

Notes on topos theory

Jon Sterling

Carnegie Mellon University

In the subjectivization of mathematical objects, the activity of a scientist centers on the artificial delineation of their characteristics into *definitions* and *theorems*. Depending on the ends, different delineations will be preferred; in these notes, we prefer to work with concise and memorable definitions constituted from a common body of semantically rich building blocks, and treat alternative characterizations as theorems.

Other texts To learn toposes and sheaves thoroughly, the reader is directed to study Mac Lane and Moerdijk's excellent and readable *Sheaves in Geometry and Logic* [8]; also recommended as a reference is the *Stacks Project* [13]. These notes serve only as a *supplement* to the existing material.

Acknowledgments I am grateful to Jonas Frey, Pieter Hofstra, Ulrik Buchholtz, Bas Spitters and many others for explaining aspects of category theory and topos theory to me, and for putting up with my ignorance. All the errors in these notes are mine alone.

1 Toposes for concepts in motion

Do mathematical *concepts* vary over time and space? This question is the fulcrum on which the contradictions between the competing ideologies of mathematics rest. Let us review their answers:

Platonism No.

Constructivism Maybe.

Intuitionism Necessarily, but space is just an abstraction of time.

Vulgar constructivism No.¹

Brouwer's radical intuitionism was the first conceptualization of mathematical activity which took a *positive* position on this question; the incompatibility of intuitionism with classical mathematics amounts essentially to the fact that they take opposite positions as to the existence of mathematical objects varying over time.

¹I mean, of course, the Markov school.

Constructivism, as exemplified by Bishop [2] takes a more moderate position: we can neither confirm nor deny the variable character of mathematical concepts. In this way, mathematics in Bishop’s sense is simultaneously the mathematics of *all* forms of variation, including the chaotic (classical) form.

This dispute has been partly trivialized under the unifying perspective of *toposes*,² which allow the scientific study of mathematical systems and their relationships, including Platonism (the *category of sets*), constructivism (the *free topos*), intuitionism (the *topos of sheaves over the universal spread*)³ and vulgar constructivism (the *effective topos*).

Toposes have both a *geometric* and a *logical* character; the geometric aspect was the first to be developed, in the form of *Grothendieck toposes*, which are universes of sets which vary continuously over some (generalized) form of space. More generally, the *logical* aspect of topos theory is emphasized in Lawvere and Tierney’s notion of an elementary topos, an abstract and axiomatic generalization of Grothendieck’s concept.

These two aspects of topos theory go hand-in-hand: whilst the laws of an elementary topos are often justified by appealing to their realization in a Grothendieck topos, it is frequently easier to understand the complicated and fully analytic definitions of objects in a Grothendieck topos by relating them to their logical counterparts. We will try and appeal to both the geometric and the logical intuitions in this tutorial where possible.

2 Presheaves and presheaf toposes

Presheaves are the simplest way to capture mathematical objects which vary over a category $\mathbb{C} : \mathbf{Cat}$.

Definition 2.1 (Presheaf). A presheaf on $\mathbb{C} : \mathbf{Cat}$ is a functor $\mathcal{F} : \mathbb{C}^{\text{op}} \rightarrow \mathbf{Set}$.

Unfolding definitions, this means that for every object $\Psi : \mathbb{C}$, we have a set $\mathcal{F}(\Psi) : \mathbf{Set}$; moreover, for any morphism $\psi : \Phi \rightarrow \Psi$, we have an induced *restriction* map $\mathcal{F}(\psi) : \mathcal{F}(\Psi) \rightarrow \mathcal{F}(\Phi)$ which preserves identities and composition. The presheaves on \mathbb{C} form a category, which we will call $\widehat{\mathbb{C}} : \mathbf{Cat}$. When the presheaf $\mathcal{F} : \widehat{\mathbb{C}}$ is understood from context, we will write $m \cdot \psi : \mathcal{F}(\Phi)$ for $\mathcal{F}(\psi)(m)$.

Definition 2.2 (Yoneda embedding). We have a full and faithful functor $\mathbf{y} : \mathbb{C} \rightarrow \widehat{\mathbb{C}}$ called the *Yoneda embedding*, which is defined as follows:

$$\mathbf{y}(\Psi) \triangleq \text{Hom}_{\mathbb{C}}(-, \Psi)$$

A functor which is isomorphic to $\mathbf{y}(\Psi)$ is called *representable* by Ψ .

²That is, if we are content to temporarily ignore the matter of *predicativity*; in practice, this can be dealt with through the notion of a “pretopos” with a large subobject classifier, or a hierarchy thereof.

³There are other options too, such as the gros topos of sheaves over the category of separable locales equipped with the open cover topology [5, 6].

Definition 2.3 (Sieve). A sieve on an object $\Psi : \mathbb{C}$ is a subfunctor of the presheaf represented by Ψ , i.e. $S \mapsto \mathbf{y}(\Psi) : \widehat{\mathbb{C}}$.⁴ In particular, a sieve on Ψ picks out *functorially* a collection of arrows ending in Ψ ; the *maximal sieve* is $\mathbf{y}(\Psi)$ itself, which chooses all arrows ending in Ψ .

2.1 Sieves and subobjects

In classical set theory, every subset $A \subseteq B$ has a *characteristic function* $\text{char}_A : B \rightarrow \mathbf{2}$, defined as follows:

$$\text{char}_A(b) \triangleq \begin{cases} 0 & \text{if } b \in A \\ 1 & \text{if } b \notin A \end{cases}$$

In toposes, there is always an object analogous to $\mathbf{2}$, called the *subobject classifier*; this object is always written Ω . Following [8], a subobject classifier is a monomorphism $\text{true} : \mathbf{1} \rightarrow \Omega$ which for any monomorphism $f : A \rightarrow B$ induces a *unique* characteristic morphism $\text{char}_f : B \rightarrow \Omega$ such that the following diagram is a pullback:

$$\begin{array}{ccc} A & \xrightarrow{\quad ! \quad} & \mathbf{1} \\ \downarrow f & & \downarrow \text{true} \\ B & \xrightarrow{\quad \text{char}_f \quad} & \Omega \end{array}$$

In the category of sets, the subobject classifier is simply the two-element set; its construction in a presheaf topos is more complicated, essentially because it must be made to respect the fact that the objects under consideration are “in motion”.

The subobject classifier in a presheaf topos The subobject classifier in a presheaf topos is defined using *sieves*:

$$\begin{aligned} \Omega(\Psi) &\triangleq \{S \mid S \mapsto \mathbf{y}(\Psi) : \widehat{\mathbb{C}}\} \\ (S \cdot \psi : \Phi \rightarrow \Psi)(X) &\triangleq \{ \phi : X \rightarrow \Phi \mid \psi \circ \phi \in S(X) \} \\ \text{true}_\Psi &\triangleq \mathbf{y}(\Psi) \end{aligned}$$

Remark 2.4. It may not be immediately clear why the subobject classifier is defined in this way: what does the collection of sieves have to do with $\mathbf{2}$ in (classical) set theory? One way to understand what is happening is to observe how Ω behaves when our base category \mathbb{C} is *chaotic*, in the sense that every two objects is connected by exactly one arrow. If \mathbb{C} is chaotic, then either $\mathbf{y}(\Psi)(\Phi) = \mathbf{1}$ for every $\Phi : \mathbb{C}$; therefore, the judgment $S \mapsto \mathbf{y}(\Psi) : \widehat{\mathbb{C}}$ comes out to mean simply $S(\Psi) \subseteq \mathbf{1}$, i.e. $S(\Psi) \in \mathcal{P}(\mathbf{1})$. Therefore, $\Omega(\Psi) = \mathcal{P}(\mathbf{1}) = \mathbf{2}$. An analogous argument can be made in case \mathbb{C} is a groupoid.

⁴Usually, an alternative definition is given in terms of “sets of arrows closed under precomposition”, but we prefer a definition with fewer moving parts. In practice it will be useful to use the alternative definition when reasoning.

Understanding Ω using Yoneda's Lemma as a weapon Much like how a candidate construction of the exponential in a functor category can be hypothesized using an insight from the Yoneda Lemma, it is also possible to apply the same technique to the construction of the subobject classifier in a presheaf category.

Based on our intention that Ω shall be a construction of the subobject classifier, we want to identify maps $\phi : X \rightarrow \Omega$ with subobjects of $X : \widehat{\mathbb{C}}$, i.e. we intend to exhibit a bijection $[X, \Omega] \cong \text{Sub}(X)$. Now, cleverly choose $X \triangleq \mathbf{y}(\mathbb{U})$; then we have $[\mathbb{U}, \Omega] \cong \text{Sub}(X)$. But the Yoneda Lemma says that $[\mathbf{y}(\mathbb{U}), \Omega] \cong \Omega(\mathbb{U})$! Therefore, we may take as a scientific hypothesis the definition $\Omega(\mathbb{U}) \triangleq \text{Sub}(\mathbf{y}(\mathbb{U}))$. It remains to show that this definition exhibits the correct properties (exercise for the reader).

3 Generalized topologies

Let us remark that so far, we have described via presheaves a notion of *variable set* which requires only functoriality; in case we are varying over a poset, this corresponds to *monotonicity* in Kripke models. We will now consider a notion of set which varies *continuously*, a property which corresponds to *local character* in Beth models.

The definition of a Grothendieck topology is quite complicated, but we will show how to understand it conceptually using the logical perspective that we alluded to in the introduction.

Definition 3.1 (Grothendieck Topology [8]). A Grothendieck topology is, for each object $\Psi : \mathbb{C}$, a collection $J(\Psi)$ of sieves on Ψ ; a sieve $S \in J(\Psi)$ is called a *covering sieve*. To be called a topology, the predicate J must be closed under the following rules:

$$\frac{\mathbf{y}(\Psi) \in J(\Psi)}{S \in J(\Psi)} \text{ maximality} \quad \frac{S \in J(\Psi) \quad \psi : \Phi \rightarrow \Psi}{S \cdot \psi \in J(\Phi)} \text{ stability}$$

$$\frac{S \in J(\Psi) \quad R \in \Omega(\Psi) \quad \forall \psi \in S(\Phi). R \cdot \psi \in J(\Phi)}{R \in J(\Psi)} \text{ transitivity}$$

The above rules seem fairly poorly-motivated at first; however, it is easy to understand their purpose when one considers the *logical* perspective. First, one should recognize that the *stability* law above is a disguised form of functoriality for J : that is, it ensures that J itself be a presheaf, namely, a subobject of Ω .

Now, every subobject induces a characteristic map into Ω , and it turns out that it will be far more informative to ignore the analytic aspects of J and focus only on the properties of its characteristic map $j : \Omega \rightarrow \Omega \triangleq \mathbf{char}_J$:

$$\begin{array}{ccc} J & \xrightarrow{!} & \mathbf{1} \\ \downarrow & & \downarrow \text{true} \\ \Omega & \xrightarrow{j \triangleq \mathbf{char}_J} & \Omega \end{array}$$

This perspective, to be developed in the next section, is justified by the fact that in a topos, subobjects are completely determined by their characteristic maps.

3.1 The logical view

Definition 3.2. For an arbitrary topos \mathcal{E} , a Lawvere-tierney operator (also called a Lawvere-Tierney topology, local operator, or nucleus) is a map $j : \Omega \rightarrow \Omega$ which exhibits the following characteristics:

$$j \circ \text{true} = \text{true} \quad (3.1)$$

$$j \circ j = j \quad (3.2)$$

$$j \circ \wedge = \wedge \circ (j \times j) \quad (3.3)$$

In other words, j is a \wedge -preserving closure operation in the internal logic of the topos. The above requirements can be rephrased equivalently in the internal language of \mathcal{E} in the following way:

$$p : \Omega \mid p \vdash j(p) \quad (3.4)$$

$$p : \Omega \mid j(j(p)) \vdash j(p) \quad (3.5)$$

$$p, q : \Omega \mid j(p) \wedge j(q) \vdash j(p \wedge q) \quad (3.6)$$

There is, however, a better-motivated way to state these laws which makes more sense from a logical perspective; in particular, axiom 3.6 can be replaced with the more intuitive internal *monotonicity* condition that j shall preserve implication:

$$p, q : \Omega \mid p \Rightarrow q \vdash j(p) \Rightarrow j(q) \quad (3.7)$$

We will restate without proof the following result from [8]:

Proposition 3.3 (Lawvere-Tierney subsumes Grothendieck). *In a presheaf topos $\widehat{\mathbb{C}}$, a subobject $J \rightarrow \Omega : \mathbb{C}$ is a Grothendieck topology iff its characteristic map $\text{char}_J : \Omega \rightarrow \Omega$ is a Lawvere-Tierney topology.*

Notation 3.4. It is common to write j for char_J to mean that Lawvere-Tierney topology which corresponds to the Grothendieck topology J .

Relating the logical and the geometric views Each of the rules for a Grothendieck topology corresponds to an intuitive logical requirement: *maximality*, i.e. the inclusion of represented functors as covers, corresponds to the requirement that our local operator shall preserve truth; *stability* corresponds to the requirement that J shall in fact be a presheaf; the *transitivity* law corresponds exactly to axioms 3.5 and 3.7, composed to form the Kleisli extension for the monad j .

Remark 3.5 (Pretopologies and coverages). There are several other ways to define some form of topology on a category, including *coverages* and *pretopologies*. In some contexts, these are allegedly easier to work with, but they tend to impose extra requirements on the category \mathbb{C} , and end up obscuring the crisp logical character of topologies and their correspondence with modal operators. From a logical perspective, the concept of a “pretopology” is essentially meaningless, so we prefer to avoid it.

3.2 Covering sieves as a poset

As we have described it, the covering sieves $J(\Psi)$ form a set, but we can also regard $J(\Psi)$ as a poset by imposing the inclusion order.

Definition 3.6 (Filtered poset). A poset $P : \mathbf{Pos}$ is *filtered* or *directed* iff it is inhabited, and every pair of elements $m, n : P$ has an upper bound $l : P$, i.e. both $m \leq l$ and $n \leq l$.

Lemma 3.7 (Covers are cofiltered). *For every $\Psi : \mathbb{C}$ the poset $J(\Psi) : \mathbf{Pos}$ is cofiltered, i.e. the opposite poset $J(\Psi)^{\text{op}} : \mathbf{Pos}$ is filtered.*

Proof. First, $J(\Psi)$ is inhabited by the maximal sieve $\mathbf{y}(\Psi)$, so it suffices to show that for any $S, T : J(\Psi)$ we can exhibit $S \cap T : J(\Psi)$ such that $S \cap T \subseteq S$ and $S \cap T \subseteq T$. The construction of this intersection is clear, but we need to show that it is still a cover.

At an analytic level, we can use the *transitivity* axiom of a Grothendieck topology to show that the intersection of two covering sieves is a covering sieve. But it is more clear and less bureaucratic to work internally, using axiom 3.3 for a Lawvere-Tierney topology which states that $j \circ \wedge = \wedge \circ (j \times j)$. Because $j : \Omega \rightarrow \Omega$ is the characteristic map of the subobject $J \rightarrow \Omega : \widehat{\mathbb{C}}$, this really means that $S \cap T \in J \Leftrightarrow S \in J \wedge T \in J$ from an internal perspective, which is precisely what we were trying to prove. \square

4 Sheaves on a site

A *site* is a category \mathbb{C} together with a topology $J \rightarrow \Omega : \widehat{\mathbb{C}}$. We will now proceed to give perspicuous definitions of what it means for a presheaf $\mathcal{F} : \widehat{\mathbb{C}}$ to be *separated* and a *sheaf* respectively.

There are many different definitions of separated presheaves and sheaves, most of which involve a number of complicated analytic conditions; we prefer to give an equivalent, simpler definition (which is usually presented as a *theorem*).

First, observe that for any sieve $S \in \Omega(\Psi)$, we have a canonical function between hom-sets $i_S^* : [\mathbf{y}(\Psi), \mathcal{F}] \rightarrow [S, \mathcal{F}]$ as follows:

$$\begin{array}{ccc}
 \mathbf{y}(\Psi) & \xrightarrow{m} & \mathcal{F} \\
 \uparrow i_S & \nearrow i_S^* m & \\
 S & &
 \end{array}$$

Definition 4.1. The presheaf $\mathcal{F} : \widehat{\mathbb{C}}$ is *separated* iff for every $S \in \Omega(\Psi)$, the induced map between hom-sets $i_S^* : [\mathbf{y}(\Psi), \mathcal{F}] \rightarrow [S, \mathcal{F}]$ is a monomorphism. \mathcal{F} is a *sheaf* iff this map is also an isomorphism.

It is worth taking a moment to cultivate some insight as to what is going on here. First, recall that by the Yoneda lemma, we can identify elements of $\mathcal{F}(\Psi)$ with natural transformations from the maximal sieve, i.e. $[\mathbf{y}(\Psi), \mathcal{F}]$. So, sheafhood for \mathcal{F} is really saying that as far as \mathcal{F} is concerned, the elements of $\mathcal{F}(\Psi)$ can be identified with natural transformations from *any* sieve that covers Ψ , not just the maximal one.

Definition 4.2 (Matching families and amalgamations). A natural transformation $m : [S, \mathcal{F}]$ for $S \in J(\Psi)$ is usually called a *matching family* for S ; then, the member of $\mathcal{F}(\Psi)$ which is determined by the sheaf-induced isomorphism (and the Yoneda lemma) is called an *amalgamation*.

Notation 4.3. As a matter of convention, we will use Latin letters m, n, \dots to range over elements of a presheaf, and German letters m, n, \dots to range over matching families for a presheaf.

5 Sheafification: the essence of toposes

For any Grothendieck topology $J \multimap \Omega : \widehat{\mathbb{C}}$, there is a canonical way to turn a presheaf on \mathbb{C} into a sheaf on the site (\mathbb{C}, J) . Later on, we will see that this is actually this can actually be taken as *definitive* without explicitly invoking Grothendieck (or Lawvere-Tierney) topologies, exposing the principal contradiction of sheaf theory as the adjunction induced by a *lex reflective subcategory* of a topos.

To begin with, we will exhibit the J -sheaves as a lex reflective subcategory of the topos $\widehat{\mathbb{C}}$ in the sense that we have a geometric morphism $\mathbf{a} : \widehat{\mathbb{C}} \rightarrow \mathbf{Sh}(\mathbb{C}, J) \dashv \iota : \mathbf{Sh}(\mathbb{C}, J) \hookrightarrow \widehat{\mathbb{C}}$, where ι is the (full and faithful) inclusion of sheaves into presheaves and \mathbf{a} is a “sheafification” functor.

5.1 The plus construction

Sheafification in presheaf toposes is obtained from the iterated application of something called “Grothendieck’s plus construction”, which we define below, fixing a presheaf $\mathcal{F} : \widehat{\mathbb{C}}$:

$$\begin{aligned} (-)^+ : \widehat{\mathbb{C}} &\rightarrow \widehat{\mathbb{C}} \\ \mathcal{F}^+(\Psi) &\triangleq \mathbf{colim}_{S:J(\Psi)^{\text{op}}} [S, \mathcal{F}] \\ (S, m) \cdot \psi : \Phi \rightarrow \Psi &\triangleq (S \cdot \psi, m \circ (\psi \circ -)) \\ (f : \mathcal{F} \rightarrow \mathcal{G})_{\Psi}^+ (S, m) &\triangleq (S, f \circ m) \end{aligned}$$

In the above colimiting construction, we write $J(\Psi)$ for the *poset* of covering sieves on Ψ .

In other words, the *plus construction* for a presheaf replaces its elements by covering sieves equipped with matching families: conceptually, this means equipping a presheaf with “formal amalgamations”. We will see that this process is remarkably well-behaved.

Lemma 5.1 (Equality in \mathcal{F}^+). *The following are equivalent:*

1. $(R, m) = (S, n) \in \mathcal{F}^+(\Psi)$
2. $(R, m) \sim (S, n)$, where \sim is the smallest equivalence relation which relates (R, m) and (S, n) when $R \subseteq S \cdot \psi : \Phi \rightarrow \Psi \in \mathcal{R}(\Phi)$, we have $m(\psi) = n(\psi) \in \mathcal{F}(\Phi)$.

3. There is a common refinement $T \subseteq R \cap S \in J(\Psi)$ such that for all $\psi : \Phi \rightarrow \Psi \in T(\Phi)$, we have $m(\psi) = n(\psi) \in \mathcal{F}(\Phi)$.

Proof. (1) \Leftrightarrow (2) is the definition of equality for a colimit in **Set**. (1) \Leftrightarrow (3) is an equivalent characterization of equality for a *filtered* colimit, i.e. a colimit of a digram whose domain is a filtered category. We have shown that $J(\Psi)^{\text{op}}$ is indeed filtered in Lemma 3.7. \square

The unit to the plus construction We can also form a unit natural transformation $\eta_+^{\mathcal{F}} : \mathcal{F} \rightarrow \mathcal{F}^+$ using the maximal sieve:

$$\eta_+^{\mathcal{F}} : \mathcal{F} \rightarrow \mathcal{F}^+ \\ (\eta_+^{\mathcal{F}})_{\Psi}(m) \triangleq (\mathbf{y}(\Psi), m \cdot -)$$

Lemma 5.2 (Executing formal amalgamations). *When $\mathcal{F} : \mathbf{Sh}(\mathbb{C}, J)$ is a sheaf, the unit $\eta_+^{\mathcal{F}} : \mathcal{F} \rightarrow \mathcal{F}^+$ is an isomorphism.*

Proof. This holds by definition; the action of the inverse to the unit is to use the sheaf structure to turn a formal amalgamation (a covering sieve together with a matching family) into the appropriate unique element of \mathcal{F} . \square

Notation 5.3. We will write $\text{glue}_{\mathcal{F}} : \mathcal{F}^+ \rightarrow \mathcal{F}$ for the inverse to $\eta_+^{\mathcal{F}}$.

Now, we reproduce a lemma from [8], giving a bit more detail.

Lemma 5.4. *For $\mathcal{F} : \mathbf{Sh}(\mathbb{C}, J)$ and $\mathcal{G} : \widehat{\mathbb{C}}$, any map $f : \mathcal{G} \rightarrow \iota\mathcal{F}$ in $\widehat{\mathbb{C}}$ factors as $f = \tilde{f} \circ \eta_+^{\mathcal{G}}$ for some unique $\tilde{f} : \mathcal{G}^+ \rightarrow \iota\mathcal{F}$:*

$$\begin{array}{ccc} \mathcal{G} & \xrightarrow{\eta_+^{\mathcal{G}}} & \mathcal{G}^+ \\ & \searrow f & \downarrow \tilde{f} \\ & & \iota\mathcal{F} \end{array} \quad (5.1)$$

Proof. Explicitly, we can construct a candidate definition of \tilde{f} using the functoriality of the plus construction and Lemma 5.2:

$$\begin{array}{ccccc} \mathcal{G}^+ & \xrightarrow{f^+} & \iota\mathcal{F}^+ & \xrightarrow{\text{glue}_{\mathcal{F}}} & \iota\mathcal{F} \\ & \searrow & \text{---} & \nearrow & \\ & & \tilde{f} & & \end{array}$$

First, we have to show that $\text{glue}_{\mathcal{F}} \circ f^+ \circ \eta_+^{\mathcal{G}} = f$; this is easily established:

$$m \xrightarrow{(\eta_+^{\mathcal{G}})_{\Psi}} (\mathbf{y}(\Psi), m \cdot -) \xrightarrow{f^+_{\Psi}} (\mathbf{y}(\Psi), f_{\Psi}(m) \cdot -) \xrightarrow{\text{glue}_{\mathcal{F}}^{\Psi}} f_{\Psi}(m)$$

Next we have to show $\text{glue}_{\mathcal{F}} \circ f^+$ is the only possible \tilde{f} such that our diagram commutes. Fix $(S, \mathfrak{m}) \in \mathcal{G}^+(\Psi)$; we will show that $\tilde{f}_{\Psi}(S, \mathfrak{m})$ *must* be $\text{glue}_{\mathcal{F}}^{\Psi}(f_{\Psi}^+(S, \mathfrak{m}))$. For all $\psi \in S(\Phi)$, we have the following:

$$\begin{aligned} \tilde{f}_{\Psi}(S, \mathfrak{m}) \cdot \psi &= \tilde{f}_{\Phi}((S, \mathfrak{m}) \cdot \psi) && \text{(Naturality)} \\ &= \tilde{f}_{\Phi}(S \cdot \psi, \mathfrak{m} \circ (\psi \circ -)) && \text{(Definition)} \end{aligned}$$

Observe that because $\psi \in S(\Phi)$, the sieve $S \cdot \psi$ is in fact the maximal sieve $\mathfrak{y}(\Phi)$. So we have:

$$\begin{aligned} \tilde{f}_{\Psi}(S, \mathfrak{m}) \cdot \psi &= \tilde{f}_{\Phi}(\mathfrak{y}(\Phi), \mathfrak{m}_{\Phi}(\psi) \cdot -) \\ &= \tilde{f}_{\Phi}(\eta_+^{\mathcal{G}}(\mathfrak{m}_{\Phi}(\psi))) && \text{(definition)} \\ &= f_{\Phi}(\mathfrak{m}_{\Phi}(\psi)) && \text{(Diagram 5.1)} \end{aligned}$$

Now, because \mathcal{F} is a sheaf, and S is a covering sieve and $f \circ \mathfrak{m} : [S, \iota\mathcal{F}]$ is a matching family, $\tilde{f}_{\Psi}(S, \mathfrak{m})$ is the *unique* element $n \in \mathcal{F}(\Psi)$ such that $n \cdot \psi = f_{\Phi}(\mathfrak{m}_{\Phi}(\psi))$. In other words $\tilde{f}_{\Psi}(S, \mathfrak{m}) = n = \text{glue}_{\mathcal{F}}^{\Psi}(S, f \circ \mathfrak{m})$, so it suffices to confirm that $f_{\Psi}^+(S, \mathfrak{m}) = (S, f \circ \mathfrak{m})$. \square

Lemma 5.5. *For any presheaf $\mathcal{G} : \widehat{\mathbb{C}}$, we have the identity $(\eta_+^{\mathcal{G}})^+ = \eta_+^{\mathcal{G}^+} : \mathcal{G}^+ \rightarrow \mathcal{G}^{++}$.*

Proof. Working pointwise, we have the following:

$$\begin{aligned} (\eta_+^{\mathcal{G}})^+(S, \mathfrak{m}) &= (S, \eta_+^{\mathcal{G}} \circ \mathfrak{m}) \\ (\eta_+^{\mathcal{G}^+})_{\Psi}(S, \mathfrak{m}) &= (\mathfrak{y}(\Psi), (S, \mathfrak{m}) \cdot -) \end{aligned}$$

We need to show that $(S, \eta_+^{\mathcal{G}} \circ \mathfrak{m}) = (\mathfrak{y}(\Psi), (S, \mathfrak{m}) \cdot -)$. By the definition of equality for formal amalgamations (Lemma 5.1), it suffices to show that for all $\psi : \Phi \rightarrow \Psi \in S(\Phi)$ we have $\eta_+^{\mathcal{G}}(\mathfrak{m}_{\Phi}(\psi)) = (S, \mathfrak{m}) \cdot \psi$. Calculate:

$$\begin{aligned} (S, \mathfrak{m}) \cdot \psi &= (S \cdot \psi, \mathfrak{m}_{\Phi} \circ (\psi \circ -)) \\ &= (\mathfrak{y}(\Phi), \mathfrak{m}_{\Phi}(\psi) \cdot -) \\ &= \eta_+^{\mathcal{G}}(\mathfrak{m}_{\Phi}(\psi)) \end{aligned}$$

\square

In the theorems below, we will follow in the broad strokes the proofs given in [8], but giving more detail and using a more “nuts-and-bolts” style based on the view of sieves as subfunctors and matching families as natural transformations.

Theorem 5.6. *When $\mathcal{F} : \widehat{\mathbb{C}}$ is any presheaf, its plus construction $\mathcal{F}^+ : \widehat{\mathbb{C}}$ is separated, i.e. for any sieve $S \in \mathbb{J}(\Psi)$ the precomposition map between hom-sets induced by $i_S : S \rightarrow \mathfrak{y}(\Psi)$ is injective:*

$$[\mathfrak{y}(\Psi), \mathcal{F}^+] \xrightarrow{i_S^*} [S, \mathcal{F}^+]$$

Proof. Fix two matching families $m_1, m_2 : [\mathbf{y}(\Psi), \mathcal{F}^+]$. We need to show that if $m_1 \circ i_S = m_2 \circ i_S : [S, \mathcal{F}^+]$, then already $m_1 = m_2 : [\mathbf{y}(\Psi), \mathcal{F}^+]$.

By the Yoneda Lemma it suffices to prove only $m_1 = m_2 \in \mathcal{F}^+(\Psi)$ (defining $m_i \triangleq m_i(\text{id}_\Psi)$), and we can rewrite our assumption to say that for all $\psi \in S(\Phi)$ we have $m_1 \cdot \psi = m_2 \cdot \psi \in \mathcal{F}^+(\Phi)$.

Let $m_i \equiv (T_i \in J(\Psi), n_i : [T_i, \mathcal{F}])$; by Lemma 5.1, our equality hypothesis is the same as to say that there exists common refinements $T_\psi \subseteq (T_1 \cdot \psi) \cap (T_2 \cdot \psi)$ such that $n_1(\psi \circ \phi) = n_2(\psi \circ \phi) \in \mathcal{F}(X)$ for all $\phi : X \rightarrow \Psi \in T_\psi(X)$.

We need to show that there exists a common refinement $T \subseteq T_1 \cap T_2$ such that for all $\psi : \Phi \rightarrow \Psi \in T(\Phi)$, we have $n_1(\psi) = n_2(\psi) \in \mathcal{F}(\Phi)$. Now we will choose $T(X) \triangleq \{\psi \circ \phi \mid \psi \in S(\Phi), \phi \in T_\psi(X)\}$; clearly $T \subseteq T_1 \cap T_2$. It remains to show the following:

1. $T \in J(\Psi)$. Applying the *transitivity* axiom for covering sieves with $S \in J(\Psi)$, it suffices to show that for all $\psi \in S(\Phi)$ we have $T \cdot \psi \in J(\Phi)$. Observe that $(T \cdot \psi)(X) \equiv \{\phi \mid \psi \circ \phi \in T(X)\}$, i.e. $T \cdot \psi = T_\psi$ which we have already assumed to be a cover.
2. For all $\psi \in T(X)$, $n_1(\psi) = n_2(\psi) \in \mathcal{F}(X)$. Unfolding things, this means that for all $\psi \in S(\Phi)$ and $\phi \in T_\psi(X)$ we must show $n_1(\psi \circ \phi) = n_2(\psi \circ \phi) \in \mathcal{F}(X)$, which we have already assumed.

□

Theorem 5.7. *When $\mathcal{F} : \widehat{\mathbb{C}}$ is a separated presheaf, its plus construction $\mathcal{F}^+ : \mathbb{C}$ is in fact a sheaf, i.e. for any covering sieve $S \in J(\Psi)$ the precomposition map between hom-sets induced by $i_S : S \rightarrow \mathbf{y}(\Psi)$ is bijective:*

$$[\mathbf{y}(\Psi), \mathcal{F}^+] \xleftarrow{i_S^*} [S, \mathcal{F}^+]$$

Proof. Fix a matching family $m_S : [S, \mathcal{F}^+]$; we need to exhibit a unique amalgamation $m_S \in \mathcal{F}^+(\Psi)$ for m_S . Because \mathcal{F}^+ is separated (Theorem 5.6), it suffices to merely exhibit *some* such amalgamation, since its uniqueness will follow.

First, observe that for each $\psi : \Phi \rightarrow \Psi \in S(\Phi)$, the matching family gives a “formal amalgamation” for some covering sieve, i.e. $m_S(\psi) \in \mathcal{F}^+(\Phi) \equiv (T_\psi, n_\psi)$ for some $T_\psi \in J(\Phi)$ and $n_\psi : [T_\psi, \mathcal{F}]$.

Our task is to choose a suitable covering sieve $R \in J(\Psi)$, together with a matching family $m_R : [R, \mathcal{F}]$ so that we can define $m_S \triangleq (R, m_R)$. Using a tactic similar to what we did in the proof of Theorem 5.6, choose:

$$R(X) \triangleq \{X \xrightarrow{\phi} \Phi \xrightarrow{\psi} \Psi \mid \psi \in S(\Phi), \phi \in T_\psi(X)\}$$

As in Theorem 5.6, we have $R \in J(\Psi)$ by the transitivity axiom. Now we need to exhibit a matching family $m_R : [R, \mathcal{F}]$; first define its components as follows:

$$m_R^X(\psi \circ \phi) \triangleq n_\psi^X(\phi)$$

Functionality To show that the components of this matching family are well-defined (functional), fix $\phi_1, \phi_2 : T_\psi(X)$ such that $\psi \circ \phi_1 = \psi \circ \phi_2$; we need to show that $n_\psi^X(\phi_1) = n_\psi^X(\phi_2) \in \mathcal{F}(X)$.

By naturality of m_S , we have $(T_\psi \cdot \phi, n_\psi \cdot \phi) = (T_{\psi \circ \phi}, n_{\psi \circ \phi}) \in \mathcal{F}^+(X)$ for all $\phi : X \rightarrow \Phi$; by Lemma 5.1, this means that there is some family of covering sieves $\mathcal{U}_{\psi, \phi} \subseteq (T_\psi \cdot \phi) \cap (T_{\psi \circ \phi}) \in J(X)$ such that $n_\psi^\Upsilon(\phi \circ \chi) = n_{\psi \circ \phi}^\Upsilon(\chi) \in \mathcal{F}(\Upsilon)$ for all $\chi : \Upsilon \rightarrow X \in \mathcal{U}_{\psi, \phi}(\Upsilon)$.

Because covering sieves are closed under intersection, we also have a cover $\mathcal{U}_{\psi, \phi_1} \cap \mathcal{U}_{\psi, \phi_2} \in J(X)$. Fixing $\chi : \Upsilon \rightarrow X \in (\mathcal{U}_{\psi, \phi_1} \cap \mathcal{U}_{\psi, \phi_2})(\Upsilon)$, calculate:

$$\begin{aligned} n_\psi^X(\phi_1) \cdot \chi &= n_\psi^\Upsilon(\phi_1 \circ \chi) && \text{(naturality of } n_\psi) \\ &= n_{\psi \circ \phi_1}^\Upsilon(\chi) && \text{(see above)} \\ &= n_{\psi \circ \phi_2}^\Upsilon(\chi) && \text{(assumption)} \\ &= n_\psi^\Upsilon(\phi_2 \circ \chi) && \text{(see above)} \\ &= n_\psi^X(\phi_2) \cdot \chi && \text{(naturality of } n_\psi) \end{aligned}$$

Because \mathcal{F} is separated, we have immediately our equality goal $n_\psi^X(\phi_1) = n_\psi^X(\phi_2) \in \mathcal{F}(X)$. Therefore, the components m_R^X are well-defined, i.e. functional.

Naturality We still have to show that this candidate definition of m_R is in fact a natural transformation. Fixing $v : X \rightarrow \Upsilon$, we have to show that the following diagram commutes:

$$\begin{array}{ccc} S(\Upsilon) & \xrightarrow{-\cdot v} & S(X) & \psi \circ \phi & \xrightarrow{-\cdot v} & \psi \circ \phi \circ v \\ m_R^\Upsilon \downarrow & & \downarrow m_R^X & m_R^\Upsilon \downarrow & & \downarrow m_R^X \\ \mathcal{F}(\Upsilon) & \xrightarrow{-\cdot v} & \mathcal{F}(X) & n_\psi(\phi) & \xrightarrow{-\cdot v} & n_\psi(\phi) \cdot v = n_\psi(\phi \circ v) \end{array}$$

In this way, the naturality of m_R reduces directly to the naturality of n_ψ . We may now define $m_S \triangleq (R, m_R)$. All that remains is to show that for all $\psi : \Phi \rightarrow \Psi \in S(\Phi)$, we have $m_S \cdot \psi = m_S(\psi) \in \mathcal{F}^+(\Phi)$.

First, observe that $R \cdot \psi = T_\psi$, and that $(m_R \cdot \psi)(\phi) = m_R(\psi \circ \phi) = n_\psi(\phi)$, whence $m_S \cdot \psi = (T_\psi, n_\psi) = m_S(\psi)$. \square

Corollary 5.8. For any presheaf $\mathcal{F} : \widehat{\mathbb{C}}$, the double application of the plus construction \mathcal{F}^{++} is a sheaf, i.e. $\mathcal{F}^{++} : \mathbf{Sh}(\mathbb{C}, J)$.

Proof. By Theorems 5.7 and 5.6. \square

5.2 Sheafification

Theorem 5.9 (Sheafification). We are now equipped to show that the J -sheaves comprise a lex reflective subcategory of the presheaf topos \mathbb{C} , defining a $\mathcal{F} \triangleq \mathcal{F}^{++}$. In other words,

1. $\iota : \mathbf{Sh}(\mathbb{C}, J) \rightarrow \widehat{\mathbb{C}}$ is full and faithful.
2. $\mathbf{a} : \widehat{\mathbb{C}} \rightarrow \mathbf{Sh}(\mathbb{C}, J)$ preserves finite limits.
3. We have the adjunction $\mathbf{a} \dashv \iota$.

Proof. Clearly ι is full and faithful, because it is the identity on morphisms. To show that \mathbf{a} preserves finite limits, it suffices to show that the plus construction preserves finite limits in $\widehat{\mathbb{C}}$. Because finite limits in a sheaf topos are formed pointwise, it suffices to show that finite limits are preserved at the level of sets. This follows from the fact that the plus construction is a filtered colimit, and that $[S, -]$ preserves limits.

Finally, to show that $\mathbf{a} \dashv \iota$, we must exhibit the following bijection of hom-sets, natural in \mathcal{F}, \mathcal{G} (we leave naturality as an exercise to the reader):

$$\frac{\mathbf{a} \mathcal{G} \rightarrow \mathcal{F} : \mathbf{Sh}(\mathbb{C}, J)}{\mathcal{G} \rightarrow \iota \mathcal{F} : \widehat{\mathbb{C}}}$$

We will construct the bijection of maps in the diagrams below. Because maps in sheaves are the same as maps in presheaves, for simplicity's sake the following diagrams will reside in $\widehat{\mathbb{C}}$, regarding each sheaf as a presheaf. The first diagram is completed using composition with the unit to the plus construction; the second diagram is completed using two applications of Lemma 5.4.

$$\begin{array}{ccc}
 \mathcal{G} & \xrightarrow{\eta_+^{\mathcal{G}}} & \mathcal{G}^+ & \xrightarrow{\eta_+^{\mathcal{G}^+}} & \mathbf{a} \mathcal{G} \\
 & \searrow \text{dashed} & & & \downarrow f \\
 & & & & \mathcal{F} \\
 f_{\#} \triangleq f \circ \eta_+^{\mathcal{G}^+} \circ \eta_+^{\mathcal{G}} & & & &
 \end{array}
 \qquad
 \begin{array}{ccc}
 \mathcal{G} & \xrightarrow{\eta_+^{\mathcal{G}}} & \mathcal{G}^+ & \xrightarrow{\eta_+^{\mathcal{G}^+}} & \mathbf{a} \mathcal{G} \\
 & \searrow f & \downarrow \tilde{f} & \swarrow \text{dashed} & \\
 & & \iota \mathcal{F} & & f_{\#} \triangleq \widetilde{(f)}
 \end{array}$$

Next we show that this actually constitutes a bijection; the first direction is almost immediate.

$$\begin{aligned}
 (f_{\#})^{\#} &= \widetilde{(f)} \circ \eta_+^{\mathcal{G}^+} \circ \eta_+^{\mathcal{G}} \\
 &= \tilde{f} \circ \eta_+^{\mathcal{G}} && \text{(Lemma 5.4)} \\
 &= f && \text{(Lemma 5.4)}
 \end{aligned}$$

The other direction will require some effort and careful manipulation, using

Lemma 5.4 and Lemma 5.5.

$$\begin{aligned}
(f^\#)_\# &= f \circ \widetilde{\eta_+^{\mathcal{G}^+}} \circ \eta_+^{\mathcal{G}} \\
&= \text{glue}_{\mathcal{F}} \circ \left(\text{glue}_{\mathcal{F}} \circ \left(f \circ \eta_+^{\mathcal{G}^+} \circ \eta_+^{\mathcal{G}} \right)^+ \right)^+ && \text{(definition of } \widetilde{f} \text{)} \\
&= \text{glue}_{\mathcal{F}} \circ \left(\text{glue}_{\mathcal{F}} \circ f^+ \circ \left(\eta_+^{\mathcal{G}^+} \right)^+ \circ \left(\eta_+^{\mathcal{G}} \right)^+ \right)^+ && \text{(functoriality)} \\
&= \text{glue}_{\mathcal{F}} \circ \left(\widetilde{f} \circ \left(\eta_+^{\mathcal{G}^+} \right)^+ \circ \left(\eta_+^{\mathcal{G}} \right)^+ \right)^+ && \text{(definition of } \widetilde{f} \text{)} \\
&= \text{glue}_{\mathcal{F}} \circ \left(\widetilde{f} \circ \eta_+^{\mathcal{A}^{\mathcal{G}}} \circ \eta_+^{\mathcal{G}^+} \right)^+ && \text{(Lemma 5.5)} \\
&= \text{glue}_{\mathcal{F}} \circ \left(f \circ \eta_+^{\mathcal{G}^+} \right)^+ && \text{(Diagram 5.1)} \\
&= \text{glue}_{\mathcal{F}} \circ f^+ \circ \left(\eta_+^{\mathcal{G}^+} \right)^+ && \text{(functoriality)} \\
&= \text{glue}_{\mathcal{F}} \circ f^+ \circ \eta_+^{\mathcal{A}^{\mathcal{G}}} && \text{(Lemma 5.5)} \\
&= f && \text{(Diagram 5.1)}
\end{aligned}$$

We have shown that $\mathbf{Sh}(\mathbb{C}, \mathbb{J})$ is a lex reflective subcategory of $\widehat{\mathbb{C}}$. □

5.3 The Fetishism of Sheaves, and the Secret Thereof

A sheaf appears, at first sight, a very trivial thing, and easily understood. Its analysis shows that it is, in reality, a very queer thing, abounding in metaphysical subtleties and topological niceties.⁵ So far as it is a presheaf enjoying certain conditions, there is nothing mysterious about it, whether we consider it from the point of view of the *espace étalé*, or as a continuously varying family of sets.

But just so soon as the collection of sheaves for a topology stands as a body in relation to all other objects in a topos, it acquires certain characteristics which at first set it apart from other collections of objects, but eventually come to be fully constitutive of the sheaf concept. I am of course referring to the characteristic of being a *lex reflective subcategory* of the presheaf topos (Theorem 5.9), which we will see is in fact coextensive with the notion of *sheaf subcategory*.

Fix a category \mathbb{E} and a lex reflective subcategory $\mathbb{F} \hookrightarrow \mathbb{E}$. In particular, we have $f^* \dashv f_* : \mathbb{F} \hookrightarrow \mathbb{E}$ arranged in the following way:

$$\begin{array}{ccc}
& f^* & \\
\mathbb{F} & \xleftarrow{\quad} & \mathbb{E} \\
& \perp & \\
& \xrightarrow{\quad} & \mathbb{E} \\
& f_* &
\end{array}$$

For convenience, we assume without loss of generality that \mathbb{F} is *replete* (closed under isomorphism).

⁵With apologies to Karl Marx [9].

Notation 5.10. For readability, we will fix the following notation for the monad and comonad of the above adjunction:

$$\begin{aligned}\circlearrowleft : \mathbb{E} &\rightarrow \mathbb{E} \triangleq f_* \circ f^* \\ \square : \mathbb{F} &\rightarrow \mathbb{F} \triangleq f^* \circ f_*\end{aligned}$$

Lemma 5.11. *With $\eta : \text{id}_{\mathbb{E}} \rightarrow \circlearrowleft$ the unit of our adjunction, we have the identity $\eta_{\circlearrowleft A} = \circlearrowleft \eta_A : \circlearrowleft A \rightarrow \circlearrowleft \circlearrowleft A$.*

Proof. Since the monad of a replete reflective subcategory is idempotent [3], its the components of its multiplication operator $\mu : \circlearrowleft \circlearrowleft \rightarrow \circlearrowleft$ are isomorphisms. Therefore, to show that $\eta_{\circlearrowleft A} = \circlearrowleft \eta_A$, it suffices to show that $\mu_A \circ \eta_{\circlearrowleft A} = \mu_A \circ \circlearrowleft \eta_A$, but this just the unit axiom of the monad. \square

Lemma 5.12. *For any two maps $l_1, l_2 : \circlearrowleft C \rightarrow f_* D$, from $l_1 \circ \eta_C = l_2 \circ \eta_C$ we can conclude $l_1 = l_2$.*

Proof. Suppose $l_1 \circ \eta_C = l_2 \circ \eta_C : C \rightarrow f_* D$; then their adjoint transposes are equal too:

$$(l_1 \circ \eta_C)^\sharp = (l_2 \circ \eta_C)^\sharp : f^* C \rightarrow D$$

Now calculate:

$$\begin{aligned}(l_i \circ \eta_C)^\sharp &= \epsilon_D \circ f^*(l_i \circ \eta_C) \\ f_*(l_i \circ \eta_C)^\sharp &= f_* \epsilon_D \circ \circlearrowleft l_i \circ \circlearrowleft \eta_C \\ &= f_* \epsilon_D \circ \circlearrowleft l_i \circ \eta_{\circlearrowleft C} && \text{(Lemma 5.11)} \\ &= f_* \epsilon_D \circ \eta_{f_* D} \circ l_i && \text{(naturality)} \\ &= l_i && \text{(adjunction)}\end{aligned}$$

Since have $f_*(l_1 \circ \eta_C)^\sharp = f_*(l_2 \circ \eta_C)^\sharp$, by the above reasoning we have $l_1 = l_2$. \square

Lemma 5.13. *If \mathbb{E} is finitely complete, then so is \mathbb{F} .*

Proof. Fix a category $\mathbb{J} : \text{Cat}$ and a diagram (functor) $D : \mathbb{J} \rightarrow \mathbb{F}$. Fix a limiting cone $(A, [\alpha_i : A \rightarrow f_* D_i]_{i:\mathbb{J}})$ in \mathbb{E} . We will show that the unit $\eta_A : A \rightarrow \circlearrowleft A$ in \mathbb{E} is an isomorphism, and therefore $A \in \mathbb{F}$ (whence \mathbb{F} is closed under finite limits).

Using the adjunction, each edge $\alpha_i : A \rightarrow f_* D_i$ of our cone in \mathbb{E} may be transposed into a unique map $\alpha_i^\sharp : f^* A \rightarrow D_i$ in \mathbb{F} . We can use this to form another cone in \mathbb{E} :

$$\left(\circlearrowleft A, \left[f_* \alpha_i^\sharp : \circlearrowleft A \rightarrow f_* D_i \right]_{i:\mathbb{J}} \right)$$

Using the universal property of the limit in \mathbb{E} , we acquire a unique map $u : \circlearrowleft A \rightarrow A$ such that for every $i : \mathbb{J}$, we have $f_* \alpha_i^\sharp = \alpha_i \circ u : \circlearrowleft A \rightarrow f_* D_i$. Observe that we have also $f_* \alpha_i^\sharp \circ \eta_A = \alpha_i : A \rightarrow f_* D_i$:

$$\begin{aligned}f_* \alpha_i^\sharp \circ \eta_A &= f_*(\epsilon_{D_i} \circ f^* \alpha_i) \circ \eta_A \\ &= f_* \epsilon_{D_i} \circ \circlearrowleft \alpha_i \circ \eta_A && \text{(reassociate)} \\ &= f_* \epsilon_{D_i} \circ \eta_{f_* D_i} \circ \alpha_i && \text{(naturality)} \\ &= \alpha_i && \text{(adjunction)}\end{aligned}$$

Now we can show that u is a two-sided inverse to η_A . To see that $u \circ \eta_A = \text{id}_A$, calculate:

$$\begin{aligned} f_* \alpha_i^\sharp &= \alpha_i \circ u && \text{(assumption)} \\ f_* \alpha_i^\sharp \circ \eta_A &= \alpha_i \circ u \circ \eta_A \\ \alpha_i &= \alpha_i \circ u \circ \eta_A && \text{(see above)} \end{aligned}$$

Now it is plausible that $u \circ \eta_A = \text{id}_A$, and it is in fact the case because of the uniqueness of mediating maps induced by limiting cones. To see that $\eta_A \circ u = \text{id}_{\bigcirc A}$, by Lemma 5.12 it suffices to show that $\eta_A \circ u \circ \eta_A = \eta_A$, and this follows directly from the other direction of the isomorphism. \square

Lemma 5.14. *If \mathbb{E} has exponentials, then so does \mathbb{F} .*

Proof. Todo. \square

Corollary 5.15. *\mathbb{F} is cartesian closed.*

Lemma 5.16. *If \mathbb{E} has a subobject classifier, then so does \mathbb{F} .*

Proof. Todo. \square

Corollary 5.17. *If \mathbb{E} is a topos, then so is \mathbb{F} .*

6 Applications and examples

We will now survey a few useful topologies.

6.1 Dense and atomic topologies

Definition 6.1 (Dense topology). The *dense topology* is defined as follows:

$$J_{\text{dense}}(\Psi) \triangleq \{ S \in \Omega(\Psi) \mid \forall \psi : \Phi \rightarrow \Psi. \exists \phi : X \rightarrow \Phi. \psi \circ \phi \in S(X) \}$$

Lemma 6.2 (Aromatherapy). *When distilled into its pure essence as a Lawvere-Tierney local operator (Definition 3.2), the dense topology corresponds **classically** to the double-negation modality $\neg\neg$.*⁶

Proof. It suffices to “compile” double-negations from the internal language of the topos into statements about sieves. Recall that $S \in J(\Psi)$ iff $\Psi \Vdash S = \top$, where $\top \triangleq \mathbf{y}(\Psi)$ is the maximal sieve. First, we unfold the meaning of $\Psi \Vdash \neg\neg S = \top$ using the Beth-Kripke-Joyal semantics of the topos as a weapon:

$$\Psi \Vdash \neg\neg S = \top \tag{6.1}$$

$$\forall \psi : \Phi \rightarrow \Psi. \neg(\Phi \Vdash \neg S \cdot \psi = \top) \tag{6.2}$$

$$\forall \psi : \Phi \rightarrow \Psi. \neg(\forall \phi : X \rightarrow \Phi. \neg(X \Vdash S \cdot \psi \cdot \phi = \top)) \tag{6.3}$$

$$\forall \psi : \Phi \rightarrow \Psi. \exists \phi : X \rightarrow \Phi. X \Vdash S \cdot \psi \cdot \phi = \top \tag{6.4}$$

$$\forall \psi : \Phi \rightarrow \Psi. \exists \phi : X \rightarrow \Phi. X \Vdash S \cdot (\psi \circ \phi) = \top \tag{6.5}$$

Now it suffices to show that $X \Vdash S \cdot (\psi \circ \phi) = \top$ iff $\psi \circ \phi \in S(X)$. (\Rightarrow) Unfolding the meaning of our assumption, we have for all $\rho : Y \rightarrow X$ that $\psi \circ \phi \circ \rho \in S(Y)$. Now choose $Y \triangleq X$ and $\rho \triangleq \text{id}_X$; therefore $\psi \circ \phi \in S(X)$. (\Leftarrow) Now suppose $\psi \circ \phi \in S(X)$. We have to show that for all $\rho : Y \rightarrow X$, then $\psi \circ \phi \circ \rho \in S(Y)$. This follows because sieves are closed under precomposition. \square

Remark 6.3. Observe the essential use of De Morgan’s duality in the passage between 6.3 and 6.4 above. The dense topology does *not* correspond to the double negation topology in a constructive metatheory; moreover, the version of the dense topology which *does* correspond to double negation often does not suffice for standard applications in a constructive metatheory, as the author observed in [14].

Applications of the dense topology The most famous use of the dense topology *qua* double negation is in modern proofs of the independence of the Continuum Hypothesis. Letting $\mathbb{P} : \mathbf{Pos}$ be a forcing poset, the sheaves on the site $(\mathbb{P}, J_{\text{dense}})$ form a boolean topos from which it is possible to obtain a model of ZFC in which the continuum hypothesis fails. The reader is referred to [8] for a summary of this construction.

⁶The definition of the dense topology that we have assumed is *not*, however, equivalent to the double negation topology in a constructive metatheory, since the equivalence relies on De Morgan duality [12, 4].

Another instance of the dense topology in practice is in its incarnation as the *atomic topology*, which can be imposed on any site which satisfies the *right Ore condition* defined below.

Definition 6.4 (Ore condition). A category \mathbb{C} satisfies the *right Ore condition* if every pair of morphisms $\psi_0 : \Phi_0 \rightarrow \Psi$ and $\psi_1 : \Phi_1 \rightarrow \Psi$ can be completed into a commutative square:

$$\begin{array}{ccc} \Phi & \dashrightarrow & \Phi_0 \\ \downarrow & & \downarrow \psi_0 \\ \Phi_1 & \xrightarrow{\psi_1} & \Psi \end{array}$$

Having pullbacks is a sufficient condition, but not a necessary one.

Definition 6.5 (Atomic topology). If \mathbb{C} satisfies the right Ore condition from Definition 6.4, then it is possible to impose the *atomic topology* on \mathbb{C} , where all inhabited sieves⁷ are covering:

$$J_{\text{at}}(\Psi) \triangleq \{ S \in \Omega(\Psi) \mid \cup_{\Phi} S(\Phi) \text{ inhabited} \}$$

If the reader prefers to work with pretopologies or coverages, they should be aware that the atomic topology is the one generated from *singleton* covering families.

Applications of the atomic topology Letting $\mathbf{Inj} : \mathbf{Cat}$ be the category of finite sets and only injective maps between them, the topos of sheaves on the site $(\mathbf{Inj}^{\text{op}}, J_{\text{at}})$ is known as the Schanuel topos, and is equivalent to the category of nominal sets [10]. Among other things, the geometry and logic of the Schanuel topos accounts for constructions which involve an *abundance* of atomic names which may be compared for equality.

It is helpful to consider what characteristics this sheaf subcategory has which $\widehat{\mathbf{Inj}}^{\text{op}}$ lacks. Because constructions within this presheaf topos must be stable under only injective maps of name contexts, it is clear that $\widehat{\mathbf{Inj}}^{\text{op}}$ justifies operations which depend on apartness of names.

In particular, defining the presheaf of “available names” $\mathbb{A} \triangleq \mathbf{y}(\{\bullet\})$ as the obvious representable functor, we can define the following natural transformation:

$$\begin{aligned} \text{test} : \mathbb{A} \times \mathbb{A} &\rightarrow \mathbf{2} \\ \text{test}_{\Psi}(\alpha, \beta) &\triangleq \begin{cases} \mathbf{t} & \text{if } \alpha = \beta \\ \mathbf{f} & \text{if } \alpha \neq \beta \end{cases} \end{aligned}$$

The above is well-defined/natural because injective maps are precisely those which preserve apartness.

⁷In classical sheaf theory, the covering sieves for the atomic topology are the *non-empty* ones; however, in a constructive metatheory, this is not enough to develop the necessary results, including Lemma 6.6.

In fact, we can do even more and use Day’s convolution to form a *separating product* and *separating function space* in $\widehat{\mathbf{Inj}}^{\text{op}}$, adding a locally monoidal-closed structure which is distinct from the standard locally cartesian closed structure of the topos [11].

At the level of the logic, this corresponds with four new connectives, $\{*, -*, \exists^*, \forall^*\}$ which are separating conjunction, separating implication, separating existential quantification, and separating universal quantification respectively.⁸ We can also define a *fresh name quantifier* $\mathcal{N}x.\phi(x)$ as $\forall^*x : \mathbb{A}.\phi(x)$, but we will see that this quantifier is not yet well-behaved.

In particular, this presheaf topos is not closed under a crucial principle, which is the “abundance” of fresh names. In particular, the principle $(\mathcal{N}x.\phi(\alpha)) \Rightarrow \phi(\alpha)$ fails to hold in the internal logic of $\widehat{\mathbf{Inj}}^{\text{op}}$. This principle, most properly understood as an instance of local character, is precisely what the sheaf subcategory $\mathbf{Sh}(\widehat{\mathbf{Inj}}^{\text{op}}, J_{\text{at}})$ is closed under.

As soon as the atomic topology has been imposed, the “freshness” quantifier becomes self-dual, in the following sense:

$$\exists^*x : \mathbb{A}.\phi(x, \alpha) = \mathcal{N}x.\phi(x, \alpha) = \forall^*x : \mathbb{A}.\phi(x, \alpha)$$

Lemma 6.6. *The atomic topology coincides with the dense topology.*

Proof. (\Rightarrow) Suppose $S \in J_{\text{at}}(\Psi)$, i.e. $S \in \Omega(\Psi)$ and $\cup_{\Phi} S(\Phi)$ *inhabited*. Fix $\psi : \Phi \rightarrow \Psi$; we have to exhibit some $\phi : X \rightarrow \Phi$ such that $\psi \circ \phi \in S(X)$. By assumption, there is some $\Upsilon : \mathbb{C}$ for which we have some $\psi' : \Upsilon \rightarrow \Psi$ such that $\psi' \in S(\Upsilon)$; by the right Ore condition (Definition 6.4), we have some $X : \mathbb{C}$ with the following property:

$$\begin{array}{ccc} X & \overset{\phi}{\dashrightarrow} & \Phi \\ \downarrow v & & \downarrow \psi \\ \Upsilon & \xrightarrow{\psi'} & \Psi \end{array}$$

Because sieves are closed under precomposition, we have $\psi' \circ v \in S(X)$; because the diagram above commutes, we therefore have $\psi \circ \phi \in S(X)$.

(\Leftarrow) Suppose $S \in J_{\text{dense}}(\Psi)$, i.e. for any $\psi : \Phi \rightarrow \Psi$ there exists some $\phi : X \rightarrow \Phi$ such that $\psi \circ \phi \in S(X)$. We have to exhibit some $\Upsilon : \mathbb{C}$ together with some $\psi' : \Upsilon \rightarrow \Psi$ such that $\psi' \in S(\Upsilon)$. Choose $\Phi \triangleq \Psi$ and $\psi \triangleq 1$; then, we have some $\phi : X \rightarrow \Psi$ such that $\phi \in S(X)$. Then choose $\Upsilon \triangleq X$ and $\psi' \triangleq \phi$. \square

Remark on constructivity One should be cautious about the numerous results in topos theory of the form “Any topos with property X is boolean” (e.g. “well-pointed”), which, far from elucidating an essential consequence of the property X , merely expose a leakage of information from the (Platonistic) external world into the topos. This kind of glitch serves only to underscore the essentially Tarskian deviation [7] which classical topos theory has inherited from old-fashioned mathematics and semantics.

⁸Be cautious about trying to naively develop this logic in the subobject lattice of the topos; Biering, Birkedal et al have shown that this construction degenerates into standard structural logic [1]. However, at the very least, the separating quantifiers can be developed in a sensible way.

To resist this deviation amounts to adopting Bishop's dictum that meaningful distinctions must be preserved; in doing so, we enter a profoundly alien world in which, for instance, the Schanuel topos is *not* boolean. We do not take a strident position on this here; our remarks are meant only to provide hope to the radical constructivist that it is possible to use these tools without incidentally committing oneself to a classical ontology.

References

- [1] B. Biering, L. Birkedal, and N. Torp-Smith. Bi-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Trans. Program. Lang. Syst.*, 29(5), Aug. 2007. 18
- [2] E. Bishop. *Foundations of Constructive Analysis*. McGraw-Hill, New York, 1967. 2
- [3] F. Borceux. *Handbook of Categorical Algebra 2 – Categories and Structures*. Cambridge Univ. Press, 1994. 14
- [4] T. Coquand. About Goodman’s theorem. *Annals of Pure and Applied Logic*, 164(4):437 – 442, 2013. 16
- [5] M. P. Fourman. Continuous Truth I: Non-constructive objects. In G. L. G. Lolli and A. Marcja, editors, *Logic Colloquium ’82*, volume 112 of *Studies in Logic and the Foundations of Mathematics*, pages 161 – 180. Elsevier, 1984. 2
- [6] M. P. Fourman. Continuous Truth II: Reflections. In L. Libkin, U. Kohlenbach, and R. de Queiroz, editors, *Logic, Language, Information, and Computation: 20th International Workshop, WoLLIC 2013, Darmstadt, Germany, August 20-23, 2013. Proceedings*, pages 153–167, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 2
- [7] J.-Y. Girard. *The Blind Spot: Lectures on Logic*. European Mathematical Society, Sept. 2011. 18
- [8] S. Mac Lane and I. Moerdijk. *Sheaves in geometry and logic : a first introduction to topos theory*. Universitext. Springer, New York, 1992. 1, 3, 4, 5, 8, 9, 16
- [9] K. Marx, E. Mandel, and B. Fowkes. *Capital: A Critique of Political Economy*. Number Volume 1 in Penguin classics. Penguin Books Limited, 2004. 13
- [10] A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, New York, NY, USA, 2013. 17
- [11] U. Schöpp and I. Stark. A dependent type theory with names and binding. In *Computer Science Logic: Proceedings of the 18th International Workshop CSL 2004*, number 3210 in Lecture Notes in Computer Science, pages 235–249. Springer-Verlag, 2004. 18
- [12] B. Spitters. The space of measurement outcomes as a spectrum for non-commutative algebras. In *Proceedings Sixth Workshop on Developments in Computational Models: Causality, Computation, and Physics, DCM 2010, Edinburgh, Scotland, 9-10th July 2010.*, pages 127–133, 2010. 16
- [13] T. Stacks Project Authors. Stacks Project. <http://stacks.math.columbia.edu>, 2017. 1
- [14] J. Sterling. Dense topology / double negation operator in a constructive metatheory? Mathematics Stack Exchange. <http://math.stackexchange.com/q/2068915>. 16