original diagram is the limit $\lim F_R X$ of the new diagram. Notice that this method, replacing maps by fibrations and then taking the ordinary pullback, is the usual means for calculating homotopy pullbacks. It happens to be the case that it is often enough to convert only one of the two maps to a fibration.

All this said, we would be better off if we could avoid replacing X in either the injective or Reedy model structures.

The homotopy limit is a compromise solution: it avoids the fibrant replacement that plagues the derived limit and is therefore more explicit and calculable, at the expense of some weakening of homotopy invariance. It has the further advantage that it is a "point-set level" functor, not a functor on homotopy categories. Here the key motivation for the homotopy limit comes from shifting attention from the injective to the projective model structure. The derived limit was $[\hom^I(*_I, FX)]$. Though this cotensor $\hom^I$ is not in fact a mapping space, it behaves rather like one. In particular note that $*_I$ is cofibrant in the injective model structure on sSet^I, and FX is by definition fibrant in the injective model structure on $\mathcal{C}^I$, and so we expect $\hom^I(*_I, FX)$ to have, as it does, a well behaved homotopy type. Suppose that instead of fibrantly replacing X in the injective model structure on $\mathcal{C}^I$, we cofibrantly replace $*_I$ in the projective model structure on sSet^I. That is, consider the cotensor $\hom^I(C(*_I), X)$, where C is cofibrant replacement in the projective model structure. Provided X is objectwise fibrant (therefore fibrant in the projective model structure), we might expect this cotensor to have a reasonable homotopy type. Indeed this is the case:

Lemma 2.11. *If the I-diagram X in $\mathcal{C}$ is objectwise fibrant, then the cotensors* $\hom^I(*_I, FX)$ *and* $\hom^I(C(*_I), X)$ *are weakly equivalent, where F is fibrant replacement in the injective model structure on $\mathcal{C}^I$ and C is cofibrant replacement in the projective model structure on sSet^I.*

The lemma is also true if we substitute the Reedy model structure fibrant replacement $F_R X$ (if it makes sense) in place of the injective model structure fibrant replacement FX.

Note that the construction $\hom^I(C(*_I), X)$ has the huge advantage that the replacement $C(*_I)$ only depends on the category I and not on the category $\mathcal{C}$ or the particular diagram X. We can therefore make such a choice of replacement once and for all. The nerve $N(I/-)$ of the overcategory is a cofibrant object in the projective model structure on sSet^I and so provides such a choice [5, Prop 14.8.9]. The definition of homotopy limit follows:

$$\begin{aligned} \operatorname{holim} : \mathcal{C}^I &\to \mathcal{C} \\ X &\mapsto \hom^I(N(I/-), X) \end{aligned}$$

We reiterate that this is a point-set level functor, and is functorial both with respect to the diagram X and with respect to the category I; this would have been difficult to arrange using the cotensor $\hom^I(*_I, FX)$ because we would need to have made a choice, compatible for all diagram categories $\mathcal{C}^I$, of a functorial fibrant replacement functor F. Instead, we make use of the simple functorial cofibrant replacement $C(*_I) = N(I/-)$.

The main disadvantage of the homotopy limit is, as one might guess from Lemma 2.11, that it is not homotopy invariant when the diagram X is not objectwise fibrant. In some categories, such as spaces, this does not present a problem (and indeed the proper behavior of holim on spaces probably accounts for its widespread use and the general lack of clarity concerning its deficiencies). As we intend to work in categories of spectra and ring spectra, though, we must correct this lack of invariance by precomposing with a functorial fibrant replacement. The resulting functor is called the corrected homotopy limit:

$$\begin{aligned} \operatorname{corholim} : \mathcal{C}^I &\to \mathcal{C} \\ X &\mapsto \hom^I(N(I/-), F_{\mathrm{obj}} X) \end{aligned}$$

Here F_{obj} is objectwise functorial fibrant replacement. According to the extent that one views objectwise fibrant replacement as a minor adjustment, one might welcome or disdain the occasional conflation of holim and corholim.

The corrected homotopy limit brings us full circle in so far as it represents the derived functor of the limit: by Lemma 2.11, the cotensor $\hom^I(*_I, FX)$ (or $\hom^I(*, F_R X)$ in the case of a Reedy diagram, either of which represent the derived limit) is weakly equivalent to the corrected homotopy limit $\hom^I(N(I/-), F_{\mathrm{obj}} X)$.

Remark 2.12. The reader may be wondering why we did not define sheaves to be presheaves satisfying a corholim, rather than a holim, condition. Indeed, in all respects that probably would have been wiser, but for reasons of convention we stick to the holim definition. We can get away with this because in the end we will restrict our attention to presheaves of fibrant objects, in which case a holim and a corholim condition amount to the same thing. Further justification for brushing off this distinction will come in the chapter on model categories and model structures on categories of presheaves; there we will see that the fibrant objects in an appropriate model structure on presheaves are objectwise fibrant and satisfy a holim condition (which in particular is then the same as a corholim condition).

We briefly describe the colimit analog of the above discussion. The ordinary colimit is, as before, the functor

$$\begin{aligned} \operatorname{colim} : \mathcal{C}^I &\to \mathcal{C} \\ X &\mapsto X \otimes_I *_I \end{aligned}$$

This functor is not homotopy invariant. It is, though, a left Quillen functor from the projective model structure on $\mathcal{C}^I$ to $\mathcal{C}$, and therefore has a total left derived functor:

$$\begin{aligned} L\operatorname{colim} : \operatorname{Ho}(\mathcal{C}^I) &\to \operatorname{Ho}(\mathcal{C}) \\ [X] &\mapsto [CX \otimes_I *_I] \end{aligned}$$

Here C is cofibrant replacement in the projective model structure.

This cofibrant replacement can be painful to calculate, so instead of replacing X we replace $*_I$. If X is objectwise cofibrant, then $CX \otimes_I *_I$ and $X \otimes_I C(*_I)$ are weakly equivalent—here CX is, as before, the cofibrant replacement in the projective model structure on $\mathcal{C}^I$, while $C(*_I)$ is the cofibrant replacement in the projective model structure on $\mathsf{sSet}^{(I^{\mathrm{op}})}$. The nerve $N(-/I)^{\mathrm{op}}$ is cofibrant in the projective model structure on $\mathsf{sSet}^{(I^{\mathrm{op}})}$ and so provides a particular choice of the latter cofibrant replacement, and thereby the definition of homotopy colimit:

$$\begin{aligned} \operatorname{hocolim} : \mathcal{C}^I &\to \mathcal{C} \\ X &\mapsto X \otimes_I N(-/I)^{\mathrm{op}} \end{aligned}$$

We can reestablish homotopy invariance by precomposing this functor with an objectwise cofibrant replacement; the result is the corrected homotopy colimit:

$$\begin{aligned} \operatorname{corhocolim} : \mathcal{C}^I &\to \mathcal{C} \\ X &\mapsto C_{\mathrm{obj}} X \otimes_I N(-/I)^{\mathrm{op}} \end{aligned}$$

The corrected homotopy colimit represents the derived functor, as desired: the tensor $CX \otimes_I *_I$ (which represents the derived colimit) is weakly equivalent to the corrected homotopy colimit $C_{\mathrm{obj}} X \otimes_I N(-/I)^{\mathrm{op}}$. Here CX is the cofibrant replacement in the projective model structure on $\mathcal{C}^I$, and C_{obj} is objectwise functorial cofibrant replacement.

Example 2.13. If the diagram I is Reedy and has fibrant constants (that is every constant I-diagram at a fibrant object is Reedy fibrant), then the colimit is a left Quillen functor from the Reedy model structure on $\mathcal{C}^I$ [5, Thm 15.10.8]. We can therefore calculate the derived colimit using a Reedy cofibrant replacement: $L\operatorname{colim} X = [C_R X \otimes_I *_I]$. If the diagram X has the form $B \xleftarrow{j} A \xrightarrow{k} C$, a Reedy cofibrant replacement $C_R X$ is a diagram $B' \xleftarrow{j'} A' \xrightarrow{k'} C'$ with an objectwise weak equivalence $C_R X \to X$ such that B', A', and C' are cofibrant and j' and k' are cofibrations. The derived limit is the pushout $\operatorname{colim} C_R X$ of this modified diagram. Indeed, replacing maps by cofibrations in such a diagram is the usual way to calculate homotopy pushouts. Note that it is often sufficient to convert one of the two maps to a cofibration.

Remark 2.14. A last important distinction between the homotopy limit and colimit and their corrected versions is that the latter depend on a choice of model structure on the underlying category $\mathcal{C}$, while the former do not.

2.3.3. *Sheaves of Orthogonal Spectra, Symmetric Spectra, and S-Modules.* It is time to bite the bullet and specify the particular categories of spectra in which we intend to work. The relevant options are symmetric spectra, orthogonal spectra, S-modules, and the categories of commutative symmetric ring spectra, commutative orthogonal ring spectra, and commutative S-algebras. We briefly review the definitions of these various categories. Along the way we describe the tensor and cotensor over simplicial sets that we needed in the definition of hocolim and holim. We refer, however, to the chapter on model categories of spectra for the notions of fibrancy and cofibrancy needed for the corrected homotopy colimit and limit.

Remark 2.15. At the end of the day, we are trying to make sense of the notion of a "sheaf of E_∞ ring spectra". By definition an E_∞ ring spectrum in a particular category $\mathcal{S}$ of spectra is an algebra in $\mathcal{S}$ over an E_∞ operad. However, provided $\mathcal{S}$ is for example symmetric or orthogonal spectra or S-modules, the category of E_∞ ring spectra in $\mathcal{S}$ is Quillen equivalent to the category of commutative monoids in $\mathcal{S}$. Therefore, we stick to these various categories of commutative monoids.

The reader may wonder, then, why the notion of E_∞ comes in at all, if the technicalities of sheaves of commutative ring spectra are best handled directly with commutative monoids in spectra. The answer is that the obstruction theory we need to actually construct such spectra uses in a fundamental way the operadic formulation of the commutativity conditions on ring spectra.

Orthogonal spectra and symmetric spectra are both examples of diagram spectra, and as such their formulations are nearly identical. We describe orthogonal spectra and then mention the modifications for symmetric spectra. An excellent reference for diagram spectra is [9] and our discussion follows the treatment there.

The basic "diagram" $\mathcal{J}$ for orthogonal spectra is the category of finite-dimensional real inner product spaces, together with orthogonal isomorphisms. From this diagram, we define the category $\mathcal{JT}$ of $\mathcal{J}$-spaces to be the category of continuous functors from $\mathcal{J}$ to based spaces, together with natural transformations between these functors.

The key observation is that $\mathcal{JT}$ is a symmetric monoidal category with product as follows:

$$\begin{aligned} &\mathcal{JT} \times \mathcal{JT} \xrightarrow{\wedge} \mathcal{JT} \\ &(X, Y) \mapsto \Big(V \xmapsto{X \wedge Y} \bigvee_{W \subset V} X(W) \wedge Y(V - W)\Big) \end{aligned}$$

Note that the wedge product $\bigvee_{W \subset V} X(W) \wedge Y(V-W)$ is topologized using the ordinary topology on subspaces of V. There is a natural commutative monoid S in $\mathcal{JT}$, namely $S(V) = S^V$; here S^V denotes the one point compactification of V. The product on S is induced by direct sum of vector spaces: $\bigvee_{W \subset V} S^W \wedge S^{V-W} \to S^V$; here $V - W$ is the orthogonal complement of W in V and the map $W \times (V - W) \to V$ is $(a, b) \mapsto i(a) + j(b)$ for $i : W \to V$ and $j : V - W \to V$ the inclusions. The reader is invited to check that this monoid really is strictly commutative, simply because direct sum of disjoint orthogonal vector subspaces is strictly commutative.

Definition 2.16. An orthogonal spectrum is a $\mathcal{J}$-space with an action by the monoid S. In other words, it is an S-module in $\mathcal{JT}$. Denote the category of orthogonal spectra by $\mathcal{JS}$.

Because S is a *commutative* monoid, the category $\mathcal{JS}$ itself has a symmetric monoidal structure with product denoted $\wedge_S$:

$$X \wedge_S Y := \mathrm{colim}_{\mathcal{JT}}(X \wedge S \wedge Y \rightrightarrows X \wedge Y)$$

This coequalizer is, of course, the usual way to define tensor products of modules in algebra. Finally we have our desired notions of ring spectra:

Definition 2.17. An orthogonal ring spectrum is a monoid in $\mathcal{JS}$. A commutative orthogonal ring spectrum is a commutative monoid in $\mathcal{JS}$.

The tensor and cotensor on the category of orthogonal spectra are particularly simple: they are both levelwise, which is to say that for X an S-module in $\mathcal{JT}$ and A a (based) space, the tensor $X \otimes A$ is given by $(X \otimes A)(V) = X(V) \wedge A$ and the cotensor X^A is given by $(X^A)(V) = X(V)^A$. Note that these tensors over topological spaces can be extended to simplicial sets via the realization functor. The cotensors on both orthogonal ring spectra and commutative orthogonal ring spectra are also levelwise. However, the tensors on these categories are rather less explicit and we do not discuss them; luckily, we only need the cotensors for our discussion of sheaves.

The definition of symmetric spectra is entirely analogous. The diagram Σ in question is the category of finite sets with isomorphisms. The category $\Sigma\mathcal{T}$ of Σ spaces is symmetric monoidal:

$$\Sigma\mathcal{T} \times \Sigma\mathcal{T} \xrightarrow{\wedge} \Sigma\mathcal{T}$$
$$(X, Y) \mapsto (X \wedge Y)(N) = \bigvee_{M \subset N} X(M) \wedge Y(N \backslash M)$$

The distinguished commutative monoid S in $\Sigma\mathcal{T}$ has $S(N) = S^N$. A symmetric spectrum is a Σ-space with an action of S. The category of symmetric spectra $\Sigma\mathcal{S}$ has a smash product, given by an appropriate coequalizer, and (commutative) symmetric ring spectra are (commutative) monoids in $\Sigma\mathcal{S}$.

Remark 2.18. Symmetric spectra are sometimes defined using the skeleton diagram Σ^{skel} whose objects are the rigid finite sets $\mathbf{n} = \{1, \ldots, n\}$, for $n \geq 0$. This variant may look more elementary but requires a rather less intuitive formula for the smash product of Σ^{skel}-spaces: $(X \wedge Y)(\mathbf{n}) = \bigvee_{m \leq n} \Sigma_{n+} \wedge_{\Sigma_m \times \Sigma_{n-m}} X(\mathbf{m}) \wedge Y(\mathbf{n} - \mathbf{m})$.

The tensor and cotensor on the category of symmetric spectra and the cotensors on symmetric ring spectra and commutative symmetric ring spectra are all levelwise, as in the case of orthogonal spectra. The tensors for (commutative) symmetric ring spectra are not levelwise, and we leave them as a mystery.

The last category of spectra we consider is the category of S-modules. S-modules are somewhat more technical than diagram spectra, and we give only the most cursory treatment, closely following EKMM [3]. Fix a universe, that is a real inner product space U isomorphic to $\mathbb{R}^\infty$. A prespectrum is an assignment to each finite dimensional subspace $V \subset U$ a based space $E(V)$ together with compatible (adjoint) structure maps $E(V) \to \Omega^{W-V} E(W)$. Denote the category of prespectra by $\mathcal{P}U$ or simply $\mathcal{P}$. A (Lewis-May-Steinberger) spectrum is a prespectrum in which all the structure maps are homeomorphisms, and the category of such is denoted $\mathcal{S}U$ or $\mathcal{S}$. The forgetful functor $\mathcal{S} \to \mathcal{P}$ has a left adjoint $L : \mathcal{P} \to \mathcal{S}$ called spectrification.

There is an external smash product of spectra $\mathcal{S}U \times \mathcal{S}U' \xrightarrow{\wedge} \mathcal{S}(U \oplus U')$. Given a pair of spectra (E, E'), the assignment $F : V \oplus V' \mapsto E(V) \wedge E'(V')$ defines a prespectrum on the decomposable subspaces of $U \oplus U'$. There is a spectrification functor here as well that produces from F a spectrum LF on the decomposable subspaces of $U \oplus U'$; there is moreover a left adjoint ψ to the restriction to such subspaces, which in turn produces our desired smash $E \wedge E' := \psi LF \in \mathcal{S}(U \oplus U')$ indexed on all finite dimensional subspaces of $U \oplus U'$.

We would like to internalize this smash product, using the space of linear isometries from $U \oplus U$ to U. If we have a linear isometry $f : U \to U'$ we can transport a spectrum $E \in \mathcal{S}U$ to a spectrum $f_* E \in \mathcal{S}U'$: define $f_* E$ to be the spectrification of the prespectrum taking $V' \subset U'$ to $E(V) \wedge S^{V' - f(V)}$, where $V = f^{-1}(V' \cap \operatorname{im} f)$. Given an A-parameter family of linear isometries, that is a map $\alpha : A \to \mathcal{I}(U, U')$, there is a spectrum $A \ltimes E \in \mathcal{S}U'$ called the twisted half smash product, which is in an appropriate sense a union of all the spectra $\alpha(a)_* E$ for $a \in A$.

Let $\mathcal{L}(j) = \mathcal{I}(U^j, U)$ be the space of all internalizing linear isometries. Given $E, F \in \mathcal{S}U$, the twisted half smash $\mathcal{L}(2) \ltimes (E \wedge F)$ is a canonical internal smash product, but it is not associative. We fix this by restricting to $\mathbb{L}$-spectra: an $\mathbb{L}$-spectrum is a spectrum E with an action $\mathcal{L}(1) \ltimes E \to E$ by the isometries $\mathcal{L}(1)$. The smash product of $\mathbb{L}$-spectra is $M \wedge_{\mathcal{L}} N := \mathcal{L}(2) \ltimes_{\mathcal{L}(1) \times \mathcal{L}(1)} (M \wedge N)$. (Here the twisted half smash over $\mathcal{L}(1) \times \mathcal{L}(1)$ is given by the expected coequalizer.) We would be done, except that the category of $\mathbb{L}$-spectra doesn't have a point-set-level unit.

We can conjure up a unit as follows. There is a natural map $\lambda : S \wedge_{\mathcal{L}} M \to M$, where S is the spectrification of the prespectrum $V \mapsto S^V$. An S-module is by definition an $\mathbb{L}$-spectrum such that λ is an isomorphism. The smash product of two S-modules is simply their smash product as $\mathbb{L}$-spectra. The category of S-modules is our desired symmetric monoidal, unital category of spectra.

S-algebras and commutative S-algebras are simply monoids and commutative monoids respectively in S-modules.

Given an S-module M and a based space X, the tensor S-module $M \otimes X$ is defined to be the spectrification of the prespectrum $V \mapsto M(V) \wedge X$. The cotensor M^X is defined to be $S \wedge_{\mathcal{L}} \phi(M)^X$, where ϕ forgets from S-modules to spectra, and the cotensor on spectra is $E^X(V) = E(V)^X$.

The cotensors on S-algebras and on commutative S-algebras are simply given by taking the cotensor in S-modules. The tensors on S-algebras and commutative S-algebras by contrast are not created in S-modules. However, the tensor on commutative S-algebras has a convenient description, as follows. Given a finite set $[n]$ and a commutative S-algebra R, define the tensor $R \otimes [n] := R^{\wedge [n]}$. Now for X a finite simplicial set, the tensor $R \otimes X$ in commutative S-algebras is the realization as a simplicial S-module of the levelwise tensor $R \otimes X_*$.

Picking one of the above three models for spectra, and feeding the cotensors back into the construction of homotopy limits, we now have a precise definition of sheaf of spectra, sheaf of ring spectra, and sheaf of commutative ring spectra.

Remark 2.19. The reader may worry that there could be a confusing difference between sheaves of ring spectra and presheaves of ring spectra that are sheaves of spectra. Luckily, this is not the case: the two notions agree because the sheaf condition is a homotopy limit condition, and this homotopy limit is built using limits and cotensors; these limits and cotensors are, in any of the above categories of ring and commutative ring spectra, simply computed in the underlying category of spectra.

We invite the reader to ruminate on the fact that cosheaves of ring spectra are very different objects from precosheaves of ring spectra that are cosheaves of spectra.

3. The Descent Spectral Sequence

Recall that the main theorem (of Hopkins and Miller) is that there exists a sheaf $\mathcal{O}^{\mathrm{top}}$ of spectra on the moduli stack $(\mathcal{M}_{ell})_{\text{ét}}$ of elliptic curves in the étale topology. Sections 1 and 2 described what it means to have such a sheaf. In particular, section 2.2 described the Grothendieck site of the moduli stack in the étale topology, while section 1.1 defined sheaves on such a site with values in a category (as presheaves satisfying a homotopy limit condition), and section 2.3 discussed the homotopy limits of spectra needed for this definition of sheaves.

Given the sheaf $\mathcal{O}^{\mathrm{top}}$, we are primarily interested in understanding its spectrum of global sections $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$—recall that this spectrum is called *TMF*. By the definition of a sheaf, information about the global sections $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ is contained in the spectra $\mathcal{O}^{\mathrm{top}}(U)$ associated to small open subsets U of the moduli stack $\mathcal{M}_{ell}$; the goal of this section is describe precisely how this local information is assembled into the desired global information. In particular, we construct a spectral sequence beginning with the sheaf cohomology of the moduli stack with coefficients in the sheafification of the presheaf of homotopy groups of $\mathcal{O}^{\mathrm{top}}$, strongly converging to the homotopy groups of the global spectrum $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$:

$$E^2 = H^q(\mathcal{M}_{ell}, \pi_p^\dagger \mathcal{O}^{\mathrm{top}}) \Longrightarrow \pi_{p-q}\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$$

In the chapter on the homotopy groups of *TMF*, we will compute the E^2 term and describe the elaborate pattern of differentials.

We begin in section 3.1 by reviewing the notions of sheaf cohomology and Cech cohomology and discussing how they are related. Then in section 3.2 we construct a spectral sequence beginning with Cech cohomology and converging to the homotopy of the global sections of a sheaf of spectra. Finally, in section 3.3, we specialize to the sheaf $\mathcal{O}^{\mathrm{top}}$, using properties of this particular sheaf to simplify the spectral sequence from the preceding section.

3.1. **Sheaf Cohomology and Cech Cohomology.** We are studying a sheaf $\mathcal{O}^{\mathrm{top}}$ of spectra on the moduli stack of elliptic curves. We can consider the homotopy groups $\pi_*(\mathcal{O}^{\mathrm{top}}(U))$ of the spectra $\mathcal{O}^{\mathrm{top}}(U)$ associated to particular objects U of the étale site of the moduli stack. These homotopy groups fit together into a good-old down-to-earth presheaf of graded abelian groups. The spectral sequence computing the homotopy groups of *TMF* will begin with the sheaf cohomology of the sheafification of this presheaf. We review sheaf cohomology, and the related Cech cohomology, in some generality.

Let $\mathcal{A}$ be an abelian category. A morphism $f : X \to Y$ in $\mathcal{A}$ is a monomorphism if $fg_1 = fg_2$ implies $g_1 = g_2$. Recall that an object $I \in \mathcal{A}$ is injective if for all maps $m : X \to I$ and all monomorphisms $X \hookrightarrow Y$, there exists an extension of m to Y. The category $\mathcal{A}$ is said to have enough injectives if for all objects $A \in \mathcal{A}$ there is a monomorphism $A \hookrightarrow I$ into an injective object I. We are interested, of course, in the category of sheaves on a site:

Note 3.1. For any site $\mathcal{C}$, the category $\mathsf{Shv}_{\mathsf{Ab}}(\mathcal{C})$ of sheaves of abelian groups on the site is an abelian category with enough injectives.

We can therefore use the usual definition of sheaf cohomology:

Definition 3.2. For $\pi \in \mathsf{Shv}_{\mathsf{Ab}}(\mathcal{C})$ a sheaf of abelian groups on a site $\mathcal{C}$, the sheaf cohomology of an object $X \in \mathcal{C}$ of the site with coefficients in π is

$$H^q(X, \pi) := H_q\big(0 \to I^0(X) \to I^1(X) \to I^2(X) \to \cdots\big)$$

where $0 \to \pi|_X \to I^0 \to I^1 \to I^2 \to \cdots$ is an injective resolution of $\pi|_X$ in $\mathsf{Shv}_{\mathsf{Ab}}(\mathcal{C}/X)$.

Sheaf cohomology has a more concrete cousin, Cech cohomology, which does not involve an abstract resolution:

Definition 3.3. Let $\pi \in \mathsf{Shv}_{\mathsf{Ab}}(\mathcal{C})$ be a sheaf of abelian groups, and let $\mathfrak{U} = \{U_i \to U\}$ be a cover in the site $\mathcal{C}$. The Cech cohomology of U with respect to the cover $\mathfrak{U}$ and with coefficients in π is

$$\check{H}^q_{\mathfrak{U}}(U, \pi) := H_q\Big(0 \to \prod \pi(U_i) \to \prod \pi(U_{ij}) \to \prod \pi(U_{ijk}) \to \cdots\Big)$$

Here U_I refers to the intersection (ie fibre product over U) of the $U_i, i \in I$, and the maps are the alternating sums of the various natural restriction maps.

Cech cohomology is computable by hand, while sheaf cohomology is evidently independent of a particular choice of cover. If the object U of $\mathcal{C}$ has an acyclic cover, then the two theories agree:

Proposition 3.4. *For $\pi \in \mathsf{Shv}_{\mathsf{Ab}}(\mathcal{C})$ a sheaf of abelian groups, and $\mathfrak{U} = \{U_i \to U\}_{i \in I}$ a cover in $\mathcal{C}$ such that $H^q(U_J, \pi) = 0$ for all $J \subset I$ and all $q \geq 1$, sheaf and Cech cohomology agree: $H^q(U, \pi) = \check{H}^q_{\mathfrak{U}}(U, \pi)$.*

The proof is the usual double complex argument: build the double complex $[I^q(\coprod_{|J|=p} U_J)]$ from an injective resolution I^* of π, and show that the two resulting spectral sequences collapse respectively to sheaf and to Cech cohomology.

Remark 3.5. We will be interested in cases where the site $\mathcal{C}$ has a global terminal object, usually denoted (confusingly) $\mathcal{C}$, and from now on we assume we are in that situation.

3.2. **The Spectral Sequence for a Sheaf of Spectra.** We begin with a sheaf of spectra $\mathcal{O}$ on the étale site $(\mathcal{M}_{ell})_{\text{ét}}$ of the moduli stack of elliptic curves, and we would like to construct a spectral sequence converging to the homotopy groups $\pi_*(\mathcal{O}(\mathcal{M}_{ell}))$. The spectral sequence is meant to start with the local data of $\mathcal{O}$; we therefore chose a cover $\{\mathcal{N}_i \to \mathcal{M}_{ell}\}$ of the moduli stack. In outline, we use this cover to build a simplicial object of the site, then apply the sheaf $\mathcal{O}$ to get a cosimplicial spectrum, from which we get a tower of spectra, which we wrap up into an exact couple, and thereby arrive at our desired spectral sequence:

$$\begin{tikzcd}
\text{cover } \mathfrak{N} = \{\mathcal{N}_i \to \mathcal{M}_{ell}\} \text{ in } (\mathcal{M}_{ell})_{\text{ét}} \arrow[d, rightsquigarrow, "e"] \\
\text{simplicial object } \coprod \mathcal{N}_i \leftleftarrows \coprod \mathcal{N}_{ij} \cdots \text{ in } (\mathcal{M}_{ell})_{\text{ét}} \arrow[d, rightsquigarrow, "d"] \\
\text{cosimplicial spectrum } \mathcal{O}^* = \{\mathcal{O}(\coprod \mathcal{N}_i) \rightrightarrows \mathcal{O}(\coprod \mathcal{N}_{ij}) \cdots\} \arrow[d, rightsquigarrow, "c"] \\
\text{tower of spectra } \cdots \to (\operatorname{Tot} \mathcal{O})^2 \xrightarrow{\phi_2} (\operatorname{Tot} \mathcal{O})^1 \xrightarrow{\phi_1} (\operatorname{Tot} \mathcal{O})^0 \arrow[d, rightsquigarrow, "b"] \\
\text{exact couple } \bigoplus \pi_{p+q}(\operatorname{Tot} \mathcal{O})^q \to \bigoplus \pi_{p+q}(\operatorname{Tot} \mathcal{O})^q \to \bigoplus \pi_{p+q}(\operatorname{corhofib} \phi_q) \arrow[d, rightsquigarrow, "a"] \\
\text{spectral sequence } E^2_{pq} = \check{H}^q_{\mathfrak{N}}(\mathcal{M}_{ell}, \pi_p \mathcal{O}) \Longrightarrow \pi_{p-q} \mathcal{O}(\mathcal{M}_{ell})
\end{tikzcd}$$

In this section we describe this chain of constructions in detail. In section 3.3 we use particular properties of the sheaf $\mathcal{O}^{\text{top}}$ to compare the E^2 term of this spectral sequence for $\pi_*(\mathcal{O}^{\text{top}}(\mathcal{M}_{ell}))$, which is the Cech cohomology of the presheaf $\pi_p \mathcal{O}^{\text{top}}$, to the sheaf cohomology of the sheafification $\pi_p^\dagger \mathcal{O}^{\text{top}}$ of the presheaf $\pi_p \mathcal{O}^{\text{top}}$.

We begin at the end (step a) and work our way back to the beginning (step e). We assume the reader is familiar with the construction of a spectral sequence from an exact couple—see McCleary [10] for a detailed presentation of the construction and Boardman [2] for a careful treatment of convergence issues. We therefore proceed to step b, building an exact couple from a tower of spectra.

3.2.1. *The Spectral Sequence of a Filtration of Spectra and of a Tower of Spectra.* We are interested in towers of spectra and their associated inverse limits. Along the way we address the slightly more intuitive situation of a filtration of spectra and its associated direct limit.

Suppose we have a filtration of spectra:

$$\begin{tikzcd}
* \arrow[r, "\phi_0"] & F_0 \arrow[r, "\phi_1"] \arrow[dl] & F_1 \arrow[r, "\phi_2"] \arrow[dl] & F_2 \arrow[r] \arrow[dl] & \cdots \\
\operatorname{corhocofib} \phi_0 & \operatorname{corhocofib} \phi_1 & \operatorname{corhocofib} \phi_2 & &
\end{tikzcd}$$

Denote the corrected homotopy colimit $\operatorname{corhocolim}_i F_i$ by F. We think of the sequence F_i as a filtration of F, and expect any spectral sequence constructed from the filtration to give information about F. Take homotopy groups of the spectra F_i and of the corrected homotopy cofibres $\operatorname{corhocofib} \phi_i$, and wrap up the resulting triangles into an exact couple. This produces a spectral sequence with $E^1_{pq} = \pi_{p+q}(\operatorname{corhocofib} \phi_q)$. This spectral sequence is a half plane spectral sequence with exiting differentials (in the sense of Boardman [2]), and converges strongly to $\pi_{p+q} F$:

$$E^1_{pq} = \pi_{p+q}(\operatorname{corhocofib} \phi_q) \underset{\text{strong}}{\Longrightarrow} \pi_{p+q} \operatorname{corhocolim}_i F_i.$$

Note 3.6. Lest there be any confusion, note that for X a spectrum, by $\pi_i X$ we mean the set of maps $\operatorname{Ho}_{\mathsf{Spec}}([S^i], [X])$ in the homotopy category of spectra between the sphere $[S^i]$ and $[X]$, not for instance homotopy classes of maps in the category of spectra from a sphere S^i to X. The above spectral sequence converges, a priori, to $\operatorname{colim}_i \pi_{p+q} F_i$; we have used implicitly the equality

$$\operatorname{colim}_i \pi_{p+q} F_i = \pi_{p+q} \operatorname{corhocolim}_i F_i.$$

In the dual picture, we begin with a tower of spectra:

$$
\begin{tikzcd}
\cdots \arrow[r] & F^2 \arrow[rr, "\phi_2"] & & F^1 \arrow[rr, "\phi_1"] & & F^0 \arrow[rr, "\phi_0"] & & * \\
& & \operatorname{corhofib}\phi_2 \arrow[ul] & & \operatorname{corhofib}\phi_1 \arrow[ul] & & \operatorname{corhofib}\phi_0 \arrow[ul] &
\end{tikzcd}
$$

Let F denote the corrected homotopy limit $\operatorname{corholim}_i F^i$. Again we take homotopy groups of the whole diagram of the F^i and $\operatorname{corhofib}\phi_i$, and wrap up the resulting triangles into an exact couple. The spectral sequence associated to this exact couple has $E^1_{pq} = \pi_{p-q}(\operatorname{corhofib}\phi_q)$. It is a half plane spectral sequence with entering differentials and converges conditionally to $\lim_i \pi_{p-q}F^i$:

$$E^1_{pq} = \pi_{p-q}(\operatorname{corhofib}\phi_q) \underset{\text{cond}}{\Longrightarrow} \lim_i \pi_{p-q}F^i.$$

In general, one must address two issues: whether or not this conditionally convergent spectral sequence in fact converges strongly, and whether or not $\lim_i \pi_{p-q}F^i = \pi_{p-q}F$, the latter of which is of course a $\lim^1$ problem.

Note 3.7. The spectral sequence of a tower has target $\lim_i \pi_{p-q}F^i$, but we are usually more interested in the homotopy $\pi_{p-q}\operatorname{corholim}_i F^i$ of the corrected homotopy limit of the tower. These are related by the Milnor exact sequence:

$$0 \to \lim^1_i \pi_{p-q-1}F^i \to \pi_{p-q}\operatorname{corholim}_i F^i \to \lim_i \pi_{p-q}F^i \to 0.$$

3.2.2. *The Realization of a Simplicial Spectrum and Tot of a Cosimplicial Spectrum.* Here we note that a simplicial spectrum leads to a filtration of spectra, and dually that a cosimplicial spectrum leads to a tower of spectra. We also identify the colimit (resp. limit) of the resulting filtration (tower) in terms of the homotopy colimit (limit) of the original (co)simplicial spectrum.

Notation 3.8. We begin by fixing some notation. The cosimplicial category Δ has objects $[0], [1], [2], \ldots$, (where we think of $[n]$ as $\{0, 1, \ldots, n\}$) and morphisms weakly order preserving maps. The bold $\boldsymbol{\Delta}$ will denote the standard cosimplicial simplicial set, whose simplicial set of n-cosimplicies $\boldsymbol{\Delta}(n)$ is the simplicial n-simplex $\Delta[n]$; this n-simplex $\Delta[n]$ has k-simplicies $\Delta[n]_k = \Delta([k], [n])$. We use the symbol Δ^n_0 to denote the full subcategory of Δ whose objects are $[0], \ldots, [n]$.

Let A be a simplicial spectrum, that is a simplicial object $A : \Delta^{\mathrm{op}} \to \mathsf{Spec}$ in the category of spectra. The realization of A is defined as follows:

$$|A| := A \otimes_{\Delta^{\mathrm{op}}} \boldsymbol{\Delta} = \operatorname{colim}\left(\coprod_{\phi: n \to m} A_m \otimes \Delta[n] \rightrightarrows \coprod_{n \geq 0} A_n \otimes \Delta[n] \right).$$

We construct a filtration whose colimit is the realization $|A|$ as follows. Recall that the n-skeleton $\operatorname{sk}_n A$ of A is the left Kan extension to Δ^{op} of the restriction of A from Δ^{op} to $(\Delta^n_0)^{\mathrm{op}}$. Intuitively, $\operatorname{sk}_n A$ consists of the simplicies of A of dimension less than or equal to n, together with the possible degeneracies of those simplicies; (here "possible" means possible in a simplicial spectrum agreeing with A through dimension n). Let $|A|_n$ denote the realization $|\operatorname{sk}_n A|$. We have the sequence

$$* \to |A|_0 \to |A|_1 \to |A|_2 \to \cdots$$

The colimit of this sequence is, by construction, the realization $|A|$.

Note 3.9. We can build the skeleta of A inductively as follows. Define the n-th latching object of A as

$$L_n A = (\operatorname{sk}_{n-1} A)_n.$$

This latching object can, roughly speaking, be thought of as the spectrum of degenerate n-simplicies of A—in general, though, the map $L_n A \to A_n$ need not be a cofibration, which means in particular that the latching object $L_n A$ may record more degenerate simplicies than are present in A itself. We have a pushout diagram [4, p.367]:

$$\begin{tikzcd} A_n \otimes \partial\Delta^n \cup_{L_n A \otimes \partial\Delta^n} L_n A \otimes \Delta^n \arrow[r] \arrow[d] & \mathrm{sk}_{n-1} A \arrow[d] \\ A_n \otimes \Delta^n \arrow[r] & \mathrm{sk}_n A \end{tikzcd}$$

Here only, by "$\otimes$" we mean the external tensor, taking a spectrum and a simplicial set and giving a simplicial spectrum.

Ignoring the terms involving $L_n A$, the pushout says that the n-skeleton $\mathrm{sk}_n A$ is built by gluing an n-simplex along its boundary onto the $(n-1)$-skeleton $\mathrm{sk}_{n-1} A$, for each "element" of the spectrum A_n of n-simplicies. But in fact, the $(n-1)$-skeleton already includes degenerate n-simplicies, and the pushout accounts for this by quotienting out the latching object $L_n A \otimes \Delta^n$.

Example 3.10. Consider the first stage $|A|_1 = |\,\mathrm{sk}_1 A|$ of the filtration of $|A|$. Ignoring degeneracies, the realization of the 1-skeleton is roughly speaking the colimit

$$\mathrm{colim}\left(A_1 \otimes \Delta^0 \sqcup A_1 \otimes \Delta^0 \rightrightarrows A_0 \otimes \Delta^0 \sqcup A_1 \otimes \Delta^1\right).$$

We can schematically picture this 1-skeleton glued together as follows:

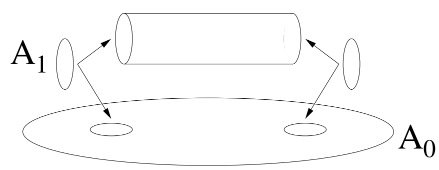

Notice that the filtration $* \to |A|_0 \to |A|_1 \to \cdots$ above has *colimit* the realization $|A|$. However, recall that the spectral sequence from section 3.2.1 converges to the homotopy groups of the *corrected homotopy colimit* of the filtration. We therefore need to address the issue of when these two agree. Conveniently, the condition ensuring this agreement can be expressed in terms of Reedy cofibrancy, as follows.

Definition 3.11. A simplicial spectrum A is Reedy cofibrant if the maps $L_n A \to A_n$ are cofibrations for all n. Roughly speaking, this is true when the degenerate simplicies of A are freely generated, and for all n the degenerate n-simplicies of A map by a cofibration into all the n-simplicies of A. A sequence of spectra $* \to X_0 \to X_1 \to \cdots$ is Reedy cofibrant if all the maps in the sequence are cofibrations.

If a simplicial spectrum A is Reedy cofibrant, then the morphisms $|A|_{n-1} \to |A|_n$ are cofibrations (see [4, p.385]), which is to say the realization sequence $* \to |A|_0 \to |A|_1 \to \cdots$ is Reedy cofibrant. In section 2.3.2, we saw that the corrected homotopy colimit represents the derived colimit and observed that in the case of a Reedy cofibrant diagram, this derived colimit is represented by the honest colimit. Altogether, when the simplicial spectrum A is Reedy cofibrant, we have a weak equivalence

$$\mathrm{corhocolim}_i\, |A|_i \underset{\mathrm{RC}}{\simeq} \mathrm{colim}_i\, |A|_i = |A|.$$

The subscript "RC" will serve as a reminder that an equivalence depends on Reedy cofibrancy.

Remark 3.12. When the sequence $* \to |A|_0 \to |A|_1 \to \cdots$ is Reedy cofibrant, all the terms in the sequence are necessarily cofibrant (though this is not, per se, part of the Reedy condition). This objectwise cofibrancy implies that the corrected homotopy colimit corhocolim $|A|_i$ agrees with the usual homotopy colimit hocolim $|A|_i$. We will not, however, need this fact.

Confusingly, there is another homotopy colimit floating around, namely the homotopy colimit of the Δ^{op}-diagram A itself. We also need to compare the realization $|A|$ to this homotopy colimit $\operatorname{hocolim}_{\Delta^{\mathrm{op}}} A$: when we are studying a cosheaf of spectra, we will care about the homotopy colimit of the associated simplicial spectrum, but the spectral sequence constructed in section 3.2.1 converges, if A is Reedy cofibrant, to the homotopy groups of the realization of this simplicial spectrum.

Recall that the homotopy colimit of the Δ^{op}-diagram A was defined as $A \otimes_{\Delta^{\mathrm{op}}} N(-/\Delta^{\mathrm{op}})^{\mathrm{op}}$. We might therefore expect a close relationship between the realization and the homotopy colimit:

Proposition 3.13 ([5, Thm 18.7.4])**.** *If the simplicial spectrum A is Reedy cofibrant, then there is a natural weak equivalence (the Bousfield-Kan map)*

$$|A| \xleftarrow{\simeq} \operatorname{hocolim}_{\Delta^{\mathrm{op}}} A.$$

This equivalence is most plausible: the reader can note that, ignoring degeneracy maps, the nerves $N(-/\Delta^{\mathrm{op}})^{\mathrm{op}}$ of the overcategories in Δ^{op} are the barycentric subdivisions of the standard simplicies of $\mathbf{\Delta}$. The Reedy cofibrancy condition will be automatically satisfied (see Lemma 3.20) in the situation we care about, and so need not concern us.

In summary, we have the following chain of equalities, relating the target of our spectral sequence to the homotopy of the homotopy colimit of our simplicial spectrum:

$$\operatorname{colim}_i \pi_* |A|_i = \pi_* \operatorname{corhocolim}_i |A|_i \underset{\mathrm{RC}}{=} \pi_* \operatorname{colim}_i |A|_i = \pi_* |A| \underset{\mathrm{RC}}{=} \pi_* \operatorname{hocolim}_{\Delta^{\mathrm{op}}} A.$$

The indicated equalities depend on A being Reedy cofibrant.

The dual, cosimplicial picture is the one that concerns us more directly. Let $B : \Delta \to \mathit{Spec}$ be a cosimplicial spectrum. The "corealization" is traditionally called the "totalization" of the cosimplicial spectrum and is defined as follows:

$$\operatorname{Tot} B := \hom^{\Delta}(\mathbf{\Delta}, B) = \lim \left(\prod_{n \geq 0} (B^n)^{\mathbf{\Delta}(n)} \rightrightarrows \prod_{\phi : n \to m} (B^m)^{\mathbf{\Delta}(n)} \right).$$

Here again $\mathbf{\Delta}$ is the cosimplicial standard simplex with $\mathbf{\Delta}(n)$ the standard simplicial n-simplex $\Delta[n]$.

The key feature of Tot of a cosimplicial spectrum is that it is the inverse limit of a tower of spectra built from the coskeleta of the cosimplicial spectrum, and this tower leads to our desired spectral sequence. We define the n-coskeleton $\operatorname{cosk}^n B$ of a cosimplicial spectrum $B : \Delta \to \mathsf{Spec}$ to be the right Kan extension to Δ of the restriction of B from Δ to Δ_0^n. Intuitively, $\operatorname{cosk}^n B$ consists of the cosimplicies of B of dimension between 0 and n; for $k > n$ the coskeleton has a k-cosimplex for every possible combination of n-cosimplicies of B that could be the image under the codegeneracy maps of a k-cosimplex of a cosimplicial spectrum agreeing with B through dimension n. Let $\operatorname{Tot}^n B$ denote the totalization $\operatorname{Tot} \operatorname{cosk}_n B$ of the coskeleton of B. We have a tower

$$\cdots \to \operatorname{Tot}^2 B \to \operatorname{Tot}^1 B \to \operatorname{Tot}^0 B \to *$$

The limit of this sequence is the full totalization $\operatorname{Tot} B$.

Remark 3.14. The above definition of Tot^n is not formally the same as the usual one in the literature, for instance as the definition in Bousfield-Kan or Goerss-Jardine, and we would like to spell out and emphasize the difference.

For a simplicial spectrum A, recall that the n-skeleton $\operatorname{sk}_n A$ and n-coskeleton $\operatorname{cosk}_n A$ of A are defined respectively as the left and right Kan extensions of $A|_{(\Delta_0^n)^{\mathrm{op}}}$ to Δ^{op}. These are both simplicial spectra in their own right. Now, somewhat unconventionally, for a cosimplicial spectrum B, we define the n-skeleton $\operatorname{sk}^n B$ and n-coskeleton $\operatorname{cosk}^n B$ to be respectively the left and right Kan extensions of $B|_{\Delta_0^n}$ to Δ. These are, of course, both cosimplicial spectra.

It is standard to define

$$|A|_n := |\operatorname{sk}_n A| = \operatorname{sk}_n A \otimes_{\Delta^{\mathrm{op}}} \mathbf{\Delta}$$
$$\operatorname{Tot}^n B := \hom^{\Delta}(\operatorname{sk}_n \mathbf{\Delta}, B)$$

This is unsettling for two reasons. First, here $\operatorname{sk}_n \mathbf{\Delta}$ does not refer to a left Kan extension of the cosimplicial object $\mathbf{\Delta}$, but to a levelwise left Kan extension of the simplicial levels of $\mathbf{\Delta}$. Second, it does not express the layers Tot^n as totalizations in their own right, and moreover entirely obscures the precise duality between the skeletal filtration and the totalization tower.

Instead we prefer

$$|A|_n := |\operatorname{sk}_n A| = \operatorname{sk}_n A \otimes_{\Delta^{\mathrm{op}}} \mathbf{\Delta}$$
$$\operatorname{Tot}^n B := \operatorname{Tot} \operatorname{cosk}^n B = \hom^{\Delta}(\mathbf{\Delta}, \operatorname{cosk}^n B)$$

We leave it to the reader to verify that this results in the same spectrum $\operatorname{Tot}^n B$ as the usual formulation.

Note 3.15. As we could for the skeleta of a simplicial spectrum, we can build the coskeleta of our cosimplicial spectrum B inductively. In a classic example of mathematical nomenclature, the duals of latching objects are called matching objects:

$$M^n B = (\operatorname{cosk}^{n-1} B)^n.$$

Readers should be warned that this indexing is not the same as that in Goerss-Jardine or Bousfield-Kan; instead we specialize the abstractly consistent scheme of Hirschhorn. The inductive pullback diagram is

$$\begin{tikzcd}
\operatorname{cosk}^n B \arrow[rr] \arrow[d] & & (B^n)^{\Delta^n} \arrow[d] \\
\operatorname{cosk}^{n-1} B \arrow[r] & (B^n)^{\partial\Delta^n} \times_{(M^n B)^{\partial\Delta^n}} (M^n B)^{\Delta^n} &
\end{tikzcd}$$

The totalization of this diagram gives a corresponding pullback for $\operatorname{Tot}^n B$ in terms of $\operatorname{Tot}^{n-1} B$.

Ignoring the terms involving matching objects, this pullback would indicate that $\operatorname{Tot}^n B$ can be seen as the pairs of maps $\phi : \Delta^n \to B^n$ and "points" $\psi \in \operatorname{Tot}^{n-1} B$ that agree as maps $\theta : \partial\Delta^n \to B^n$; here, ψ determines θ by the coface maps of B. This idea is illustrated in the following example. More precisely, though, the matching terms in the pullback account for the fact that the $(n-1)$-coskeleton already contains a collection of potential n-cosimplicies.

Example 3.16. We consider the first stage of the Tot tower. Ignoring codegeneracy issues, the totalization of the 1-coskeleton is, roughly, the limit

$$\lim\left((B^0)^{\Delta^0} \times (B^1)^{\Delta^1} \rightrightarrows (B^1)^{\Delta^0} \times (B^1)^{\Delta^0}\right),$$

which is to say a 0-cosimplex, together with a path of 1-cosimplicies agreeing at the ends with the cofaces of the 0-cosimplex:

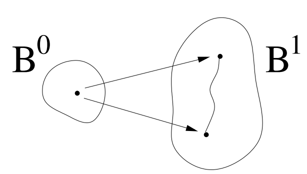

By construction the limit of the Tot tower is $\operatorname{Tot} B$. There is the pesky issue of whether this limit is the same as the corrected homotopy limit of the tower—recall that, in the absence of $\lim^1$ problems, the spectral sequence of the tower has target the homotopy groups of the corrected homotopy limit. The condition on the cosimplicial spectrum B that tethers the limit and the corrected homotopy limit is, as expected, Reedy fibrancy.

Definition 3.17. A cosimplicial spectrum B is Reedy fibrant if the maps $B^n \to M^n B$ are fibrations. A tower of spectra $\cdots \to Y^1 \to Y^0 \to *$ is Reedy fibrant if the maps in the tower are all fibrations.

When a cosimplicial spectrum B is Reedy fibrant, the maps $\mathrm{Tot}^n B \to \mathrm{Tot}^{n-1} B$ are fibrations, so the Tot tower is Reedy fibrant. We saw in section 2.3.2 that the corrected homotopy limit represents the derived limit; in the case of a Reedy fibrant diagram, the honest limit also represents this derived limit, and the limit and corrected homotopy limit agree. When the cosimplicial spectrum B is Reedy fibrant, we therefore have a weak equivalence

$$\mathrm{corholim}_i \mathrm{Tot}^i B \underset{\mathrm{RF}}{\simeq} \lim_i \mathrm{Tot}^i B = \mathrm{Tot}\, B.$$

Remark 3.18. When the tower $\cdots \to \mathrm{Tot}^1 B \to \mathrm{Tot}^0 B \to *$ is Reedy fibrant, all the terms in the tower are fibrant, even though this is not explicitly part of the Reedy condition. This objectwise fibrancy implies that the corrected homotopy limit $\mathrm{corholim}\, \mathrm{Tot}^i B$ is equal to the homotopy limit $\mathrm{holim}\, \mathrm{Tot}^i B$, though we do not need to consider the latter homotopy limit.

We have a tower of spectra whose limit is the totalization of our cosimplicial spectrum, and a spectral sequence with target the homotopy of this totalization (provide we have Reedy fibrancy, and no $\lim^1$ problem). However, in the end we will be interested in the homotopy limit of the cosimplicial spectrum (not its totalization), because that homotopy limit will carry the global homotopical information in a sheaf of spectra. Therefore, we need to compare the totalization of the tower to the homotopy limit of the cosimplicial diagram itself:

Proposition 3.19 ([5, Thm 18.7.4])**.** *If the cosimplicial spectrum B is Reedy fibrant, then there is a natural weak equivalence (again called the Bousfield-Kan map)*

$$\mathrm{Tot}\, B \xrightarrow{\simeq} \mathrm{holim}_\Delta B.$$

The totalization is $\hom^\Delta(\boldsymbol{\Delta}, B)$, while the homotopy limit is $\hom^\Delta(N(\Delta/-), B)$. We already remarked that, glossing over degeneracy issues, the nerves of the undercategories $N(\Delta/-)$ are the barycentric subdivisions of the standard simplicies of $\boldsymbol{\Delta}$, and so this weak equivalence is unsurprising. The cosimplicial spectra coming from our sheaves of spectra will always be Reedy fibrant—see Lemma 3.20.

In summary, we have the following chain relating the target of the spectral sequence of the Tot tower to the homotopy of the homotopy limit of our cosimplicial spectrum:

$$\lim_i \pi_* \mathrm{Tot}^i B \underset{\lim^1}{\rightsquigarrow} \pi_* \mathrm{corholim}_i \mathrm{Tot}^i B \underset{\mathrm{RF}}{=} \pi_* \lim_i \mathrm{Tot}^i B = \pi_* \mathrm{Tot}\, B \underset{\mathrm{RF}}{=} \pi_* \mathrm{holim}_\Delta B.$$

Here the first arrow refers to the Milnor short exact sequence, and the indicated equalities depend on B being Reedy fibrant.

3.2.3. *Cosheaves and Simplicial Spectra, and Sheaves and Cosimplicial Spectra.* In the last section, we constructed a filtration of spectra out of a simplicial spectrum, and a tower of spectra out of a cosimplicial spectrum. In either case we have an associated exact couple and therefore spectral sequence. In this section, given a cosheaf (resp. sheaf) of spectra, we build a simplicial (resp. cosimplicial) spectrum, and we describe in detail the E^1 and E^2 terms of the resulting spectral sequence.

Let $\mathcal{C}$ be a site, for instance the étale site of the moduli stack of elliptic curves, and let $\mathfrak{U} = \{U_i \to U\}_{i \in I}$ be a cover in $\mathcal{C}$. Assuming we have coproducts in $\mathcal{C}$, this cover yields a simplicial object in $\mathcal{C}$:

$$\mathbf{U}. := \left(\coprod_i U_i \leftleftarrows \coprod_{i,j} U_{ij} \Lleftarrow \coprod_{i,j,k} U_{ijk} \cdots \right).$$

A precosheaf of spectra on $\mathcal{C}$ is a covariant functor $\mathcal{G} : \mathcal{C} \to \mathsf{Spec}$; similarly of course we may have precosheaves with values in other categories. If we apply such a precosheaf $\mathcal{G}$ to the simplicial object $\mathbf{U}.$, we get a simplicial spectrum:

$$\mathcal{G}(\mathbf{U}.) = \left(\mathcal{G}(\coprod U_i) \leftleftarrows \mathcal{G}(\coprod U_{ij}) \Lleftarrow \cdots \right).$$

A cosheaf of spectra is a precosheaf $\mathcal{G}$ such that for all objects U of the site $\mathcal{C}$ and for all covers $\mathfrak{U}$ of U, the map

$$\mathcal{G}(U) \xleftarrow{\simeq} \operatorname{hocolim}_{\Delta^{\mathrm{op}}} \mathcal{G}(\mathbf{U}.)$$

is a weak equivalence.

The spectral sequence associated to the the simplicial spectrum $\mathcal{G}(\mathbf{U}.)$ has the form

$$E^1_{pq} = \pi_{p+q}(\operatorname{corhocofib}(|\mathcal{G}(\mathbf{U}.)|_{q-1} \to |\mathcal{G}(\mathbf{U}.)|_q)) \Longrightarrow \pi_{p+q} \operatorname{corhocolim}_i |\mathcal{G}(\mathbf{U}.)|_i.$$

In order to better identify both the E^1 term and the target of this spectral sequence, we need to know that simplicial spectra built from precosheaves are well behaved:

Lemma 3.20. *Let $\mathcal{G}$ be a precosheaf on a site $\mathcal{C}$ with values in cofibrant objects of a (model) category $\mathcal{D}$; suppose moreover that $\mathcal{G}$ preserves coproducts. Let $\mathfrak{U} = \{U_i \to U\}$ be a cover in $\mathcal{C}$, and $\mathbf{U}. = (\coprod U_i \leftleftarrows \coprod U_{ij} \overset{\leftarrow}{\leftleftarrows} \cdots)$ the associated simplicial object of $\mathcal{C}$. The simplicial object $\mathcal{G}(\mathbf{U}.)$ of $\mathcal{D}$ is Reedy cofibrant. Similarly, if $\mathcal{F}$ is a presheaf on $\mathcal{C}$ with values in the fibrant objects of $\mathcal{D}$, and if $\mathcal{F}$ takes coproducts to products, the associated cosimplicial object $\mathcal{F}(\mathbf{U}.)$ is Reedy fibrant.*

We assume our precosheaf $\mathcal{G}$ takes values in cofibrant spectra and preserves coproducts. The associated simplicial spectrum $\mathcal{G}(\mathbf{U}.)$ is therefore Reedy cofibrant, and the maps $|\mathcal{G}(\mathbf{U}.)|_{q-1} \to |\mathcal{G}(\mathbf{U}.)|_q$ are cofibrations between cofibrant spectra. The corrected homotopy cofibres of these maps, which appear in the E^1 term of the spectral sequence, are then simply ordinary cofibres. We can identify the cofibre of $|\mathcal{G}(\mathbf{U}.)|_{q-1} \to |\mathcal{G}(\mathbf{U}.)|_q$ explicitly as the q-fold suspension of the q-simplicies $\coprod_{|Q|=q} \mathcal{G}(U_Q)$ of the simplicial spectrum $\mathcal{G}(\mathbf{U}.)$:

$$\operatorname{cofib}\left(|\mathcal{G}(\mathbf{U}.)|_{q-1} \to |\mathcal{G}(\mathbf{U}.)|_q\right) = \Sigma^q \left(\coprod_{|Q|=q} \mathcal{G}(U_Q) \right).$$

For example, in picture 3.10 the cofibre of the inclusion of the 0-skeleton is evidently the 1-fold suspension of the 1-simplicies; (note that the 0-skeleton includes degeneracies of the 0-simplicies). The E^1 term of the spectral sequence is therefore $E^1_{pq} = \pi_p(\coprod_{|Q|=q} \mathcal{G}(U_Q))$. Tracing the d^1 differential from the exact couple through to this description of the E^1 term, we find the E^2 term of the spectral sequence is $E^2_{pq} = \check{H}^{\mathfrak{U}}_q(U, \pi_q \mathcal{G})$—the Cech homology of U with respect to the cover $\mathfrak{U}$ with coefficients in the precosheaf of abelian groups $(\pi_p \mathcal{G})(V) := \pi_p(\mathcal{G}(V))$.

Now suppose $\mathcal{G}$ is a cosheaf of cofibrant spectra. The simplicial spectrum $\mathcal{G}(\mathbf{U}.)$ is Reedy cofibrant, and the results of section 3.2.2 identify the target of the spectral sequence:

$$\operatorname{corhocolim}_i |\mathcal{G}(\mathbf{U}.)|_i \simeq \operatorname{colim}_i |\mathcal{G}(\mathbf{U}.)|_i = |\mathcal{G}(\mathbf{U}.)| \simeq \operatorname{hocolim}_{\Delta^{\mathrm{op}}} \mathcal{G}(\mathbf{U}.) \simeq \mathcal{G}(U).$$

The last equivalence here comes from the cosheaf condition. Altogether then, for such a cosheaf of spectra we have a strongly convergent spectral sequence

$$E^2_{pq} = \check{H}^{\mathfrak{U}}_q(U, \pi_p \mathcal{G}) \underset{\text{strong}}{\Longrightarrow} \pi_{p+q} \mathcal{G}(U).$$

The reader can imagine how the dual story progresses. Given a presheaf $\mathcal{F}$ of spectra, we have a cosimplicial spectrum

$$\mathcal{F}(\mathbf{U}.) = \left(\mathcal{F}(\coprod U_i) \rightrightarrows \mathcal{F}(\coprod U_{ij}) \overset{\rightarrow}{\rightrightarrows} \cdots \right).$$

The spectral sequence associated to this cosimplicial spectrum has the form:

$$E^1_{pq} = \pi_{p-q}(\operatorname{corhofib} \operatorname{Tot}^q \mathcal{F}(\mathbf{U}.) \to \operatorname{Tot}^{q-1} \mathcal{F}(\mathbf{U}.)) \Longrightarrow \lim_i \pi_{p-q} \operatorname{Tot}^i \mathcal{F}(\mathbf{U}.).$$

Suppose the presheaf $\mathcal{F}$ takes values in fibrant spectra, and takes coproducts to products. The cosimplicial spectrum $\mathcal{F}(\mathbf{U}.)$ is then Reedy fibrant, and the maps $\operatorname{Tot}^q \mathcal{F}(\mathbf{U}.) \to \operatorname{Tot}^{q-1} \mathcal{F}(\mathbf{U}.)$ in the Tot tower are fibrations between fibrant spectra. The corrected homotopy fibres appearing in the E^1 term of the spectral sequence are therefore ordinary fibres, which are explicitly identifiable.

The fibre of the map $\mathrm{Tot}^q\,\mathcal{F}(\mathbf{U}.) \to \mathrm{Tot}^{q-1}\,\mathcal{F}(\mathbf{U}.)$ is, up to homotopy, the q-fold loop space of the q-cosimplicies $\prod_{|Q|=q} \mathcal{F}(U_Q)$ of our cosimplicial spectrum:

$$\mathrm{fib}(\mathrm{Tot}^q\,\mathcal{F}(\mathbf{U}.) \to \mathrm{Tot}^{q-1}\,\mathcal{F}(\mathbf{U}.)) \simeq \Omega^q(\prod_{|Q|=q} \mathcal{F}(U_Q)).$$

For example, for a cosimplicial spectrum B, the map $\mathrm{Tot}^1 B \to \mathrm{Tot}^0 B$ is, roughly speaking, projecting to the 0-cosimplicies B^0; the fibre over a given 0-cosimplex is an end-fixed path space in the 1-cosimplicies B^1, which (provided the 1-cosimplicies B^1 are connected) has the homotopy type of the loop space ΩB^1—consider picture 3.16. The E^1 term of the spectral sequence associated to the cosimplicial spectrum $\mathcal{F}(\mathbf{U}.)$ is now $E^1_{pq} = \pi_p(\prod_{|Q|=q} \mathcal{F}(U_Q))$. The corresponding E^2 term is $E^2_{pq} = \check{H}^q_{\mathfrak{U}}(U, \pi_p\mathcal{F})$, the Cech cohomology with coefficients in the presheaf $(\pi_p\mathcal{F})(V) := \pi_p(\mathcal{F}(V))$.

Provided there is no $\lim^1$ problem, that is if $\lim^1_i \pi_r \,\mathrm{Tot}^i\,\mathcal{F}(\mathbf{U}.) = 0$, the target of the spectral sequence is $\lim_i \pi_{p-q}\,\mathrm{Tot}^i\,\mathcal{F}(\mathbf{U}.) = \pi_{p-q}\,\mathrm{corholim}_i\,\mathrm{Tot}^i\,\mathcal{F}(\mathbf{U}.)$. Finally suppose that our presheaf $\mathcal{F}$ is actually a sheaf taking values in fibrant spectra. The associated cosimplicial spectrum is Reedy fibrant, and using the results of section 3.2.2 this allows us to further identify the target spectrum:

$$\mathrm{corholim}_i\,\mathrm{Tot}^i\,\mathcal{F}(\mathbf{U}.) \simeq \lim_i \mathrm{Tot}^i\,\mathcal{F}(\mathbf{U}.) = \mathrm{Tot}\,\mathcal{F}(\mathbf{U}.) \simeq \mathrm{holim}_\Delta\,\mathcal{F}(\mathbf{U}.) \simeq \mathcal{F}(U).$$

The last equivalence is the sheaf condition on $\mathcal{F}$. The spectral sequence, at long last, is

$$E^2_{pq} = \check{H}^q_{\mathfrak{U}}(U, \pi_p\mathcal{F}) \underset{\mathrm{cond,no}\lim^1}{\Longrightarrow} \pi_{p-q}\mathcal{F}(U).$$

3.3. **The spectral sequence for $\pi_* TMF$.** We specialize the spectral sequence of Section 3.2 to the particular sheaf $\mathcal{O}^{\mathrm{top}}$ of (fibrant) spectra on the moduli stack $(\mathcal{M}_{ell})_{\text{ét}}$ of elliptic curves in the étale topology. We use particular properties of this sheaf and the results of Section 3.1 to identify the E^2 term as a sheaf cohomology of the moduli stack. Moreover, we address the $\lim^1$ problem for the relevant tower of spectra, thereby pinning down the target of the spectral sequence as the spectrum of global sections $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$.

For a cover $\mathfrak{U} = \{U_i \to \mathcal{M}_{ell}\}$ of $\mathcal{M}_{ell}$ in the site $(\mathcal{M}_{ell})_{\text{ét}}$, and $\mathbf{U}.$ the associated simplicial object of $(\mathcal{M}_{ell})_{\text{ét}}$, the spectral sequence associated to $\mathcal{O}^{\mathrm{top}}$ and $\mathfrak{U}$ has the form

$$E^2_{pq} = \check{H}^q_{\mathfrak{U}}(\mathcal{M}_{ell}, \pi_p\mathcal{O}^{\mathrm{top}}) \underset{\mathrm{cond}}{\Longrightarrow} \lim_i \pi_{p-q}\,\mathrm{Tot}^i\,\mathcal{O}^{\mathrm{top}}(\mathbf{U}.).$$

We can remove the dependence of the E^2 term on the particular cover $\mathfrak{U}$ by restricting attention to covers of the moduli stack by affine schemes.

Proposition 3.21. *Suppose the cover $\mathfrak{U} = \{U_i \to \mathcal{M}_{ell}\}$ is by affine schemes U_i. In this case, for any collection of indices $J = \{i_1, \ldots, i_j\}$, the value of the presheaf $\pi_p\mathcal{O}^{\mathrm{top}}$ on U_J is the same as the corresponding value of the sheafification of $\pi_p\mathcal{O}^{\mathrm{top}}$:*

$$\pi_p\mathcal{O}^{\mathrm{top}}(U_J) = (\pi_p^\dagger\mathcal{O}^{\mathrm{top}})(U_J).$$

It follows from this proposition that there is an equality of Cech cohomology groups:

$$\check{H}^q_{\mathfrak{U}}(\mathcal{M}_{ell}, \pi_p\mathcal{O}^{\mathrm{top}}) = \check{H}^q_{\mathfrak{U}}(\mathcal{M}_{ell}, \pi_p^\dagger\mathcal{O}^{\mathrm{top}}).$$

The same affine condition on the cover gets us the rest of the way to sheaf cohomology:

Proposition 3.22. *Suppose again the cover $\mathfrak{U}$ is by affine schemes. Then for all J, the intersection U_J is acyclic for the sheaf $\pi_p^\dagger\mathcal{O}^{\mathrm{top}}$, that is*

$$H^i(U_J, \pi_p^\dagger\mathcal{O}^{\mathrm{top}}) = 0, \quad i > 0.$$

By the general Proposition 3.4, the acyclicity of U_J for the sheaf $\pi_p^\dagger\mathcal{O}^{\mathrm{top}}$ implies that in this case Cech and sheaf cohomology agree:

$$\check{H}^q_{\mathfrak{U}}(\mathcal{M}_{ell}, \pi_p^\dagger\mathcal{O}^{\mathrm{top}}) = H^q(\mathcal{M}_{ell}, \pi_p^\dagger\mathcal{O}^{\mathrm{top}}).$$

Having identified the E^2 term, we reconsider the target of the spectral sequence. Both the convergence properties of the spectral sequence and the potential $\lim^1$ term are controlled by the same finiteness condition on the differentials—see Boardman [2]:

Proposition 3.23. *Let $E^1_{pq} = \pi_{p-q}(\mathrm{corhofib}\,\phi_q) \Longrightarrow \lim_i \pi_{p-q} F^i$ be the spectral sequence associated to the tower of spectra $\cdots \to F^2 \xrightarrow{\phi_2} F^1 \xrightarrow{\phi_1} F^0 \xrightarrow{\phi_0} *$. If for all p and q there are only finitely many r such that the differential originating at E^r_{pq} is nonzero, then the spectral sequence convergences strongly and $\lim^1_i \pi_{p-q-1} F^i = 0$.*

The spectral sequence associated to the sheaf $\mathcal{O}^{\mathrm{top}}$ evaluated on an affine cover of the moduli stack $\mathcal{M}_{ell}$ has the feature that there are only finitely many nonzero differentials originating at any term.

Corollary 3.24. *For $\mathbf{U}.$ the simplicial object of $(\mathcal{M}_{ell})_{\acute{e}t}$ associated to an affine cover $\mathfrak{U}$ of the moduli stack $\mathcal{M}_{ell}$, the $\lim^1$ term in the Milnor sequence for $\mathrm{Tot}^* \mathcal{O}^{\mathrm{top}}(\mathbf{U}.)$ vanishes:*

$$\lim^1_i \pi_{p-q-1} \mathrm{Tot}^i \mathcal{O}^{\mathrm{top}}(\mathbf{U}.) = 0.$$

The target $\lim_i \pi_{p-q} \mathrm{Tot}^i \mathcal{O}^{\mathrm{top}}(\mathbf{U}.)$ of our spectral sequence is therefore equal to $\pi_{p-q}\,\mathrm{corholim}_i \mathrm{Tot}^i \mathcal{O}^{\mathrm{top}}(\mathbf{U}.)$, which is, in turn, equal to $\pi_{p-q}\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$. The spectrum *TMF* is by definition this spectrum $\mathcal{O}^{\mathrm{top}}(\mathcal{M}_{ell})$ of global sections of the sheaf $\mathcal{O}^{\mathrm{top}}$.

Though we have restricted our attention to the moduli stack $\mathcal{M}_{ell}$ of smooth elliptic curves, the sheaf $\mathcal{O}^{\mathrm{top}}$ extends to the Deligne-Mumford compactification $\overline{\mathcal{M}}_{ell}$. The process we have described also provides a spectral sequence for the homotopy of the spectrum $\mathcal{O}^{\mathrm{top}}(\overline{\mathcal{M}}_{ell})$ of global sections over the compactification—this spectrum is denoted *Tmf*. Altogether then, we have reached the end of our road:

Proposition 3.25. *There are strongly convergent spectral sequences*

$$E^2_{pq} = H^q(\mathcal{M}_{ell}, \pi^\dagger_p \mathcal{O}^{\mathrm{top}}) \Longrightarrow \pi_{p-q}\, TMF,$$
$$E^2_{pq} = H^q(\overline{\mathcal{M}}_{ell}, \pi^\dagger_p \mathcal{O}^{\mathrm{top}}) \Longrightarrow \pi_{p-q}\, Tmf.$$

As the spectrum *tmf* is by definition the connective cover of *Tmf*, the second spectral sequence gives in particular a means of computing the homotopy groups of *tmf*.

Appendix. Degenerate simplicies and codegenerase cosimplicies

We hope we are not alone in thinking that cosimplicial objects, totalization, matching objects, and Reedy fibrancy appear at first more abstruse than the dual notions of simplicial objects, realization, latching objects, and Reedy cofibrancy. We think this is in part due to some missing terminology, which we advertise here. We restrict attention to pointed model categories $\mathcal{C}$ that behave like the category of based spaces in the sense that fibrations are categorical epimorphisms, and the image of a cofibration $P \hookrightarrow Q$ is isomorphic to its source P.

Let A be a simplicial object in such a model category $\mathcal{C}$. The object A_n is called, of course, the *n-simplicies* of A. This object receives degeneracy maps from the k-simplicies A_k for $k < n$. We would like to build an object $\mathrm{Dgnt}_n A$, the *degenerate n-simplicies*, that contains precisely the *targets* of these degeneracy maps in A. In order to do this, we first use a left Kan extension to glue together the possible n-simplicies that could be targets of degeneracy maps from the k-simplicies A_k for $k < n$—this forms the latching object $L_n A$, which we think of as the *possible degenerate n-simplicies* of A. This latching object was described in section 3.2.2, where we noted that "possible" refers to n-simplicies that may exist in simplicial objects agreeing with A in levels less than n.

There is a natural map $\lambda_n : L_n A \to A_n$ from the n-th latching object to the n-simplicies. Define the object $\mathrm{Dgnt}_n A$ of degenerate n-simplicies of A to be the fibre of the cofibre (that is the image) of the map λ_n. The degenerate n-simplicies $\mathrm{Dgnt}_n A$ are the possible degenerate simplicies that actually occur in A. (Note that in an arbitrary model category, it would be wiser to define the degenerate n-simplicies to be the cofibre of the fibre (that is the coimage) of λ_n; however, in topological contexts coimage rarely coincides with our intuition about the target of a map, and we scuttle the coimage formulation.)

Recall that A is Reedy cofibrant if $\lambda_n : L_n A \to A_n$ is a cofibration. This occurs precisely when the map $L_n A \to \mathrm{Dgnt}_n A$ is an isomorphism and the map $\mathrm{Dgnt}_n A \to A_n$ is a cofibration—the first condition says that all possible degeneracies actually occur, which is to say the degenerate simplicies are freely generated, while the second condition says that the degenerate simplicies include into all simplicies by a cofibration. This pair of conditions is a convenient mnemonic for Reedy cofibrancy of simplicial diagrams.

Now let B be a cosimplicial object in the category $\mathcal{C}$. The object B^n is called the *n-cosimplicies* of B. The n-cosimplicies map by codegeneracy maps to the k-cosimplicies, for $k < n$. We would like to build an object that encodes information about the *sources* of these codegeneracy maps—we will call the resulting object $\mathrm{Codgns}^n B$, the *codegenerase n-cosimplicies*. Note well that these are not the "codegenerate n-cosimplicies", a term which would refer to n-cosimplicies in the target of codegeneracy maps, and this object also has nothing to do with coface maps. In order to build this codegenerase object, we first use a right Kan extension to assemble the possible n-cosimplicies that could be sources of codegeneracy maps to the k-cosimplicies B^k for $k < n$—this forms the matching object $M^n B$, which we think of as the *possible codegenerase n-cosimplicies* of B. Matching objects were defined in section 3.2.2; the "possible" here refers to cosimplicies that could appear in cosimplicial objects agreeing with B below level n. (Pedantically speaking, we could emphasize the left versus right Kan extension by thinking of $L_n A$ as the "copossible degenerate" n-simplicies and $M^n B$ as the "possible codegenerase" n-cosimplicies, but we draw a line before the term "copossible".)

There is a natural map $\mu^n : B^n \to M^n B$ from the n-cosimplicies to the n-th matching object. Define the object $\mathrm{Codgns}^n B$ of codegenerase n-cosimplicies of B to be the fibre of the cofibre (that is the image) of the map μ^n. (Note that we do use the image here, not as one might expect the coimage.) The codegenerase n-cosimplicies $\mathrm{Codgns}^n B$ is the object of possible codegenerase cosimplicies that actually occur in B.

The cosimplicial object B is Reedy fibrant if $\mu^n : B^n \to M^n B$ is a fibration. This happens exactly when the map $B^n \to \mathrm{Codgns}^n B$ is a fibration and the map $\mathrm{Codgns}^n B \to M^n B$ is an isomorphism—the first condition is that the n-cosimplicies map by a fibration onto the codegenerase cosimplicies, and the second condition is that all possible codegenerase cosimplicies occur, which is to say that the codegenerase cosimplicies are cofreely generated. This provides a convenient perspective on the meaning of Reedy fibrancy for cosimplicial diagrams.

References

[1] Tibor Beke. Sheafifiable homotopy model categories. *Math. Proc. Cambridge Philos. Soc.*, 129(3):447–475, 2000.

[2] J. Michael Boardman. Conditionally convergent spectral sequences. In *Homotopy invariant algebraic structures (Baltimore, MD, 1998)*, volume 239 of *Contemp. Math.*, pages 49–84. Amer. Math. Soc., Providence, RI, 1999.

[3] A. D. Elmendorf, I. Kriz, M. A. Mandell, and J. P. May. *Rings, modules, and algebras in stable homotopy theory*, volume 47 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1997. With an appendix by M. Cole.

[4] Paul G. Goerss and John F. Jardine. *Simplicial homotopy theory*, volume 174 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1999.

[5] Philip S. Hirschhorn. *Model categories and their localizations*, volume 99 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2003.

[6] S. Hollander. A homotopy theory for stacks. Preprint (2001), math.AT/0110247.

[7] M. J. Hopkins. Complex oriented cohomology theories and the language of stacks. Course notes (1999).

[8] J. Lurie. Higher topos theory. Preprint (2006).

[9] M. A. Mandell, J. P. May, S. Schwede, and B. Shipley. Model categories of diagram spectra. *Proc. London Math. Soc. (3)*, 82(2):441–512, 2001.

[10] John McCleary. *A user's guide to spectral sequences*, volume 58 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.

[11] M. Shulman. Homotopy limits and colimits and enriched homotopy theory. Preprint (2006), math.AT/0610194.