

10 Sheaves on a Site

In which we turn to various powerful generalizations of the sheaf notion, moving beyond the topological case to consider sheaves on a site (Grothendieck toposes).

The chapter begins by exploring *Grothendieck topologies*, moving beyond the usual topological notion of a cover to a more general categorical setting. Using the resulting notion of a *site*—a category together with a choice of Grothendieck topology—the chapter then introduces the notion of a *Grothendieck topos* in terms of sheaves on a site. The rest of the chapter is devoted to developing a particular angle on such sheaves when specialized to posets, and to considering further examples. The chapter concludes by dwelling with the *idea* of such Grothendieck toposes and looking ahead to the even more general notion of *elementary toposes*, via a *Lawvere-Tierney topology*.

10.1 Revisiting Covers: Toward General Sheaves

Sheaf theory emerged as the investigation of the global consequences of properties that are defined locally; and in the classical definition of a sheaf, this notion of *local* ultimately derives from the underlying topology. In the classical definition of a sheaf on a topological space X , sheaves were basically an association of information to the open sets of X , satisfying a gluing axiom specified in terms of a pointwise covering—where, of course, by “pointwise covering” we mean that for $U \subseteq X$, we have that open subsets $\{U_i\}_{i \in I}$ cover U iff $\bigcup_i U_i = U$, where every point that is in U comes from some U_i .

In the appendix, we consider some matters addressing why, in the usual topological definition of sheaves, *open* sets are always emphasized; but, as mentioned there, sheaves were historically defined over closed sets before they were defined over open sets. Considerations of this sort—regarding the objects used to define topologies—is one way to prompt a more careful consideration of questions surrounding *covers*. The objects of study in general (point-set) topology are collections of open sets that are stable (closed) under arbitrary union and finite intersection, where such collections are used above all to study continuity, especially those properties of spaces which are invariant under continuous deformation. In short, as we have seen, a topology is basically a structure that enables us to define objects locally and then look at how these objects are continuously transformed into one another. However, while it has a good deal of power for its simplicity, general topology is not useful for all purposes. For one thing, many interesting structures simply do not fit into this framework or satisfy the relevant axioms. For instance, attempting to treat all the substrings

of a string as open sets would not work, since such a collection will not be stable under union. But it seems sensible to think of the substrings {"Groth", "thend", "ndieck"} as being some sort of covering of the string "Grothendieck," since we can join together the substrings along the overlapping parts ("th" and "nd") to yield the original string.

While general topology is, despite the name, perhaps insufficiently *general*, it remains one of the cornerstones of modern mathematics because of the immense utility and wide applicability of many of its notions and the prevalence of situations where topological intuitions come in handy. There are contexts in mathematics where some form of topological intuition appears relevant or useful, yet where no description of the situation in terms of a space (or any other common ways, e.g., in terms of "locales") can be given.

Moreover, from the perspective of the *association of information* to the open sets of a topological space, we know that the definition of a presheaf, for its part, has a straightforward generalization beyond the topological case to the use of an arbitrary (small) category \mathbf{C} . These sorts of considerations naturally lead one to wonder whether the concept of a sheaf—presented thus far in terms of topological covers—can also be extended beyond the usual topological settings, admitting a definition on more general "topologies." As you can imagine, given this preamble, the answer is yes!

Grothendieck was responsible for a more general notion of a "topology" on a category, one that can now be defined for any category and which we will be able to use to extend the sheaf notion. This notion largely arose from a desire to study entities—in particular, from algebraic geometry, arising in the study of cohomology theories—that appeared to exhibit properties that *looked like* those of a sheaf, yet where the underlying categories involved were not the lattice of open sets of a topological space. The cohomology theory is built on a structure that *seems* like a space, insofar as it comes with some cover-like structure, yet which has features that cannot be found in topology. Grothendieck gave the right description of such contexts, using categories equipped with a general notion of coverings, allowing for a more unified account.

To begin to appreciate how this works, first notice that the open neighborhoods of a topological space are really just topological maps $U \rightarrow X$ that are monic (cancellable on the left).¹⁶⁷ Grothendieck was led to think that perhaps, in place of the neighborhoods U in a topological space X , we could use more general maps into X , where these are not necessarily monic. In this manner, the old-style pointwise covers by open sets, mentioned in the beginning of this section, would come to be replaced with a "covering" by a family of maps that simply satisfies certain conditions or axioms. Using this to consider an arbitrary category, the points of the usual pointwise covering approach will then vanish, leaving only some abstract "open sets," related no longer by inclusion arrows but by arbitrary arrows. The "topology" making the objects behave like the usual open sets is then entirely captured by the specification of the generalized covering in terms of abstract conditions on such families of maps.

A Grothendieck topology can initially be thought of—where, really, we are currently describing what is called a pretopology—as a rule for specifying when certain objects of a category "ought to" cover another object of the category, but purifying (by axiomatizing)

167. Recall that monic maps, or monomorphisms, are the category-theoretic generalization of the notion of injective maps from sets; see definition 65 (chapter 3).

the usual topological notion of an open cover. As such, we are basically putting a structure on a category that allows its components to be decomposed, demanding that the objects “behave like” the open sets of a topological space, except that rather than looking at intersections, we can look at pullbacks, and unions don’t even have to come into play. While this is already a powerful generalization, we can actually go further, dropping the assumption that our category has pullbacks. The Grothendieck topology that emerges will effectively tell us about when abstract morphisms act as one might hope any cover would, formalizing the notion of being *locally the case*, without confining us to the notions or assumptions from the usual topological setting. A major difference between this new “topology” and the old notion of open covers from general topology is that our “open sets” (now objects in a category) need not be stable under union (i.e., coproducts need not exist in the category), nor even under intersection (i.e., pullbacks need not exist).¹⁶⁸

Via this wider notion of a Grothendieck topology—where, instead of the usual inclusion relation between open sets, we can consider arbitrary arrows in a category—we can make use of many topological intuitions and devices in situations where there does not appear to be any topology (in the traditional sense) at play. *That*—and not just generality for its own sake!—is part of what is behind the real power of all this. The fundamental takeaway here is that there is no need to restrict ourselves to topological spaces to do sheaf theory: as long as the category gives information *similar to* open covers (which similarity is formalized with the notion of a Grothendieck topology), we can generalize the constructions involved in creating a sheaf beyond the category of open sets of a topological space. Using Grothendieck topologies, we develop a more powerful notion of a sheaf defined on a (small) category, disposing of the condition that the base be formed of open sets and the lattice formed by inclusions arrows between these sets.¹⁶⁹

Before diving into the relevant definitions and constructions, here’s a last way of thinking of all this. Suppose you have an (arbitrary) presheaf. You may want to say that if you know the value of that presheaf on some family of objects X_i —let’s not assume anything else of those objects—then that “ought to” be enough to tell you the value of the presheaf on some X (not necessarily a space), which we would thus like to think of as being “covered” by the family X_i . In other words, you want to know which families of arrows into X are such that the presheaf data valued locally on that family *extends* to presheaf data on X along those arrows. But your domain category might not be topological at all, and so the objects might not present as “open sets” in the strict sense. By developing the notions discussed above, we will be able to address these needs. The important notion of a *site* will emerge as a category together with a choice of Grothendieck topology—using this, we will be able to define the all-important notion of a *Grothendieck topos* in terms of sheaves on a site.

168. While Grothendieck’s use of the term “topology” for his own notion is entirely sensible, we should emphasize that it is not literally a topology in the usual sense. A Grothendieck topology, strictly speaking, has nothing to do with—or, rather, does not rely on—open sets or closed sets in the usual topological sense (i.e., as treated in chapter 4).

169. Sheaf theory cast in terms of this notion of Grothendieck topologies was first described in Artin, Grothendieck, and Verdier (1972, SGA 4 Exposé I–IV).

10.2 Grothendieck Toposes

10.2.1 Grothendieck Topologies: First Pass

The first way of commonly approaching this new sort of “covering” is by considering how the construction can be accomplished in any category \mathbf{C} assumed to have pullbacks. We will use the ingredients about to be introduced to define what is technically a *pretopology*. As the name suggests, one can think of it as “on the way toward” the still more general notion toward which we are building.

For each object c of \mathbf{C} , let’s start by considering a set S of indexed families of morphisms to c , that is,

$$S = \{f_i : c_i \rightarrow c \mid i \in I\}.$$

If you think of an arrow into c as a “perspective” on c , taking such an S is like bunching together a number of perspectives on, or ways of seeing, c . The poet Wallace Stevens has a poem “Thirteen Ways of Looking at a Blackbird,” which consists of thirteen short parts, each of which present a distinct “perspective” on a blackbird. For intuition, S is the poem which comprises thirteen (so index set $I = 13$) ways of looking at $c = \text{blackbird}$. Suppose next that for each object c of \mathbf{C} we assign a collection

$$K(c) = \{S, S', S'', \dots\}$$

consisting of a selection of certain of the families of morphisms (each of the same form as S). On the “perspective” way of seeing things, this is like selecting certain “poems” about our blackbird—as if to say, “I am interested in *these* particular families of ways of seeing c .” But suppose we make you abide by a rule: you can only make such a selection provided it meets certain (three) conditions. Provided it does meet those conditions, then your collection $K(c)$ is allowed, and we agree to call it a *covering* and the families in it *covering families* (or *covers* of c). Such a $K(c)$ then makes up our “(pre)topology.”

Definition 236 Let \mathbf{C} be a category that has pullbacks. Then a *Grothendieck pretopology* (or *basis* (for a Grothendieck topology)) on \mathbf{C} is a function K that assigns to each $c \in \text{Ob}(\mathbf{C})$ a collection $K(c)$ of families of morphisms into c , as above, where this satisfies three conditions:

1. Whenever an arrow $f : c' \xrightarrow{\cong} c$ in \mathbf{C} is an isomorphism, then the family $\{f : c' \xrightarrow{\cong} c\}$ consisting of just that arrow is in the collection $K(c)$.¹⁷⁰
2. If $\{f_i : c_i \rightarrow c\}$ is a covering family (i.e., is in $K(c)$) and $g : b \rightarrow c$ is any morphism in \mathbf{C} , then the family of pullbacks $\{\pi_2 : c_i \times_c b \rightarrow b\}_{i \in I}$ exists and is a covering family (i.e., is in $K(b)$).¹⁷¹
3. If $\{f_i : c_i \rightarrow c\}$ is a covering family (i.e., is in $K(c)$) and for each i we also have a family $\{g_{ij} : b_{ij} \rightarrow c_i\}_{j \in J_i}$ that is a covering family (i.e. is in $K(c_i)$), then the family of composites $\{f_i \circ g_{ij} : b_{ij} \rightarrow c_i \rightarrow c\}_{i \in I, j \in J_i}$ is also a covering family (i.e., is in $K(c)$).¹⁷²

170. Think of this as the “solipsist” stipulation: if you’re going to collect different “poems” about c , you can only do so as long as you agree to “let c tell her own story from her own perspective (or from the perspective of anyone who essentially shares her perspective).”

171. Attempt to finish the analogy for yourself (this is an important one!): Given a “poem” with a bunch of perspectives on c , and then given another isolated perspective on c (from the vantage of some b), . . .

172. Try formulating for yourself the right analogy here as well.

To formulate what it means to have a sheaf for the coverings comprising such a pretopology, we can actually rehash the classical (topological) definition of a sheaf. We just have to make one change. Recall from the definition of a sheaf on a topological space the stipulation that for each open cover $\{U_i\}_{i \in I}$ of some U , every family of elements $\{x_i \in P(U_i)\}_{i \in I}$ matching on the intersections $U_i \cap U_j$ for all i and j can be glued together into a unique element $x \in P(U)$. The same sort of thing works for a definition using covering families of an object c ; all we need to do is replace intersection $c_i \cap c_j$ by the pullback $c_i \times_c c_j$, that is,

$$\begin{array}{ccc} c_i \times_c c_j & \xrightarrow{h_{ij}} & c_j \\ v_{ij} \downarrow & & \downarrow f_j \\ c_i & \xrightarrow{f_i} & c. \end{array}$$

Now, to define a sheaf, we need a *presheaf* to work with. But we can of course just use a functor $P : \mathbf{C}^{op} \rightarrow \mathbf{Set}$, where we only assume of \mathbf{C} that it has pullbacks. Then, applying P to the diagrams in \mathbf{C} of the above sort, we would get associated diagrams living in \mathbf{Set} :

$$\begin{array}{ccc} P(c_i \times_c c_j) & \xleftarrow{P(h_{ij})} & P(c_j) \\ P(v_{ij}) \uparrow & & \uparrow P(f_j) \\ P(c_i) & \xleftarrow{P(f_i)} & P(c). \end{array}$$

We then just follow the lead given by the topological definition. The matching or gluing condition for a sheaf will thus read: if $\{x_i \in P(c_i)\}_{i \in I}$ is a family of elements that match, in the sense that $P(v_{ij})(x_i) = P(h_{ij})(x_j)$ for all i, j , then that family determines a unique element $x \in P(c)$ such that

$$P(f_i)(x) = x_i$$

for all $i \in I$. And this is equivalent to requiring that the arrow e is an equalizer

$$P(c) \xrightarrow{e} \prod_i P(c_i) \rightrightarrows \prod_{i,j} P(c_i \times_c c_j),$$

where $e(x) = (P(f_i)(x))_{i \in I}$, for every covering family $\{c_i \xrightarrow{f_i} c\}_{i \in I}$.

Assuming a category \mathbf{C} has pullbacks, we will thus have a way of talking about sheaves on it! No topology (in the usual sense) needed!

The definition of pretopologies via covering families just sketched requires that pullbacks exist in the underlying category. We can just as well provide a more general definition that does not rely on their existence, and this is done not just for the sake of giving the most general definition but also because there are certain categories of interest that do not have pullbacks. This generalization is typically accomplished by replacing the indexed families S of morphisms into an object with the *sieves* they generate, a construction the reader is in fact already familiar with (under another name). By using sieves, we will see that we can dispense (in principle) with the assumption that \mathbf{C} has pullbacks.

But in fact, without yet getting into sieves, we can point out that the three conditions given above are not created equal: it is really the second condition in the definition that is most decisive, as far as sheaves are concerned. Following Johnstone (2002), we might thus isolate the second condition and define the following notion.

Definition 237 Let \mathbf{C} be a category. Then a *coverage* on \mathbf{C} is a function T that assigns to each $c \in \text{Ob}(\mathbf{C})$ a collection $T(c)$ of families of morphisms $\{f_i : c_i \rightarrow c\}_{i \in I}$ into c (these are called *T-covering families*, or just *covering families*, when T is understood), where this satisfies:

- If $\{f_i : c_i \rightarrow c\}_{i \in I}$ is a covering family and $g : b \rightarrow c$ is any morphism with codomain c , then there exists a covering family $\{h_j : b_j \rightarrow b\}_{j \in J}$ such that each composite $g \circ h_j$ factors through some f_i , in the sense that

$$\begin{array}{ccc} b_j & \xrightarrow{h_j} & b \\ k \downarrow & & \downarrow g \\ c_i & \xrightarrow{f_i} & c. \end{array}$$

Remark 238 The traditional way of defining a *Grothendieck topology*, as we shall see, is to use three conditions (similar to those found in the definition of a Grothendieck pretopology—see definition 236). But, really, the second condition in 236 is decisive for the introduction of the sheaf notion, while the other conditions (usually called “saturation conditions” or “closure properties”) are conditions that are often met in practice (and so, typically assumed), yet fundamentally do not affect the sheaf notion. In particular, the collection of families of morphisms for which a given functor satisfies the associated sheaf axiom will typically have a number of closure properties (essentially encoded by the notion of a sieve), which are thus commonly added in the definition itself.

However, following Johnstone (2002), one can take the approach of isolating this essential condition and giving the definition of a coverage, as in 237, and then reserving the name *Grothendieck coverage* for a coverage that also satisfies the two other saturation conditions (where a Grothendieck coverage is the same thing as what is traditionally called, and that we will call, a Grothendieck topology).¹⁷³ Grothendieck himself originally considered coverages that also obey the additional saturation conditions.

Finally, observe that a Grothendieck pretopology differs from the above in that it assumes the category \mathbf{C} has pullbacks and replaces the condition in the definition 237 by the stronger condition:

If $\{f_i : c_i \rightarrow c\}$ is a covering family and $g : b \rightarrow c$ is any morphism with codomain c , then the family of pullbacks $\{\pi_2 : c_i \times_c b \rightarrow b\}_{i \in I}$ exists and is a covering family.

This condition is satisfied in many examples; however, strictly speaking, it is not needed for defining sheaves—the more general condition given in 237 will be enough.

Using this notion of coverage, we can already define the notion of a site, and then define sheaves for this, just as one would expect.

Definition 239 A *site* will mean a pair (\mathbf{C}, T) consisting of a category \mathbf{C} equipped with a coverage T . (A *small site* is a site for which the underlying category is small.)

Definition 240 For \mathbf{C} a category, a functor $F : \mathbf{C}^{op} \rightarrow \mathbf{Set}$ is said to satisfy the *sheaf axiom* for a family of morphisms $\{f_i : c_i \rightarrow c\}_{i \in I}$ if, whenever we have a family of elements $s_i \in$

173. Johnstone 2002 renamed things to avoid the “emotional baggage” associated with the language of a “topology.”

$F(c_i)$ that are *compatible*—in the sense that whenever $g : b \rightarrow c_i$ and $h : b \rightarrow c_j$ satisfy $f_i \circ g = f_j \circ h$ (for any i, j , not necessarily distinct), we have $F(g)(s_i) = F(h)(s_j)$ —then there exists a unique $s \in F(c)$ such that $F(f_i)(s) = s_i$ for each $i \in I$.¹⁷⁴

Then, given T a coverage on \mathbf{C} , F is said to be a T -sheaf provided it satisfies the sheaf axiom for every T -covering family.

Again, the collection of all families of morphisms for which a functor satisfies the sheaf axiom, as above, will generally have some further “closure properties.” On account of this, it is typical to add such conditions to the single condition defining a coverage, obtaining an expanded definition. We will now expand on the covering families approach via the notion of a *sieve*, which will lead to the original (and common) definition of a Grothendieck topology. Moreover, as we will see—and as the name “pretopology (or basis for a Grothendieck topology)” suggests—a pretopology on a category can ultimately be used to generate a Grothendieck topology (though different pretopologies may generate the same topology).


10.2.2 Sieves

Definition 241 A sieve S on an object c of \mathbf{C} is a family of morphisms in \mathbf{C} , all of which have codomain c , that satisfies

$$f \in S \text{ implies } f \circ g \in S,$$

whenever such a composition makes sense.

In other words, a sieve S on c is a collection of arrows with codomain c , where this is “closed under composition on the right” (or “closed under precomposition”). A generic picture of a sieve might look something like:



In the ordinary sense of the term, a sieve is of course a device used to strain, sift or sort some material, filtering out undesired pieces from a larger input. In a similar fashion, if we start with all arrows into c , a sieve (in our sense) on c is thus a collection of arrows that sifts or filters certain arrows from that set of all arrows into c : if there is an arrow $b \rightarrow c$ going through to c , then whenever there is an arrow $b' \rightarrow b$, the arrow $b' \rightarrow c$ “finds its way” through to c as well. In slogan form (explaining the term “sieve”):

If b goes through the sieve, then so too does anything “smaller than” b .

174. By the way, notice how the definition simplifies in the case where \mathbf{C} has pullbacks: specifically, one need only check the compatibility of s_i and s_j on the pullback of f_i along f_j , instead of on arbitrary pairs of functions (g, h) , as stipulated by this more general definition.

This defining condition of a sieve being closed under pre-composition is sometimes called a “saturation condition.” Moreover, given a collection of morphisms landing in c , it can always be saturated in this way, leaving us with a sieve on c .

Here is another important saturation feature of the notion of a sieve. The definition implies that a sieve S is such that for any $b \rightarrow c$ in S , any path to b followed by the path from b to c will itself be a path in S . In particular, then, if S is a sieve on c and $h : b \rightarrow c$ is any arrow to c , we will have that

$$h^*(S) = \{g \mid \text{codom}(g) = b \text{ and } h \circ g \in S\}$$

is a sieve on b . This is another closure or saturation condition, in that it says that sieves are “closed under pulling back.” It is straightforward to show that $h^*(S)$ is itself a sieve.

This is a powerful notion that makes sense in any category. But to get a better handle on this notion, it may again help to specialize to posets. Recall that a *downset* is a subset $D \subseteq P$ for which, for all $x \in D$ and for all $y \in P$, if $y \leq x$, then $y \in D$. Moreover, recall that for each element p of the poset, the downset generated by p —its *principal downset*—is

$$\downarrow p := \{q \in \mathcal{P} \mid q \leq p\}.$$

Applying the definition of a sieve in the context of posets, we can thus say that a subset $S \subseteq P$ will be a sieve on p if $q \leq p$ for each $q \in S$ and $r \in S$ if $r \leq q$ for some q . In other words, a subset $S \subseteq \mathcal{P}$ is a sieve on p provided S belongs to $\mathcal{D}(\downarrow p)$, the collection of all downsets of the principal downset $\downarrow p$. Moreover, in a poset, regarded as a category, recall that there is an arrow $q \rightarrow p$ precisely when $q \leq p$, and there is *at most* one arrow from q to p . Accordingly, in such a setting, we can identify the underlying families of morphisms $S = \{f_i : c_i \rightarrow c\}_{i \in I}$, consisting of arrows into c closed under precomposition, with a family of *elements* $\{c_i \mid c_i \leq c \text{ for all } i \in I\}$ that is downward closed.

In a moment, we will use this formulation for posets to present the main definitions in terms of posets. But let us first make a few other, more general observations. On account of the intimate connection between principal downsets and representable functors—as explored in chapter 6—it should come as no surprise to the reader that, in the more general case, each sieve S on an object c can be identified with a subpresheaf $S \subseteq \text{Hom}_{\mathbf{C}}(-, c) := Y_c$ of the representable functor (at least in settings where \mathbf{C} is locally small).

Proposition 242 A sieve S on an object c of a category \mathbf{C} can be exhibited as a subfunctor of the representable hom-functor $Y_c := \text{Hom}_{\mathbf{C}}(-, c)$, via the following specification:

$$\begin{aligned} b &\longmapsto \{f \in \text{Hom}_{\mathbf{C}}(b, c) \mid f \in S\} \\ (a \xrightarrow{g} b) &\longmapsto g^*, \text{ such that } g^* : f \mapsto f \circ g. \end{aligned}$$

In other words, given a sieve S on c , defining $Q(b) = \{f \mid f : b \rightarrow c \text{ and } f \in S\} \subseteq \text{Hom}_{\mathbf{C}}(b, c)$ will produce a functor $Q : \mathbf{C}^{op} \rightarrow \mathbf{Sets}$ that is a subfunctor of Y_c . Conversely, via the Yoneda embedding $\mathbf{y} : \mathbf{C} \rightarrow \mathbf{Set}^{\mathbf{C}^{op}}$ taking an object $c \in \mathbf{C}$ to Y_c , we can further consider that if we have a subfunctor $Q \subseteq Y_c$, then the set

$$S = \{f \mid \text{for some object } b, f : b \rightarrow c \text{ and } f \in Q(b)\}$$


will form a sieve on c .

Altogether, since we can pass from S to Q and from Q to S in reciprocal fashion, this informs us of the fact that:

$$\text{sieve on } c = \text{subfunctor of } Y_c.$$

This already gives us a more category-theoretic way of formulating the notion of a sieve, as compared to the definition that takes a sieve to be some *set* of arrows satisfying certain constraints. A third characterization of sieves is also in this more category-theoretic spirit.

Proposition 243 A sieve $S = \{c_i \xrightarrow{f_i} c \mid i \in I\}$ on an object c of a category \mathbf{C} can be seen as a subcategory of the slice (comma) category \mathbf{C}/c , where the objects are just the arrows $c_i \xrightarrow{f_i} c$ of S , and the morphisms are given by arrows u making triangles of the following form commute:



In short, a sieve has three (equivalent) characterizations:

- as a set of arrows into some object (where these arrows satisfy a condition of closure under composition on the right);
- as a subfunctor of the representable hom-functor;
- as a subcategory of the slice category.

Before putting this notion of a sieve to work, let us make a simple observation that will allow us to define one other notion that will be of use to us. Observe that a sieve on some object clearly does not need to include *all* the arrows into that object. However, taking all arrows into an object will indeed amount to a sieve, and we reserve a special name for such a sieve.

Definition 244 For an object c of \mathbf{C} , the set

$$M_c = \{f \mid \text{codomain}(f) = c\}$$

of *all* arrows into c will be a sieve, one that (for reasons that should be clear) is called the *maximal sieve on c* .

Observe that, specialized to posets, the maximal sieve on $p \in P$ is just given by taking $\downarrow p$, the principal downset itself.

Finally, relating this back to the earlier discussions using families of morphisms into c : observe that any family $\{f_i : c_i \rightarrow c\}_{i \in I}$ generates a sieve on c , by taking all those morphisms with codomain c which factor through at least one of the f_i . And it is straightforward to demonstrate that given a family $\{f_i : c_i \rightarrow c\}_{i \in I}$ of morphisms of \mathbf{C} and F a presheaf on \mathbf{C} , F satisfies the sheaf axiom for $\{f_i : c_i \rightarrow c\}_{i \in I}$ iff it satisfies the sheaf axiom for the sieve generated by that family.

10.2.3 Grothendieck Topologies and Sites Defined

Equipped with the notion of a sieve, we are now in a position to define the following concept of a Grothendieck topology.

Definition 245 A *Grothendieck topology* (or *Grothendieck coverage*) on a category \mathbf{C} is a function J that assigns to each object c of \mathbf{C} a collection $J(c)$ of sieves on c , in such a way that

1. *identity cover* (or *maximality axiom*): the maximal sieve M_c is in $J(c)$;
2. *stability under change of base* (or *stability axiom*): if $S \in J(c)$ and $h: b \rightarrow c$ is any morphism with codomain c , then $h^*(S)$ is in $J(b)$;¹⁷⁵
3. *local character condition* (or *transitivity axiom*): if $S \in J(c)$ and R is another sieve on c such that, for each f in S (e.g., $f: b \rightarrow c$), the sieve $f^*(R)$ is in $J(b)$, then R belongs to $J(c)$.

At first glance, this might seem unmotivated or difficult to parse. But observe that the conditions on a Grothendieck topology are rather intuitive: (1) is just an inclusion of representable functors as covers; (2) is really a disguised way of insisting that J itself be functorial (a presheaf); and (3) requires that things are transitive.

Terminologically, if $S \in J(c)$, we say that S is a *covering sieve*, or that S *covers* c (or, sometimes, for explicitness, that S is a J -*cover* of c). We will also say that a sieve S on c *covers an arrow* $f: b \rightarrow c$ if $f^*(S)$ covers b . In other words, S covers c iff S covers the identity arrow on c . In these terms, the three axioms for a Grothendieck topology given above can equivalently be formulated in terms of a relation between the sieves and the arrows of the category (this is called the *arrow form* of the definition in Mac Lane and Moerdijk 1994):

- (1a) *identity*: if S is a sieve on c and f is in S , then S covers f ; informally, this says that any “open set” covers itself, or (in terms of sets) that any set is covered by all its possible subsets.
- (2a) *stability*: if S covers an arrow $f: b \rightarrow c$, it also covers the composition $f \circ g$, for any arrow $g: a \rightarrow b$; informally, this says that coverings “pull back.”
- (3a) *transitivity*: if S covers an arrow $f: b \rightarrow c$ and R is a sieve on c which covers all arrows of S , then R covers f ; informally, this ensures that a cover of a cover will be a cover.

A further feature that one might wish to impose on such covering sieves in fact already follows from the three main axioms (in either form): namely, that any two covers have a common *refinement*, that is, $J(c)$ is closed under finite intersections. Explicitly, we mean

- (4) if $R, S \in J(c)$, then $R \cap S \in J(c)$,

or in arrow form:

- (4a) if R and S both cover $g: b \rightarrow c$, then $R \cap S$ covers g .

Equipped with the notion of a Grothendieck topology, we can define the following:

Definition 246 A *site* will mean a pair (\mathbf{C}, J) consisting of a category \mathbf{C} equipped with a Grothendieck topology J .¹⁷⁶

175. Observe that, in the presence of the other two conditions, this essentially amounts to a simplification of the coverage condition in definition 237.

176. In such contexts, \mathbf{C} is typically assumed to be small.

Both because it offers a more immediately accessible version of the above definition and because we will make particular use of this specialized version in what follows, let us also give a version of the definition specialized to posets.

Definition 247 (*Grothendieck topology specialized to posets*) A Grothendieck topology J on a poset \mathcal{P} is a map

$$p \mapsto J(p)$$

assigning to each element $p \in \mathcal{P}$ a collection $J(p)$ of sieves on p —where, as indicated above, $S \subseteq \mathcal{P}$ is a sieve on p provided $q \leq p$ for each $q \in S$ and if $r \leq q$ for some $q \in S$, then also $r \in S$ —that satisfies the following conditions:

1. the maximal sieve $\downarrow p$ is an element of $J(p)$;
2. if $S \in J(p)$ and $q \leq p$, then $S \cap \downarrow q \in J(q)$;
3. if $S \in J(p)$ and R is a sieve on p such that $R \cap \downarrow q \in J(q)$ for each $q \in S$, then also $R \in J(p)$.¹⁷⁷

Earlier it was mentioned that the new definition of a Grothendieck topology in terms of sieves, as in definition 245, has some advantages, including the advantage of allowing us to dispense with the assumption that the underlying category \mathbf{C} has pullbacks. Even if the individual morphisms in \mathbf{C} cannot be pulled back, a sieve can *always* be pulled back along a morphism of \mathbf{C} —definition 245(2) tells us how this should go. This is ultimately the case because each sieve on an object c of the category can be identified with a subpresheaf of the representable hom-functor Y_c , and the presheaf category $\mathbf{Set}^{\mathbf{C}^{op}}$ in which this lives itself has pullbacks. Being able to identify sieves with subpresheaves of the representable hom-functor Y_c will be especially useful to us in defining sheaves for such topologies built from covering sieves. Another reason for the new definition in terms of sieves is that—referring back to the notion of a Grothendieck pretopology—two different pretopologies may yield the exact same sheaves, so there was a lingering imprecision in the definition of a Grothendieck pretopology. It is partly in order to overcome this imprecision that we move from focusing on covering families to covering sieves.

On the other hand, in practice, it is often the case that covering families are easier to specify and work with than the sieves that they generate; especially when the underlying category \mathbf{C} has pullbacks, it also appears to be simpler to work with the generating covering families, at least when assessing whether or not a particular functor is a sheaf.

Before coming back to some of these niceties, let us look at some examples of Grothendieck topologies. There will generally be a number of Grothendieck topologies that could be attached to a category. The next two examples represent two important extremes.

177. With this, it becomes especially evident that J in fact must be a functor $\mathcal{P}^{op} \rightarrow \mathbf{Set}$, once we define $J(q \xrightarrow{\leq} p)(S) = S \cap \downarrow q$ for each $S \in J(p)$. If we have $r \leq q \leq p$ in \mathcal{P} and $S \in J(p)$, then

$$J(r \xrightarrow{\leq} p)(S) = S \cap \downarrow r = (S \cap \downarrow q) \cap \downarrow r = J(r \xrightarrow{\leq} q)J(q \xrightarrow{\leq} p)(S).$$

Defining things in this way, the stability axiom ((2) above) then just expresses that $J \in \mathbf{PreSh}(\mathcal{P})$. However, observe that, in general, as a functor, J itself is not a sheaf—for, given a cover $\{c_i \xrightarrow{f_i} c\}_{i \in I}$ of c and a compatible family of covers $\{x_i \in J(c_i)\}_{i \in I}$, there may in fact be a number of covers of c that extend this latter family.

Example 248 The *minimal topology*—also going under the names *trivial*, *indiscrete*, or *coarse topology*— J_{ind} on a category \mathbf{C} is the one in which the *only* sieve covering an object c is the maximal sieve M_c . In other words, we are declaring that only sieves of the form $245(1)$ —that is, $\text{Hom}_{\mathbf{C}}(-, c)$ —are covering sieves. In particular, when dealing with \mathcal{P} a poset, since we know that representable functors are the principal downsets, the indiscrete Grothendieck topology on \mathcal{P} is just given by taking the principal downsets, that is,

$$J_{ind}(p) = \{\downarrow p\}.$$

This topology is the *coarsest* (smallest) of all topologies that can be put on \mathbf{C} .

Example 249 The *maximal topology*—also *discrete topology*—on \mathbf{C} is such that *every sieve* is declared a covering sieve. Specialized to \mathcal{P} a poset, the discrete Grothendieck topology on \mathcal{P} will be given by

$$J_{dis}(p) = \mathcal{D}(\downarrow p).$$

This is clearly the *finest* (biggest) topology on \mathbf{C} . In practice, this topology can sometimes lead to rather nonintuitive covers. As the names suggest, the maximal and the minimal topologies provide the two extremes on the spectrum of possible topologies.

Example 250 As is commonly known in differential geometry, the category **Man** of smooth manifolds and their smooth maps does not have all pullbacks—in particular, pullbacks of manifolds are not generally manifolds. (It has a number of other deficiencies, from a categorical perspective.) In general, one needs to impose some structure on the maps in order to ensure the existence of pullbacks.

But if we take for covering families the families of “open embeddings” $\{f_\alpha : U_\alpha \hookrightarrow M\}_{\alpha \in A}$ such that the family of open sets $\{f_\alpha(U_\alpha)\}_{\alpha \in A}$ is an open cover of M , this gives us a Grothendieck topology on **Man**—and using such a site, one can observe that the representable functors are sheaves.

The notion of a Grothendieck topology and its covering sieves is sometimes said to be a vast generalization of the usual notions of a topological space and its covers.¹⁷⁸ For this to be a valid perspective, we should at the very least be able to check that the latter notion is indeed captured, as a particular case, by the former.

Example 251 By viewing the lattice (poset) $\mathcal{O}(X)$ of open subsets of a topological space X as a category, a sieve on U is just a family $S = \{U_i\}_{i \in I}$ of open subsets of U satisfying the condition that $W \subseteq V \in S$ implies that $W \in S$. One then says that S covers U provided U is contained in the union of the open sets of S . The usual topological notion of a cover is thus recovered by taking as our Grothendieck topology that specified by

$$J_{\mathcal{O}(X)}(U) = \left\{ S \in \mathcal{D}(\downarrow U) \mid \bigcup S = U \right\},$$

where, as before, $\mathcal{D}(\downarrow U)$ is the collection of all the downsets of $\downarrow U$. You can verify that the standard open cover definition in fact satisfies the axioms for a Grothendieck topology,

178. This is not exactly the right way to think about it, since it is not at all a straightforward abstraction from the topological notions. But, as the next example shows, it does indeed embrace the usual topological notions.

and that the above Grothendieck topology corresponds to the usual topological notion of an open cover.¹⁷⁹

In greater generality, given a *complete Heyting algebra* H (that is, a Heyting algebra with arbitrary joins \bigvee and meets)—or a *frame*, introduced below—we can define the Grothendieck topology J_H by taking

$$\{c_i\}_{i \in I} \in J_H(c) \text{ iff } \bigvee_{i \in I} c_i = c.$$

This is sometimes called the *canonical topology on H* .

Example 252 (*The dense topology*) The dense topology J_{dense} —also called the double-negation topology, for reasons we will see shortly—can be defined for an arbitrary category \mathbf{C} in terms of collections of sieves $J_{dense}(c)$ of the form:

$$S \in J_{dense}(c) \text{ iff for any } f : b \rightarrow c \text{ there exists a } g : a \rightarrow b \text{ such that } f \circ g \in S.$$

It is useful to look at this in the special case when we are dealing with a poset \mathcal{P} . In general, for a poset \mathcal{P} , a subset $D \subseteq \mathcal{P}$ is said to be *dense* if, for every $p \in P$, there is a $q \leq p$ with $q \in D$. The subset $D \subseteq \{q \in P \mid q \leq p\}$ is then said to be *dense below p* if for every $p' \leq p$, there exists a $q \leq p'$ with $q \in D$. Observe that since $D \subseteq \{q \in P \mid q \leq p\}$, D is already a sieve. If D is dense below p , we can call it a *dense sieve (on p)*. The collection of dense sieves will supply us with a topology J_{dense} on a poset by taking

$$J_{dense}(p) = \{D \mid q \leq p \text{ for all } q \in D, \text{ and } D \text{ is a sieve dense below } p\}.$$

Observe that another way of saying this is

$$J_{dense}(p) = \{D \in \mathfrak{D}(\downarrow p) \mid \downarrow p \subseteq \uparrow D\},$$

that is, $J_{dense}(p)$ will consist of a selection of sets D from the collection of all downsets of the principal downset $\downarrow p$, in such a way that the principal downset itself is contained by $\uparrow D = \{x \in \mathcal{P} \mid x \geq d \text{ for some } d \in D\} = \bigcup_{d \in D} \uparrow d$, the upper set generated by such D .

It may be helpful to prove that, defined thus, such objects indeed supply us with a Grothendieck topology.

Exercise 22 Prove that such dense sieves form a Grothendieck topology.

Solution

Proof. We need to check the three requirements:

1. the maximal sieve (principal downset) $\downarrow p$ is an element of $J(p)$;
2. if $S \in J(p)$ and $q \leq p$, then $S \cap \downarrow q \in J(q)$ (stability);
3. if $S \in J(p)$ and R is any sieve on p such that $R \cap \downarrow q \in J(q)$ for each $q \in S$, then also $R \in J(p)$ (transitivity).


179. However, there is a subtlety here. As we know, in the usual context of a topological space, an open cover of U is typically described as a family $\{U_i \mid i \in I\}$ of open subsets of U whose union $\bigcup_i U_i$ is equal to U . Note that such a family is not necessarily a sieve. However, it does *generate* a sieve, via taking the collection of all those open $V \subseteq U$ where $V \subseteq U_i$ for some U_i . This motivates a more careful consideration of the notion of *generating a covering sieve*, something that can be described in the more general context of an arbitrary category (with or without pullbacks), via a *basis* for a Grothendieck topology. See below and Mac Lane and Moerdijk (1994, III.2) for details.

Requirement 1 is basically immediate: certainly $\downarrow p$ itself is (trivially) in the collection of downsets of $\downarrow p$, and $\downarrow p \subseteq \uparrow(\downarrow p)$. Thus, $\downarrow p \in J_{dense}(p)$.


Requirement 2: let $D \in J_{dense}(p)$ and consider $q \leq p$. But $\downarrow p \subseteq \uparrow D$ (since $D \in J_{dense}(p)$), so if $r \in \downarrow q$, we will have that $r \in \downarrow p$; hence, there will have to be an $s \in D$ with $r \leq s$. But then $s \in D \cap \downarrow q$, as $q \geq r$, from which we must have $\downarrow q \subseteq \uparrow(D \cap \downarrow q)$, which means $D \cap \downarrow q \in J_{dense}(q)$.

Requirement 3: let $D \in J_{dense}(p)$ and $E \in \mathcal{D}(\downarrow p)$ so that $E \cap \downarrow q \in J_{dense}(q)$ for $q \in D$. Thus, $\downarrow p \subseteq \uparrow D$ and $\downarrow q \subseteq \uparrow(E \cap \downarrow q)$ for $q \in D$. If we now consider $r \in \downarrow p$, we can show that $r \in \uparrow E$ as follows: $\downarrow p \subseteq \uparrow D$, so there exists a $q \in D$ with $q \leq r$; we also have that $\downarrow q \subseteq \uparrow(E \cap \downarrow q)$; thus, there exists an $s \in (E \cap \downarrow q)$ with $s \leq q$, which means that $r \in \uparrow E$, with $s \leq q \leq r$. This shows that $E \in J_{dense}(p)$. \square

For concreteness, consider the following poset



Its associated downset completion poset $\mathcal{D}(\mathcal{P})$ is then



J_{dense} is then given on the objects of \mathcal{P} as follows:

$$J_{dense}(x) = \{\{x\}\}$$

$$J_{dense}(y) = \{\{y\}\}$$


$$J_{dense}(z) = \{\{x\}, \{x, z\}\}$$

$$J_{dense}(w) = \{\{x, y\}, \{x, y, z\}, \{w, x, y, z\}\}.$$

Consider, for instance, why $\{x, z\}$ does not belong to $J_{dense}(w)$. While both x and z are below w , $D = \{x, z\}$ is not a sieve dense below w . In particular, given that y is an element in

the poset below w , we know that if $\{x, z\}$ were dense below w , then there must exist some element $q \in \{x, z\}$ that gets below y in the poset. But there is no such element, except for y itself, which is not in $\{x, z\}$.

Intuitively, one can picture how these four J_{dense} assignments of dense below downsets altogether supply the pieces of something that acts to “cover” the poset,



An important special case of the dense topology is the following.

Example 253 (*The atomic topology*) The atomic topology J_{atom} is specified by saying that $S \in J_{atom}(c)$ whenever the sieve S is nonempty. Then axiom (2) in the definition of a Grothendieck topology will be satisfied if we make the assumption (sometimes called the *right Ore condition*)¹⁸⁰ that any two morphisms $f : d \rightarrow c$ and $g : e \rightarrow c$, both with codomain c , can be completed to a commutative square

$$\begin{array}{ccc}
 \bullet & \dashrightarrow & d \\
 \vdots & & \downarrow f \\
 e & \xrightarrow{g} & c.
 \end{array}$$

Specialized to posets, again, the atomic topology on \mathcal{P} can be defined only when \mathcal{P} is downward directed, where in general a (nonempty) $N \subseteq \mathcal{P}$ is *downward directed* if for every $x, y \in N$ there exists a lower bound $z \in N$.¹⁸¹ In such a case, the atomic topology is defined by taking


$$J_{atom}(p) = \mathcal{D}(\downarrow p) \setminus \{\emptyset\}.$$

To get a slightly better handle on this, consider the following poset that is *not* downwards directed, for which the stability axiom for the candidate J_{atom} would then fail.¹⁸²

180. Actually, it turns out that we do not even need this assumption to define the atomic topology—we can just define it as the smallest Grothendieck topology that contains all the nonempty sieves, that is, as the intersection of all Grothendieck topologies with the property that all nonempty sieves cover. See Caramello (2012) for details.

181. In other words, the existence of (finite) meets will suffice to guarantee that a poset is downward directed.

182. This example is derived from Lindenhovius (2014), a nice resource for Grothendieck topologies applied to posets and examples of sheaves on the resulting sites.



Here, $\downarrow a \in J_{atom}(c)$, and as $b \leq c$, we ought to have that $\downarrow a \cap \downarrow b \in J_{atom}(c)$, that is, that $\emptyset \in J_{atom}(c)$. But if it is really the atomic topology, this cannot be.

Conversely, if a poset is downward directed, then one can verify that the stability axiom will necessarily hold.

If a poset \mathcal{P} is downward directed, the dense topology is nothing other than the atomic topology (which is why we referred to it as a “special case” of the dense topology). To see this, assume \mathcal{P} is downward directed, and let S be a nonempty sieve on an element p . Then, as $S \in \mathcal{D}(\downarrow p)$, there is an $x \leq p$. Now let $q \in \downarrow p$. By downward directedness, there exists an $r \leq q$, $r \leq x$, from which we have that $r \in S$, and so $q \in \uparrow S$ —altogether implying that $S \in J_{dense}(p)$. Going the other way: supposing $S \in J_{dense}(p)$, we must have $p \in \downarrow p \subseteq \uparrow S$; but this can be the case only if $S \neq \emptyset$. Thus, $S \in J_{atom}(p)$.

A few other examples of Grothendieck topologies will be given throughout this chapter. But with the notion of Grothendieck topology, and thus a site, we are already in a position to define a sheaf in greater generality.

10.2.4 Sheaves on a Site


Sheaves on a site can be defined in a very similar way to sheaves on a topological space, according to the “compatible families of sections can be uniquely patched together” model. Suppose we have a site (\mathbf{C}, J) , and of course, a presheaf is simply a functor $P : \mathbf{C}^{op} \rightarrow \mathbf{Set}$. Let S be a (J -covering) sieve of $c \in \mathbf{C}$. We first define the following notions:

Definition 254 A *matching family* (or *compatible family*) for S of elements of P is a function that assigns to each arrow f of S an element $x_f \in P(\text{dom}(f))$ in such a way that given any arrow $g \in \mathbf{C}$ with $\text{cod}(g) = \text{dom}(f)$, we have

$$P(g)(x_f) = x_{f \circ g}.$$

Observe that this makes sense, as $f \circ g$ will automatically be an element of S , by virtue of the fact that S is a sieve.

Recall the alternate formulation of a sieve $S = \{c_i \xrightarrow{f_i} c \mid i \in I\}$ on an object c of a category \mathbf{C} from proposition 243, in terms of a subcategory of the slice category \mathbf{C}/c , where the objects are just the arrows f_i of S , and the morphisms are given by arrows u making triangles of the following form commute:



In these terms, a *matching family* of elements of the presheaf P with respect to the sieve S would then be defined as a collection of elements $\{s_i \in P(c_i) \mid i \in I\}$, one for each arrow of the sieve S , such that for any arrow $u : c_i \rightarrow c_j$ in S for which the above triangle commutes, the function $F(u)$ sends s_j onto s_i .

Finally, note that since a covering sieve S of c can be construed as a subfunctor of the representable functor Y_c , when we insist on seeing the sieve S in this way, then a matching family $\{x_f\}_{f \in S}$ is just a natural transformation

$$\begin{aligned} x : S &\rightarrow P \\ f &\mapsto x_f, \end{aligned}$$

where the assignment of f to x_f is the component of the natural transformation at $\text{dom}(f) \in \mathbf{C}$. If we want to emphasize this perspective, where a matching family of a presheaf P with respect to a sieve S is regarded as a natural transformation $S \rightarrow P$, we will sometimes denote the collection of matching families of P with respect to S by $\text{Nat}(S, P)$ or $\text{Match}(S, P)$.

Definition 255 An *amalgamation* for a matching family $\{x_f\}_{f \in S}$ for S , with S a sieve on c , is a unique element $x \in P(c)$ such that

$$P(f)(x) = x_f \text{ for all } f \in S.$$

Making use of both of these notions—matching families and amalgamations—lets us define when a presheaf is a sheaf with respect to a site.

Definition 256 P is a *sheaf* (with respect to J)—or *J-sheaf*—iff each matching family $\{x_f\}_{f \in S}$ with respect to any J -covering sieve $S \in J(c)$ of any object c in \mathbf{C} has a *unique amalgamation*.

In terms of diagrams, a presheaf P will be a sheaf if the diagram

$$P(c) \xrightarrow{e} \prod_{f \in S} P(\text{dom}(f)) \xrightarrow[p]{q} \prod_{f, g \in S, \text{cod}(g) = \text{dom}(f)} P(\text{dom}(g))$$

is an equalizer for each object $c \in \mathbf{C}$ and each cover $S \in J(c)$. Here, e is the map that takes $x \in P(c)$ to $\{P(f)(x)\}_{f \in S}$, while the products range over all composable pairs f, g with $f \in S$ (thus also $f \circ g \in S$), so that $p(\{x_f\}_{f \in S})_{f, g} = x_{f \circ g}$ and $q(\{x_f\}_{f \in S})_{f, g} = P(g)(x_f)$.

Using the fact that a sieve S on c is the same thing as a subfunctor of Y_c , we get another way of presenting the definition. We just saw how a matching family taking $f \mapsto x_f$ for $f \in S$ is the same thing as specifying a natural transformation $S \rightarrow P$. Thus, that the matching family $\{x_f\}_{f \in S}$ has a unique amalgamation is equivalent to requiring that $S \rightarrow P$ can be uniquely extended as in

$$\begin{array}{ccc} S & \xrightarrow{\phi} & P \\ \downarrow & \nearrow & \\ Y_c & & \end{array}$$

And thus, a presheaf P is a sheaf iff, for every covering sieve S of c , any natural transformation $\phi : S \rightarrow P$ has a unique extension to a morphism $Y_c \rightarrow P$, in the sense that given the diagram

$$\begin{array}{ccc} S & \xrightarrow{\phi} & P \\ \downarrow & & \\ Y_c & & \end{array}$$

there is always exactly one extension of this to the following commutative diagram:

$$\begin{array}{ccc}
 S & \xrightarrow{\phi} & P \\
 \downarrow & \nearrow & \\
 Y_c & &
 \end{array}$$

But what does this mean? This is just to say that P is a sheaf precisely when for every covering sieve S on c (where this S is regarded as a functor), the inclusion $i_S : S \hookrightarrow Y_c$ induces an isomorphism

$$\text{Nat}(Y_c, P) \xrightarrow{\cong} \text{Nat}(S, P).$$

But since the natural transformations are morphisms in the presheaf category, by the Yoneda lemma this $\text{Nat}(Y_c, P) := \text{Hom}_{\text{PreSh}(\mathbf{C})}(Y_c, P)$ is the same as $P(c)$. Moreover, $\text{Nat}(S, P)$ are just the matching families of P with respect to S . Thus, the above just says that a presheaf P on \mathbf{C} is a J -sheaf iff the map

$$\begin{aligned}
 \kappa_S : P(c) &\rightarrow \text{Match}(S, P) \\
 x &\mapsto x \circ i_S
 \end{aligned}$$

is bijective for any object $c \in \mathbf{C}$ and any covering sieve $S \in J(c)$.

Definition 257 (*Alternate definition of sheaf on a site*) A sheaf on a site (\mathbf{C}, J) is a presheaf $P : \mathbf{C}^{op} \rightarrow \mathbf{Set}$ such that for every object c of \mathbf{C} and every covering sieve $S \in J(c)$, each morphism $S \rightarrow P$ in $\mathbf{Set}^{\mathbf{C}^{op}}$ has exactly one extension to a morphism $Y_c \rightarrow P$.

Covering sieves can be generated by covering families, and we saw already, in definition 236, how such covering families are used to define the notion of a basis (pretopology) of a Grothendieck topology.¹⁸³ Moreover, in terms of a coverage, the Grothendieck topology generated by a coverage will be the smallest collection of sieves containing it which also happens to be closed under the maximal and transitivity conditions (1 and 3 in the definitions). In section 10.2.1, we then saw how to provide a definition of a sheaf in terms of a basis K . We used the notions of a matching family and an amalgamation, now with respect to a basis K ; let’s gather these definitions together again.

Definition 258 Given a presheaf P , a basis (pretopology) K , and a family of morphisms $R = \{f_i : c_i \rightarrow c \mid i \in I\} \in K(c)$, we call a family of elements $\{x_i \in P(c_i)\}_{i \in I}$ a *matching family* for R iff

$$P(\pi_1)(x_i) = P(\pi_2)(x_j) \quad \forall i, j \in I,$$

where $\pi_1 : c_i \times_c c_j \rightarrow c_i$ and $\pi_2 : c_i \times_c c_j \rightarrow c_j$ are the projections from the pullback, as in

$$\begin{array}{ccc}
 c_i \times_c c_j & \xrightarrow{\pi_2} & c_j \\
 \pi_1 \downarrow & & \downarrow f_j \\
 c_i & \xrightarrow{f_i} & c.
 \end{array}$$

183. We can also define a basis (pretopology) even when \mathbf{C} does not have pullbacks. To do so, replace the second condition in 236 with the condition from the definition of coverage (definition 237).

Definition 259 Given a matching family $\{s_i\}_{i \in I}$ for $R = \{f_i : c_i \rightarrow c\}_{i \in I} \in K(c)$, we define an *amalgamation* for that matching family as an element $s \in P(c)$ such that

$$P(f_i)(s) = s_i \quad \forall i \in I.$$

A sheaf is then defined as it was in section 10.2.1. But now note that given a basis (pretopology) K , this will *generate* a Grothendieck topology J by taking

$$S \in J(c) \text{ iff } \exists R \in K(c) \text{ s.t. } R \subseteq S.$$

Often, in practice, it is more convenient to verify the sheaf condition for the sieve generated by a covering family in terms of the covering family itself (and the associated sheaf definition).

Proposition 260 (*Sheaf with respect to a basis*) Given a site (\mathbf{C}, J) , a presheaf P on \mathbf{C} will be a sheaf for J iff for any family of morphisms $R = \{f_i : c_i \rightarrow c\}_{i \in I} \in K(c)$ in the basis (pretopology) K , any matching family $\{x_i\}_{i \in I}$ has a unique amalgamation.

In the particular cases where \mathbf{C} has pullbacks, things are simplified and the above proposition can be expressed by saying that a presheaf P on a category \mathbf{C} is a sheaf for J iff for any family of morphisms $R = \{f_i : c_i \rightarrow c\}_{i \in I} \in K(c)$ in the basis, the diagram

$$P(c) \xrightarrow{e} \prod_i P(c_i) \rightrightarrows \prod_{i,j} P(c_i \times_c c_j),$$

is an equalizer.

Moreover, using a coverage and appealing to the definition 240, if a presheaf P satisfies the sheaf condition with respect to a coverage, then it satisfies the sheaf condition with respect to the Grothendieck topology generated by it.¹⁸⁴

If this is your first time seeing these definitions (and even if it isn't!), they probably remain very abstract. We will try to get a better handle on things by first specializing things to posets, and then considering how these definitions look in the case of the corresponding definition of $\mathbf{2}$ -enriched J -sheaves on a poset. But, before doing this, it is worth observing that when working with a site consisting of the atomic topology (defined in example 253), sheaves can be given an especially simple description:

Proposition 261 A presheaf P is a sheaf for the atomic topology on \mathbf{C} iff for any morphism $f : b \rightarrow c$ and any $y \in P(b)$, if $P(a \xrightarrow{g} b)(y) = P(a \xrightarrow{h} b)(y)$ for all diagrams

$$a \begin{array}{c} \xrightarrow{g} \\ \xrightarrow{h} \end{array} b \xrightarrow{f} c$$

with $f \circ g = f \circ h$, then $y = P(b \xrightarrow{f} c)(x)$ for a unique $x \in P(c)$.¹⁸⁵

10.2.5 Simpler Version with Posets

As we have said many times now, by regarding a poset \mathcal{P} as a category, there will be an arrow (and at most one arrow) $q \rightarrow p$ precisely when $q \leq p$, and we can identify indexed families S of arrows into p with a collection of *elements* “below” p in the poset. Then,

184. For further discussion of some subtleties here, having to do with bases (pretopologies) and the like, we refer the reader to Johnstone (2002, C.2.1).

185. A proof of this fact can be found in Mac Lane and Moerdijk (1994, 127).

taking a poset \mathcal{P} with J a Grothendieck topology, and $F : \mathcal{P}^{op} \rightarrow \mathbf{Set}$ a presheaf, we can define the following:

Definition 262 Let $p \in \mathcal{P}$ and $S \in J(p)$. Then a *matching family* for the cover S is a family $\{x_q \mid q \in S\} \in \prod_{q \in S} F(q)$ such that

$$F(r \leq q)(x_q) = x_r$$

for each $r, q \in S$ with $r \leq q$.

Definition 263 An element $x \in F(p)$ for which we have

$$F(q \leq p)(x) = x_q$$

for each $q \in S$ will be called an *amalgamation*.

Using the previous definitions, we can then define a J -sheaf as follows.

Definition 264 (*J-sheaf for posets*) F will be a J -sheaf provided for each $p \in \mathcal{P}$, each $S \in J(p)$, and each matching family $\{x_q\}_{q \in S}$ there is a *unique* amalgamation $x \in F(p)$.

It will help to see some examples, working with particular Grothendieck topologies.

Example 265 Take for J the indiscrete (minimal) topology on \mathcal{P} :

$$J_{ind}(p) = \{\downarrow p\}.$$

Suppose we have a presheaf $F \in \mathbf{Set}^{\mathcal{P}^{op}}$ and let $p \in \mathcal{P}$. But with J_{ind} as our topology, we know that, for each p , there is only *one* cover of p , namely that given by the principal downset $\downarrow p$ itself. Now, suppose that we have a matching family $\{x_q\}_{q \in \downarrow p}$ for $\downarrow p$. Then, since in particular p itself is in $\downarrow p$, we must have

$$F(q \leq p)(x_p) = x_q.$$

But this just means that x_p is an amalgamation of the matching family. In fact, it must be unique as well. For, suppose $x \in F(p)$ were another amalgamation. Then, by definition, we would need in particular to have

$$x = F(p \leq p)(x) = x_p.$$

Thus, our amalgamation x_p is unique.

In the previous example of a sheaf, notice how we have basically just described what it is to be a presheaf. The more or less trivial nature of the previous example—where it seems that being a sheaf with respect to the given topology is the same as being a presheaf on the original poset—is not an accident. In fact, it reflects a general result, one that also calls back to earlier special results about being able to treat presheaves on a poset as the same thing as sheaves on the poset (once it has been equipped with an appropriate topology). Namely, with the site (\mathcal{P}, J_{ind}) , where J_{ind} is the indiscrete topology, sheaves are the same thing as presheaves—that is, we have that $\mathbf{Sh}(\mathcal{P}, J_{ind})$ is precisely $\mathbf{Set}^{\mathcal{P}^{op}}$.

Example 266 Suppose \mathcal{P} is downward directed, and take the atomic topology J_{atom} on it. The reader should convince themselves that the sheaves on such a site are constant (up to isomorphism), essentially recovering what it is to be a *set*—that is, that $\mathbf{Sh}(\mathcal{P}, J_{atom}) \simeq \mathbf{Set}$.

The next, more extended, example will give us an opportunity to present a new angle on how to think about sheaves. In building to this construction, which will be developed over the course of the next few pages, we will need to appeal to the following general facts, and some new definitions.

Proposition 267 Let $\{J_i \mid i \in I\}$ be a set of Grothendieck topologies on \mathbf{C} . Then

1. $\bigcap_{i \in I} J_i$ is a topology and is the infimum of the J_i ;
2. there exists a topology $\bigvee_{i \in I} J_i$ which is the supremum of the J_i .¹⁸⁶

The Grothendieck topologies on \mathbf{C} are partially ordered by inclusion. Focusing on a poset category in particular, the set of Grothendieck topologies on a poset \mathcal{P} can be partially ordered (pointwise) as follows: for any two Grothendieck topologies J, K on \mathcal{P} , we define $J \leq K$ provided $J(p) \subseteq K(p)$ for all $p \in \mathcal{P}$. Then, the poset of Grothendieck topologies on a poset in fact forms a complete lattice, with the infimum $\bigwedge_{i \in I} J_i$ of each collection $\{J_i\}_{i \in I}$ of Grothendieck topologies on \mathcal{P} just given by pointwise intersection ($\bigwedge_{i \in I} J_i$)(p) = $\bigcap_{i \in I} J_i(p)$.

Returning to the more general account,

Proposition 268 For any presheaf F , there will exist a unique *largest* topology J_M for which F is a sheaf.¹⁸⁷

Putting this last proposition together with proposition 267 (1), we can deduce that *there is a unique largest topology for which each of a given class of presheaves is a sheaf*. One class of presheaves that are naturally of special interest is the representable functors. This result allows to ask about the largest topology for which, in particular, all the representable functors are sheaves. Such a topology is called the *canonical topology*.

Definition 269 For a category \mathbf{C} , we know that the Yoneda embedding provides us, for each object c of \mathbf{C} , with the functor $Y_c = \text{Hom}_{\mathbf{C}}(-, c)$. We define the *canonical topology* to be the finest (largest) Grothendieck topology such that every presheaf of the form Y_c is a sheaf—that is, where all the representable functors are sheaves.¹⁸⁸

Definition 270 We call a Grothendieck topology J *subcanonical* whenever every representable presheaf is itself a sheaf (with respect to J), that is, if Y_c is a J -sheaf for each $c \in \mathbf{C}$. In particular, then, subcanonical topologies are in general smaller (coarser) than the canonical topology (hence the name), which is the finest (largest) of the subcanonical topologies.

In addition to supplying us with further examples of Grothendieck topologies, we are discussing such matters in order to build toward a particular construction and result, the

186. This is lemma 0.34 in Johnstone (2014). This proof of the first item is immediate by definition; and for the second item, one just applies the first result to the set of upper bounds for the J_i .

187. This is lemma 0.35 in Johnstone (2014); a proof of this can also be found in Johnstone (2014).

188. It is important to realize that the canonical topology is not the indiscrete topology, even though representable functors are involved in the determination of both. Recall that the indiscrete topology was the one for which the *only* sieves covering an object $c \in \mathbf{C}$ are things of the form $\text{Hom}_{\mathbf{C}}(-, c)$; moreover, for this topology, sheaves were the same as presheaves. The canonical topology, by contrast, is the largest of the topologies requiring that *all* representable functors be sheaves. This means, in particular, that this may involve other functors as well—just that all the representable ones are sheaves.

full force of which may not become fully evident for a few pages. Recall from chapter 6, section 6.5, how the $\mathbf{2}$ -enriched presheaves ($\mathbf{2}$ -presheaves, for short) on a poset \mathcal{P} emerged the same thing as the poset $(\mathcal{D}(\mathcal{P}), \subseteq) := \mathbf{Down}(\mathcal{P})$ formed from the collection $\mathcal{D}(\mathcal{P})$ of all downsets of \mathcal{P} , ordered by inclusion, that is,

$$\mathbf{Down}(\mathcal{P}) \cong \mathbf{2}\text{-PreSh}(\mathcal{P}),$$

a fact discussed in the context of the $\mathbf{2}$ -enriched version of the Yoneda-embedding

$$\mathcal{P} \xrightarrow{y=\downarrow} \mathbf{2}\text{-PreSh}(\mathcal{P}) \cong \mathbf{Down}(\mathcal{P}).$$

This so-called downset completion of \mathcal{P} is not just a poset, but a *lattice* (where the meet is just set intersection and the join is union). As it turns out, $\mathbf{Down}(\mathcal{P})$ is not just a lattice, but it is a special one, where finite meets distribute over arbitrary joins (it is what is called a “frame”).

Definition 271 A *frame* is a complete lattice in which finite meets distribute over arbitrary joins, that is, in which the following distributivity law

$$a \wedge \left(\bigvee_{i \in I} b_i \right) = \bigvee_{i \in I} (a \wedge b_i)$$

holds, where I is an arbitrary indexing set, and a, b_i are elements of the lattice.

Maps between frames are given by *frame homomorphisms*, where these are functions f that both preserves all joins (including the empty join 0 or “bottom”), that is,

$$f \left(\bigvee_{i \in I} b_i \right) = \bigvee_{i \in I} f(b_i)$$

for any $\{b_i\}$ in the lattice, and preserves binary (finite) meets (including the empty meet 1 or “top”), that is,

$$f(a \wedge b) = f(a) \wedge f(b), \quad f(\top) = \top \tag{10.1}$$

for any a, b in the lattice.

Frames together with their frame homomorphisms assemble into a category, **Frm**.

While frames are rather general entities, observe that every topology X (in the usual sense of topology discussed in chapter 4) forms a frame, in the sense that the lattice of open sets $\mathcal{O}(X)$ is a complete lattice which moreover satisfies the distributivity law, with \wedge becoming set-theoretic intersection and \bigvee the union.¹⁸⁹ Moreover, it can be shown that every frame is in fact a (complete) Heyting algebra—that is, that a complete lattice satisfies the distributivity law above iff it is a Heyting algebra—and object-wise they are identical; the distinction between them has to do only with the relevant homomorphisms, as a frame homomorphism is not necessarily a Heyting algebra homomorphism (one that preserves the Heyting arrow). Thus, while “frame” and “complete Heyting algebra” mean the same thing when considering *objects* (i.e., frames, cHas), they become distinct, and must be treated as such, when we involve the morphisms.

Now recall that a meet-semilattice is a poset in which all finite (hence also empty) meets (greatest lower bounds) exist. A homomorphism f between meet-semilattices S, T

189. More explicitly, one can show that the law holds in $\mathcal{O}(X)$ by using the fact that it holds in the Boolean algebra $\mathbb{P}(X)$, and then considering that the inclusion $\mathcal{O}(X) \rightarrow \mathbb{P}(X)$ preserves finite meets and all joins.

is just a map that preserves all finite (including empty) meets (see equation 10.1). This data assembles into a category, **Meet**, the category of meet-semilattices. One can observe that the homomorphisms serving as the morphisms of this category will automatically be monotone; yet observe also that in general the converse need not hold—a monotone map between meet-semilattices need not be a homomorphism in this category (i.e., need not preserve all finite meets). Definitionally, a meet-semilattice that has *all* joins, and where these are distributive, is nothing other than a frame.

Before relating **Meet** and **Frm**, and discussing the point of all this, let us record another important definition that will be of some use to us.

Definition 272 Let \mathcal{P} be a meet-semilattice. A family $\{b_i\}_{i \in I}$ of elements in \mathcal{P} is said to be *join distributive* (or have a *distributive join*) if

- its join $\bigvee_{i \in I} b_i$ exists in \mathcal{P} , and
- for every $a \in \mathcal{P}$,

$$a \wedge \left(\bigvee_{i \in I} b_i \right) = \bigvee_{i \in I} (a \wedge b_i)$$

for I an arbitrary indexing set.¹⁹⁰


Then, a map $f : P \rightarrow Q$ between two meet-semilattices is said to *preserve distributive joins* provided whenever $\{b_i\}_{i \in I}$ is join distributive in P , then $\{f(b_i)\}_{i \in I}$ is join distributive in Q , and

$$f \left(\bigvee_{i \in I} b_i \right) = \bigvee_{i \in I} f(b_i).$$

Again, one can observe that morphisms that preserve distributive joins in the above sense are automatically monotone. If we take as our morphisms those that preserve distributive joins (and also finite meets), we get the category **Meet**_{∨*dist*}, which of course forms a subcategory of **Meet**.

Join distributivity might seem like a natural condition, but distributive lattices are in fact rather special—and, historically, it is curious to note that it was originally believed, quite wrongly, that *every* lattice was distributive. Here is a simple example of a lattice that does *not* have this property, so that you have something more concrete to hang on to as we press forward with abstractions.

Example 273 The finite lattice depicted below is *not* distributive:



190. In case it helps, for finite joins, note that this just reduces to the usual

$$a \wedge (b \vee d) = (a \wedge b) \vee (a \wedge d).$$

Proof. Observe that $r \wedge (p \vee q) = r$. However, $(r \wedge p) \vee (r \wedge q) = \perp$ (the bottom of the lattice). Thus,

$$r \wedge (p \vee q) \neq (r \wedge p) \vee (r \wedge q),$$

making the lattice fail to distributive over joins. □

Now, as a frame is a lattice with a certain structure—so that, in particular, it is a meet-semilattice (and a join-semilattice, for that matter)—by forgetting about the join-structure of a frame, we should just recover the data of a meet-semilattice. Doing so just describes a forgetful functor

$$U : \mathbf{Frm} \rightarrow \mathbf{Meet}$$

that lets us view each frame as a meet-semilattice and each frame morphism as a meet-semilattice morphism. **Frm** is of course also a subcategory—but not a *full* one—of the category **Meet** of meet-semilattices. But **Frm** in fact forms a *full* subcategory of **Meet**_{∨*dist*} (by just taking those of the meet-semilattices that are frames). As we go forward, unpacking what is going on here, we will need the following general category-theoretic notion:

Definition 274 A (full) subcategory **C** of a category **D** is said to be a *reflective subcategory* when its inclusion $\mathbf{C} \rightarrow \mathbf{D}$ has a left adjoint $L : \mathbf{D} \rightarrow \mathbf{C}$, called the *reflector*. The unit of the adjunction has components $r_D : D \rightarrow L(D)$ for $D \in \mathbf{D}$, where these are called the *reflections*.

The reflector left adjoint to the inclusion typically acts as a sort of “completion” operator. In the frame theory literature, one can find the result that **Frm** is a reflective subcategory of **Meet**.¹⁹¹ Effectively, this informs us that for every meet-semilattice S there exists a frame F_S together with an embedding $i_S : S \hookrightarrow F_S$ in **Meet**—meaning that i_S is an injection that, being a map in **Meet**, moreover preserves finite meets—where this has the universal property that for any other frame H and any other morphism $g : S \rightarrow H$ in **Meet** there will exist a unique morphism $\bar{g} : F_S \rightarrow H$ in **Frm** that extends g along i_S , in the sense that $\bar{g} \circ i_S = g$.

$$\begin{array}{ccc} S & \xrightarrow{g} & H \\ & \searrow i_S & \uparrow \bar{g} \\ & & F_S \end{array}$$

Now, given U the forgetful functor, we would expect there to be an associated “free functor” going in the other direction, left adjoint to U . This functor is in fact given by the embedding that takes each element s of a meet-semilattice to its principal downset $\downarrow s$, altogether taking meet-semilattices to the downset frame. As was already seen in chapter 6, the embedding $\downarrow (-) : \mathcal{P} \rightarrow \mathbf{Down}(\mathcal{P})$ is monotone, and the construction can be carried out for any poset \mathcal{P} . With this embedding, arbitrary infima are preserved (just as the analogous Yoneda embedding will in general *preserve all limits* that exist in **C**) and finite suprema (joins) are preserved. While this obtains for any poset, when $\mathcal{P} = S$ happens to be a meet-semilattice, this same embedding will yield a meet-semilattice morphism. To appreciate this, first observe that, trivially, $\downarrow(\top) = S$, the top of $\mathbf{D}(S)$. Moreover, taking any $a, b \in S$, for each $x \in S$ we will have $x \in \downarrow(a \wedge b)$ iff $x \leq a \wedge b$ iff $x \leq a$ and $x \leq b$ iff $x \in \downarrow(a)$

191. See, for instance, Johnstone (1986).

and $\downarrow(b)$ iff $x \in \downarrow(a) \cap \downarrow(b)$ —which shows that $\downarrow(a \wedge b) = \downarrow(a) \cap \downarrow(b)$. In fact, the embedding $\downarrow(-)$ gives us precisely the universal “frame completion,” as described in the previous paragraph.

While the embedding $\downarrow(-) : \mathcal{P} \rightarrow \mathbf{Down}(\mathcal{P})$ must preserve all meets that exist in \mathcal{P} , it need not preserve arbitrary suprema. To appreciate this latter failure, it suffices to consider the poset of natural numbers (together with $\{\infty\}$) under the natural ordering. Moreover, when \mathcal{P} again happens to be a meet-semilattice, the embedding will *not* yield a morphism of $\mathbf{Meet}_{\vee dist}$ —for, if \mathcal{P} has distributive joins, these are not necessarily going to be preserved by $\downarrow(-)$. In the usual setting of frame theory, this problem is addressed by construing \mathbf{Frm} as a *full* reflective subcategory of $\mathbf{Meet}_{\vee dist}$, which effectively informs us how for every meet-semilattice (i.e., poset for which all finite meets exist) S , there will exist a universal frame completion in $\mathbf{Meet}_{\vee dist}$ involving $i_S : S \hookrightarrow F_S$, where F_S is a frame and the map i_S is an injection that both preserves finite meets and also preserves any distributive joins that exist in S .

The idea here is to proceed by using an equivalence relation on the frame $\mathbf{Down}(\mathcal{P})$ to generate a quotient, denoted $\mathbf{Down}_{\vee dist}(\mathcal{P})$, whose elements are all those downsets that are in fact *closed under the taking of distributive joins*; then the same construction of the universal frame completion in an appropriate category can be exhibited as a situation where the embedding $\mathcal{P} \hookrightarrow \mathbf{Down}(\mathcal{P})$ factors through $\mathbf{Down}_{\vee dist}(\mathcal{P})$, as in

$$\begin{array}{ccc}
 \mathcal{P} & \xrightarrow{y=\downarrow} & \mathbf{Down}(\mathcal{P}) \\
 & \searrow & \downarrow \uparrow \\
 & & \mathbf{Down}_{\vee dist}(\mathcal{P}).
 \end{array}$$

More explicitly, given a meet-semilattice \mathcal{P} , we entertain the following binary relation R on $\mathbf{Down}(\mathcal{P})$: if $\{b_j\}_{j \in J}$ is a family with distributive join in \mathcal{P} , then $\bigcup_j \downarrow b_j$ will be regarded as R -related to $\downarrow(\bigvee_j b_j)$. Then the *congruence* (equivalence relation on the frame) generated by R can be used to get a quotient of $\mathbf{Down}(\mathcal{P})$, denoted $\mathbf{Down}_{\vee dist}(\mathcal{P})$, the elements of which are those downsets *closed under distributive joins*. In other words, whenever $\{b_i\}_{i \in I}$ is a family with distributive join in \mathcal{P} , the elements b_i of which are in $D \in \mathbf{Down}_{\vee dist}(\mathcal{P})$, then $\bigvee_i b_i$ will also be in D . This congruence allows us to give the universal frame completion in the category of meet-semilattices (with join-distributive-preserving maps) as a factorization of the embedding $\mathcal{P} \hookrightarrow \mathbf{Down}(\mathcal{P})$ through $\mathbf{Down}_{\vee dist}(\mathcal{P})$.

So we consider just those downsets that are closed under distributive joins by forming the congruence on $\mathbf{Down}(\mathcal{P})$; and as a quotient, this situation in fact assembles into a Galois connection (or adjunction)

$$\mathbf{Down}_{\vee dist}(\mathcal{P}) \xleftarrow{\perp} \mathbf{Down}(\mathcal{P}).$$

The surjection here is left adjoint to the embedding, so it preserves suprema/joins (as well as finite meets, in fact), by the LAPC result of proposition 175; while the embedding, as a right adjoint, preserves infima/meets, by RAPL.

Now we come to the purpose of all this. Just as $\mathbf{2}$ -presheaves on a poset \mathcal{P} are essentially the same as the downsets in $\mathbf{Down}(\mathcal{P})$, we can show that the $\mathbf{2}$ -sheaves on the so-called canonical topology (the largest topology for which all representable presheaves are sheaves) recover precisely those of the downsets that are closed under distributive joins.

Thus, the quotient frame $\mathbf{Down}_{\vee dist}(\mathcal{P}) \hookrightarrow \mathbf{Down}(\mathcal{P})$ exhibited above emerges as a special case of the more general $\mathbf{Sh}(\mathcal{P}, J) \hookrightarrow \mathbf{PreSh}(\mathcal{P})$, and the factorization sketched above will be seen as a special case of a more general process of sheafification of presheaves on a meet-semilattice.

That the downset embedding does not necessarily preserve distributive joins that exist in \mathcal{P} can accordingly be seen as a special case of a more general matter. In general, we know that the Yoneda embedding \mathbf{y} will preserve all *limits* that exist in \mathbf{C} . However—and this is the relevant point— \mathbf{y} need not behave as well with respect to *colimits*. To address this, the idea is that one could attempt to restrict the codomain of the embedding, the presheaf category $\mathbf{PreSh}(\mathbf{C})$, by requiring that (at least certain) colimits that happen to exist in \mathbf{C} get preserved in moving to $\mathbf{PreSh}(\mathbf{C})$. The category $\mathbf{PreSh}(\mathbf{C})$ in fact has many nice properties—so the aim of such a restriction of the codomain, then, would ultimately be to describe something that, while treating colimits in the desired fashion, did not restrict so much as to give up any of those nice properties of $\mathbf{PreSh}(\mathbf{C})$.

By making use of the notions of Grothendieck topologies and sheaves on a site, we can make better sense of the underlying situation. In very broad strokes, the idea is this: for each Grothendieck topology J on \mathbf{C} , we can consider the sheaves $\mathbf{Sh}(\mathbf{C}, J)$ on the site, and the resulting inclusion functor

$$\mathbf{Sh}(\mathbf{C}, J) \hookrightarrow \mathbf{PreSh}(\mathbf{C})$$

in fact admits a left adjoint which moreover preserves finite limits. As a left adjoint, it will automatically preserve all colimits. By bringing the Yoneda embedding $\mathbf{C} \rightarrow \mathbf{PreSh}(\mathbf{C})$ into this setup, we can consider the subcanonical topologies (where all representable presheaves are sheaves), and form the “corestriction” of the Yoneda embedding to an embedding $\mathbf{C} \rightarrow \mathbf{Sh}(\mathbf{C}, J)$, exhibiting the category \mathbf{C} as a subcategory of the category of sheaves on the site (\mathbf{C}, J) . By taking the *finest* topology for which all the representable presheaves are sheaves—that is, using the canonical topology for our J —we will be left with a universal sort of solution to the underlying problem: we will have the smallest category $\mathbf{Sh}(\mathbf{C}, J)$ into which the category \mathbf{C} may be embedded using a corestriction of the Yoneda embedding. Altogether, this can be thought of as a powerful generalization of the story just sketched of frames and “frame completions”—and such a generalization not only gives us a nice and concrete way of thinking about sheaves in terms of properties of lattices (like distributivity), but it will also allow us to put to work some earlier definitions from the context of $\mathbf{2}$ -enriched category theory and to provide an explicit description of an additional Grothendieck topology. It will be worth the reader’s while to grasp the particular connection at the heart of this generalization. But we need a few more definitions first.¹⁹²

In what follows, we will take \mathcal{P} to be a poset with finite and empty meets, i.e., a meet-semilattice. We view this as a category in the usual way: regard the underlying poset as a category \mathcal{P} , and then stipulate that it has finite limits (and a terminal object). We now

192. In his insightful paper, Isar Stubbe covers the main ideas behind the observations of the last few paragraphs and those that follow on the next few pages (see Stubbe 2005). The main observation concerning how the “frame completion” situation discussed above can be seen as a special case of the more general sheafification of presheaves on a poset equipped with the canonical Grothendieck topology (involving some sort of covers that dealt with distributivity) was something I had stumbled on independently, but Stubbe’s paper is a very nice and thorough presentation of the core relevant facts, and includes an explicit description of the canonical topology in question, so the account that follows leans on that paper and refers the reader there for further details.

describe the particular Grothendieck topology J on \mathcal{P} that will be the most relevant for our purposes, that is, the topology according to which the resulting $\mathbf{2}$ -sheaves on \mathcal{P} paired with the topology J will emerge as nothing other than those downsets of \mathcal{P} that are closed under distributive joins.

Definition 275 We define the Grothendieck topology $J_{\vee dist}$ by regarding this as a function assigning to each $p \in \mathcal{P}$ the collection of *distributive covers* of p in \mathcal{P} , where this means what you might expect, namely

$$J_{\vee dist}(p) = \{ \{p_i\}_{i \in I} \mid \{p_i\} \text{ is a family in } \mathcal{P} \text{ that has } p \text{ as its distributive join} \}.$$

Note that a family $\{p_i\}_{i \in I}$ in \mathcal{P} has distributive join provided (1) its join $\bigvee_i p_i$ exists, and (2) for every $q \leq \bigvee_i p_i$ in \mathcal{P} ,

$$q = \bigvee_i (q \wedge p_i).$$

In other words, for such a topology, $\bigvee_i p_i = p \in \mathcal{P}$ and $\forall q \in \mathcal{P}$, the (join-)distributivity law

$$q \wedge \left(\bigvee_i p_i \right) = \bigvee_i (q \wedge p_i)$$

holds.¹⁹³

Since we have been specializing matters to posets (lattices), where the relevant maps are ultimately maps between posets, the idea is now to reconstrue the sheaf definition in its $\mathbf{2}$ -enriched version, in which setting it becomes rather straightforward to discern what a sheaf is, namely the presheaves that are “ J -continuous.”

Definition 276 (2-enriched sheaf) A $\mathbf{2}$ -sheaf on a site (\mathcal{P}, J) is a presheaf $\phi : \mathcal{P}^{op} \rightarrow \mathbf{2}$ for which for every $p \in \mathcal{P}$ and every $\{p_i\}_{i \in I} \in J(p)$, if $\phi(p_i) = 1$ for all $i \in I$, then $\phi(p) = 1$.

The poset of $\mathbf{2}$ -sheaves is then denoted $\mathbf{2}\text{-Sh}(\mathcal{P}, J)$; this is a subset of $\mathbf{2}\text{-PreSh}(\mathcal{P})$.

In the isomorphism $\mathbf{2}\text{-PreSh}(\mathcal{P}) \cong \mathbf{Down}(\mathcal{P})$ first introduced in chapter 6, the isomorphism takes any $\phi \in \mathbf{2}\text{-PreSh}(\mathcal{P})$ to $\phi^{-1}(1) \in \mathbf{Down}(\mathcal{P})$; and, given a downset D of \mathcal{P} , one defines $\phi_D : \mathcal{P}^{op} \rightarrow \mathbf{2}$ by setting $\phi_D(x) = 1$ precisely when $x \in D$. For any element $p \in \mathcal{P}$, there will be a representable $\mathbf{2}$ -presheaf

$$\phi_p : \mathcal{P}^{op} \rightarrow \mathbf{2}$$

that makes $q \mapsto 1$ iff $p \leq q$. In this way, the representable presheaves act as the “characteristic maps” of the principal downsets of \mathcal{P} , and the Yoneda embedding taking each $p \mapsto \phi_p$ is essentially the same as the inclusion of the elements of the poset into its downsets. The above $\mathbf{2}$ -enriched sheaf definition just tells us that ϕ (a $\mathbf{2}$ -presheaf) is a sheaf if whenever ϕ sends every element in a cover to “true” (i.e., 1), then it sends the element covered by those elements to “true” as well. This gives us a particularly easy way of thinking about a sheaf as involving a gluing together of its local assignments.

In general, as indicated earlier, for a category \mathbf{C} and for every Grothendieck topology J , the inclusion functor $\mathbf{Sh}(\mathbf{C}, J) \hookrightarrow \mathbf{PreSh}(\mathbf{C})$ from the sheaves on the site (\mathbf{C}, J) to the

193. Once we allow that joins exist, notice that this is effectively the stability condition (condition 2) in the general definition of a Grothendieck topology. The other conditions making $J_{\vee dist}$ count as a Grothendieck topology are evident.

presheaves on \mathbf{C} has a left adjoint that preserves finite limits. Since a topology is subcanonical when all representable presheaves are sheaves, that a given topology J is subcanonical amounts to stipulating that the Yoneda embedding $\mathbf{y} : \mathbf{C} \rightarrow \mathbf{PreSh}(\mathbf{C})$ can be “co-restricted” to an embedding $\mathbf{C} \rightarrow \mathbf{Sh}(\mathbf{C}, J)$ revealing \mathbf{C} as a subcategory of $\mathbf{Sh}(\mathbf{C}, J)$. Now, with a subcanonical topology J , a meet-semilattice (or poset) \mathcal{P} can be embedded in the poset $\mathbf{2-Sh}(\mathcal{P}, J)$ by simply sending $p \in \mathcal{P}$ to $\phi_p \in \mathbf{2-Sh}(\mathcal{P}, J)$, making this embedding a “corestriction” of the Yoneda embedding. It is straightforward to see that the topology $J_{\vee dist}$ will be subcanonical, and that the following result obtains:

Proposition 277 $\mathbf{2-Sh}(\mathcal{P}, J_{\vee dist}) \cong \mathbf{Down}_{\vee dist}(\mathcal{P})$.

Thus, in taking

$$\begin{array}{ccc} \mathcal{P} & \xrightarrow{\mathbf{y}} & \mathbf{PreSh}(\mathcal{P}) \\ & \searrow & \downarrow \begin{array}{c} s \dashv \vdash i \\ \Downarrow \end{array} \\ & & \mathbf{Sh}(\mathcal{P}, J_{\vee dist}) \end{array}$$

to $\mathbf{2}$ -enriched category theory, we have

$$\begin{array}{ccccc} \mathcal{P} & \xrightarrow{\mathbf{y} \downarrow} & \mathbf{2-PreSh}(\mathcal{P}) & \xrightarrow{\cong} & \mathbf{Down}(\mathcal{P}) \\ & \searrow & \downarrow \begin{array}{c} s \dashv \vdash i \\ \Downarrow \end{array} & & \downarrow \begin{array}{c} \dashv \vdash \\ \Downarrow \end{array} \\ & & \mathbf{2-Sh}(\mathcal{P}, J_{\vee dist}) & \xrightarrow{\cong} & \mathbf{Down}_{\vee dist}(\mathcal{P}), \end{array}$$

and the desired frame completion is just given by this corestriction $\mathcal{P} \hookrightarrow \mathbf{2-Sh}(\mathcal{P}, J_{\vee dist})$, which acts to sheafify the $\mathbf{2}$ -presheaves along the topology of distributive covers on \mathcal{P} . This map will preserve not just all meets but all distributive joins that exist in \mathcal{P} .

The category $\mathbf{Sh}(\mathcal{P}, J_{\vee dist})$ is in fact the *smallest* corestriction of the Yoneda embedding for which \mathcal{P} can still be considered a subcategory. This can be seen by showing that the topology $J_{\vee dist}$ is the *finest* (largest) of the subcanonical topologies on \mathcal{P} , in turn demonstrating universality in the frame completion. The overall idea here is that the larger the topology, in general, the fewer sheaves there will be on the associated site—after all, the larger the topology, the more covers there are for which one must verify the sheaf condition is met, and so the more chances for there to be a failure to meet the condition! Since we know that, in general, the canonical topology gives us the largest of the subcanonical topologies, if our topology $J_{\vee dist}$ can be shown to be the canonical topology on \mathcal{P} , in considering the sheaves on the resulting site, we will be left with the universal solution—in the sense of the *least* (smallest) category $\mathbf{Sh}(\mathcal{P}, J)$ in which \mathcal{P} can be embedded via “corestricting” the Yoneda embedding.

Proposition 278 $J_{\vee dist}$ is the *finest* subcanonical topology on \mathcal{P} , making it the *canonical* topology.

Proof. $J_{\vee dist}$ is subcanonical. Let J' be another subcanonical topology on \mathcal{P} . But if $\{p_i\}_{i \in I} \in J'(p)$, then we must have that the join exists and $p = \bigvee_i p_i$, since if $\{p_i\}_{i \in I} \in J'(x)$, then $p_i \leq p$ for all p_i , making p an upper bound for its covering family; and if p were not the *least* upper bound, then there would have to exist a $q < p \in \mathcal{P}$ such that $p_i \leq q$ for all p_i , but since J' is subcanonical, the representable ϕ_q is a sheaf, and this entails that $\phi_q(p) = 1$ (as $\phi_q(p_i) = 1$ for all p_i and ϕ_q is a sheaf), yet this contradicts the assumption that $q < p$.

The join of such a covering family is, moreover, necessarily distributive. Thus, $J'(p) \subseteq J_{\vee dist}(p)$ for every p . By definition, this makes $J_{\vee dist}$ finer than any other subcanonical topology.¹⁹⁴ \square

This topology of distributive covers $J_{\vee dist}$ on a meet-semilattice \mathcal{P} (qua category), as the canonical Grothendieck topology, is in fact the canonical Grothendieck topology even for more general sheaves, landing in **Set** (as opposed to **2**-enriched sheaves).¹⁹⁵

We leave it to the reader to continue to ponder the impacts of thinking about sheaves and sheafification in terms of the topology of distributive covers and maps that preserve all distributive joins. Working with particularly simple lattices should help to uncover some of the significance of this special way of thinking about the more general result, and the process of obtaining a sheaf as one that engages distributivity—or, rather, as one of adding in all the joins, while keeping the meets, and then demanding that distributive joins are preserved.

10.2.6 Grothendieck Toposes

Returning to a more general approach, the sheaves on a site (\mathbf{C}, J) form a category, where the maps are the natural transformations (between presheaves). As such, this category of J -sheaves, which we denote as $\mathbf{Sh}(\mathbf{C}, J)$, forms a full subcategory of the presheaf category

$$\mathbf{Sh}(\mathbf{C}, J) \hookrightarrow \mathbf{Sets}^{\mathbf{C}^{op}}.$$

Definition 279 A *Grothendieck topos* is a category which is equivalent to the category $\mathbf{Sh}(\mathbf{C}, J)$ of sheaves on some site (\mathbf{C}, J) —that is, a category \mathbf{E} is a Grothendieck topos if there exists a site (\mathbf{C}, J) such that \mathbf{E} is equivalent to $\mathbf{Sh}(\mathbf{C}, J)$.

On account of this, Grothendieck toposes also sometimes go under the name *sheaf toposes*. Let us first look at some (mostly) trivial examples.

Example 280 In the previous section, in example 265, we already saw how if we take the site consisting of a category \mathbf{C} with the indiscrete (trivial) topology J_{ind} , then J_{ind} -sheaves are the same thing as presheaves on \mathbf{C} . Thus, $\mathbf{Sh}(\mathbf{C}, J_{ind})$ reduces to the presheaf category $\mathbf{Set}^{\mathbf{C}^{op}}$. In particular, then, for any (small) category \mathbf{C} , by taking J as the indiscrete topology, any category of the form $\mathbf{Set}^{\mathbf{C}^{op}}$ —that is, any presheaf category—will be a Grothendieck topos. This gives us a large stock of examples.

Example 281 Specializing the previous example, one can verify that $\mathbf{Set}^{\mathbf{1}^{op}} = \mathbf{Set}$ is a Grothendieck topos, equivalent to the category of sheaves on the terminal category $\mathbf{1}$.

To see this more explicitly, we can specialize this to the one-point topological space $\mathbf{1} = \{*\}$ with its unique topology, the topology in which all subsets are open. Observe that for the space $\mathbf{1}$ —or just the category with one object and one morphism, the identity morphism—there are just two opens, \emptyset and $\{*\}$ itself. For any sheaf $F \in \mathbf{Sh}(\mathbf{1})$, being a sheaf requires that for the empty set, which is covered by the empty collection of subsets, the only matching family for the empty cover of \emptyset is the trivial empty tuple, that is, $F(\emptyset) = \{()\}$; but then the only information supplied by the sheaf F is the set $F(\{*\})$. As

194. Our proof of this mimics Stubbe (2005).

195. This is the substance of theorem 1 in Stubbe (2005); a proof can be found there.

such, a set may be regarded as a sheaf on a point. $\mathbf{Set} = \mathbf{Sh}(\{*\})$ accordingly plays the role of a point in topos theory, and is thus referred to as the *punctual topos*.

Example 282 For a topological space X , the Grothendieck topos $\mathbf{Sh}(\mathcal{O}(X), J_{\mathcal{O}(X)})$ of sheaves on the site $(\mathcal{O}(X), J_{\mathcal{O}(X)})$ (as defined in example 251) is the same as the usual category $\mathbf{Sh}(X)$ of sheaves on X .

Example 283 Recall the Sierpiński space \mathbb{S} , first introduced in exercise 4 (chapter 4). This was formed from a two-point set $\{0, 1\}$, with the collection of open sets given by taking just one open point, that is, $\{\emptyset, \{1\}, \{0, 1\}\}$. This open set category can accordingly be displayed by

$$\emptyset \rightarrow \{1\} \rightarrow \{0, 1\}.$$

A presheaf P on such an open set category just amounts to the specification of

$$P(\{0, 1\}) \rightarrow P(\{1\}) \rightarrow P(\emptyset),$$

that is, three sets together with two functions between them. But recall from theorem 129 (chapter 5) that when P is a sheaf, we must have that $P(\emptyset) \cong \{1\}$, that is, that $P(\emptyset)$ is a set with just one element. This means that the category of sheaves on the Sierpiński space reduces to the data of two sets $F(\{0, 1\}), F(\{1\})$ and one function $F(\{0, 1\}) \rightarrow F(\{1\})$ between them, where $\{0, 1\}$ is the whole space and $\{1\}$ is the sole open singleton. In other words, the objects (sheaves) of this category are just functions between sets, and the morphisms are necessarily natural transformations between such objects (sheaves). Altogether, this informs us that the category of sheaves on the Sierpiński space \mathbb{S} is nothing other than the arrow category $\mathbf{Set}^{\rightarrow}$ of \mathbf{Set} .¹⁹⁶ $\mathbf{Sh}(\mathbb{S})$ is accordingly a Grothendieck topos, called the *Sierpiński topos*.

A Grothendieck topos has certain desirable properties—in particular, it is complete and cocomplete. To appreciate this, it suffices to realize that given a site (\mathbf{C}, J) , the presheaf category $\mathbf{Set}^{\mathbf{C}^{op}}$ is complete and cocomplete; one computes limits and colimits pointwise. One can show that if a category is complete and cocomplete, then so too is the reflective subcategory. $\mathbf{Sh}(\mathbf{C}, J)$ is a (full) reflective subcategory of the presheaf category $\mathbf{Set}^{\mathbf{C}^{op}}$, and hence inherits all limits and colimits. Such properties can help us rule out certain categories as being Grothendieck toposes.

Example 284 \mathbf{FinSet} is *not* a Grothendieck topos. In fact, various toposes arising from finite sets—like $\mathbf{FinSet}^{\mathbf{C}^{op}}$ (with \mathbf{C} a fixed finite category)—will not count as one either, since they lack infinite colimits.¹⁹⁷

Before moving on to some more explicit and exciting examples, we will take the opportunity to look more closely at the important adjoint “sheaf functor” associated to the inclusion of presheaves into sheaves on a site—a functor that sends each presheaf P to the sheaf that most closely “approximates” it, and exhibits the sheaf topos as a reflective subcategory.

196. See definition 20 (chapter 1) for the definition of the arrow category.

197. There are also prominent non-Grothendieck toposes that are treated in topos-theoretical approaches to nonstandard analysis. How can something be a topos but not a Grothendieck topos? We will look at that shortly.

10.2.7 The Plus Construction

In the previous two sections, we have repeatedly appealed to the fact that the inclusion functor

$$\iota : \mathbf{Sh}(\mathbf{C}, J) \hookrightarrow \mathbf{Sets}^{\mathbf{C}^{op}},$$

taking J -sheaves to their underlying presheaves, admits a left adjoint

$$\mathbf{a} : \mathbf{Set}^{\mathbf{C}^{op}} \rightarrow \mathbf{Sh}(\mathbf{C}, J),$$

which sends an arbitrary presheaf to a sheaf. This \mathbf{a} is called the *associated sheaf functor*. In other words, we will have

$$\mathrm{Hom}_{\mathbf{Sh}(\mathbf{C}, J)}(\mathbf{a}(P), F) \cong \mathrm{Hom}_{\mathbf{PreSh}(\mathbf{C})}(P, \iota(F)).$$

To show such an adjunction, what do we need? Well, first of all we would of course need to show that such a functor indeed exists, and then show that it is indeed left adjoint to the stated inclusion. While this adjunction is of some importance in its own right—and we return to it in the next chapter—this section will be mainly devoted to constructing the functor in question. Such a functor is valuable in that it effectively generalizes, to sites, the functor $\Gamma\Lambda$ introduced in equation 8.1, and gives us a canonical way of upgrading a presheaf to a sheaf.

The process involved in constructing the functor \mathbf{a} in question is carried out in two main steps, the first of which involves converting a presheaf into a *separated presheaf*, using something called the *plus construction*, denoted $(-)^+$. The plus construction turns out to be a functor $(-)^+ : \mathbf{Set}^{\mathbf{C}^{op}} \rightarrow \mathbf{Set}^{\mathbf{C}^{op}}$ in its own right, but we will also show that for any presheaf P , applying this functor to get P^+ will leave us with a presheaf that is moreover separated, that is,

$$\mathbf{PreSh}(\mathbf{C}) \xrightarrow{(-)^+} \mathbf{SepPreSh}(\mathbf{C}).$$

While a separated presheaf is not yet necessarily a sheaf, it gets us “one step closer” to a sheaf. As it turns out, it can be shown that P^+ is in fact a sheaf provided P is already separated. Thus, for any presheaf (including those that are not already separated), by applying the plus construction to the presheaf twice—this is the second of the two main steps—we will be sure to end up with a sheaf, since regardless of whether or not P was itself separated, P^+ will be separated and applying the plus construction again to this P^+ will leave us with a sheaf. This composite $((-)^+)^+$ indeed gives us the desired sheafification left adjoint functor \mathbf{a}

$$\begin{array}{ccc} & \mathbf{a} & \\ & \curvearrowright & \\ \mathbf{PreSh}(\mathbf{C}) & \xrightarrow{(-)^+} \mathbf{SepPreSh}(\mathbf{C}) & \xrightarrow{(-)^+} \mathbf{Sh}(\mathbf{C}, J). \end{array}$$

But what does an explicit description of this plus construction look like? The first thing to realize is this process takes an arbitrary presheaf and ultimately tells us something about how it behaves with respect to covering sieves, so we will need to bring covering sieves into the mix. The relevant functor can be defined in a number of (equivalent) ways, but fundamentally one can think of this construction as taking a presheaf and replacing its elements by covering sieves that are themselves equipped with matching families, adjoining unique amalgamations (gluings) to every matching family.

Given P a presheaf on a category \mathbf{C} , for each object $c \in \mathbf{C}$, we will define

$$P^+(c) = \varinjlim_{S \in \mathbf{Cov}(c)^{op}} \text{Match}(S, P),$$

where the colimit is taken over all covering sieves of c , ordered by reverse inclusion, and where $\text{Match}(S, P)$ is the set of matching families for the cover S of c . Let us unpack this.

First of all, observe that the covering sieves of c form a poset, with the natural order inherited from the inclusion order on sieves: $S \leq S'$ iff $S(d) \subseteq S'(d)$ for all $d \rightarrow c$. But we know that we can also regard covering sieves as subfunctors of the representable functor $Y_c = \text{Hom}(-, c)$, and so it is also natural to consider the poset of covering sieves ordered by reverse inclusion (refinement). This category of all covering sieves of c ordered by reverse inclusion ($J(c), \subseteq^{op}$) will be denoted $\mathbf{Cov}(c)^{op}$, where an inclusion $S \subseteq S'$ of covering sieves for c is a morphism $S' \rightarrow S$ in the category $\mathbf{Cov}(c)^{op}$. Moreover, we know from the induced condition (4) of definition 245 that the intersection of any two covering sieves on an object will be a covering sieve on that object—that is, any two covering sieves have a common refinement. In terms of $\mathbf{Cov}(c)^{op}$, the way of highlighting this refinement property is to say that $\mathbf{Cov}(c)^{op}$ is *filtered*.

Definition 285 A category \mathbf{J} is *filtered* (or *directed*) if

- it is not empty;
- for every pair of objects $j, j' \in \mathbf{J}$, there exists an object k and two morphisms $f : j \rightarrow k$ and $f' : j' \rightarrow k$ in \mathbf{J} ; and
- for every two parallel arrows $u, v : i \rightarrow j$ in \mathbf{J} , there exists an object k and an arrow $w : j \rightarrow k$ such that $w \circ u = w \circ v$.

This is a generalization of the notion of a (upward) directed poset, from order theory, where that means a poset that is inhabited (nonempty) and for which every finite subset has an upper bound.

A *filtered colimit* is a colimit of a functor $F : \mathbf{J} \rightarrow \mathbf{C}$, where \mathbf{J} is a filtered category. And in particular, a colimit over a filtered poset \mathcal{P} is the same as the colimit over a *cofinal* subset Q of that poset, where this means that for every element $p \in \mathcal{P}$, there exists an element $q \in Q$ with $p \leq q$.

Proposition 286 $\mathbf{Cov}(c)^{op}$ is *filtered*.

Proof. To see this, observe:

- Y_c is always a covering sieve, so $J(c)$ is nonempty.
- We are using the reverse order, so the requirement amounts to saying that any two covering sieves must have a common refinement, which we know is the case.
- Again reversing the arrows in the requirement, and using pullbacks, one can easily see that this is also true.

□

Thus, $P^+(c)$ will be a filtered colimit, which matters as we are taking a colimit over (finer and finer) covering sieves.

Now suppose $R = \{c_i \xrightarrow{f_i} c \mid i \in I\}$ is a covering sieve of c , for an object $c \in \mathbf{C}$. We know that a sieve can be regarded as a functor, specifically a subfunctor $i_R : R \hookrightarrow \text{Hom}_{\mathbf{C}}(-, c)$ of

the representable functor on c , which is in turn an object of $\mathbf{PreSh}(\mathbf{C})$. In other words, we have a functor

$$i_c : \mathbf{Cov}(c)^{op} \rightarrow \mathbf{PreSh}(\mathbf{C})$$

taking our covering sieve R into the presheaf category. But this inclusion in fact induces a map of natural transformations, for any presheaf F ,

$$\mathbf{Nat}(\mathbf{Hom}_{\mathbf{C}}(-, c), F) \rightarrow \mathbf{Nat}(R, F). \tag{10.2}$$

By the Yoneda lemma, each morphism in $\mathbf{Nat}(\mathbf{Hom}_{\mathbf{C}}(-, c), F)$ just corresponds to an element $x \in F(c)$. How about the other side? What is a natural transformation $\sigma : R \Rightarrow F$? Well, it will just amount to a matching family of elements $\{x_i \in F(c_i) \mid i \in I\}$! We can appreciate this by simply unpacking the definition of a natural transformation. As such, $\sigma : R \Rightarrow F$ will take each arrow $f_i : c_i \rightarrow c$ of the sieve R to an element $\sigma_{c_i}(f_i) \in F(c_i)$. Moreover, for any commutative triangle in the sieve,

$$\begin{array}{ccc} c_i & \xrightarrow{u} & c_j \\ & \searrow f_i & \swarrow f_j \\ & c & \end{array}$$

we will have that σ_{c_i} takes $f_i \mapsto x_i \in F(c_i)$, σ_{c_j} takes $f_j \mapsto x_j \in F(c_j)$, and $F(u)$ takes $x_j \mapsto x_i$. And this just informs us—viewing the sieve in terms of its slice category description—that each morphism from R to F amounts to a matching family of elements on R . Moreover, we know that F will be a sheaf precisely when the map in equation 10.2 is an isomorphism.

Altogether, we have a composite functor

$$\begin{array}{ccc} \mathbf{Cov}(c)^{op} & \xrightarrow{i_c} & \mathbf{PreSh}(\mathbf{C}) \xrightarrow{\mathbf{Nat}(-, F)} \mathbf{Set}, \\ & \searrow & \nearrow \\ & & \mathbf{Nat}(-, F) \circ i_c \end{array}$$

where this takes a covering sieve R to the set of natural transformations from R to F , where each such morphism is the same as a matching family of elements for R . For $f \in R$, this functor will then map $f \mapsto x_f$, where x_f is a component of the natural transformation $R \rightarrow F$ at $\text{dom}(f) \in \mathbf{C}$.

Now, given a presheaf P , $P^+(c)$ will be exactly the colimit of the composite functor just described. Given an $R \in \mathbf{Cov}(c)$, we will designate the associated morphism of the colimit by

$$x_R^P : \mathbf{Nat}(R, P) \rightarrow P^+(c).$$

Recall that one way of reading the sheaf condition on a presheaf P is as saying: corresponding to each matching family $\{x_f \in P(\text{dom}(f))\}_{f \in R}$ for the sieve R on c , there is a unique element $x \in P(c)$ such that $x_f = P(f)(x)$ for all $f \in R$. In this way, we can start to appreciate that this composite functor gives us a presheaf that sends a covering sieve $R \hookrightarrow \mathbf{Hom}(-, c)$ to the set $\mathbf{Nat}(R, P)$ —or, alternatively, to the set $\text{Match}(R, P)$ of matching families for P on R . The difference between this presheaf and the original $P(c)$ can be seen as supplying us with a measure of the degree to which P fails to obey the sheaf conditions at c .

Now, any morphism $f : c' \rightarrow c$ will induce (by pulling back) a functor $\mathbf{Cov}(c) \rightarrow \mathbf{Cov}(c')$. How does this work? Suppose given a morphism $f : c' \rightarrow c$ in \mathbf{C} . Then each sieve R on c determines, by pulling back,

$$\begin{array}{ccc}
 R_f & \longrightarrow & \text{Hom}(-, c') \\
 f_R \downarrow & & \downarrow \text{Hom}(-, f) \\
 R & \longrightarrow & \text{Hom}(-, c),
 \end{array}$$

a sieve on c' , where we are here denoting the sieve $f^*(R)$ by $R_f := \{g \mid f \circ g \in R\}$. And we can take

$$\begin{array}{ccc}
 \text{Nat}(R_f, P) & \longleftarrow & \text{Nat}(\text{Hom}(-, c'), P) \\
 \text{Nat}(f_R, P) \uparrow & & \uparrow \\
 \text{Nat}(R, P) & \longleftarrow & \text{Nat}(\text{Hom}(-, c), P),
 \end{array}$$

which is the same (using the Yoneda lemma) as

$$\begin{array}{ccc}
 \text{Nat}(R_f, P) & \longleftarrow & P(c') \\
 \text{Nat}(f_R, P) \uparrow & & \uparrow \\
 \text{Nat}(R, P) & \longleftarrow & P(c).
 \end{array}$$

Altogether, we have a functor

$$\begin{aligned}
 \mathbf{Cov}(c) &\rightarrow \mathbf{Cov}(c') \\
 R &\mapsto R_f,
 \end{aligned}$$

and this induces a unique morphism $P^+(c) \rightarrow P^+(c')$, from which we can deduce the functoriality of P^+ .

Let us see how this works in more detail. Given a morphism $f : c' \rightarrow c$ in \mathbf{C} as above, as we range through the R in $\mathbf{Cov}(c)$, the induced morphisms

$$\begin{array}{ccc}
 \text{Nat}(R, P) & \xrightarrow{\text{Nat}(f_R, P)} & \text{Nat}(R_f, P) & \xrightarrow{x_{R_f}^P} & P^+(c') \\
 & \searrow & & \nearrow & \\
 & & & & x_{R_f}^P \circ \text{Nat}(f_R, P)
 \end{array}$$

actually form a co-cone on the diagram defining $P^+(c)$, that is, on

$$\begin{array}{ccc}
 \text{Nat}(R, P) & \xrightarrow{\text{Nat}(f_R, P)} & \text{Nat}(R_f, P) & \xrightarrow{x_{R_f}^P} & P^+(c) \\
 & \searrow & & \nearrow & \\
 & & & & x_R^P
 \end{array}$$

But by definition, $P^+(c)$ is the colimit, and so it is the universal co-cone:

$$\begin{array}{ccc}
 \text{Nat}(R, P) & \xrightarrow{\text{Nat}(f_R, P)} & \text{Nat}(R_f, P) \\
 \searrow x_R^P & & \swarrow x_{R_f}^P \\
 & P^+(c) & \\
 \swarrow x_{R_f}^P \circ \text{Nat}(f_R, P) & \downarrow P^+(f) & \searrow x_{R_f}^P \\
 & P^+(c') &
 \end{array}$$

And this supplies us with a unique morphism $P^+(f) : P^+(c) \rightarrow P^+(c')$ making the diagram commute, in the sense that it satisfies

$$P^+(f) \circ x_R^P = x_{R_f}^P \circ \text{Nat}(f_R, P).$$

But then the functoriality of P^+ is basically immediate from this uniqueness condition in the definition of $P^+(f)$. In particular, given another morphism $g : c'' \rightarrow c'$ of \mathbf{C} , we can see that

$$\begin{aligned} P^+(f \circ g) \circ x_R^P &= x_{R_{f \circ g}}^P \circ \text{Nat}((f \circ g)_R, P) \\ &= x_{(R_f)_g}^P \circ \text{Nat}(g_R, P) \circ \text{Nat}(f_R, P) \\ &= P^+(g) \circ x_{R_f}^P \circ \text{Nat}(f_R, P) \\ &= P^+(g) \circ P^+(f) \circ x_R^P. \end{aligned}$$

But, using the uniqueness condition in the definition of $P^+(f \circ g)$ inherited from the universality of the colimit construction, this just tells us that

$$P^+(f \circ g) = P^+(g) \circ P^+(f)$$

whenever $\text{cod}(g) = \text{dom}(f)$. It is also evident that $P^+(\text{id}_c) = \text{id}_{P^+(c)}$. Altogether, this tells us that the P^+ we defined is in fact itself a presheaf (functor).

Before going forward, let us reconsider the description of $P^+(c)$ once again. Suppose we have two covers $R, S \in \mathbf{Cov}(c)$. Since finite intersections of elements of $\mathbf{Cov}(c)$ are in $\mathbf{Cov}(c)$, and since for any two covers we have a common refinement $Q \subseteq R \cap S$, with $Q \in \mathbf{Cov}(c)$, the induced maps $S \rightarrow Q$ and $R \rightarrow Q$ in our relevant category $\mathbf{Cov}(c)^{op}$

$$\begin{array}{ccc} & R & \\ & \downarrow & \\ S & \longrightarrow & Q \end{array}$$

allow us to form the pullback diagram

$$\begin{array}{ccc} S \cap R & \longrightarrow & R \\ \downarrow & & \downarrow \\ S & \longrightarrow & Q. \end{array}$$

In terms of presheaves, the diagram of covering sieves then becomes

$$\begin{array}{ccc} & \text{Match}(R, P) & \\ & \uparrow & \\ \text{Match}(S, P) & \longleftarrow & \text{Match}(Q, P), \end{array}$$

the colimit of which is then the pushout (the categorical dual of the pullback)

$$\begin{array}{ccc} \text{Match}(S, P) & \coprod_{\text{Match}(Q, P)} & \text{Match}(R, P) \longleftarrow \text{Match}(R, P) \\ & \uparrow & \uparrow \\ \text{Match}(S, P) & \longleftarrow & \text{Match}(Q, P). \end{array}$$

And this $\text{Match}(S, P) \coprod_{\text{Match}(Q, P)} \text{Match}(R, P) \cong \text{Match}(S \cap R, P)$ just amounts to equivalence classes of matching families. Specifically, two matching families will belong to the same equivalence class if their restriction along a common refinement agrees. In other words, $P^+(c)$ consists of equivalence classes $[\{x_f\}_{f \in S}]$ —where, as you would expect, the $\{x_f\}_{f \in S}$ are such that $x_f \in P(c')$ and for any $g : c'' \rightarrow c'$, we have $P(g)(x_f) = x_{f \circ g}$. On this setup, two families, $(S, \{x_f\}_{f \in S})$ and $(S', \{y_h\}_{h \in S'})$, are then taken to be equivalent— $(S, \{x_f\}_{f \in S}) \sim (S', \{y_h\}_{h \in S'})$ —precisely when there is a common refinement $Q \subseteq S \cap S'$ on which the restrictions of $\{x_f\}_{f \in S}$ and $\{y_h\}_{h \in S'}$ agree, that is, where $x_i = y_i$ for every $i \in Q$.

Now, given a morphism $l : c' \rightarrow c$ in \mathbf{C} , there will then be an induced restriction map $P^+(c) \rightarrow P^+(c')$ between the presheaves, given by

$$P^+(c) \rightarrow P^+(c')$$

$$[\{x_f\}_{f \in S}] \mapsto P^+(l)([\{x_f\}_{f \in S}]) = [\{x_{l \circ \hat{f}}\}_{\hat{f} \in l^*(S)}].$$

One can verify that this map preserves equivalence classes and is well defined. It is also straightforward to check the functor conditions, ensuring that this P^+ is itself a presheaf (functor).

Altogether, we have seen (in more than one way) how the construction of the presheaf P^+ works, and that it is a presheaf.

10.2.7.1 $(-)^+$ is a functor as well We can actually show that $(-)^+ : \mathbf{PreSh}(\mathbf{C}) \rightarrow \mathbf{PreSh}(\mathbf{C})$, the process we used to construct the presheaf P^+ from a presheaf P , is itself functorial. To see this, let us first define the functor. Again observe that a natural transformation $\alpha : F \Rightarrow G$ between two presheaves F, G will induce a natural transformation

$$\text{Nat}(-, F) \circ i_c \Rightarrow \text{Nat}(-, G) \circ i_c,$$

which will supply us with a unique factorization

$$\alpha_c^+ : F^+(c) \rightarrow G^+(c)$$

between the colimits, making is so that for each covering sieve R , we have

$$\alpha_c^+ \circ x_R^F = x_R^G \circ \text{Nat}(R, \alpha).$$

We can show that α^+ is natural as follows. If we have $f : c' \rightarrow c$ in \mathbf{C} , then using the definition of a colimit, we need only establish that

$$G^+(f) \circ \alpha_c^+ \circ x_R^F = \alpha_{c'}^+ \circ F^+(f) \circ x_R^F$$

for each $R \in \mathbf{Cov}(c)$. But the left-hand side of this equation

$$\begin{aligned} &= G^+(f) \circ x_R^G \circ \text{Nat}(R, \alpha) \\ &= x_{R_f}^G \circ \text{Nat}(f_R, G) \circ \text{Nat}(R, \alpha) \\ &= x_{R_f}^G \circ \text{Nat}(R_f, \alpha) \circ \text{Nat}(f_R, F) \\ &= \alpha_{c'}^+ \circ x_{R_f}^F \circ \text{Nat}(f_R, F) \\ &= \alpha_{c'}^+ \circ F^+(f) \circ x_R^F, \end{aligned}$$

as desired.

Having defined $(-)^+$, we can show that it is functorial. Given another natural transformation $\beta : G \Rightarrow H$, one can show—in effectively the same way as we showed P^+ to be functorial, using the uniqueness condition in the definition of $(-)^+$ applied to the composite $\beta \circ \alpha$ —that $(\beta \circ \alpha)_c^+ = \beta_c^+ \circ \alpha_c^+$. We can also construct the natural transformation

$$\eta : \text{id} \Rightarrow (-)^+,$$

by constructing, for each presheaf F , a natural transformation $\eta_F : F \Rightarrow F^+$. On objects $c \in \mathbf{C}$, set

$$\begin{aligned} \eta_F^c : F(c) &\rightarrow F^+(c) \\ s &\mapsto x_{Y_c}^F(\hat{s}), \end{aligned}$$

where $\hat{s} : Y_c \Rightarrow F$ is just the natural transformation associated to F (using the Yoneda lemma). In particular, given $c' \in \mathbf{C}$,

$$\begin{aligned} \hat{s}_{c'} : \text{Hom}(c', c) &\rightarrow F(c') \\ f &\mapsto F(f)(s), \end{aligned}$$

which completes the definition of η_F^c . Finally, one can show that η_F and η are natural (which we leave to the reader),¹⁹⁸ altogether leaving us with a functor $(-)^+$.

10.2.7.2 When is P^+ a sheaf? The point of all this, we suggested, was to generate a sheaf! Running the plus construction on an arbitrary presheaf P , will P^+ be a *sheaf*? Well, not always! But it does get us “closer” to having a sheaf, in a precise sense to be unpacked. Observe how if P^+ were a sheaf, we could show that for any matching family there exists an amalgamation, where this amalgamation is moreover unique. Thus, if something fails to give us a sheaf, what may have gone wrong is that there may not be an amalgamation for a matching family, or there may be “too many.”

The important sense in which a presheaf may be “closer” to being a sheaf is that it is *separated*, where fundamentally a separated presheaf is one that satisfies the uniqueness part of the sheaf amalgamation condition but not necessarily the existence part. Referring back to the equalizer diagram defining a sheaf in definition 256, for a general presheaf P , if P fails to be a sheaf, it will thus fail to be an equalizer for the diagram

$$P(c) \xrightarrow{e} \prod_{f \in S} P(\text{dom}(f)) \begin{matrix} \xrightarrow{p} \\ \xrightarrow{q} \end{matrix} \prod_{f, g \in S, \text{cod}(g) = \text{dom}(f)} P(\text{dom}(g)).$$

Since P^+ may fail to be a sheaf, it can thus fail to be an equalizer. However—and this is the point!—even if it fails to be a sheaf, we will see that P^+ will always be separated, where this can be captured more formally by saying that the map e above must still be injective.

Definition 287 A presheaf F is *separated for* $Q \in \mathbf{Cov}(c)$ if

$$F(c) \rightarrow \prod_{f \in Q} F(\text{dom}(f))$$

is injective, that is, if a section of F is entirely determined by its restrictions along the elements of a cover. F is *separated (period)* if it is separated for every covering sieve Q .


198. See, for instance, Borceux (1994) for details.

Observe how, in terms of the map in equation 10.2, we could have said that a presheaf P is separated if $\text{Hom}_{\text{Set}^{\text{cop}}}(Y_c, P) \rightarrow \text{Hom}_{\text{Set}^{\text{cop}}}(S, P)$ is a monomorphism for each covering sieve.

Thus, a presheaf is separated if a matching family can have *at most one* amalgamation. Before proving the main result, let us first appreciate, with an example, how the plus construction does not always leave us with a sheaf.

Example 288 Recall, from chapter 8, section 8.2, the constant presheaf P on X a topological space consisting of two elements, endowed with the discrete topology (i.e., every set is open). In other words, we have a presheaf P on X that assigns any set (or abelian group) with more than one element to each of the four open sets and the identity map to each of the nine restriction maps (five plus the four trivial self-maps), for example, $P(\{0, 1\}) = P(\{0\}) = P(\{1\}) = P(\emptyset) = \mathbb{Z}$. Observe that this presheaf is not only not a sheaf but in particular it is not separated (precisely for the empty cover of the empty set).

What is P^+ ? I claim that it is



Observe what has changed. While we have $P^+(U) = \mathbb{Z}$ whenever $U \neq \emptyset$, this construction forces $P^+(\emptyset) = \{*\}$, the terminal object. This is because, for the empty cover, there is exactly one matching family.

At this point, we should make two observations: First, observe that P^+ is still not yet a sheaf. For, let $U_1 = \{0\}$ and $U_2 = \{1\}$ be the disjoint opens of X . Together, $\{U_1, U_2\}$ covers X . But then $s_1 \in P^+(U_1)$ and $s_2 \in P^+(U_2)$ will both have, for their restriction to $P^+(\emptyset)$, the unique element of $P^+(\emptyset) = \{*\}$. But this means that $\{s_1, s_2\}$ is actually now a matching family for the cover $\{U_1, U_2\}$ of $U_1 \cup U_2 = \{0, 1\}$. Yet, as we saw in chapter 8, section 8.2, if s_1 and s_2 are taken to be different elements of \mathbb{Z} , they cannot have an amalgamation in P^+ . This brings us to the second observation: Notice how, in running this plus construction, we may introduce matching families—and thus, families that ought to have amalgamations—that were not there in the original presheaf P (and so we need not have amalgamations for them).

Again, the plus construction does not necessarily leave us with a sheaf, but it does take us in the right direction, as running this $(-)^+$ once gets us “halfway” toward a sheaf—in the precise sense that it leaves us with a separated presheaf. Here is the main result:

Proposition 289 Given P an arbitrary presheaf, P^+ is always a separated presheaf.

Proof. To show that P^+ is separated, it will suffice to show that the map e in

$$F(c) \xrightarrow{e} \prod_{h \in Q} F(\text{dom}(h))$$

is injective. In particular, consider two elements $\{x_f\}_{f \in S}$ and $\{y_g\}_{g \in R}$ in $P^+(c)$ for which $P^+(h)(\{x_f\}_{f \in S}) = P^+(h)(\{y_g\}_{g \in R})$ for all collections of morphisms $h: c' \rightarrow c$ in Q , with Q a cover of c . One must then show that the matching families $\{x_f\}_{f \in S}$ and $\{y_g\}_{g \in R}$ are in fact the same. The proof is left to the reader. \square

Thus, the plus construction takes an arbitrary presheaf and will make it a separated presheaf. But what happens if the presheaf on which we run the plus construction is *already* separated? In that case, running the plus construction will actually leave us with a sheaf!

Proposition 290 If a presheaf P is already separated, then P^+ will be a sheaf.

Proof. Again left to the reader. \square

But by putting these last two results together, we have that for *any* presheaf F , $(F^+)^+$ will be a sheaf! In other words, we can functorially upgrade any presheaf on a site to a sheaf by just two applications of the plus construction! While sometimes one can get away with just running it once, in order to take an arbitrary presheaf (separated or not) to a sheaf, it will always suffice to apply it just twice.

Using the double application of the plus functor as our definition of the sheaf functor \mathbf{a} , we can now prove the following:

Theorem 291 \mathbf{a} is left adjoint to the inclusion functor,

$$\mathbf{Sh}(\mathbf{C}, J) \xrightleftharpoons[\iota]{\mathbf{a}} \mathbf{PreSh}(\mathbf{C}).$$

Proof. The adjunction amounts to saying that to any map $P \rightarrow \iota(F)$ there will correspond a unique map $\mathbf{a}(P) \rightarrow F$. But consider the natural transformation $\eta: P \rightarrow P^+$. On components $c \in \mathbf{C}$, this is defined as

$$\begin{aligned} \eta_c : P_c &\rightarrow P^+(c) \\ x &\mapsto \eta_c(x) = \{P(f)(x) \mid f \in M_c\}, \end{aligned}$$

using M_c the maximal sieve. As we know, applying η twice gives us a map from P to $(P^+)^+ = \mathbf{a}(P)$. We would like to show that any map from P to a sheaf F will factor uniquely through the map $P \xrightarrow{\eta \circ \eta} \mathbf{a}(P)$, in the sense that

$$\begin{array}{ccccc} P & \xrightarrow{\eta} & P^+ & \xrightarrow{\eta} & (P^+)^+ \\ & \searrow & & & \vdots \\ & & & & F. \end{array}$$

But the same map η is being applied twice, so it would be enough to show that for

$$\begin{array}{ccc} P & \xrightarrow{\eta} & P^+ \\ & \searrow \beta & \vdots \alpha \\ & & F \end{array}$$

there is a unique map α making the diagram commute.

Take an element $\{x_f\}_{f \in R} \in P(c)$ for a cover R of c . Then, for a map $g : c' \rightarrow c$ in R , we will have that $\eta_{c'}(x_g) = \{P(h)(x_g) \mid h \in M_{c'}\}$. But since $\{x_f\}_{f \in R}$ is a matching family, we will have $P(g)(\{x_f\}_{f \in R}) = \{x_{g \circ f'} \mid f' \in g^*(R)\}$. As $g \in R$, $g^*(R)$ will be equal to the maximal sieve $M_{c'}$, giving $\eta_{c'}(x_g) = P(g)(\{x_f\}_{f \in R})$, which will hold for all $g \in R$.

Now, if there indeed existed the map α , as desired, it would preserve the above equality, in the sense that there would exist a unique $\alpha(\{x_f\}_{f \in R}) \in F(c)$ satisfying

$$F(g)(\alpha(\{x_f\}_{f \in R})) = \alpha([P(g)(\{x_f\}_{f \in R})]) = \alpha(\eta_{c'}(x_g)) = \beta(x_g)$$

for all $g \in R$. Since F is a sheaf and $\{\beta(x_g) \mid g \in R\}$ is a matching family, there will exist such a unique element $\alpha(\{x_f\}_{f \in R}) \in F(c)$ satisfying the above equation.

Thus, given a map $P \rightarrow \iota(F)$, this will uniquely determine a map $h : \mathbf{a}(P) \rightarrow F$ making

$$\begin{array}{ccc} P & \xrightarrow{\eta \circ \eta} & \iota \mathbf{a}(P) \\ & \searrow & \downarrow \iota(h) \\ & & \iota(F) \end{array}$$

commute, and telling us that $\eta \circ \eta$ is in fact the unit of the adjunction. □

As \mathbf{a} is a left adjoint, by LAPC, \mathbf{a} preserves colimits—thus, all small colimits will exist in $\mathbf{Sh}(\mathbf{C}, J)$, since they exist in $\mathbf{PreSh}(\mathbf{C})$. In effect, we thus ensure that the “nice” features of the presheaf category (topos) carry over to our sheaf topos. In fact, we could show that the functor $(-)^+$ also preserves finite limits of presheaves; using this, and running $(-)^+$ twice, it further follows that $\mathbf{a} : \mathbf{PreSh}(\mathbf{C}) \rightarrow \mathbf{Sh}(\mathbf{C}, J)$ preserves finite limits as well. These facts will become important in the next chapter.

Example 292 Let us return to example 288. What if we run $(-)^+$ once more, on P^+ itself? Then we should have


$$\begin{array}{ccccc} & & P^+(P^+({0, 1})) = \mathbb{Z} \oplus \mathbb{Z} & & \\ & \swarrow \pi_1 & \downarrow & \searrow \pi_2 & \\ P^+(P^+({0})) = \mathbb{Z} & & \mathbf{0} & & P^+(P^+({1})) = \mathbb{Z} \\ & \searrow \mathbf{0} & \downarrow & \swarrow \mathbf{0} & \\ & & P^+(P^+(\emptyset)) = \{*\} & & \end{array}$$

which is indeed a sheaf.

For a related example, again taking X the discrete topological space on a two-element set, readers should convince themselves that for *any* set-valued presheaf P on X , $P^+(X)$ will be equal to the pullback of

$$\begin{array}{ccc}
 & & P(\{1\}) \\
 & & \downarrow \rho_{\emptyset}^1 \\
 P(\{0\}) & \xrightarrow{\rho_{\emptyset}^0} & P(\emptyset).
 \end{array}$$

For instance, suppose the presheaf P has been given by $P(X) = \emptyset$ and $P(\{1\}) = P(\{0\}) = P(\emptyset) = A$, where A is some set with (at least) two elements and where the restriction maps $\rho_{\emptyset}^{\{1\}}$ and $\rho_{\emptyset}^{\{0\}}$ are both the identity map. In other words, we have




This is, of course, not a sheaf. To figure out how P^+ is determined, let us first consider what the plus construction does over the whole space X . As mentioned above, $P^+(X)$ should be the pullback of


$$\begin{array}{ccc}
 & & P(\{1\}) = A \\
 & & \downarrow \rho_{\emptyset}^1 = \text{id} \\
 P(\{0\}) = A & \xrightarrow{\rho_{\emptyset}^0 = \text{id}} & P(\emptyset) = A
 \end{array}$$

which set is canonically isomorphic to A itself. Thus, $P^+(X) = A$. What about for the rest of the opens? P^+ will not change anything for the two nontrivial open subsets of X . However, as in the earlier example with abelian groups, $P^+(\emptyset)$ will be forced to be $\{*\}$, the terminal object in **Set**, that is, a singleton set. This is again on account of the fact that \emptyset is covered by the empty cover; and for the empty cover (a covering by no sets at all), there is exactly one matching family, and so any two sections of $F(\emptyset)$ will agree on this cover. In more detail, recall how products over the empty index set reduce to $\{*\}$, so the relevant equalizer diagram for this cover forces $F^+(\emptyset) = \{*\}$.

Altogether, then, this leaves us with



Again, P^+ is not yet a sheaf. But if we run $(-)^+$ once more, on P^+ itself, then $P^{++}(X)$ is again the same as the pullback, now of



With this P^{++} , we will indeed have a sheaf.

Altogether, observe how, in general, P, P^+ , and $P^+(P^+)$ may all be different—and given a presheaf we may truly need to run the plus construction twice to end up with a sheaf. On the other hand, for certain presheaves, it will suffice to run the plus construction once, as the following example demonstrates.

Example 293 Recall the presheaf B of all bounded continuous functions on a topological space, for instance,

$$B(U) := \{\text{bounded continuous functions } U \rightarrow \mathbb{R}\},$$

defined on $X = \mathbb{R}$. This is clearly a presheaf, with the usual restriction maps. Moreover, it is separated. However, as you were asked to contemplate in chapter 8, section 8.2, this is not a sheaf. For, consider the open cover of \mathbb{R} given by $\{U_n = (-n, n)\}_{n \in \mathbb{N}}$. And take the identity function $f_n(x) = x$. This will give us a sequence of functions, bounded on each U_n , that moreover agree on intersections—that is, a matching family. However, the only possible amalgamation of $\{f_n\}$ to all of $X = \mathbb{R}$ is $f(x) = x$, which is unbounded. Hence, we have a matching family with no amalgamation. The moral here was that even if we have a collection of bounded functions whose domain of definition cover all of \mathbb{R} and which are well behaved on intersections, we cannot be sure that we will have a function that is *bounded* and defined on all of \mathbb{R} , when these are glued together. There are many functions that are locally bounded yet not globally bounded—hence, you cannot expect to be able to glue together locally bounded pieces, since such a gluing might leave you with an unbounded function.

If there is some matching family that does not have an amalgamation, the plus construction effectively adds it. One can turn this presheaf B of all bounded continuous functions into a sheaf—specifically, the sheaf of all continuous functions—with just one application of the plus construction to B . One way to appreciate this is to first realize that there is a map $B^+(U) \rightarrow C(U)$, since each matching family of bounded functions $f_i : U_i \rightarrow \mathbb{R}$ on a cover U_i

of U can be joined to supply a unique continuous function $f : U \rightarrow \mathbb{R}$, and equivalent matching families will yield the same function f . Going the other way, $C(U) \rightarrow B^+(U)$, every continuous function $f : U \rightarrow \mathbb{R}$ can be got from a matching family of bounded functions $f_n : U_n \rightarrow \mathbb{R}$, where $U_n = \{x \in U : |f(x)| < n\}$, and f_n is the restriction of f to such U_n .

10.3 A Few More Examples

Let us now take a break from the more abstract results and look at a few more examples of sheaves, some of which have the added benefit of exhibiting the utility of the more general notion of Grothendieck topologies.

Example 294 Recall from the n -coloring graph discussions how, in the case of the category of undirected, connected graphs, we can define a subgraph G of a graph H as a graph such that $\text{Edges}(G) \subseteq \text{Edges}(H)$, and in that context, to define a cover of a graph G it sufficed to specify a family of subgraphs $\{G_i \hookrightarrow G \mid i \in I\}$ satisfying the condition that

$$\bigcup_{i \in I} \text{Edges}(G_i) = \text{Edges}(G).$$

The present example will also work with such notions and the implied site on the category of undirected, connected graphs.¹⁹⁹

The well-known chess “ n -queens problem” presents the task of finding configurations of n queens on an $n \times n$ chessboard that satisfy the condition that none of them can attack one another, where two queens can *attack* each other, or are *in conflict*, in the usual chess sense (i.e., if they are placed on the same row, column or diagonal). We can encode this data by working over the subgraphs of the complete graph K_n with n nodes. Each node i , $1 \leq i \leq n$, is assigned a variable q_i , corresponding to a placement of a queen in the i -th column. In the concrete case of $n = 4$, we then have the set

$$F = \{F(q_1), F(q_2), F(q_3), F(q_4)\},$$

where the $F(q_i)$ are valued in the set $\{1, 2, 3, 4\}$, each of the values $u_i \in F(q_i)$ corresponding to the *row* where the queen in the i -th column has been placed. Initially, we can think of value assignments such as $(1, 2, 4, 3)$ in terms of subgraphs ordered by inclusion. For instance, for this last assignment, we take $(1, 0, 0, 0)$ to correspond to a chosen node (labeled 1); $(1, 2, 0, 0)$ for the chosen node together with the edge from 1 to the node labeled 2; $(1, 2, 4, 0)$ for the chosen node 1 and its edge to 2 together with additional edges from 2 to the node labeled 4 (and then from 4 to 1); and then $(1, 2, 4, 3)$ for the complete graph K_4 .


What a given ordered set of numbers represents then is a given *configuration* of queens, where such assignments correspond to a particular subgraph of the complete graph. For instance, the selection of 1 in the “stalk” over q_1 would represent the placement of a queen in the first row of the first column—or, in terms of the underlying subgraphs, the assignment of data to the specified node labeled 1 in a subgraph. Then, for instance, if 2 were selected in the stalk over q_3 , yielding $(1, 0, 2, 0)$, this would correspond to a configuration with a queen in the first row of the first column and another queen in the second row of the third column. But since ultimately exactly one queen will be placed in each column (up to 4 in our current case), we can restrict our attention to the subgraphs corresponding to

199. The idea for this example, and many of the details of its presentation, comes from Srinivas (1993).

sequences that proceed stepwise from q_1 through q_4 . At the i -th step, we add a queen to the i -th column, ensuring that it does not attack any of the $i - 1$ queens that have already been placed. In other words, in constructing the sheaf over such subgraphs ordered by inclusion, we first consider the restricted set

$$\{c_{ij} \mid (q_i \neq q_j) \wedge (|q_i - q_j| \neq |i - j|) \forall i, j \in \{1, 2, 3, 4\}, i \neq j\}$$

which exhibits the set of constraints describing the problem (ruling out configurations such as (1, 2, 4, 3)). In the case of selecting 1 in the stalk at q_1 , then, the constraint demands that for such a selection, only 3 or 4 could be selected in the stalk at q_2 . In turn, in the case of selecting 3, this could be extended no further; while in the case of selecting 4, we could extend one more step, assigning the value 2 to the stalk at q_3 —however, such an assignment could not then be extended over the complete graph K_4 , for the constraint prohibits any further assignment to q_4 .²⁰⁰ On the other hand, one of the (two possible) global sections is given by selecting 2 at the stalk at q_1 , then 4 at q_2 , 1 at q_3 , and 3 at q_4 . Since the selection of an integer $1, \dots, 4$ over the nodes of the subgraphs of K_4 stands for a queen placement that respects the constraint, we might represent the data of such an assignment thus:



In constructing a sheaf, we make use of the constraint and consider only a subset of the possible sequences of q_i , so that, for instance, (1, 1, 0, 0), (1, 2, 0, 0) are not allowed. More explicitly, a site for the sheaf is built on the poset of subgraphs of the complete graph K_n on n nodes, using the covering introduced at the outset. A configuration of queens


200. Recalling the discussion from chapter 9, such assignments could thus be seen to represent (strictly) local sections.

corresponding to a given graph G is then just a function that assigns to each node in G a position on an $n \times n$ board. In other words, we let $\text{Sub}(K_n)$ be the poset of subgraphs of K_n , $N(G)$ and $E(G)$ the nodes and edges of a given (sub)graph G , and *Chessboard* the set of all pairs $\langle x, y \rangle$, where x and y range between 1 and n . But since we are not interested in all possible (arbitrary) configurations of queens, we can use **valid**(c, e) a predicate describing the special situations in which the positions in a given configuration (specified by a map c) of a pair of queens connected by the edge e do not present a conflict. In this way, we can ultimately define a contravariant functor

$$\mathbf{Config} : \text{Sub}(K_n)^{op} \rightarrow \mathbf{Set}$$

that assigns the set of valid configurations of queens to each subgraph, where a valid configuration is one in which none of the queens conflict along any of the edges of the given subgraph. In a little more detail, the functor acts to assign to each subgraph of the complete graph the set of all pairwise compatible (along each component edge) arrangements of queens on the chessboard. It is easy to see how the functor acts contravariantly with respect to inclusions of subgraphs, and the sheaf conditions are met, for the covering components will in each case intersect in at least one node, and at such nodes the queen represented there will automatically prevent conflicts with subconfigurations.

Continuing to confine our attention to $n=4$, we display the configuration data of the n -queen sheaf, where the thick arrows trace out the two “global sections” (which, as we should expect, give the two solutions to the n -queen problem when $n=4$), and the two grayed-out paths highlight two of the local sections which cannot be extended to global sections:



The full sheaf diagram is fairly straightforward and is left to the reader.

Example 295 *Self-similar groups* show up in a number of areas and are of particular use in modeling various phenomena involving some self-similarity, for example, certain “self-similar” melodies in musical composition. For self-similar groups (such as the well-known Basilica group, presented below) acting on a set (or infinite rooted binary tree, etc.), and given some subset of that set, we can form an interesting general construction called the *Schreier graph*, which is basically a generalization of a Cayley graph, but whose geometry can be much more interesting than that found in a Cayley graph. Cayley graphs are like “pictures” of a group, the geometric counterpart to an otherwise deeply algebraic entity (a group).²⁰¹ Similarly, Schreier graphs display, in a picture, how the cosets of a subgroup of G act in relation to the overall group G , specifically how the cosets “tile” the overall group. For this reason, Schreier would first refer to his graphs as *Nebengruppenbilder*, basically “coset pictures.”


In abstract algebra, one often strives to obtain, for a given group G , a characterization of its subgroups (together with their properties). In terms of the associated geometrical approach, such tasks amount to finding characterizations of a given group G ’s Schreier graphs (together with their properties). In constructing the Schreier graphs with various actual self-similar groups in particular, one quickly observes that, in passing from the graph on one fixed level to the graph on the next level, some copying, rotating, and gluing are

201. More formally, a Cayley graph is a group G that has been geometrically realized as a G -torsor, where the latter is a space that G acts on *transitively* and *freely*.

involved; one moreover observes that yet another layer of self-similarity and restriction emerges in this passages of levels. These sorts of things lead one to wonder whether there is a sheaf lurking here.

One can indeed construct a sheaf here, however there are some difficulties and the construction requires some cleverness, specifically in the construction of the covering sieves for the objects of the category associated to the graphs to which Schreier graphs are to be assigned. First, we will better motivate and describe the basic objects, Schreier graphs.

Let us begin by considering a very familiar sort of object, an infinite rooted binary tree T (part of which is displayed):



An *automorphism* of T is a bijective map from the nodes to the nodes which preserves adjacencies, that is, if the nodes are joined by an edge, then the nodes they are mapped to under the map are again joined by an edge.²⁰² If a and b are automorphisms of T , the notation ab means first apply the a transformation, then b . Now, since automorphisms are by definition bijective maps, they have inverses. The *inverse* of an automorphism x is an automorphism y such that $xy = e (= yx)$, where e is the identity/unit map. Of course, if we let G be a set of automorphisms of T together with their inverses, such that for each $x, y \in G$, the products xy and yx are also in G , then we obtain a *group*. We can guarantee this closure under multiplication by taking a set of automorphisms, say a and b , together with their inverses (which for the moment, we assume, for simplicity, are just a and b themselves), and then let G be the set of all finite products of these automorphisms. We then say that G is *generated* by the set $\{a, b\}$, and a and b are called the *generators*. Whenever a group is generated by a finite set of elements, in this case automorphisms, we call it a *finitely generated group*.

We could also define automorphisms via *rules*, for instance, for w an arbitrary binary string, we might have:

$$\begin{aligned} a(0w) &= 1.e(w), & a(1w) &= 0.e(w); \\ b(0w) &= 0.e(w), & b(1w) &= 1.a(w), \end{aligned}$$


where $e(w)$ means apply the identity map (“do nothing”) to the suffixed string w . Products of a and b can also be expressed with rules, for example, $ba(0w) = a(0.e(w)) = 1.ee(w)$. Observe that in the above rules, the rules for a uses only e while the rules for b uses a and e . In other words, the set $\{a, b, e\}$ of automorphisms is described by a set of self-referential or *self-similar* rules.

202. The definition is far more general and works for any graph Γ , but we confine our attention to binary trees like T .

Definition 296 Let G be a group of automorphisms of T .²⁰³ Then G is called a *self-similar group* if for each $g \in G$, each $x \in \{0, 1\}$, and each binary string w , there is a $y \in \{0, 1\}$ and an $h \in G$ such that

$$g(xw) = y.h(w)$$

Another particularly suggestive way of describing automorphisms of T is via *automata*. For instance, the automaton for the $\{a, b, e\}$ group from above looks like:



The idea then is that if we feed, for instance, the string 10110 in to the automaton, starting at b , then b will send 10110 along to a , keeping the 1 in the first position the same; then a will switch the second letter 0 to 1 and send the new string along to e ; e in turn will preserve the subsequent three letters as they were, ultimately spitting out the string 11110.

Now, in general, whenever we have a group G generated by a finite set \mathcal{G} , acting on a set X , and given M some subset of X , there is a Schreier graph. For concreteness, in illustrating this notion, we confine our attention to X being the nodes of a tree T , M the set of nodes at some fixed level k of T , and G as a self-similar group acting on T . Roughly, the idea is that for each element of the subset M of T , we first draw a node labeled by this element. We then connect the nodes m_i, m_j with a directed edge assigned some label $s \in \mathcal{G}$ whenever $m_j = sm_i$. This graph is obviously *connected* if for any two nodes in M there is some group element (which need not be a single generator, but might be some combination or finite product of generators from \mathcal{G}) which carries us from m_i to m_j . In this case, we say that G acts transitively on M .


In more generality, given a group G with generating set \mathcal{G} , and any subgroup H , the set of left cosets G/H is a set on which G will act transitively, which entails that we can construct Schreier graphs for G acting on G/H . This will result in a graph Γ whose vertex set is G/H and in which two cosets Hg and Hg' are connected with a directed edge labelled with a generator $a \in \mathcal{G}$ iff $Hga = Hg'$. In the special case where H is the trivial subgroup, the Schreier graph coincides with the usual Cayley graph. If G is a self-similar group acting transitively on each level of T , let H be the subgroup of G containing all those elements which fix a node at level k . Then the Schreier graph for G acting on this G/H will be the graph for G acting on T when M is held to be the set of nodes at level k . More formally, a *Schreier graph* is a (connected and rooted) $2n$ -regular (each of whose vertices has degree $2n$) graph Γ the edges of which are colored (using n different colors), and where for each vertex there is exactly one incoming and one outgoing edge of a given color attached to it.

203. This definition could be extended to groups of more general automorphisms, in particular to rooted n -ary trees or graphs.


We give a concrete illustration of this with the so-called Basilica group B . Let B be the group finitely generated by the set of automorphisms consisting of a , b , and e , acting on the rooted binary tree T . These generators are described by the following self-similar rules:

$$\begin{aligned} a(0w) &= 1.b(w), & a(1w) &= 0.e(w); \\ b(0w) &= 0.a(w), & b(1w) &= 1.e(w). \end{aligned}$$


The automaton encoding these rules is given below:




We know then that given the group B acting on the set of nodes of T an infinite rooted binary tree, we can look at the subset $M \subseteq T$ where M is the set of nodes at some fixed level k . For concreteness, let us consider $k \leq 5$. Then, for each element of M , we will draw a node. If $m_i = sm_j$ for some $s \in \mathcal{B}$ (where \mathcal{B} is the generating set of the group acting on M), that is, if from any node we can get to any other node via some (combination of) action in the group, we say that B acts transitively on M or that the induced Schreier graph is *connected* at each level k . Concretely, for our Basilica group B acting on T , and M as just described, we have:




$k = 1$. To get the next level, duplicate, flip, and glue along dashed arrow




$k = 2$. To get back to level 1, reconnect the duplicates along dashed arrows



$k = 3$. To get back to level 2, reconnect the duplicates along dashed arrows



$k = 4$



The sorts of manipulations involved here prompt the notion that such Schreier structures over a graph might form a sheaf. However, while assigning Schreier graphs to the subgraphs of a given graph (or 1-simplicial complex), endowed with discrete topology (every subset open), describes a *presheaf*, it does not give us a sheaf. Yet, by using the more general notion of a Grothendieck topology, we can in fact describe things so as to produce a sheaf on an appropriate site. We will sketch how this works.²⁰⁴

Before developing this, we mention a few important features of Schreier graphs. Schreier graphs are graphs provisioned with a distinguished base point—the *root*. Moreover, Schreier graphs are *regular*, that is, each vertex has the same degree. For finite generating sets, the degree of vertices in a Schreier graph is always even. Since the action of G on G/H is transitive, Schreier graphs are connected. Furthermore, the quotient map construction $p : G \rightarrow G/H$ that takes the cosets of H and collapses them will induce a covering of Schreier graphs $p : G \rightarrow \Gamma$. Finally, if we denote by $L(G)$ the lattice of subgroups of G , and by $\Lambda(G)$ the corresponding space of Schreier graphs of G , we can observe that the space $L(G)$ is a lattice ordered by inclusion, while $\Lambda(G)$ is a lattice ordered by coverings. This latter claim means that $\Gamma \leq \Delta$ if there exists a covering map $p : \Gamma \rightarrow \Delta$ that takes the root of Γ to the root of Δ . Then, the map $f : L(G) \rightarrow \Lambda(G)$ that sends a subgroup $H \leq G$ to its Schreier graph is a lattice isomorphism sending an ordered pair $H \leq H'$ to the covering map $p : \Gamma \rightarrow \Gamma'$, where Γ and Γ' are the Schreier graphs of H and H' . The inverse map $f^{-1} : \Lambda(G) \rightarrow L(G)$ just takes Γ a Schreier graph of G to the subgroup H , which is the stabilizer of Γ under the action of G (i.e., the set of elements of G which fix the root of Γ).

204. I discovered a construction of a Schreier graph sheaf in an earlier draft of this book, but later discovered Cannizzo (2014) had done something similar. Moreover, I subsequently found an error in my own earlier approach, so the following account, including many details, largely derives from Cannizzo (2014).


In short, a Schreier graph is a connected and rooted $2n$ -regular graph Γ whose edges are labeled by the generators of a group—we can think of the edges as coming in n different “colors,” after assigning a color to each of the generators in the generating set \mathcal{G} and painting the edges of Γ accordingly—and that is such that, for each vertex of Γ , there is exactly one incoming and one outgoing edge of a given color attached to it.²⁰⁵ A Schreier structure Σ on Γ , a $2n$ -regular rooted graph, is just a labeling of its edges by the generators of the free group $\mathbb{F}_n = \langle a_1, \dots, a_n \rangle$, whereby Γ is made into a Schreier graph. In other words, a Schreier structure is just a map

$$\Sigma : E_0(\Gamma) \rightarrow A,$$


where $E_0(\Gamma)$ specifies a choice of orientation attached to each edge $(x, y) \in \Gamma$, and where for each $x \in \Gamma$ and each $1 \leq i \leq n$, there is, attached to x , exactly one incoming edge labeled with a_i (from A) and one outgoing edge labeled with a_i . Finally, we can define the space of Schreier graphs $\Lambda(G)$ of a given (finitely generated) group G by defining $\Lambda_r(G)$ as the set of isomorphism classes of r -neighborhoods (defined below) centered at the roots of the Schreier graphs in $\Lambda(G)$, and then setting $\Lambda(G) = \varprojlim \Lambda_r(G)$. We can, moreover, identify the space of Schreier graphs $\Lambda(G)$ and the lattice of subgroups $\mathcal{L}(G)$ of G .

It has been known for some time that every $2n$ -regular graph admits a Schreier structure. Suppose we have Γ some $2n$ -regular graph. We then want to describe the space of Schreier structures $Schr(\Gamma)$ over Γ . We can show that the association of Schreier structures to subgraphs of a given graph Γ supplies the data of a presheaf. However, by assigning the Schreier graphs to the subgraphs of a graph endowed with the discrete topology (in which every subset is open), or realized as a simplicial complex, one cannot make this a sheaf.

Suppose you place the discrete topology on a graph Γ , that is, the natural metric topology. Consider the graph Γ



The two edges $(v_0 v_1)$ and $(v_1 v_2)$ together form an open cover of the entire graph G . Their intersection is the sole vertex v_1 . We can then assign a Schreier structure to the edge $(v_0 v_1)$ by first directing the edge from v_0 to v_1 and then labeling this with a generator a ; similarly, we can assign a Schreier structure to the edge $(v_1 v_2)$ by first directing the edge from v_2 to v_1 and then also labeling this with the generator a , so that the assignments are to



On the overlap $v_1 = (v_0 v_1) \cap (v_1 v_2)$, the assigned Schreier structures trivially agree. However, the sheaf gluing axiom then requires that there exists a gluing $s \in \Lambda(\Gamma)$ such that its restriction to each edge equals a . But both edges are directed towards v_1 and have the same label, so the gluing condition cannot be satisfied and we do not have a sheaf.

But by reworking what it means for a family of subsets of a graph to serve as an open cover, specifically using a Grothendieck topology, we can arrive at a viable description of a sheaf of Schreier structures on such graphs. We begin by considering only sets that are

205. Technically, this graph will then be a Schreier graph of the free group generated by the group’s generating set, or its Cayley graph.

r -neighborhoods, where this means sets $U \subseteq \Gamma$ of the form

$$U = \{x \in \Gamma \mid \rho(K, x) \leq r\},$$

where K is an arbitrary subset of Γ and ρ the standard metric on graphs. We say then that a subset $K \subseteq U$ of an r -neighborhood $U \subseteq \Gamma$ in a graph Γ is an r -basis of U provided the r -neighborhood of K (inside of Γ) is equal to U . This lets us define a cover as follows:

Definition 297 Let Γ be a graph with the usual graph metric. A system of r -neighborhoods $\{U_i\}_{i \in I}$ is said to be a *cover* of Γ provided there exist bases $K_i \subseteq U_i$ covering Γ in the ordinary sense, that is,

$$\bigcup_{i \in I} K_i = \Gamma.$$

While the open sets of covers thus defined will overlap in a nice way, rectifying our previous problem, allowing the assignment of Schreier structures to sets to be a sheaf, there is an issue in that the defined collection of r -neighborhoods is not actually a *topology*, since it is not closed under intersections. This is where the fully abstract notion of a Grothendieck topology comes to the rescue. This will, in turn, enable us to define a proper sheaf of Schreier structures on a graph.

For Γ a graph, we proceed by defining a category \mathbf{C}_Γ associated to Γ , that has

- objects: the 1-neighborhoods in Γ ; and
- morphisms: inclusions.

Let us now define a Grothendieck topology J_1 on \mathbf{C}_Γ thus: for U an object of \mathbf{C}_Γ , we let $J_1(U)$, the collection of covering sieves of U , be the collection of all sets of inclusions $\{U_i \hookrightarrow U\}_{i \in I}$ subject to the requirement that each U_i is a 2-neighborhood and that the system $\{U_i\}_{i \in I}$ covers U in the sense of the definition above (taking $r=2$). With the resulting Grothendieck topology, the functor assigning to a given 1-neighborhood U the set of Schreier structures over U of Γ is in fact a sheaf.

Theorem 298 Let Γ be a graph and let (\mathbf{C}_Γ, J) be its associated site. Then the contravariant functor $Schr : \mathbf{C}_\Gamma \rightarrow \mathbf{Set}$ that assigns to a 1-neighborhood U in \mathbf{C}_Γ the set of Schreier structures $Schr(U)$ over U is in fact a sheaf.²⁰⁶

Proof. The assignment of Schreier structures to a given graph is a contravariant assignment, so the functor $Schr$ is a presheaf. Take U an object, that is, a 1-neighborhood, in $\mathbf{C}(\Gamma)$, and $S \in J(U)$ a covering sieve that consists of inclusions $\{U_i \hookrightarrow U\}_{i \in I}$. We can check that it is a sheaf by verifying the two usual sheaf conditions.

First, uniqueness: suppose $\Sigma, \Sigma' \in Schr(U)$ are two Schreier structures on U such that their restrictions to U_i agree for all $i \in I$. Suppose, moreover, that $\Sigma \neq \Sigma'$. This inequality means that there will have to exist an edge $(xy) \subseteq U$ that has different labels in Σ and Σ' . The definition of covering sieves stipulates that the point x belongs to the basis of some U_i , and since U_i is a neighborhood of this basis, it must also contain the point y and thus also the edge (x, y) connecting them. But this is a contradiction, so we get that, provided there are such gluings, there can only be one.

206. This theorem, and its proof, follows Cannizzo (2014) very closely.

Second, gluing: suppose that $\{\Sigma_i \in \text{Schr}(U_i)\}_{i \in I}$ is a system of Schreier structures with the property that, for all $i, j \in I$, the restrictions of Σ_i and Σ_j to the fibered product $U_i \times_U U_j$ agree. Take an arbitrary $x \in U$. If the set of all edges incident with x (called $\text{star}(x)$) is contained in the 1-basis of one of the U_i , then the Schreier structure of this $\text{star}(x)$ that comes from the Σ_i can be extended to the entire neighborhood U . On the other hand, if x lies on the boundary of the 1-basis K_i of some U_i , making $\text{star}(x)$ contain a point y not in K_i , then this point y will have to be contained in the 1-basis K_j of some U_j . Such an edge (x, y) will then belong to the fibered product $U_i \times_U U_j$, and this fibered product is itself a 1-neighborhood.²⁰⁷ Thus, a Schreier structure can be extended to all of U . \square

Fundamentally, the theorem says that the functor assigning Schreier structures to subgraphs of a given graph can be regarded as a sheaf, provided the subgraph is furnished with an r -neighborhood structure, and provided one then defines an appropriate Grothendieck topology.

The next two constructions are left to the reader.

Example 299 In general, for a topological group G equipped with the prodiscrete topology, the category of continuous G -sets or $\mathbf{Cont}(G)$ —having for objects the continuous actions of G on discrete sets and for morphisms the equivariant maps—forms a Grothendieck topos; this category is in fact equivalent to the topos of sheaves $\mathbf{Sh}(\mathbf{C}, J_{At})$, where \mathbf{C} consists of the (nonempty) transitive actions of $\mathbf{Cont}(G)$ and J_{At} is the atomic topology on this subcategory.

A cellular automaton (CA) is a sort of machine that takes in some input and produces some output.²⁰⁸ They are excellent devices for modeling complex behaviors with minimal building blocks. A CA consists, roughly, of the following data:

- An *alphabet* A : the set of symbols we allow, for example, $\{0, 1\}$, $\{\text{blue, orange, white}\}$, and so on; individual elements of this alphabet set are called *states* or *colors*.
- A *group* G , called the *universe*, typically displayed in the form of an ordered grid of cells (i.e., tessellation with some basic shape, e.g., squares); elements of the universe are called *cells*.
- A notion of *neighborhood*, for example, the eight cells immediately surrounding a given cell in a 2-D tessellation with squares;
- A *local transition function* or *rule* that provides instructions for moving from one assignment of elements of the alphabet to the lattice of cells to another; these instructions for producing the output state of a given cell are given in terms of the input states of a finite number of a cell’s neighboring cells, which might include the cell itself.

We consider A^G , the set consisting of all maps from the universe G to the alphabet set A :

$$A^G = \prod_{g \in G} A = \{x: G \rightarrow A\}.$$

207. For $U \hookrightarrow W$ and $V \hookrightarrow W$ two morphisms in \mathbf{C}_Γ , the fibered product $U \times_W V$ is specifically the 1-neighborhood of the intersection of the maximal bases of U and of V , that is, $U_1(U_{\max} \cap V_{\max})$. See Cannizzo (2014) for details.

208. Much of the basic background theory concerning CA’s can be found in Ceccherini-Silberstein and Coornaert (2010), a book that would serve as an excellent reference for anyone desiring to learn more about cellular automata.

Elements of this set A^G are called *configurations*. In other words, a configuration is a way of attaching an element of the alphabet to each member of the underlying group (i.e., to each cell in the universe), where each cell is assigned one state at a time. Informally, a configuration is a specific coloring of all the cells of the universe. There is a natural action of the group on the set of configurations, called the *shift action*. In brief, a CA is a self-mapping of the set of configurations given by a system of local rules which commutes with this shift action. Since G acts continuously on the set A^G (equipped with the prodiscrete topology)²⁰⁹ it would appear that we are dealing with the category of continuous G -sets, $\mathbf{Cont}(G)$, whose objects are the continuous actions of G on the discrete sets and whose arrows are the equivariant maps between them; one could then construct a Grothendieck topos here by constructing a site out of the open subgroups and using the atomic topology on the full subcategory of $\mathbf{Cont}(G)$ that consists of the nonempty transitive actions. We leave readers to explore this on their own.

Example 300 For those who know about combinatorial game theory, try to construct a sheaf out of data of the *thermographs* of games.²¹⁰

10.4 Philosophical Pass: The Idea of Toposes

Box 10.1

The Idea of Toposes

This chapter offered a first look at Grothendieck toposes, where these are categories equivalent to a category of sheaves on some site; moreover, every Grothendieck topos can be shown to arise in this way. When Grothendieck first introduced the essential concepts, it was in order to be able to treat cohomology within algebraic geometry—in this manner, toposes were able to play an important role in problems with a more arithmetic bent. But ultimately, this powerful concept of a Grothendieck topos allows us to deal with situations or theories wherein seemingly topological notions would appear to be useful or relevant, but where the ordinary topological spaces are not present. Grothendieck toposes accordingly allow us to study a variety of “space-like” situations, to regard space as something ubiquitous throughout mathematics, and even to have a unified framework we can use to study arithmetic and geometry. As Grothendieck stressed (see promenade 13 of Grothendieck 1986 in particular), the concept of a Grothendieck topos in a sense joins together the continuous and structures that would appear to be thoroughly algebraic and discrete.

From one perspective, the story of Grothendieck toposes can be seen as a decisive step in the story of the metamorphosis of the concept of space. The associated notion of a Grothendieck topology in terms of systems of covers is an important purification of what is arguably fundamental to the spatiality of a space—and the more general notion allows us to capture this even in situations quite distant from literal topological spaces. The essence of what is “spatial”

209. The prodiscrete topology is where A^G has the product topology, each factor A of A^G given the discrete topology (all subsets of A open).

210. See Siegel (2013), Berlekamp (1982), Berlekamp and Wolfe (2012), or Guy, Conway, and Berlekamp (1982) for details on combinatorial game theory and the notion of “temperature” and thermographs; the reader can also visualize some of these ideas, and perform some computations that would be too difficult to do by hand, with the open-source program *CGSuite*.

about space is arguably to be found wherever there is a structure where localization makes sense. Grothendieck topologies, and the sites built out of them, allow us to speak about this. When we can define something on covers and check agreement on intersections of such covers, it makes sense to speak of localization. And, fundamentally, the idea of covering sieves is one that enables us to examine the relations between the local and global features of some system of objects. When we can uniquely glue data assigned to covers, we can talk about sheaves on such a system. By considering all the sheaves on a site—and so a Grothendieck topos—we are recapturing structures that behave like sheaves on a topological space, even in situations where there may not be a literal topological space (and so, where there was not already a notion of what sorts of assignments of data to that space are “sheaf-like”).

In this framework, the characteristic relations between local and global features of some system of objects emerge as more fundamental than any sort of *points*; the concept of a site accordingly starts to tease out, in a particularly powerful fashion, this idea that one can examine something spatial, and things defined on such a space, without looking at its *points*. The focus shifts away from points and behavior on points to how a localization or covering system can be made to define a “variable structure” that varies over that system of covers.

The usual (point-set) notion of a topology was ultimately designed to capture and understand two things: locally defined phenomena and continuous transformations. If Grothendieck’s notion of a site, built out of a transformed notion of a topology, supplies us with a more general notion of a local system, consideration of the sheaves on a site provides some sort of continuous or continuously variable perspective on a category that may in principle be quite “nontopological.” The notion of *geometric morphisms* between toposes, which we cover in the next chapter, further extends the generalized notion of continuous transformations. With the notion of a Grothendieck topos, we ultimately have a framework in which the (discrete) combinatorial diagrammatic approach (via categories) is paired harmoniously with (continuous) concepts native to topological spaces.

Furthermore, the introduction of Grothendieck toposes can also be seen as motivated from observations that a number of important properties of topological spaces themselves can be reformulated as certain invariant properties of their associated category of sheaves. In certain cases, as we have already seen, a topological space X can even be recovered from its associated category of sheaves $\mathbf{Sh}(X)$; and regarding a topological space in terms of its associated category of sheaves has various merits (since the latter has a rather rich categorical structure, as we will better appreciate in the next chapter).

It is often observed how attempting to obtain a topos appropriate for a particular sort of math by constructing the domain of variation (via a site), and then considering the category of sheaves over that site, gives rise to the perspective of *variable set theory*, wherein the standard constant universe of sets is replaced by something like a universe of continuously variable sets. As Johnstone (2014, xvii) claims, this is in fact the very essence of the topos-theoretic view of things:

it consists in the rejection of the idea that there is a fixed universe of “constant” sets within which mathematics can and should be developed, and the recognition that the notion of “variable structure” may be more conveniently handled within a universe of *continuously variable* sets than by the method, traditional since the rise of abstract set theory, of considering separately a domain of variation (i.e., a topological space) and a succession of constant structures attached to the points of this domain. In the words of Lawvere, “Every notion of constancy is relative, being derived perceptually or conceptually as a limiting

case of variation, and the undisputed value of such notions in clarifying variation is always limited by that origin.”

When dealing with an ordinary topological space X , the variable set (indexed by the opens of X) just yields a sheaf of sets on X . Instead of using a topological space X , we can insist on the perspective of X as a parameter space, yielding the category of variable sets with parameters in X . In this way, the space withdraws into the background, providing the form of variability or cohesion for the objects in the foreground. (This perspective ultimately leads to variable topological algebra, which has deep connections with arithmetic.) By enforcing different axioms over and above those specified by topos theory in general (on which more below), one can obtain different “universes” of variable sets within which to do certain kinds of math.

Grothendieck toposes can also be seen as *unifying* in their ability to house disparate mathematical constructions, useful not only for studying relationships between different mathematical theories, but also for transferring information and ideas across different situations. (Readers are encouraged to consult the work of Olivia Caramello, who has developed the idea of Grothendieck toposes as “bridges” for unifying different theories and transferring information between them.) There can also be a number of different relevant sites for a given Grothendieck topos, which sometimes appears to amount to letting us look at the same theory or situation in different ways.

The definition of a Grothendieck topos can be applied to any category of sheaves associated to a site, and in particular we saw that this notion of topos encompasses any presheaf (or variable set) category $\mathbf{Set}^{C^{op}}$ (just set J as the minimal topology). In particular, we looked at $\mathbf{Set}^{op} = \mathbf{Set}$. But the notion of a Grothendieck topos in fact depends on the assumed model of set theory, and this suggests the need for even greater generality. Lawvere and Tierney first introduced the notion of an *elementary topos* in part to provide a characterization of categories that resemble Grothendieck toposes from one perspective, but that can be defined strictly by elementary axioms that are independent of set theory. In the next chapter, we develop the story, pursuing some of these further generalizations and giving some examples that put the underlying ideas to work.

This is a section of [doi:10.7551/mitpress/12581.001.0001](https://doi.org/10.7551/mitpress/12581.001.0001)

Sheaf Theory through Examples

By: Daniel Rosiak

Citation:

Sheaf Theory through Examples

By: Daniel Rosiak


DOI: [10.7551/mitpress/12581.001.0001](https://doi.org/10.7551/mitpress/12581.001.0001)

ISBN (electronic): 9780262370424

Publisher: The MIT Press

Published: 2022

The open access edition of this book was made possible by generous funding and support from Arcadia – a charitable fund of Lisbet Rausing and Peter Baldwin, and MIT Press Direct to Open



The MIT Press

© 2022 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.

Subject to such license, all rights are reserved.



The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in LaTeX by the author.

Library of Congress Cataloging-in-Publication Data

Names: Rosiak, Daniel, author.

Title: Sheaf theory through examples / Daniel Rosiak.

Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.

Identifiers: LCCN 2021058949 | ISBN 9780262542159 (paperback)

Subjects: LCSH: Sheaf theory.

Classification: LCC QA612.36 .R67 2022 | DDC 514/.224—dc23/eng20220521

LC record available at <https://lccn.loc.gov/2021058949>