



UNIVERSITY OF
CAMBRIDGE

Classifying Toposes

Geno Racklin Asher

gnr21@cam.ac.uk

University of Cambridge

Contents

Introduction	3
1 Elementary Toposes	4
1.1 The Category of Sets	4
1.2 Examples of Elementary Toposes	6
1.3 The Lattice of Subobjects	8
2 Grothendieck Toposes	10
2.1 Sieves, Sites and Sheaves	10
2.2 Grothendieck Toposes are Elementary Toposes	12
2.3 The Associated Sheaf Functor	13
2.4 Geometric Morphisms	15
3 Categorical Logic	18
3.1 Models of Theories	18
3.2 Kripke-Joyal Semantics	21
3.3 An Application to Algebraic Geometry	22
4 Classifying Toposes	24
4.1 Geometric Logic	24
4.2 The Classifying Topos of a Theory	26
4.3 Examples of Classifying Toposes	29
Conclusion	32
Bibliography	32

Introduction

The notion of a “topos” has its roots in two separate mathematical disciplines: On the one hand, an *Elementary Topos* can be viewed as a generalised mathematical universe; that is, as a foundation for higher-order intuitionistic logic rather than the classical set-theoretic setting. From this perspective, a topos is simply a category sufficiently rich in categorical structure that it admits an “internal language” which can be reasoned about according to the deduction rules of intuitionistic logic, and we have a natural notion of models of theories inside such categories.

The second viewpoint – *Grothendieck Toposes* – arose from Grothendieck’s work in the 1960s reshaping algebraic geometry with the invention of scheme theory. This class of toposes is a generalisation of the category of sheaves on a topological space, and thus can be analysed via arguments of a geometric nature. Placing these “geometric” constraints on toposes imposes structure on their internal logic; in particular, we can show that the theories whose models are preserved by *geometric morphisms* between Grothendieck toposes (that is, functors which preserve the geometric structure) are those which admit axiomatisations whose axioms are all of a specific form. We call such theories *geometric theories*.

In this essay, we will introduce the basic theory of elementary and Grothendieck toposes, and show how logic can be interpreted in these categories. We will use the machinery we have built to show that for any geometric theory \mathbb{T} , there is a Grothendieck topos $\mathbf{Set}[\mathbb{T}]$ which contains a *universal model* of \mathbb{T} inside, in the sense that any model of \mathbb{T} in a Grothendieck topos can be viewed as the inverse image of this model along some geometric morphism. We call $\mathbf{Set}[\mathbb{T}]$ the *classifying topos* of \mathbb{T} .

The first section of the essay is devoted to Elementary Toposes. We consider mathematically-useful features of set theory as categorical properties of \mathbf{Set} , and use these features to define a topos. After exploring some key examples, we build up some important machinery, in particular showing that for any object in a topos, its category of subobjects forms a Heyting algebra. The material covered in this section broadly follows [1, Ch. I & IV].

Section 2 introduces Grothendieck toposes, following the material found in [2, Ch. 1.1]; we first define the objects of interest, in particular the notion of a *Grothendieck site* as a generalisation of a topological space. We then define the category of sheaves on a site, and give a proof based on [3, C2.2.6] that this category is an elementary topos. We then prove, following [1, Ch. III.5], the existence of an analogue to the “sheafification functor” found in algebraic geometry. Finally, we define a suitable notion of morphism between Grothendieck toposes, and prove an important equivalence between the category of such morphisms and a class of so-called *flat* functors.

In Section 3, we explore the internal logic of a topos. First, approximately following [2, Ch. 1.3], we define models of theories in toposes, and prove that the semantics of such models is intuitionistically sound. We next define the *Kripke-Joyal semantics*, which offers a way to determine validity of formulae in the *internal* language of a topos, and show that it admits a more concrete interpretation of “truth” in a topos. In the final subsection, based on [4], we offer an application of these logical constructions to algebraic geometry, showing how we can translate intuitionistically-valid statements from commutative algebra into theorems relating to sheaves of rings and modules in a scheme-theoretic setting.

Finally, Section 4 introduces the titular Classifying Topos. Drawing material from [2, Ch. 2] and [1, Ch. X], we first introduce geometric logic, showing that theories of this form are preserved under geometric morphisms between Grothendieck toposes. We then define classifying toposes, and construct a syntactic category associated to a geometric theory \mathbb{T} in order to prove that every such theory has a classifying topos $\mathbf{Set}[\mathbb{T}]$. We then explore some simple examples of classifying toposes, before finishing with a discussion of the Zariski topos, which we prove is the classifying topos for the theory of local rings.

1 Elementary Toposes

When studying axiomatic set theory as a foundation for mathematics, a natural question arises: *Why sets?* In other words: what are the properties of set theory which make it a convenient foundation for the rest of mathematics? More generally, are there other constructions with similar useful properties which might serve as alternative “universes” for mathematical logic? In this chapter we will explore these questions from a category-theoretic perspective; the mathematical “universes” we seek shall motivate our definition of an *elementary topos*.

1.1 The Category of Sets

Let us begin by recording some of the mathematically useful features and constructions in set theory, and see if they can be readily translated into categorical language as properties of the category **Set**.

- (i) *Subsets*: given a set S and a formula $\varphi(x)$ we can form the subset $\{x \in S \mid \varphi(x)\} \subset S$. In particular, the ability to form the set $\{x \in S \mid f(x) = g(x)\}$ for some functions f, g allows us to find “sets of solutions” to equations. In categorical terms this corresponds with the construction of equalisers in **Set**.
- (ii) *Products & singletons*: given two sets A, B , we can form their cartesian product $A \times B = \{(a, b) \mid a \in A, b \in B\}$; i.e. we can form the categorical product in **Set**. Moreover, any singleton set $\mathbf{1} = \{*\}$ is a terminal object so in fact **Set** has all finite products. Note that (i) and (ii) together imply that **Set** has all finite limits.
- (iii) *Disjoint unions & the empty set*: similarly to (ii), we can show that **Set** has finite coproducts. Indeed, we can form the binary coproduct of two sets A, B as their disjoint union $A \sqcup B$; and the empty set $\mathbf{0} = \emptyset$ is the unique initial object of **Set**, so the claim follows.
- (iv) *Quotients*: given an equivalence relation \sim on a set A , we can form the quotient A/\sim . Specifically, if we have functions $f, g : S \rightarrow A$, then we can construct their coequaliser as the quotient of A by the smallest equivalence relation containing $f(x) \sim g(x) \forall x \in S$. As above, (iii) and (iv) together imply that **Set** has all finite colimits.
- (v) *Functions*: given two sets A, B , we can form the set of functions, or *exponential*, $B^A = \{f : A \rightarrow B\}$. In particular, **Set** is cartesian closed, since for any set C we have a bijection $\mathbf{Hom}(A \times C, B) \cong \mathbf{Hom}(C, B^A)$ given by *currying*, and this bijection is natural in all three factors. Note that this defines an adjunction $(- \times A) \dashv (-)^A$ with counit the *evaluation map*

$$\begin{aligned} e : B^A \times A &\rightarrow B \\ (f, a) &\mapsto f(a). \end{aligned} \tag{1}$$

- (vi) *Power sets*: given a set X we can form its power set $PX = \{S \mid S \subset X\}$. This construction extends to a contravariant functor $P : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$, which sends a function $f : A \rightarrow B$ to its “inverse image function”

$$\begin{aligned} f^{-1} : PB &\rightarrow PA \\ S &\mapsto \{x \in A \mid f(x) \in S\}. \end{aligned} \tag{2}$$

Moreover, each subset S of X (i.e. each element of PX) has a *characteristic function*

$$\begin{aligned} \varphi_S : X &\rightarrow \mathbf{2} = \{0, 1\} \\ x &\mapsto \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{if } x \notin S \end{cases} \end{aligned} \tag{3}$$

such that the subset $S = \{x \in X \mid \varphi_S(x) = 1\}$ can be defined as the pullback along φ_S of the map $t : \mathbf{1} \rightarrow \{0, 1\}$ sending $* \mapsto 1$. More generally, we have the following definition:

Definition 1.1: Let \mathcal{C} be a category. If \mathcal{C} has an object Ω and a monomorphism $t : \mathbf{1} \rightarrow \Omega$, such that for any monomorphism $S \rightarrow X$ there exists a unique characteristic morphism $\varphi_S : X \rightarrow \Omega$ which forms a pullback square

$$\begin{array}{ccc} S & \longrightarrow & \mathbf{1} \\ \downarrow & & \downarrow t \\ X & \overset{\varphi_S}{\dashrightarrow} & \Omega \end{array} \quad (4)$$

then we say that t is a *subobject classifier* for \mathcal{C} . (We sometimes informally refer to the object Ω as a subobject classifier).

Note that if \mathcal{C} is locally small and has a subobject classifier, then it is well-powered and we write $\mathbf{Sub}_{\mathcal{C}}(X)$ for the (small) poset of subobjects of $X \in \text{ob } \mathcal{C}$; we note $\mathbf{Sub}_{\mathcal{C}}(X) \cong \mathbf{Hom}_{\mathcal{C}}(X, \Omega)$.

Now we might at this stage choose to define a topos as a category possessing all of the above “nice” properties. We shall instead take a more minimalist definition and then show that the remaining properties still hold – thus less effort will be required to verify that a given category is a topos.

Definition 1.2: An *elementary topos* is a cartesian closed category with all finite limits and a subobject classifier.

We note some initial properties of toposes:

Lemma 1.3: Elementary toposes are balanced.

Proof: Suppose $m : A \rightarrow B$ is both monic and epic. Let $!_B : B \rightarrow \mathbf{1}$ denote the unique morphism. Since m is monic, it is the equaliser of its characteristic morphism $\text{char}(m)$ with $t \circ !_B : B \rightarrow \Omega$. But since m is epi, we have $\text{char}(m) = t \circ !_B$, whence m is the equaliser of $\text{char}(m)$ with itself, i.e. an isomorphism. \square

Lemma 1.4: An elementary topos \mathcal{E} has *power objects*; that is, there is a contravariant functor $P : \mathcal{E}^{\text{op}} \rightarrow \mathcal{E}$ such that $\mathbf{Hom}_{\mathcal{E}}(B \times A, \Omega) \cong \mathbf{Hom}_{\mathcal{E}}(A, PB)$ naturally in A and B .

Proof: We can take P to be $\Omega^{(-)}$; then the desired natural isomorphism follows from the definition of the exponential. \square

Lemma 1.5: Elementary toposes have images; that is, every morphism f in a topos factors uniquely as $f = me$ for a monomorphism m and epimorphism e , and we call m the *image* of f .

Proof: See [1, Prop. IV.6.1]. \square

Proposition 1.6: Elementary toposes have finite colimits.

Proof (Sketch): See [1, Ch.IV.5]; the key idea of the proof is to note that P is right adjoint to P^{op} , and appeal to the crude monadicity theorem to show that P creates finite limits in \mathcal{E}^{op} , i.e. finite colimits in \mathcal{E} . \square

The structure of toposes thus defined will, in Section 3, allow us to interpret the language of set theory in terms of these categorical constructions; the curious reader might wonder whether this structure allows us to build *bona fide* models of Zermelo-Fraenkel set theory from particular toposes. Shulman’s paper [5] goes into some detail on the precise connections between properties of toposes and axioms of set theory, in particular giving a classification of those toposes which can be viewed as the category of sets of some model of ZFC.

1.2 Examples of Elementary Toposes

Now that we have defined a topos, let us consider some simple examples of toposes. For a great wealth of additional examples, see [1, Ch. I.1] or [6, Ch. 5.2].

Example 1.7: By the discussion in Section 1.1, we see that **Set** satisfies our definition. In fact, we claim that the full subcategory **Finset** of finite sets is also an elementary topos. Indeed, equalisers and finite products of finite sets are themselves finite, so we see that **Finset** has all finite limits. Moreover, if A and B are finite sets then the exponential B^A is also finite, so again **Finset** is cartesian closed. Finally, the subobject classifier $\mathbf{1} \mapsto \{0, 1\}$ of **Set** is finite and satisfies the universal property of a subobject classifier in **Finset**. So we conclude that **Finset** is indeed a topos.

Example 1.8: Given an arbitrary (small) category \mathcal{C} , we claim that the functor category $\mathcal{E} = [\mathcal{C}^{\text{op}}, \mathbf{Set}]$, known as the category of *presheaves* (of sets) on \mathcal{C} , is a topos. We note first that if P is such a functor and $x \in PC$ then given $f : D \rightarrow C$ we have $P(f)(x) \in PD$, and we write $x \cdot f$ for $P(f)(x)$.

Now, to show \mathcal{E} has finite limits we notice that we can take limits “pointwise”: for example, given $F, G : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ we can define $(F \times G)(C) = F(C) \times G(C) \forall C \in \text{ob } \mathcal{C}$. So we inherit all limits from **Set** in this way.

To form the exponential, we note first that if Q^P exists then $\mathbf{Hom}(R \times P, Q) \cong \mathbf{Hom}(R, Q^P)$. In particular, if R is a representable functor $R = \mathbf{Hom}(-, C) = \mathbf{y}(C)$ (where \mathbf{y} is the Yoneda embedding) then by the Yoneda Lemma

$$\begin{aligned} Q^P(C) &\cong \mathbf{Hom}(\mathbf{y}(C), Q^P) \\ &\cong \mathbf{Hom}(\mathbf{y}(C) \times P, Q). \end{aligned} \tag{5}$$

Taking this as the definition of Q^P gives a well-defined functor $\mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$, and it remains to check that $(-)^P$ thus defined is indeed right adjoint to $- \times P$. Define the evaluation map $e : Q^P \times P \rightarrow Q$ by

$$e_C(\theta, y) = \theta_C(1_C, y) \in Q(C) \tag{6}$$

for $\theta : \mathbf{y}(C) \times P \rightarrow Q$ and $y \in P(C)$. We claim this is the counit of an adjunction: given a natural transformation $\phi : R \times P \rightarrow Q$ we must show there exists a (unique) $\phi' : R \rightarrow Q^P$ making the following diagram commute:

$$\begin{array}{ccc} R \times P & & \\ \phi' \times 1_P \downarrow \text{dashed} & \searrow \phi & \\ Q^P \times P & \xrightarrow{e} & Q \end{array} \tag{7}$$

We can define, for each $C \in \text{ob } \mathcal{C}$ and $u \in RC$ a natural transformation $\phi'_C(u)$ as follows:

$$\begin{aligned} (\phi'_C(u))_D &: \mathbf{Hom}_e(D, C) \times P(D) \rightarrow Q(D) \\ (f, x) &\mapsto \phi_D(u \cdot f, x) \end{aligned} \tag{8}$$

Naturality of ϕ' follows immediately; and we check that the diagram (7) commutes:

$$\begin{aligned}
e_C(\phi'_C(u), y) &= \phi'_C(u)(1_C, y) \\
&= \phi_C(u, y).
\end{aligned} \tag{9}$$

So \mathcal{E} is indeed cartesian closed.

Finally, we must verify that \mathcal{E} has a subobject classifier. Suppose Ω is a subobject classifier for the category. By the Yoneda Lemma and the definition of a subobject classifier, we have an isomorphism

$$\Omega(C) \cong \mathbf{Hom}_{\mathcal{E}}(\mathbf{y}(C), \Omega) \cong \mathbf{Sub}_{\mathcal{E}}(\mathbf{y}(C)) \tag{10}$$

for each $C \in \text{ob } \mathcal{C}$. But a subobject of a presheaf P is, up to isomorphism, a *subfunctor* of P , i.e. a functor Q such that $QC \subset PC$ for all C , and such that $Q(f)$ is a restriction of $P(f)$ for all morphisms f . So we define

$$\Omega(C) = \{S \mid S \text{ a subfunctor of } \mathbf{y}(C)\}, \tag{11}$$

and for a morphism $f : C \rightarrow D$ and subfunctor $S \rightarrow \mathbf{y}(C)$ take $\Omega(f)(S)$ to be the pullback of this subobject along $\mathbf{y}(f) : \mathbf{y}(D) \rightarrow \mathbf{y}(C)$. Finally, we define the monomorphism $t : \mathbf{1} \rightarrow \Omega$ by

$$t_C(*) = \mathbf{y}(C) \text{ for each } C \in \text{ob } \mathcal{C}, \tag{12}$$

where $\mathbf{1}(C) = \{*\}$ for every C . Now suppose $R \rightarrow F$ is a subobject. We want to construct a characteristic map $\varphi : F \rightarrow \Omega$. We set

$$\varphi_C(x)(D) = \{f : D \rightarrow C \mid x \cdot f \in R(D)\} \tag{13}$$

and it is immediate that

$$\varphi_C(x) = \mathbf{y}(C) \iff x \in R(C), \tag{14}$$

so that the square in the definition of a subobject classifier is indeed a pullback. To check that φ is unique, suppose θ were some other natural transformation satisfying the pullback condition (14). Then for any $x \in F(C)$ and $f : D \rightarrow C$, we would have

$$\begin{aligned}
x \cdot f \in R(D) &\iff \theta_D(x \cdot f) = t_D \\
&\iff \theta_C(x) \cdot f = t_D \text{ (naturality of } \theta) \\
&\iff f \in \theta_C(x)(D).
\end{aligned} \tag{15}$$

So in fact $\theta = \varphi$, and the proof that \mathcal{E} is a topos is complete.

Example 1.9: Let \mathcal{E} be an elementary topos, and C an object of \mathcal{E} . The *slice category* \mathcal{E}/C has as its objects morphisms with codomain C , and its morphisms f are commutative diagrams

$$\begin{array}{ccc}
A & \xrightarrow{f} & B \\
& \searrow & \swarrow \\
& & C
\end{array} \tag{16}$$

\mathcal{E}/C can be shown to be a topos. For a proof of this fact, see [6, Thm. 5.8.1] - we will show the construction of finite limits here. \mathcal{E}/C clearly has as its terminal object the identity $\mathbf{1} : C \rightarrow C$. We can construct the product of $f : X \rightarrow C$ and $g : Y \rightarrow C$ as the pullback

$$\begin{array}{ccc}
X \times_C Y & \dashrightarrow & Y \\
\downarrow & & \downarrow f \\
X & \xrightarrow{g} & C
\end{array}
\tag{17}$$

so we see that \mathcal{E}/C has finite products. To construct the equaliser of two morphisms of \mathcal{E}/C we can simply take their equaliser in \mathcal{E} (with the composite morphism to C):

$$\begin{array}{ccccc}
E & \longrightarrow & X & \rightrightarrows & Y \\
& & \downarrow & & \swarrow \\
& & & & C
\end{array}
\tag{18}$$

So we see that indeed \mathcal{E}/C has finite limits. The subobject classifier and exponentials are constructed similarly from suitable pullback diagrams.

1.3 The Lattice of Subobjects

We finish the chapter with a discussion of $\mathbf{Sub}_{\mathcal{E}}(X)$, for a topos \mathcal{E} and an object $X \in \text{ob } \mathcal{E}$. In the case where $\mathcal{E} = \mathbf{Set}$, the poset of subobjects of some set X is just the power set PX , ordered by inclusion. In this setting, we have logical operations on sets (unions, intersections, etc.) which allow us to define additional structure on this poset. We hope to generalise this to the topos-theoretic case, using the rich structure of a topos to define similar operations on subobjects. In particular, we will show:

Theorem 1.10: For any object C of \mathcal{E} , $\mathbf{Sub}_{\mathcal{E}}(C)$ is a Heyting algebra.

Proof: We can construct the intersection of two subobjects X, Y as their pullback over C , and similarly we can construct their union as the pushout over $X \cap Y$.

$$\begin{array}{ccc}
X \cap Y & \longrightarrow & X \\
\downarrow & & \downarrow \\
Y & \longrightarrow & X \cup Y \\
& & \searrow \text{dashed} \\
& & C
\end{array}
\tag{19}$$

That these satisfy the universal properties of the greatest lower bound and least upper bound follow from the universal properties of the pullback and pushout (note also that we can view $X \cup Y$ as the image of the morphism $X + Y \rightarrow C$). $\mathbf{Sub}_{\mathcal{E}}(C)$ clearly has a top and a bottom element (namely, the identity 1_C and the unique map $0 \rightarrow C$) so we see that $\mathbf{Sub}_{\mathcal{E}}(C)$ has the structure of a bounded lattice. It remains to construct an implication operator \Rightarrow such that for all subobjects X, Y, Z of C ,

$$(X \cap Y) \leq Z \quad \text{iff} \quad X \leq (Y \Rightarrow Z) \tag{20}$$

Let $\varphi_{Y \cap Z}, \varphi_Y : C \rightrightarrows \Omega$ be the characteristic morphisms of $Y \cap Z$ and Y respectively, as subobjects of C . Define $Y \Rightarrow Z$ as their equaliser

$$(Y \Rightarrow Z) \longrightarrow C \begin{array}{c} \xrightarrow{\varphi_{Y \cap Z}} \\ \xrightarrow{\varphi_Y} \end{array} \Omega. \tag{21}$$

Now let $\chi : X \rightarrow C$ be any subobject. We have a diagram of pullbacks

$$\begin{array}{ccccc}
X \cap Y & \twoheadrightarrow & Y & \longrightarrow & \mathbf{1} \\
\downarrow & & \downarrow & & \downarrow t \\
X & \xrightarrow{\chi} & C & \xrightarrow{\varphi_Y} & \Omega
\end{array} \tag{22}$$

so we see that $\varphi_Y \circ \chi$ is the characteristic morphism of $X \cap Y$ as a subobject of X . Similarly $\varphi_{Y \cap Z} \circ \chi$ is the characteristic morphism for $X \cap (Y \cap Z)$. Thus

$$\begin{aligned}
X \leq (Y \Rightarrow Z) & \text{ iff } \chi \text{ factors through the equaliser (21)} \\
& \text{ iff } \varphi_Y \circ \chi = \varphi_{Y \cap Z} \circ \chi \\
& \text{ iff } X \cap Y = X \cap (Y \cap Z) \\
& \text{ iff } X \cap Y = (X \cap Y) \cap Z \\
& \text{ iff } X \cap Y \leq Z.
\end{aligned} \tag{23}$$

So $\mathbf{Sub}_{\mathcal{E}}(C)$ indeed has the structure of a Heyting algebra. \square

Remark 1.11: We will see later in Example 3.7 that the subobject classifier Ω is in fact an “internal” Heyting algebra, in the sense that for each of the Heyting algebra operations there are corresponding morphisms on Ω which obey the rules of a Heyting algebra written as commuting diagrams. In this sense, Ω “encodes” the logical structure of \mathcal{E} , and we can view the structure on $\mathbf{Sub}_{\mathcal{E}}(C) \cong \mathbf{Hom}_{\mathcal{E}}(C, \Omega)$ as induced by composition from the internal structure on Ω .

Now consider the fact that any morphism $f : X \rightarrow Y$ in \mathcal{E} induces a *pullback* functor $f^* : \mathbf{Sub}_{\mathcal{E}}(Y) \rightarrow \mathbf{Sub}_{\mathcal{E}}(X)$ sending a subobject $S \twoheadrightarrow Y$ to its pullback along f .

As a helpful example, consider the case where $\mathcal{E} = \mathbf{Set}$, and f is the projection function $\pi : X \times Y \rightarrow X$. In this case, $\pi^* : PX \rightarrow P(X \times Y)$ sends $S \mapsto \{(x, y) \mid x \in S, y \in Y\}$. Then we can show that π^* has a left adjoint given by

$$T \mapsto \{x \mid \exists y \in Y (x, y) \in T\} \tag{24}$$

and a right adjoint given by

$$T \mapsto \{x \mid \forall y \in Y (x, y) \in T\}. \tag{25}$$

We will generalise this example to arbitrary toposes:

Proposition 1.12: The pullback functor f^* has a left adjoint $\exists_f : \mathbf{Sub}_{\mathcal{E}}(X) \rightarrow \mathbf{Sub}_{\mathcal{E}}(Y)$.

Proof: Consider a subobject $s : S \twoheadrightarrow X$; we define $\exists_f(S)$ to be the image of the composite $f \circ s$ (by Lemma 1.5). Now, given a subobject $r : R \twoheadrightarrow Y$ such that $S \leq f^*(R)$, then $f \circ s$ factors through r , whence so does $\exists_f(S)$. So we have $\exists_f(S) \leq R$. Conversely if $\exists_f(S) \leq R$, then we know $f \circ s$ factors through r , and thus by definition of the pullback $S \leq f^*(R)$. \square

We want to show that f^* also has a right adjoint: to prove this, we’ll first consider that pullback along f similarly induces a functor $\widehat{f}^* : \mathcal{E}/Y \rightarrow \mathcal{E}/X$.

Lemma 1.13: $\widehat{f}^* : \mathcal{E}/Y \rightarrow \mathcal{E}/X$ has a right adjoint $\Pi_f : \mathcal{E}/X \rightarrow \mathcal{E}/Y$.

Proof: Since \mathcal{E}/Y is itself a topos, and we have an isomorphism $\mathcal{E}/X \cong (\mathcal{E}/Y)/(f)$, we may reduce to the case where $Y = \mathbf{1}$, so $\mathcal{E}/\mathbf{1} \cong \mathcal{E}$. Then pullback along the unique $X \rightarrow \mathbf{1}$ is just the product functor $C \mapsto (\pi_C : C \times X \rightarrow X)$.

Let $h : D \rightarrow X$ be an object in \mathcal{E}/X . Note that composition with h induces a morphism $h^X : D^X \rightarrow X^X$. Morphisms $\pi_C \rightarrow h$ correspond to $t : C \times X \rightarrow D$ such that $ht = \pi_C$. Transposing

across the exponential adjunction, these correspond to $t' : C \rightarrow D^X$ such that $h^X \circ t' : C \rightarrow X^X$ is equal to $j_X \circ !_C$, where j_X is the transpose of 1_X . But these in turn correspond to morphisms $t'' : C \rightarrow \Pi h$, where Πh is defined to be the pullback of j_X along h^X .

$$\begin{array}{ccc}
C & \overset{t''}{\dashrightarrow} & \Pi h & \longrightarrow & D^X \\
& & \downarrow & & \downarrow h^X \\
& & \mathbf{1} & \xrightarrow{j_X} & X^X.
\end{array} \tag{26}$$

So we see that $h \mapsto \Pi h$ gives the desired right adjoint. \square

Proposition 1.14: $f^* : \mathbf{Sub}_{\mathcal{E}}(Y) \rightarrow \mathbf{Sub}_{\mathcal{E}}(X)$ has a right adjoint, denoted \forall_f .

Proof: Π_f is a right adjoint, so preserves the terminal object and monomorphisms in \mathcal{E}/X . So it restricts to a functor on subobjects of the terminal object X , i.e. Π_f restricts to a functor $\forall_f : \mathbf{Sub}_{\mathcal{E}}(X) \rightarrow \mathbf{Sub}_{\mathcal{E}}(Y)$ which is right adjoint to f^* . \square

The suggestive use of the symbols \exists and \forall is no coincidence; we will see in Section 3.1 that these adjoints can be used to interpret logical quantifiers in the internal language of a topos. But for now, we note simply that the existence of these adjoints means that f^* preserves limits and colimits, and thus commutes with all the operations defined above; that is, f^* is a Heyting algebra homomorphism.

So we have seen that the subobject lattices possess a good deal of structure, and moreover that this structure is preserved by pullback along morphisms of \mathcal{E} .

2 Grothendieck Toposes

Having defined our categories of interest, let's now hone in on a specific class of examples. We'll focus our attention on a class of toposes which can be reasoned about geometrically; namely, the categories of *sheaves* on a (generalised) space. This class of examples should be familiar to those with a background in scheme-theoretic algebraic geometry, as it generalises the notion of a sheaf (of sets) ubiquitous in the study of schemes. When we define a notion of “logic internal to a topos”, we will see that this family of toposes will admit a geometric interpretation of our semantics, which can then, for example, be applied to give simple proofs of facts in algebraic geometry.

As in the case of elementary toposes, we shall start by defining the structures in question and proving their essential properties. We will then exhibit some useful constructions, which will arise as generalisations of constructions in algebraic geometry.

2.1 Sieves, Sites and Sheaves

We begin by recalling that if X is a topological space, we define the category of *presheaves* on X to be the functor category $[\mathcal{O}(X)^{\text{op}}, \mathbf{Set}]$, where $\mathcal{O}(X)$ is the poset of open subsets of X ordered by inclusion. If F is a presheaf on X , we have a *restriction map* $F(U) \rightarrow F(V)$ induced by each inclusion of open subsets $V \subset U$, and we call elements of $F(U)$ *sections*. A *sheaf* F on X is then a presheaf satisfying the following additional property:

If $(U_i \mid i \in I)$ is an open cover of U (that is, $\bigcup_{i \in I} U_i = U$), and $(s_i \mid i \in I)$ is a family of elements of $F(U_i)$ such that for every pair (i, j) , the restrictions of s_i and s_j to $U_i \cap U_j$ agree, then there exists a unique section $s \in F(U)$ such that the restriction of s to each U_i equals s_i .

We then define the category $\mathbf{Sh}(X)$ to be the full subcategory of $[\mathcal{O}(X)^{\text{op}}, \mathbf{Set}]$ of presheaves satisfying the sheaf condition above.

We seek to generalise this notion to define sheaves on general categories \mathcal{C} , rather than this specific case where $\mathcal{C} = \mathcal{O}(X)$, the poset of open subsets of a topological space; we will prove that the category of sheaves is in fact a topos. To do so, we need to be able to generalise the notion of an “open cover” of an object in a category \mathcal{C} . Our goal, then, is to define topologies on general categories whose open covers can be reasoned about in a similar way to the usual case. Consider the following properties of open covers:

- (i) $\{V \mid V \text{ is an open subset of } U\}$ forms an open cover of U ;
- (ii) If the family $(U_i \mid i \in I)$ is an open cover of U , and $V \subset U$, then $(U_i \cap V \mid i \in I)$ is an open cover of V ;
- (iii) If $(U_i \mid i \in I)$ is an open cover of U , and $(V_j \mid j \in J)$ is such that for each i , $(V_j \cap U_i \mid j \in J)$ is an open cover of U_i , then $(V_j \mid j \in J)$ is an open cover of U as well.

We will use these three properties as the basis for our generalisation of the notion of “open cover”, considering first the following definition:

Definition 2.1: Let \mathcal{C} be a category, and $C \in \text{ob } \mathcal{C}$. A *sieve* on C is a family of morphisms in \mathcal{C} with common codomain C , such that for any $f \in S$, we have $f \circ g \in S$ for any morphism g in \mathcal{C} composable with f . If $T \subset S$, we say that T *generates* S if any $f \in S$ factors through some morphism in T .

If \mathcal{C} is a poset, then sieves are just downward-closed subsets: in the example where $\mathcal{C} = \mathcal{O}(X)$, we can identify open covers of a subset U with sieves on U such that $\bigcup_{f \in S} \text{dom}(f) = U$. So we can reformulate our rules for open covers into statements about sieves, motivating our next batch of definitions:

Definition 2.2: Let \mathcal{C} be a category. A *Grothendieck topology* J on \mathcal{C} specifies, for each $C \in \text{ob } \mathcal{C}$, a collection $J(C)$ of sieves on C satisfying:

- (i) (*Maximality*) For any object C , the *maximal sieve* $\mathfrak{m}_C = \{f \mid \text{cod}(f) = C\}$ belongs to $J(C)$.
- (ii) (*Stability under pullback*) For any $f : D \rightarrow C$ in \mathcal{C} and any $S \in J(C)$, the sieve $f^*(S) = \{g : E \rightarrow D \mid f \circ g \in S\}$ belongs to $J(D)$.
- (iii) (*Transitivity*) Whenever S is a sieve on C and $T \in J(C)$, if $f^*(S) \in J(\text{dom}(f))$ for all $f \in T$ then $S \in J(C)$.

The sieves $S \in J(C)$ are referred to as *J-covering sieves* for C .

Definition 2.3: A *site* is a pair (\mathcal{C}, J) , where \mathcal{C} is a category and J a Grothendieck topology on \mathcal{C} .

Definition 2.4: A *morphism of sites* $F : (\mathcal{C}, J) \rightarrow (\mathcal{D}, K)$ is a finite-limit-preserving functor $F : \mathcal{C} \rightarrow \mathcal{D}$ such that for each $S \in J(C)$, the image of S under F generates a K -covering sieve on $F(C)$.

Finally, given a site (\mathcal{C}, J) we can specify a sheaf condition for presheaves $\mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$:

Definition 2.5: Let (\mathcal{C}, J) be a site. A presheaf $P : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ is called a *sheaf* on (\mathcal{C}, J) if the following condition holds:

If $S \in J(C)$ is a J -covering sieve and $(s_f \in P(\text{dom}(f)) \mid f \in S)$ is a family of sections of P such that $s_f \cdot g = s_{f \circ g}$ for any $f \in S$ and any g composable with f , then there exists a unique $s \in P(C)$ such that $s \cdot f = s_f$ for every $f \in S$. We call $(s_f \mid f \in S)$ a *matching family*, and s its *amalgamation*.

We denote the full subcategory of $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$ of sheaves on (\mathcal{C}, J) by $\mathbf{Sh}(\mathcal{C}, J)$.

Definition 2.6: A category \mathcal{E} is called a *Grothendieck Topos* if it is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$ for some site (\mathcal{C}, J) .

We'll see in Section 2.2 that this terminology is justified, i.e. that $\mathbf{Sh}(\mathcal{C}, J)$ is indeed a topos. First though, let's see a couple of examples of Grothendieck topologies.

Example 2.7: The simplest example of a Grothendieck topology on any category \mathcal{C} is the *trivial topology*, whose only covering sieves are the maximal sieves \mathfrak{m}_C for each $C \in \text{ob } \mathcal{C}$. In this case, the sheaf condition becomes trivial, such that every presheaf $P : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ is a sheaf. Therefore if J is the trivial topology on \mathcal{C} , then $\mathbf{Sh}(\mathcal{C}, J) = [\mathcal{C}^{\text{op}}, \mathbf{Set}]$.

In particular, if we take $\mathcal{C} = \mathbf{1}$ the trivial one-object category, then we see that $\mathbf{Set} \simeq [\mathbf{1}, \mathbf{Set}]$ is a Grothendieck topos.

Example 2.8: Let \mathcal{C} be a *frame* (i.e. a complete Heyting algebra). The *canonical topology* on \mathcal{C} is defined by taking $J(C)$ to be the collection of sieves on C generated by families $\{D_i \leq C \mid i \in I\}$ such that $\bigvee_{i \in I} D_i = C$.

We note that if $\mathcal{C} = \mathcal{O}(X)$ for some topological space X , then the canonical topology on \mathcal{C} is induced by the topology on X (with sieves generated by open covers). In particular, the sheaf condition of Definition 2.5 is identical under this topology to the condition defined at the start of Section 2.1. So we have $\mathbf{Sh}(\mathcal{O}(X), J) = \mathbf{Sh}(X)$, the category of sheaves on X .

We'll see more examples of sites in our later discussions of Classifying Toposes; for now we will continue proving important properties of the corresponding categories of sheaves.

2.2 Grothendieck Toposes are Elementary Toposes

Let's now prove that Grothendieck Toposes indeed possess all the structure of an Elementary Topos, following the direction of proof of [3, C2.2.6]. Throughout, we assume (\mathcal{C}, J) is a small site.

Proposition 2.9: $\mathbf{Sh}(\mathcal{C}, J)$ is complete and the inclusion $i : \mathbf{Sh}(\mathcal{C}, J) \hookrightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ preserves small limits.

Proof: Since limits in $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$ are taken pointwise (see Example 1.8), it follows from the definition of an amalgamation of a matching family that such amalgamations commute with limits. In other words, given $F = \lim_{i \in I} F_i$, a limit of a small diagram of sheaves in $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$, a matching family $(x_f \mid f \in S)$ for F is just the limit of a diagram of matching families $(x_{i,f} \mid f \in S)$ for F_i . Let x_i be the amalgamations of these families guaranteed by the sheaf condition for F_i , then $x = \lim_{i \in I} x_i$ is the unique amalgamation for $(x_f \mid f \in S)$. So it follows that $\mathbf{Sh}(\mathcal{C}, J)$ is closed under (pointwise) limits, and that these limits coincide with those in $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$. \square

Proposition 2.10: $\mathbf{Sh}(\mathcal{C}, J)$ is cartesian closed.

Proof: We know by Proposition 2.9 that $\mathbf{Sh}(\mathcal{C}, J)$ is closed under finite products, so it suffices to show that for sheaves F, G in $\mathbf{Sh}(\mathcal{C}, J)$, the presheaf exponential F^G is itself a sheaf. In fact, we will show that F^G is a sheaf for *any* presheaf G . So suppose $S \in J(C)$, and $(\alpha_f \mid f \in S)$ is a matching family for F^G . Recall that we defined

$$F^G(C) = \mathbf{Hom}(\mathbf{y}(C) \times G, F) \quad (27)$$

in Example 1.8. We must find a unique $\alpha : \mathbf{y}(C) \times G \rightarrow F$ such that $\alpha \cdot f = \alpha_f \forall f \in S$.

Given any $f : D \rightarrow C$ and $y \in G(C)$, consider the family $(\alpha_{f \cdot g}(1, y \cdot g) \mid g \in f^*S)$. This is a matching family for F since α_f is a matching family. So it has a unique amalgamation $\beta_{f,y} \in$

$F(C)$ since F is a sheaf. Then defining $\alpha(f, y) = \beta_{f, y}$ gives a unique amalgamation for the original matching family $(\alpha_f \mid f \in S)$, as required. \square

Proposition 2.11: $\mathbf{Sh}(\mathcal{C}, J)$ has a subobject classifier.

Proof: Recall from Example 1.8 that we defined the subobject classifier Ω for $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$ to be

$$\Omega(C) = \{S \mid S \text{ a subfunctor of } \mathbf{y}(C)\}. \quad (28)$$

But subfunctors of $\mathbf{y}(C)$ correspond bijectively to *sieves* on C , where, given a subfunctor S , we view $S(D)$ as the set of morphisms in the sieve with domain D . So we could just as well have defined

$$\begin{aligned} \Omega(C) &= \{S \mid S \text{ a sieve on } C\} \\ \Omega(f)(S) &= f^*S \end{aligned} \quad (29)$$

to the same effect. We define a sieve S on C to be *J-closed* if, whenever $f : D \rightarrow C$ satisfies $f^*S \in J(D)$ then $f \in J(C)$. Then

$$\Omega_J(C) = \{S \mid S \text{ a } J\text{-closed sieve on } C\} \quad (30)$$

defines a subfunctor of Ω .

Now given a sheaf P and a subfunctor $Q \vDash P$, then any matching family for Q has an amalgamation in P ; Q is a sheaf iff all such amalgamations lie in Q . That is, whenever $x \in P(C)$ and $S \in J(C)$, if $x \cdot f \in Q(D)$ for all $f : D \rightarrow C$ belonging to S , then $x \in Q(C)$. In other words, Q is a sheaf iff the sieve

$$\{f : D \rightarrow C \mid x \cdot f \in Q(D)\} \quad (31)$$

is *J-closed* for every $x \in P(C)$. But (31) is precisely the characteristic map (13) for Q evaluated at x , so Q is a sheaf iff its characteristic map factors through $\Omega_J \vDash \Omega$.

So it suffices to show that Ω_J is a sheaf, then it will be the subobject classifier for $\mathbf{Sh}(\mathcal{C}, J)$ along with the map

$$\begin{aligned} t : \mathbf{1} &\rightarrow \Omega_J \\ t_C(*) &= \mathbf{m}_C. \end{aligned} \quad (32)$$

Indeed, if $(R_f \mid f \in S)$ is a matching family of *J-closed* sieves, then the sieve

$$R = \{g : D \rightarrow C \mid R_{gh} = \mathbf{m}_E \text{ for all } (h : E \rightarrow D) \in g^*S\} \quad (33)$$

is the unique *J-closed* sieve on C satisfying $f^*R = R_f$ for all f , so the matching family has a unique amalgamation. \square

Theorem 2.12: $\mathbf{Sh}(\mathcal{C}, J)$ is an elementary topos.

Proof: Immediate from the preceding propositions. \square

2.3 The Associated Sheaf Functor

An essential construction in the setting of algebraic geometry is that of the *sheafification functor* $[\mathcal{O}(X)^{\text{op}}, \mathbf{Set}] \rightarrow \mathbf{Sh}(X)$, which associates to each presheaf on a topological space a corresponding sheaf.

In this section, we aim to generalise this construction by showing that the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \hookrightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ has a left adjoint \mathbf{a} , known as the *associated sheaf functor*.

We start with the following definition:

Definition 2.13: Let (\mathcal{C}, J) be a site. A presheaf $P : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ is *separated* if, for $S \in J(C)$, any matching family $(s_f \mid f \in S)$ has *at most one* amalgamation $s \in P(C)$.

This is a weaker condition than the sheaf condition which requires that such an amalgamation exist. We shall find a functor $P \mapsto P^+$ which sends presheaves to separated presheaves, and separated presheaves to sheaves. Then we will see that defining $\mathbf{a}(P) = P^{++}$ yields the desired adjoint.

Definition 2.14: Let S be a J -covering sieve. We write $\text{Match}(S, P)$ for the set of matching families for S . Then we define

$$P^+(C) = \lim_{\rightarrow S \in J(C)} \text{Match}(S, P) \quad (34)$$

taking the colimit over the set of J -covering sieves, partially ordered by reverse inclusion.

The elements of $P^+(C)$ can be viewed as equivalence classes of matching families $(x_f \mid f \in S)$, where $(x_f \mid f \in S) \sim (y_f \mid f \in R)$ iff there is $T \subset S \cap R$ such that $x_f = y_f \ \forall f \in T$. Then P^+ can be viewed as a presheaf $\mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$, where if $g : D \rightarrow C$ we have

$$P^+(g)(x_f \mid f \in S) = (x_{gf'} \mid f' \in g^*S), \quad (35)$$

noting that this action on matching families descends to a well-defined action on equivalence classes.

Moreover, the assignment $P \mapsto P^+$ is itself functorial, since a map $P \rightarrow Q$ of presheaves induces a corresponding map $P^+ \rightarrow Q^+$. We also have a canonical map $\eta : P \rightarrow P^+$ defined by

$$\eta_C(x) = (x \cdot f \mid f \in \mathfrak{m}_C). \quad (36)$$

We now show that P^+ has the required ‘‘sheafification’’ properties:

Proposition 2.15: Let P be a presheaf.

- (i) P^+ is separated;
- (ii) If P is separated, P^+ is a sheaf.

Proof:

- (i) Suppose $\mathbf{x} = (x_f \mid f \in S)$, $\mathbf{y} = (y_g \mid g \in R) \in P^+(C)$, such that

$$\mathbf{x} \cdot h = \mathbf{y} \cdot h \quad \text{for all } h \in T. \quad (37)$$

This means that for each $h \in T$, there is $V_h \subset h^*S \cap h^*R$ such that $x_{hv} = y_{hv}$ for all $v \in V_h$. Then we can define a sieve $V = \{hv \mid h \in T, v \in V_h\} \subset S \cap R$.

But $V \in J(C)$ by transitivity, and $(x_f \mid f \in V) = (y_f \mid f \in V)$ so in fact $\mathbf{x} = \mathbf{y}$. So P^+ is separated.

- (ii) Suppose we have a matching family $(\mathbf{x}_f \mid f \in R)$ in P^+ . For each $f \in R$, we thus have $\mathbf{x}_f = (x_{f,g} \mid g \in S_f)$, itself a matching family in P . Let $T = \{f \circ g \mid f \in R, g \in S_f\}$, a covering sieve by transitivity. Then we define $\mathbf{y} \in P^+(C)$ by

$$\mathbf{y}_{f \circ g} = x_{f,g}. \quad (38)$$

\mathbf{y} is well-defined (independent of the choice of factorisations $f \circ g$), since the $(x_{f,g} \mid g \in S_f)$ are matching families; and \mathbf{y} is itself a matching family, because $(\mathbf{x}_f \mid f \in S)$ is matching - so it is indeed an element of $P^+(C)$. Then by definition, \mathbf{y} is an amalgamation for \mathbf{x}_f , and it must be unique by (i). So P^+ is a sheaf. □

Finally, we can prove the existence of the associated sheaf functor.

Theorem 2.16: The inclusion $i : \mathbf{Sh}(\mathcal{C}, J) \hookrightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ has a left adjoint.

Proof: Consider P^{++} . By Proposition 2.15, two applications of $(-)^+$ to any presheaf produce a J -sheaf, so this functor factors through i to give a functor $\mathbf{a} : [\mathcal{C}^{\text{op}}, \mathbf{Set}] \rightarrow \mathbf{Sh}(\mathcal{C}, J)$.

We have a canonical morphism $P \rightarrow i(\mathbf{a}(P)) = P^{++}$ given by the composition of $\eta_P : P \rightarrow P^+$ with $\eta_{P^+} : P^+ \rightarrow P^{++}$; we must show that this morphism is the unit of an adjunction, and thus universal in the sense that any map $\varphi : P \rightarrow F$ into a sheaf F factors through the canonical morphism. To do so, we will show that η itself is universal (and hence, that a composition of two copies of η is): we must find $\varphi^+ : P^+ \rightarrow F$ such that $\varphi = \varphi^+ \circ \eta$.

Indeed, given $(x_f \mid f \in S) \in P^+(C)$, we have that $(\varphi(x_f) \mid f \in S)$ is a matching family for F and thus has a unique amalgamation $y \in F(C)$. But then if we define $\varphi^+(x_f \mid f \in S) = y$, we see that whenever $x \in P(C)$, $\varphi(x)$ is the amalgamation of $(\varphi(x \cdot f) \mid f \in S)$, i.e. $\varphi(x) = \varphi^+(\eta(x))$. \square

2.4 Geometric Morphisms

Equipped with the notion of Grothendieck topos, a natural question now presents itself: is there a corresponding notion of morphisms between Grothendieck toposes? To answer this question in general, we first return to our motivating example; that is, the category of sheaves $\mathbf{Sh}(X)$ on a topological space X . Given a map $f : X \rightarrow Y$ between two topological spaces, there are two maps induced on their categories of sheaves:

- An *direct image functor* $f_* : \mathbf{Sh}(X) \rightarrow \mathbf{Sh}(Y)$, defined by

$$f_* F(V) = F(f^{-1}(V)); \quad (39)$$

- An *inverse image functor* $f^* : \mathbf{Sh}(Y) \rightarrow \mathbf{Sh}(X)$, where $f^* F$ is defined to be the sheafification of the presheaf given by

$$U \mapsto \lim_{\rightarrow V \supset f(U)} F(V) \quad (40)$$

Interestingly, we find [7, 2.7.1] that these induced maps $\mathbf{Sh}(X) \rightleftarrows \mathbf{Sh}(Y)$ form an adjunction ($f^* \dashv f_*$)! We will let this fact motivate our definition of morphisms between Grothendieck toposes.

Definition 2.17: Let \mathcal{E} and \mathcal{F} be Grothendieck toposes. A *geometric morphism* $f : \mathcal{E} \rightarrow \mathcal{F}$ is a pair of adjoint functors $(f^* : \mathcal{F} \rightarrow \mathcal{E}) \dashv (f_* : \mathcal{E} \rightarrow \mathcal{F})$, such that the left adjoint f^* preserves finite limits. f^* and f_* are referred to respectively as the *inverse image* and *direct image* of f .

Definition 2.18: A *geometric transformation* $f \rightarrow g$ between two geometric morphisms is a natural transformation $f^* \rightarrow g^*$. We write $\mathbf{Geom}(\mathcal{E}, \mathcal{F})$ for the category of geometric morphisms $\mathcal{E} \rightarrow \mathcal{F}$ and geometric transformations between them.

We have already seen one nontrivial example of a geometric morphism:

Example 2.19: Let (\mathcal{C}, J) be a small site. There is a geometric morphism $\mathbf{Sh}(\mathcal{C}, J) \hookrightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ whose direct image is the inclusion functor and whose inverse image is the associated sheaf functor \mathbf{a} defined in Section 2.3 (for a proof that \mathbf{a} indeed preserves finite limits, refer to remarks at the end of [1, Ch. III.5]).

In order to construct more interesting examples of geometric morphisms, we seek to better understand the structure of the category $\mathbf{Geom}(\mathcal{E}, \mathcal{F})$: we will start with the special case $\mathcal{F} = [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ then generalise to arbitrary Grothendieck toposes.

Proposition 2.20: Let \mathcal{C} be a small category, and \mathcal{E} a cocomplete and locally small category. Then for any $A : \mathcal{C} \rightarrow \mathcal{E}$, the functor $\underline{\mathbf{Hom}}_{\mathcal{E}}(A, -) : \mathcal{E} \rightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ defined by

$$\underline{\mathbf{Hom}}_{\mathcal{E}}(A, E)(C) = \mathbf{Hom}_{\mathcal{E}}(A(C), E) \quad (41)$$

has a left adjoint, which we denote by

$$(-) \otimes_{\mathcal{E}} A : [\mathcal{C}^{\text{op}}, \mathbf{Set}] \rightarrow \mathcal{E}. \quad (42)$$

Proof: Given $P : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$, we write $\int P$ for the *category of elements* of P , whose objects are pairs (C, x) with $x \in P(C)$, and whose morphisms $(C, x) \rightarrow (D, y)$ are those $f : C \rightarrow D$ such that $y \cdot f = x$. Let $\pi_P : \int P \rightarrow \mathcal{C}$ be the projection functor. Then we can define our “generalised tensor product” as the colimit

$$P \otimes_{\mathcal{E}} A = \varinjlim (A \circ \pi_P). \quad (43)$$

We claim that this functor is left adjoint to $\underline{\mathbf{Hom}}_{\mathcal{E}}(A, -)$. Indeed, a natural transformation $\alpha : P \rightarrow \underline{\mathbf{Hom}}_{\mathcal{E}}(A, E)$ is given by a family $(\alpha_C : P(C) \rightarrow \mathbf{Hom}(A(C), E))$ indexed by $\text{ob } \mathcal{C}$, satisfying the naturality condition. But then we can equivalently view α as a family $(\alpha_C(x) : A(C) \rightarrow E)$ of morphisms in \mathcal{E} , indexed by $C \in \text{ob } \mathcal{C}$ and $x \in P(C)$: i.e. the objects of $\int P$. In this interpretation, the naturality condition is precisely the statement that for each $f : C \rightarrow D$ in \mathcal{C} , the diagram

$$\begin{array}{ccc} A\pi_P(C, x) & \xrightarrow{f} & A\pi_P(D, y) \\ & \searrow \alpha_C(x) & \swarrow \alpha_D(y) \\ & E & \end{array} \quad (44)$$

commutes. But (44) is the statement that the family $(\alpha_C(x) \mid (C, x) \in \int P)$ form a cocone under the diagram $A \circ \pi_P$, and thus induce a unique map

$$P \otimes_{\mathcal{E}} A = \varinjlim (A \circ \pi_P) \rightarrow E, \quad (45)$$

whence we have a natural bijection $\mathbf{Hom}(P, \underline{\mathbf{Hom}}_{\mathcal{E}}(A, E)) \cong \mathbf{Hom}(P \otimes_{\mathcal{E}} A, E)$. \square

If we think of the Hom-tensor adjunction exhibited here as a generalisation of the familiar tensor product seen in commutative algebra, we can also transfer the notion of a flat module to this setting; recall that an R -module M is flat if the functor $(-) \otimes_R M$ is (left) exact.

Definition 2.21: Let $A : \mathcal{C} \rightarrow \mathcal{E}$ be a functor. We say that A is *flat* if the functor $(-) \otimes_{\mathcal{E}} A : [\mathcal{C}^{\text{op}}, \mathbf{Set}] \rightarrow \mathcal{E}$ preserves finite limits. The full subcategory of $[\mathcal{C}, \mathcal{E}]$ consisting of flat functors is denoted $\mathbf{Flat}(\mathcal{C}, \mathcal{E})$.

The notion of flat functors turns out to be intrinsically linked to the notion of geometric morphisms; indeed, we have the following theorem:

Theorem 2.22 (Diaconescu’s Equivalence): Let \mathcal{E} be a Grothendieck topos, and \mathcal{C} a small category. Then there is an equivalence of categories $\mathbf{Geom}(\mathcal{E}, [\mathcal{C}^{\text{op}}, \mathbf{Set}]) \simeq \mathbf{Flat}(\mathcal{C}, \mathcal{E})$.

Proof: Suppose $A : \mathcal{C} \rightarrow \mathcal{E}$ is a flat functor. Then we can define a geometric morphism $\tau(A) : \mathcal{E} \rightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ whose inverse image is given by $\tau(A)^* = (-) \otimes_{\mathcal{E}} A$ and whose direct image is given by $\tau(A)_* = \underline{\mathbf{Hom}}_{\mathcal{E}}(A, -)$. This is indeed a geometric morphism since $(-) \otimes_{\mathcal{E}} A$ preserves finite limits by flatness of A , and the functors are adjoint by Proposition 2.20. Natural

transformations $A \rightarrow B$ induce transformations $(-) \otimes_{\mathcal{C}} A \rightarrow (-) \otimes_{\mathcal{C}} B$, so we can extend this construction to a functor $\tau : \mathbf{Flat}(\mathcal{C}, \mathcal{E}) \rightarrow \mathbf{Geom}(\mathcal{E}, [\mathcal{C}^{\text{op}}, \mathbf{Set}])$.

Conversely suppose $f : \mathcal{E} \rightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ is a geometric morphism. Then we can define a functor $\rho(f) = f^* \circ \mathbf{y} : \mathcal{C} \rightarrow \mathcal{E}$. Then we have

$$\begin{aligned} \underline{\mathbf{Hom}}_{\mathcal{E}}(\rho(f), F)(C) &= \mathbf{Hom}(f^* \circ \mathbf{y}(C), F) \\ &\cong \mathbf{Hom}(\mathbf{y}(C), f_* F) \\ &\cong f_* F(C) \end{aligned} \tag{46}$$

So in particular, $\rho(f)$ is flat and $\tau\rho(f) \cong f$.

Likewise if $A : \mathcal{C} \rightarrow \mathcal{E}$ is flat, then

$$\rho\tau(A)(C) = \mathbf{y}(C) \otimes_{\mathcal{C}} A \cong A(C) \tag{47}$$

So we see that ρ and τ are the components of an equivalence of categories. \square

In fact, in the case where \mathcal{C} has finite limits, this class of functors has a more explicit description:

Lemma 2.23: Let \mathcal{C} be a (small) category with finite limits, and \mathcal{E} a Grothendieck topos. Then $F : \mathcal{C} \rightarrow \mathcal{E}$ is flat if and only if it preserves finite limits.

Proof: See [1, Ch. VII.9]. \square

We now seek to generalise the equivalence of Theorem 2.22 to categories of sheaves, rather than just presheaves.

Definition 2.24: A family $(f_i : D_i \rightarrow C \mid i \in I)$ is said to be *epimorphic* if whenever $g, h : C \rightarrow B$ such that $gf_i = hf_i$ for all i , then $g = h$.

We say that a functor $A : \mathcal{C} \rightarrow \mathcal{E}$ is *J-continuous* if it maps *J*-covering sieves into epimorphic families in \mathcal{E} ; we denote the full subcategory of $\mathbf{Flat}(\mathcal{C}, \mathcal{E})$ of *J*-continuous flat functors by $\mathbf{Flat}_J(\mathcal{C}, \mathcal{E})$.

Lemma 2.25: Let $f : \mathcal{E} \rightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ be a geometric morphism. Then f factors through $\mathbf{Sh}(\mathcal{C}, J) \hookrightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$ if and only if $f^* \circ \mathbf{y}$ is *J*-continuous.

Proof: See [1, Lemma VII.7.3]. \square

We have an immediate corollary:

Corollary 2.26: The equivalence of Theorem 2.22 restricts to an equivalence

$$\mathbf{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}, J)) \simeq \mathbf{Flat}_J(\mathcal{C}, \mathcal{E}). \tag{48}$$

As a first application of Diaconescu's equivalence, let's show that any morphism of sites $\phi : (\mathcal{C}, J) \rightarrow (\mathcal{D}, K)$ induces a corresponding geometric morphism $\mathbf{Sh}(\phi) : \mathbf{Sh}(\mathcal{D}, K) \rightarrow \mathbf{Sh}(\mathcal{C}, J)$. Indeed, consider the functor $A_\phi = \mathbf{a} \circ \mathbf{y} \circ \phi$. Since ϕ , \mathbf{a} and \mathbf{y} all preserve finite limits, so does A_ϕ , thus it is flat by Lemma 2.23. So it induces a geometric morphism $\mathbf{Sh}(\mathcal{D}, K) \rightarrow [\mathcal{C}^{\text{op}}, \mathbf{Set}]$.

ϕ is cover-preserving, and $\mathbf{a} \circ \mathbf{y}$ sends *J*-covering sieves to epimorphic families by Lemma 2.25, so we see that A_ϕ is *K*-continuous and thus by Corollary 2.26 we have a geometric morphism between the categories of sheaves, whose inverse image is $(-) \otimes_{\mathcal{C}} A_\phi$. The direct image of this geometric morphism is then $\underline{\mathbf{Hom}}_{\mathcal{E}}(A_\phi, -)$, and we have

$$\begin{aligned}
\underline{\mathbf{Hom}}_{\mathcal{E}}(\mathbf{ay}\phi, F)(C) &= \mathbf{Hom}_{\mathcal{E}}(\mathbf{ay}\phi(C), F) \\
&\cong \mathbf{Hom}_{[\mathcal{D}^{\text{op}}, \mathbf{Set}]}(\mathbf{y}\phi(C), F) \\
&\cong F \circ \phi(C)
\end{aligned} \tag{49}$$

so we can instead take the direct image to be $\mathbf{Sh}(\phi)_* = (-) \circ \phi$.

In the special case where the sites in question are the open subsets of two topological spaces endowed with the canonical topology, and where $\phi = f^{-1}$ is the map on open sets induced by a continuous map of topological spaces, we see that the induced geometric morphism $\mathbf{Sh}(\phi)$ is precisely the one described by the construction at the start of this section.

3 Categorical Logic

Let's return for the moment to the full generality of elementary toposes. We now seek to show that such categories are indeed a "suitable universe" in which mathematics can be done. To do so, we must show that toposes have an appropriate logic: i.e. semantics and model-theoretic constructions similar to those usually defined in a set theory context.

We shall operate in the language of intuitionistic predicate logic, for it will turn out that such a logical system can be represented categorically within a topos. Loosely, a language Σ consists of:

- A collection of *sorts* A, B, \dots (e.g. in the language of categories, we might have one sort for objects and another for morphisms);
- Function symbols $f : A_1 \cdots A_n \rightarrow B$;
- Relation symbols $R \subset A_1 \cdots A_n$;
- Logical connectives $\wedge, \vee, \neg, \Rightarrow$ and the symbols \perp for false and \top for true;
- Quantifiers \forall and \exists ;
- Variables x, y, z, \dots ;

composed in the usual ways. Note that in this multi-sorted paradigm, function and relation symbols have specified sorts as their (co)domains; constant symbols are just 0-ary function symbols, denoted $c : \mathbf{1} \rightarrow A$.

3.1 Models of Theories

Let's begin by transferring some standard model-theoretic notions into the realm of category theory. In the usual setting, a Σ -structure is a collection of sets (one for each sort of the language), together with functions and subsets corresponding to the function and relation symbols. We have a categorical analogue to this:

Definition 3.1: Let \mathcal{E} be a category with finite products, Σ a language. A Σ -structure M in \mathcal{E} consists of:

- An object A^M of \mathcal{E} for each sort A of the language;
- For each function symbol $f : A_1 \cdots A_n \rightarrow B$, a corresponding morphism $f^M : A_1^M \times \cdots \times A_n^M \rightarrow B^M$;
- For each relation symbol $R \subset A_1 \cdots A_n$, a corresponding subobject $R^M \hookrightarrow A_1^M \times \cdots \times A_n^M$.

We sometimes abuse notation and omit the superscript M for convenience, when it is clear we are referring to a specific Σ -structure rather than to the terms of the language.

Definition 3.2: A *homomorphism* $h : M \rightarrow N$ of Σ -structures in \mathcal{E} is a collection of morphisms $h_A : A^M \rightarrow A^N$ (one for each sort A of Σ), such that:

- For each function symbol $f : A_1 \cdots A_n \rightarrow B$, the following diagram commutes:

$$\begin{array}{ccc}
A_1^M \times \dots \times A_n^M & \xrightarrow{h_{A_1} \times \dots \times h_{A_n}} & A_1^N \times \dots \times A_n^N \\
\downarrow f^M & & \downarrow f^N \\
B^M & \xrightarrow{h_B} & B^N
\end{array} \tag{50}$$

- For each relation symbol $R \subset A_1 \cdots A_n$, there is a morphism $R^M \rightarrow R^N$ making the following diagram commute:

$$\begin{array}{ccc}
R^M & \dashrightarrow & R^N \\
\downarrow & & \downarrow \\
A_1^M \times \dots \times A_n^M & \xrightarrow{h_{A_1} \times \dots \times h_{A_n}} & A_1^N \times \dots \times A_n^N
\end{array} \tag{51}$$

We denote the category of Σ -structures in \mathcal{E} (and homomorphisms between them) by $\Sigma\text{-str}(\mathcal{E})$.

Defining Σ -structures in this way suggests a natural choice for the “internal” language of a particular category:

Definition 3.3: Let \mathcal{E} be a category with finite products. The *internal language* $\Sigma_{\mathcal{E}}$ of \mathcal{E} has:

- A sort $\lceil C \rceil$ for each object C of \mathcal{E} ;
- A function symbol $\lceil f \rceil : \lceil C_1 \rceil \cdots \lceil C_n \rceil \rightarrow \lceil D \rceil$ for each morphism $f : C_1 \times \dots \times C_n \rightarrow D$;
- A relation symbol $\lceil R \rceil \subset \lceil C_1 \rceil \cdots \lceil C_n \rceil$ for each subobject $R \hookrightarrow C_1 \times \dots \times C_n$.

There is an obvious $\Sigma_{\mathcal{E}}$ -structure in \mathcal{E} , assigning to each sort, function symbol and relation symbol its corresponding object, morphism or subobject; this structure is known as the *tautological* $\Sigma_{\mathcal{E}}$ -structure.

Given a theory \mathbb{T} in a language Σ , we want to be able to define models of \mathbb{T} in \mathcal{E} ; but to do so requires having a sensible interpretation of logical formulae in Σ , and a notion of what it means for a Σ -structure M to *satisfy* such formulae. This is where the categorical richness of toposes shines - indeed, the constructions we exhibited in Section 1.3 will allow us to define these notions in an intuitionistically-sound way.

Definition 3.4: Let \mathcal{E} be an elementary topos, Σ a language, M a Σ -structure in \mathcal{E} , and t a term in Σ . We define the *interpretation* t^M of t in \mathcal{E} inductively, as follows:

If t is a term with free variables $(x_i : A_i)$ (i.e. x_i is of sort A_i) then:

- if $t = x_i$, then t^M is the projection $\pi_i : A_1 \times \dots \times A_n \rightarrow A_i$;
- if $t = f(t_1, \dots, t_n)$ for terms $t_i : C_i$ then we define t^M to be the composite

$$A_1 \times \dots \times A_n \xrightarrow{\langle t_1^M, \dots, t_n^M \rangle} C_1 \times \dots \times C_n \xrightarrow{f^M} B \tag{52}$$

If φ is a formula with free variables $(x_i : A_i)$, then we define the *interpretation* of φ in M to be a subobject $\{\vec{x} \mid \varphi\} \hookrightarrow A_1 \times \dots \times A_n$, again inductively:

- If φ is $(s = t)$ for terms s, t of sort B then $\{\vec{x} \mid \varphi\}$ is the equaliser of $s^M, t^M : A_1 \times \dots \times A_n \rightarrow B$;
- If φ is \perp then $\{\vec{x} \mid \varphi\}$ is the initial subobject $\mathbf{0} \hookrightarrow A_1 \times \dots \times A_n$; similarly if φ is \top then $\{\vec{x} \mid \varphi\}$ is just $B \hookrightarrow B$.
- If φ is $R(t_1, \dots, t_n)$ for a relation symbol $R \subset B_1 \cdots B_n$, then $\{\vec{x} \mid \varphi\}$ is the pullback

$$\begin{array}{ccc}
\{\vec{x} \mid \varphi\} & \xrightarrow{\quad} & R^M \\
\downarrow & & \downarrow \\
A_1 \times \cdots \times A_n & \xrightarrow{\langle t_1^M, \dots, t_n^M \rangle} & B_1 \times \cdots \times B_n
\end{array} \tag{53}$$

- If φ is $\psi \wedge \chi$, $\{\vec{x} \mid \varphi\}$ is the intersection of subobjects $\{\vec{x} \mid \psi\} \cap \{\vec{x} \mid \chi\}$ as defined in Section 1.3; likewise we define the interpretation of $\psi \vee \chi$, $\neg\psi$, and $\psi \Rightarrow \chi$ as the corresponding operations on the Heyting algebra $\mathbf{Sub}_{\mathcal{E}}(A_1 \times \cdots \times A_n)$.
- If φ is $\exists y \psi(\vec{x}, y)$ then $\{\vec{x} \mid \varphi\}$ is defined to be $\exists_{\pi} \{(\vec{x}, y) \mid \psi\}$, where π is the projection onto the first n factors, and $\exists_{\pi} : \mathbf{Sub}_{\mathcal{E}}(A_1 \times \cdots \times A_n \times B) \rightarrow \mathbf{Sub}_{\mathcal{E}}(A_1 \times \cdots \times A_n)$ is as defined in Proposition 1.12. Likewise, if φ is $\forall y \psi(\vec{x}, y)$ then $\{\vec{x} \mid \varphi\}$ is the subobject $\forall_{\pi} \{(\vec{x}, y) \mid \psi\}$ of $A_1 \times \cdots \times A_n$.

Remark 3.5: If \mathcal{E} is complete (as is the case for a Grothendieck topos) then $\mathbf{Sub}_{\mathcal{E}}(X)$ has arbitrary intersections, and so we can interpret infinitary conjunctions analogously to the finite case. Similarly if \mathcal{E} is cocomplete (as is again the case for Grothendieck toposes) then we can likewise interpret infinitary disjunctions. This will prove useful in our later discussions of geometric logic.

We now have enough machinery to define validity of a formula by a Σ -structure, and thus to define the notion of a model of a theory \mathbb{T} in a topos.

Definition 3.6: Let \mathcal{E} be an elementary topos, Σ a language, and M a Σ -structure in \mathcal{E} . A formula $\varphi(\vec{x})$ is said to be *valid* in M if the subobject $\{\vec{x} \mid \varphi\} \rightarrow A_1 \times \cdots \times A_n$ is the identity, and we write $M \models \varphi$. Note that if φ is a *sentence* (i.e. it has no free variables) then $\{\vec{x} \mid \varphi\}$ is a subobject of $\mathbf{1}$.

If \mathbb{T} is a theory in Σ , we say M is a *model* of \mathbb{T} if all the axioms of \mathbb{T} are valid in M ; we write $\mathbb{T}\text{-mod}(\mathcal{E})$ for the full subcategory of $\Sigma\text{-str}(\mathcal{E})$ whose objects are the models of \mathbb{T} in \mathcal{E} .

Let's consider some simple examples of models in toposes:

Example 3.7:

- (i) A model of a theory \mathbb{T} (as defined above) in \mathbf{Set} corresponds precisely to the usual notion of a model in intuitionistic predicate logic.
- (ii) If $\mathcal{E} = \mathbf{Sh}(X)$ is the category of sheaves on a topological space, then a model of the theory of rings is just a *sheaf of rings* on X , familiar from algebraic geometry.
- (iii) In any elementary topos \mathcal{E} , the subobject classifier Ω is a model of the theory of Heyting algebras via the isomorphism $\mathbf{Sub}_{\mathcal{E}}(C) \cong \mathbf{Hom}_{\mathcal{E}}(C, \Omega)$. Indeed, under this isomorphism each of the Heyting algebra operations on $\mathbf{Sub}_{\mathcal{E}}(C)$ corresponds to a morphism $\mathbf{Hom}_{\mathcal{E}}(C, \Omega \times \Omega) \rightarrow \mathbf{Hom}_{\mathcal{E}}(C, \Omega)$, which by the Yoneda lemma must be induced by composition from a morphism $\Omega \times \Omega \rightarrow \Omega$. These morphisms then satisfy the axioms of a Heyting algebra.

Now, an important question still needs to be resolved - is this interpretation logically sound? That is to say, does our definition of validity in the internal logic of a topos agree with the notion of provability in intuitionistic logic? Let's show that this is indeed the case.

Theorem 3.8: Let M be a model of a theory \mathbb{T} in a topos \mathcal{E} . If $\varphi(\vec{x})$ is a formula in Σ that is intuitionistically provable from the axioms of \mathbb{T} , then $M \models \varphi$.

Proof: We must check that each of the deduction rules of intuitionistic logic (see [8, Ch. 2.1] for a full list) is sound in this semantics. Since $\mathbf{Sub}_{\mathcal{E}}(X)$ is a Heyting algebra for every object X ,

we can apply the completeness theorem for the Heyting semantics of intuitionistic *propositional* logic [9, Thm. 2.4.7] to immediately deduce that all the propositional deduction rules are sound (since the interpretations of propositional connectives were defined in terms of the Heyting algebra structure of $\mathbf{Sub}_{\mathcal{E}}(X)$). So it remains to check that the deduction rules for quantifiers and equality are sound. Indeed, the soundness of the \exists -introduction and elimination rules is the statement that the following holds in $\mathbf{Sub}_{\mathcal{E}}(X)$:

$$\{\vec{x} \mid \exists y \varphi(\vec{x}, y)\} \leq \{\vec{x} \mid \psi(\vec{x})\} \quad \text{iff} \quad \{(\vec{x}, y) \mid \varphi(\vec{x}, y)\} \leq \{(\vec{x}, y) \mid \psi(\vec{x})\} \quad (54)$$

which follows immediately from the adjunction $(\exists_{\pi} \dashv \pi^*)$. Likewise the soundness of the \forall -introduction and elimination rules is the statement that

$$\{\vec{x} \mid \varphi(\vec{x})\} \leq \{\vec{x} \mid \forall y \psi(\vec{x}, y)\} \quad \text{iff} \quad \{(\vec{x}, y) \mid \varphi(\vec{x})\} \leq \{(\vec{x}, y) \mid \psi(\vec{x}, y)\} \quad (55)$$

which is again an immediate consequence of the adjunction $(\pi^* \dashv \forall_{\pi})$. Finally, $(s = s)$ is always valid since the equaliser of x^M with itself is an isomorphism; moreover if $(s = t)$ is valid then $s^M = t^M$ so the interpretation of a formula involving s is the same as the interpretation of the same formula with t substituted for s . Symmetry and transitivity follow similarly. \square

3.2 Kripke-Joyal Semantics

We now have a way to understand validity of formulae in Σ -structures in a topos. When the language in question is the *internal* language $\Sigma_{\mathcal{E}}$ of the topos (and in particular when we are reasoning about the tautological $\Sigma_{\mathcal{E}}$ structure), these semantics give us a way to define what it means for a formula to be “true in the internal logic of \mathcal{E} ”. We shall see that this notion has quite a concrete interpretation.

In the usual set-theoretic setting, we interpret variables in formulae by substituting them for elements of the Σ -structure in question. We don’t quite have that luxury when dealing with categorical logic (as we no longer have elements at our disposal), but we can replace the use of “elements of A ” with *generalised elements* of A , that is, morphisms $\alpha : E \rightarrow A$. In particular, we can extend our definition of validity in \mathcal{E} as follows:

Definition 3.9: Let \mathcal{E} be a topos, and $\varphi(x)$ be a formula in $\Sigma_{\mathcal{E}}$ with a free variable x of sort $\ulcorner A \urcorner$ (recall that A is an object of \mathcal{E}). Let $\alpha : E \rightarrow A$ be a generalised element. We say that E *satisfies* $\varphi(\alpha)$, and write $E \models \varphi(\alpha)$, if α factors through the subobject $\{x \mid \varphi(x)\} \twoheadrightarrow A$.

This notion of satisfaction in the internal language is known as the *Kripke-Joyal semantics*. We note immediately that a formula $\varphi(x)$ is valid in the tautological $\Sigma_{\mathcal{E}}$ -structure if and only if $\varphi(\alpha)$ is satisfied for all generalised elements α (since these specify the subobject $\{x \mid \varphi(x)\}$ uniquely).

Note also that if $f : D \rightarrow E$ and $E \models \varphi(\alpha)$, then $D \models \varphi(\alpha \circ f)$. This “monotonicity” property should be reminiscent of the *persistence property* for Kripke models [9, Def. 2.5.1] in the usual setting for intuitionistic logic. Indeed, we can think of the Kripke-Joyal semantics as a sort of “higher-dimensional” generalisation of Kripke semantics with a topos replacing the usual poset - this generalisation in fact allows for a semantic interpretation of higher-order logic in a topos [3, D4.1], though this goes beyond the scope of this essay.

Defining satisfaction of formulae in terms of generalised elements in this way allows us to give a characterisation, in terms of the structure of formulae, of exactly when a formula is satisfied (similarly to the usual description of Kripke semantics):

Theorem 3.10: Let \mathcal{E} be a topos, and $\varphi(x), \psi(x)$ formulae of $\Sigma_{\mathcal{E}}$ with x free of sort $\ulcorner A \urcorner$. Let $\alpha : E \rightarrow A$. Then:

- $E \models (\varphi \wedge \psi)(\alpha)$ iff $E \models \varphi(\alpha)$ and $E \models \psi(\alpha)$;

- $E \models (\varphi \vee \psi)(\alpha)$ iff there are $p : C \rightarrow E$, $q : D \rightarrow E$ with $p + q : C + D \rightarrow E$ epimorphic, such that $C \models \varphi(\alpha \circ p)$ and $D \models \psi(\alpha \circ q)$.
- $E \models (\varphi \Rightarrow \psi)(\alpha)$ iff for any morphism $p : D \rightarrow E$ such that $D \models \varphi(\alpha \circ p)$, also $D \models \psi(\alpha \circ p)$;
- $E \models \neg\varphi(\alpha)$ iff for any $p : D \rightarrow E$ such that $D \models \varphi(\alpha \circ p)$, then $D \cong \mathbf{0}$.
- If \mathcal{E} is complete, then $E \models \bigwedge_{j \in J} \varphi_j(\alpha)$ iff $E \models \varphi_j(\alpha)$ for each $j \in J$.
- If \mathcal{E} is cocomplete, then $E \models \bigvee_{j \in J} \varphi_j(\alpha)$ iff there is an epimorphic family $(p_j : C_j \rightarrow E)_{j \in J}$ such that $C_j \models \varphi_j(\alpha \circ p_j)$ for each $j \in J$.

If instead $\varphi(x, y)$ has free variables x of type $\lceil A \rceil$ and y of type $\lceil B \rceil$, then:

- $E \models (\exists y \varphi(x, y))(\alpha)$ iff there exists an epimorphism $p : D \twoheadrightarrow E$ and a generalised element $\beta : D \rightarrow B$ such that $D \models \varphi(\alpha \circ p, \beta)$;
- $E \models (\forall y \varphi(x, y))(\alpha)$ iff for every object D , every $p : D \rightarrow E$ and every generalised element $\beta : D \rightarrow B$ we have $D \models \varphi(\alpha \circ p, \beta)$.

Proof (Sketch): The criteria for the connectives are essentially reformulations of the definitions of the corresponding Heyting algebra operations; likewise the criteria for \exists and \forall follow from the definitions of \exists_π and \forall_π as adjoints to the pullback functor π^* . We'll prove the criterion for \forall in full detail: the proofs of the remaining conditions are of a similar flavour, and can be found in [1, Thm. IV.6.1].

Suppose α factors through $\{x \mid \varphi(x) \vee \psi(x)\}$ via $\beta : E \rightarrow \{x \mid \varphi(x) \vee \psi(x)\}$. Let $p : C \rightarrow E$ and $q : D \rightarrow E$ be the pullbacks along α of the subobjects corresponding to φ and ψ :

$$\begin{array}{ccc}
C & \longrightarrow & \{x \mid \varphi(x)\} & & \{x \mid \psi(x)\} & \longleftarrow & D \\
p \downarrow & & \downarrow & & \downarrow & & \downarrow q \\
E & \xrightarrow{\alpha} & A & & A & \xleftarrow{\alpha} & E
\end{array} \tag{56}$$

Since pullbacks preserve coproducts, $p + q$ is the pullback along β of the epimorphism $\{x \mid \varphi(x)\} + \{x \mid \psi(x)\} \twoheadrightarrow \{x \mid \varphi(x) \vee \psi(x)\}$, thus epi. And since $\alpha \circ p$ and $\alpha \circ q$ factor through $\{x \mid \varphi(x)\}$ and $\{x \mid \psi(x)\}$ respectively, we see that $C \models \varphi(\alpha \circ p)$ and $D \models \psi(\alpha \circ q)$.

Conversely suppose p and q satisfy the conditions of the theorem. Then the squares in (56) commute (but might not be pullbacks), so we get a commuting square

$$\begin{array}{ccc}
C + D & \longrightarrow & \{x \mid \varphi(x)\} + \{x \mid \psi(x)\} \\
p + q \downarrow & & \downarrow \\
E & \xrightarrow{\alpha} & A
\end{array} \tag{57}$$

Since $p + q$ is epi, the image of α is the same as that of the overall composite $C + D \rightarrow A$, and thus factors through the image of $\{x \mid \varphi(x)\} + \{x \mid \psi(x)\} \rightarrow A$; namely, α factors through $\{x \mid \varphi(x) \vee \psi(x)\}$. \square

Thus we have a concrete interpretation of satisfaction of formulae in a topos, and we can extend this description to formulae with more than one free variable inductively [1, Ch. VI.6] to get similar statements to Theorem 3.10.

3.3 An Application to Algebraic Geometry

We can apply the notions of the internal language of a topos, and the semantics therein, to the case where the topos in question is $\mathbf{Sh}(X)$, the (Grothendieck) topos of sheaves on a topological space. In particular, we will show that we can use the concrete interpretations of satisfaction in the internal

language to “sheafify” theorems of commutative algebra to theorems in algebraic geometry (as long as they can be proved intuitionistically). We start by observing that the Kripke-Joyal semantics allow us to translate formulae in the internal language into statements about open sets:

Definition 3.11: Let $U \subset U' \subset X$ be open subsets, and \mathcal{F} a sheaf on X . For a formula φ in the internal language and $s \in \mathcal{F}(U')$, we write $U \vDash \varphi(s)$ to mean $s|_U \in \{x \mid \varphi(x)\}(U)$.

In particular, the semantics of Theorem 3.10 become the following *sheaf semantics*:

Proposition 3.12: Let $U \subset X$ an open set. Then:

$$\begin{aligned}
U \vDash s = t & \quad \text{iff } s|_U = t|_U \in \mathcal{F}(U); \\
U \vDash \perp & \quad \text{iff } U = \emptyset; \\
U \vDash \varphi \wedge \psi & \quad \text{iff } U \vDash \varphi \text{ and } U \vDash \psi; \\
U \vDash \bigwedge_{j \in J} \varphi_j & \quad \text{iff for all } j \in J, U \vDash \varphi_j \\
U \vDash \varphi \vee \psi & \quad \text{iff there is an open cover } \bigcup_i U_i = U \text{ such that for all } i: \\
& \quad U_i \vDash \varphi \text{ or } U_i \vDash \psi; \\
U \vDash \bigvee_{j \in J} \varphi_j & \quad \text{iff there is an open cover } \bigcup_i U_i = U \text{ such that for all } i: \\
& \quad U_i \vDash \varphi_j \text{ for some } j \in J; \\
U \vDash \varphi \Rightarrow \psi & \quad \text{iff for all open } V \subset U, \text{ if } V \vDash \varphi \text{ then } V \vDash \psi; \\
U \vDash \forall s : \mathcal{F}. \varphi(s) & \quad \text{iff for all open } V \subset U \text{ and } s \in \mathcal{F}(V), V \vDash \varphi(s); \\
U \vDash \exists s : \mathcal{F}. \varphi(s) & \quad \text{iff there exists an open cover } \bigcup_i U_i = U \text{ such that for all } i: \\
& \quad \text{there exists } s_i \in \mathcal{F}(U_i) \text{ such that } U_i \vDash \varphi(s_i).
\end{aligned}$$

Proof: See [1, Thm. VI.7.1] for details; the proof follows essentially the same direction as that of Theorem 3.10, as before by considering the definition of $\{x \mid \varphi(x)\}$ in terms of the Heyting algebra structure on subobjects. \square

Via the sheaf semantics, common objects in algebraic geometry have simple interpretations in the internal language. For example, the statement that “ \mathcal{R} is a sheaf of rings” translates in the internal language to “ \mathcal{R} is a ring”. As an example, the distributivity axiom

$$X \vDash \forall x, y, z : \mathcal{R}. x \cdot (y + z) = (x \cdot y) + (x \cdot z) \quad (58)$$

translates according to Proposition 3.12 into the statement that for every open $U \subset X$ and all elements $x, y, z \in \mathcal{R}(U)$, $x \cdot (y + z) = (x \cdot y) + (x \cdot z) \in \mathcal{R}(U)$. But this is precisely the statement of distributivity for sheaves of rings. Likewise, the statement that \mathcal{F} is a sheaf of \mathcal{R} -modules over a sheaf of rings \mathcal{R} translates to the internal statement that “ \mathcal{F} is a module over the ring \mathcal{R} ”.

Now, if we are given a theorem in commutative algebra which can be proved intuitionistically (that is, using only constructive arguments), then we can translate it via the sheaf semantics into a theorem relating to sheaves of rings and modules which holds automatically by soundness. Let’s work through an example of exactly this process to provide a short proof of a familiar theorem in algebraic geometry.

Lemma 3.13: Let X be a scheme, and \mathcal{F} a sheaf of \mathcal{O}_X -modules.

$$X \vDash \lceil \mathcal{F} \text{ is finitely generated as an } \mathcal{O}_X\text{-module} \rceil \quad (59)$$

if and only if \mathcal{F} is of *finite type* as a sheaf of \mathcal{O}_X -modules.

Proof: Recall that \mathcal{F} is of finite type if there is an open cover $(U_i)_{i \in I}$ of X such that for each U_i there is a surjection $\mathcal{O}_X^n|_{U_i} \twoheadrightarrow \mathcal{F}|_{U_i}$. The internal statement can be written

$$X \models \bigvee_{n \geq 0} \exists x_1, \dots, x_n : \mathcal{F}. \forall y : \mathcal{F}. \exists a_1, \dots, a_n : \mathcal{O}_X. y = \sum_{i=1}^n a_i x_i; \quad (60)$$

applying Proposition 3.12 immediately yields that these are equivalent. \square

Lemma 3.14: Let R be a ring, and let

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0 \quad (61)$$

be an exact sequence of R -modules. If A and C are both finitely generated, then so is B .

Proof: Suppose $A = \langle a_1, \dots, a_n \rangle$ and $C = \langle c_1, \dots, c_m \rangle$. Since g is surjective, there exist elements $b_1, \dots, b_m \in B$ such that $g(b_i) = c_i$ for each i . We claim that B is generated by $f(a_1), \dots, f(a_n), b_1, \dots, b_m$.

Indeed, suppose $x \in B$. There exist $r_1, \dots, r_m \in R$ such that $g(x) = r_1 c_1 + \dots + r_m c_m$. So

$$x - (r_1 b_1 + \dots + r_m b_m) \in \ker(g) = \operatorname{im}(f), \quad (62)$$

and thus there is an element $a \in A$ such that $f(a) = x - (r_1 b_1 + \dots + r_m b_m)$. But then $a = s_1 a_1 + \dots + s_n a_n$ for some $s_i \in R$, whence

$$x = \sum_{i=1}^n s_i f(a_i) + \sum_{j=1}^m r_j b_j \quad (63)$$

as required. \square

Note that the proof of Lemma 3.14 at no point invoked the law of the excluded middle, the axiom of choice, nor any other non-constructive principle. That is, the lemma holds intuitionistically. By the recipe above, we immediately deduce the following:

Proposition 3.15: Let X be a scheme and let

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F}'' \longrightarrow 0 \quad (64)$$

be an exact sequence of \mathcal{O}_X -modules. If \mathcal{F} and \mathcal{F}'' are both of finite type, then so is \mathcal{F}' .

We have seen how we can “sheafify” intuitionistically-valid theorems in commutative algebra to prove statements in algebraic geometry - in Blechschmidt’s paper [4, Part II] the author offers a large number of additional examples applying the recipe above.

4 Classifying Toposes

In this final chapter, we introduce the classifying topos of a geometric theory \mathbb{T} ; we will first describe the class of *geometric theories* whose models are preserved by geometric morphisms. We will then show that we can associate to any such theory \mathbb{T} a topos $\mathbf{Set}[\mathbb{T}]$ which “classifies” models of \mathbb{T} in any Grothendieck topos. Throughout, we will restrict our attention to Grothendieck toposes.

4.1 Geometric Logic

In Section 2.4 we defined the notion of a geometric morphism between two Grothendieck toposes \mathcal{F} and \mathcal{E} . We showed that these morphisms preserve the “geometric” structure of a Grothendieck topos. But to what extent do geometric morphisms preserve the “logical” structure we defined in Section 3.1? More precisely, given a language Σ or a theory \mathbb{T} , when can we be sure that Σ -structures and models of \mathbb{T} are preserved by a geometric morphism $f : \mathcal{F} \rightarrow \mathcal{E}$? In this section we aim to answer this question.

First, we show that Σ -structures are always preserved under geometric morphisms.

Proposition 4.1: Let Σ be a language and $F : \mathcal{E} \rightarrow \mathcal{F}$ be a finite-limit-preserving functor. Then F induces a functor $\Sigma\text{-str}(\mathcal{E}) \rightarrow \Sigma\text{-str}(\mathcal{F})$. In particular, if $f : \mathcal{F} \rightarrow \mathcal{E}$ is a geometric morphism, then it induces a functor

$$f^* : \Sigma\text{-str}(\mathcal{E}) \rightarrow \Sigma\text{-str}(\mathcal{F}). \quad (65)$$

Proof: Suppose M is a Σ -structure in \mathcal{E} . We can define a Σ -structure FM in \mathcal{F} by setting $A^{FM} = F(A^M)$ for each sort A (and likewise for each function and relation symbol). Then if $h : M \rightarrow N$ is a homomorphism, we can define Fh by $(Fh)_A = F(h_A)$ for each sort A ; since F preserves finite limits, it is immediate that the diagrams in Definition 3.2 of a Σ -structure homomorphism commute, so that $Fh : FM \rightarrow FN$ is a homomorphism. So we indeed have a functor $\Sigma\text{-str}(\mathcal{E}) \rightarrow \Sigma\text{-str}(\mathcal{F})$.

The final part of the proposition follows since, for any geometric morphism f , the inverse image f^* preserves finite limits. \square

Now suppose \mathbb{T} is a theory defined over Σ . For a geometric morphism $f : \mathcal{F} \rightarrow \mathcal{E}$ to similarly induce a functor $\mathbb{T}\text{-mod}(\mathcal{E}) \rightarrow \mathbb{T}\text{-mod}(\mathcal{F})$, we will require additional constraints on the syntactic structure of \mathbb{T} to ensure that the interpretation of each of the axioms is preserved by f^* .

Definition 4.2: A formula φ in Σ is *geometric* if it can be obtained from atomic formulae by finitary conjunction, potentially infinitary disjunction, and existential quantification. A formula of the form

$$\forall x_1 \cdots \forall x_n [\varphi(x_1, \dots, x_n) \Rightarrow \psi(x_1, \dots, x_n)], \quad (66)$$

where φ and ψ are geometric formulae, is called a *geometric implication*. We say that a theory \mathbb{T} over Σ is a *geometric theory* if it admits an axiomatisation all of whose axioms are geometric implications.

We note immediately that many familiar theories admit an axiomatisation by geometric implications: for example, the theories of groups, rings, and local rings are all geometric theories.

Proposition 4.3: Suppose M is a Σ -structure in \mathcal{E} , $f : \mathcal{F} \rightarrow \mathcal{E}$ a geometric morphism, and $\varphi(\vec{x})$ a geometric formula in Σ . Then

$$\{\vec{x} \mid \varphi(\vec{x})\}^{f^*M} = f^*\left(\{\vec{x} \mid \varphi(\vec{x})\}^M\right) \quad (67)$$

Proof: By induction over the structure of φ . First note that since f^* preserves monos, it descends to a map $\mathbf{Sub}_{\mathcal{E}}(C) \rightarrow \mathbf{Sub}_{\mathcal{F}}(f^*C)$. Since f^* preserves finite limits (and hence products, equalisers and pullbacks), as well as the top and bottom elements of the subobject lattice, it follows from the definitions of the interpretations of terms and atomic formulae in Σ (Definition 3.4) that these are preserved under f^* . So it remains to check conjunction, disjunction and existential quantification.

Since the intersection of two subobjects is given by a pullback, it is preserved by f^* and hence the interpretation of $\varphi \wedge \psi$ is preserved. Likewise the union can be given as a pushout: since f^* is a left adjoint, it preserves colimits and hence the interpretation of $\bigvee_i \varphi_i$ is also preserved. Finally, f^* preserves epi-mono factorisations of morphisms, and thus for any $\alpha : C \rightarrow D$ we have a commuting square

$$\begin{array}{ccc}
\mathbf{Sub}_{\mathcal{E}}(C) & \xrightarrow{f^*} & \mathbf{Sub}_{\mathcal{F}}(f^*C) \\
\exists_{\alpha} \downarrow & & \downarrow \exists_{f^*\alpha} \\
\mathbf{Sub}_{\mathcal{E}}(D) & \xrightarrow{f^*} & \mathbf{Sub}_{\mathcal{F}}(f^*D)
\end{array} \tag{68}$$

and hence f^* preserves the interpretation of $\exists y \varphi(\vec{x}, y)$. \square

Finally, we can show that models of geometric theories are indeed preserved by geometric morphisms.

Theorem 4.4: Let \mathbb{T} be a geometric theory and $f : \mathcal{F} \rightarrow \mathcal{E}$ a geometric morphism. Then $f^* : \Sigma\text{-str}(\mathcal{E}) \rightarrow \Sigma\text{-str}(\mathcal{F})$ restricts to a functor $\mathbb{T}\text{-mod}(\mathcal{E}) \rightarrow \mathbb{T}\text{-mod}(\mathcal{F})$.

Proof: Suppose M is a model of \mathbb{T} in \mathcal{E} . We need to show that f^*M is a model of \mathbb{T} in \mathcal{F} , that is, that all the axioms of \mathbb{T} are satisfied in f^*M . So consider an axiom of \mathbb{T} of the form $\forall \vec{x} [\varphi(\vec{x}) \Rightarrow \psi(\vec{x})]$. That this axiom is valid in M means that

$$\{\vec{x} \mid \varphi(\vec{x})\}^M \leq \{\vec{x} \mid \psi(\vec{x})\}^M \tag{69}$$

as subobjects of $A_1 \times \dots \times A_n$. But f^* is order-preserving as a map on subobjects, so we have

$$f^*(\{\vec{x} \mid \varphi(\vec{x})\}^M) \leq f^*(\{\vec{x} \mid \psi(\vec{x})\}^M). \tag{70}$$

Thus the axiom is valid in f^*M by Proposition 4.3. \square

4.2 The Classifying Topos of a Theory

At last, we can define the eponymous topos of this essay:

Definition 4.5: Let \mathbb{T} be a geometric theory over Σ . A *classifying topos* for \mathbb{T} is a Grothendieck topos $\mathbf{Set}[\mathbb{T}]$ such that for any Grothendieck topos \mathcal{E} , there is an equivalence of categories

$$\mathbf{Geom}(\mathcal{E}, \mathbf{Set}[\mathbb{T}]) \simeq \mathbb{T}\text{-mod}(\mathcal{E}) \tag{71}$$

natural in \mathcal{E} , in the sense that for any geometric morphism $f : \mathcal{E} \rightarrow \mathcal{F}$ there is a commuting square (up to natural isomorphism) in \mathbf{Cat} :

$$\begin{array}{ccc}
\mathbf{Geom}(\mathcal{E}, \mathbf{Set}[\mathbb{T}]) & \xrightarrow{\simeq} & \mathbb{T}\text{-mod}(\mathcal{E}) \\
-\circ f \downarrow & & \downarrow f^* \\
\mathbf{Geom}(\mathcal{F}, \mathbf{Set}[\mathbb{T}]) & \xrightarrow{\simeq} & \mathbb{T}\text{-mod}(\mathcal{F}).
\end{array} \tag{72}$$

We make some immediate remarks on this definition. First, we note that the hypothesis that \mathbb{T} is a geometric theory is essential for the naturality condition to make sense: otherwise, we can't be sure that a geometric morphism f descends to a functor between the categories of \mathbb{T} -models.

Consider also that if $\mathbf{Set}[\mathbb{T}]$ is a classifying topos for \mathbb{T} , then the identity geometric morphism $\mathbf{Set}[\mathbb{T}] \rightarrow \mathbf{Set}[\mathbb{T}]$ corresponds under (71) to a model $U_{\mathbb{T}}$ of \mathbb{T} in $\mathbf{Set}[\mathbb{T}]$, known as the *universal model* of \mathbb{T} . This model is so named because, if M is any model of \mathbb{T} in a Grothendieck topos \mathcal{E} , the naturality square (72) implies that $M \cong f^*U_{\mathbb{T}}$, where f is the geometric morphism corresponding to M under the equivalence (71). In other words:

Every model of \mathbb{T} in every Grothendieck topos can be obtained by taking the inverse image of the universal model along some geometric morphism!

Now, we seek to prove that every geometric theory \mathbb{T} has a classifying topos. To do so, we must build a suitable site whose category of sheaves satisfies the universal property of Definition 4.5.

In the usual setting the *Lindenbaum-Tarski algebra* associated to \mathbb{T} fulfils a similar role - the structure of this Heyting algebra “encodes” information about provability of formulae from \mathbb{T} . We will define our site as a generalisation of this idea, constructing a category rather than a poset in terms of the syntax of formulae in the language Σ , subject to an equivalence relation built from provability under \mathbb{T} .

Definition 4.6: Let \mathbb{T} be a geometric theory over a language Σ . The *syntactic category* of \mathbb{T} , denoted $\mathcal{C}_{\mathbb{T}}$, has as objects equivalence classes of geometric formulae $\varphi(\vec{x})$ in Σ , where two formulae are equivalent if they are identical up to renaming of variables; we write $[\vec{x}.\varphi(\vec{x})]$ for the equivalence class of $\varphi(\vec{x})$. The morphisms $[\vec{x}.\varphi(\vec{x})] \rightarrow [\vec{y}.\psi(\vec{y})]$ are equivalence classes $[\theta]$ of geometric formulae $\theta(\vec{x}, \vec{y})$ such that

$$\mathbb{T} \vdash \ulcorner \theta \text{ is a function } \{\vec{x} \mid \varphi(\vec{x})\} \rightarrow \{\vec{y} \mid \psi(\vec{y})\} \urcorner, \quad (73)$$

where two such formulae θ, θ' are equivalent if $\mathbb{T} \vdash \forall \vec{x}, \vec{y} [\theta(\vec{x}, \vec{y}) \Leftrightarrow \theta'(\vec{x}, \vec{y})]$.

The identity morphism on $[\vec{x}.\varphi(\vec{x})]$ is given by

$$1_{[\vec{x}.\varphi(\vec{x})]} = [\varphi(\vec{x}) \wedge \psi(\vec{y}) \wedge x_1 = y_1 \wedge \dots \wedge x_n = y_n]. \quad (74)$$

As for composition, if we have $[\theta] : [\vec{x}.\varphi(\vec{x})] \rightarrow [\vec{y}.\psi(\vec{y})]$ and $[\gamma] : [\vec{y}.\psi(\vec{y})] \rightarrow [\vec{z}.\zeta(\vec{z})]$ then we define

$$[\gamma] \circ [\theta] = [\exists \vec{y} \theta(\vec{x}, \vec{y}) \wedge \gamma(\vec{y}, \vec{z})]. \quad (75)$$

Definition 4.7: We define a Grothendieck topology $J_{\mathbb{T}}$ on $\mathcal{C}_{\mathbb{T}}$: for a sieve S on $[\vec{x}.\varphi(\vec{x})]$, we say that $S \in J_{\mathbb{T}}$ if and only if

$$\mathbb{T} \vdash \forall \vec{y} \bigvee_{[\theta] \in S} \exists \vec{x} \theta(\vec{x}, \vec{y}); \quad (76)$$

Intuitively, we can understand this condition as a syntactic version of the statement that the union of the images of the “functions” in S cover $\{\vec{x} \mid \varphi(\vec{x})\}$. the resulting site $(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})$ is known as the *syntactic site*.

Proposition 4.8: For any Grothendieck topos \mathcal{E} , there is an equivalence of categories

$$\mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E}) \simeq \mathbb{T}\text{-mod}(\mathcal{E}), \quad (77)$$

natural in \mathcal{E} .

Proof: We prove first that $\mathcal{C}_{\mathbb{T}}$ has finite limits. We can define an terminal object $[\cdot.\top]$ and binary products

$$[\vec{x}.\varphi(\vec{x})] \times [\vec{y}.\psi(\vec{y})] = [(\vec{x}, \vec{y}) . \varphi(\vec{x}) \wedge \psi(\vec{y})]; \quad (78)$$

the equaliser of $[\theta], [\zeta] : [\vec{x}.\varphi(\vec{x})] \rightarrow [\vec{x}.\psi(\vec{x})]$ is given by $[\vec{x} . \varphi(\vec{x}) \wedge \exists \vec{y} [\theta(\vec{x}, \vec{y}) \wedge \zeta(\vec{x}, \vec{y})]]$. So $\mathcal{C}_{\mathbb{T}}$ indeed has finite limits. Moreover, there is a Σ -structure $M_{\mathbb{T}}$ inside $\mathcal{C}_{\mathbb{T}}$, defined by

$$\begin{aligned} A^{M_{\mathbb{T}}} &= [\cdot.\top] \text{ (with } x \text{ a free variable of sort } A\text{);} \\ f^{M_{\mathbb{T}}} &= [y = f(\vec{x})] \text{ for function symbols } f; \\ R^{M_{\mathbb{T}}} &= [\vec{x}.R(\vec{x})] \text{ for relation symbols } R. \end{aligned} \quad (79)$$

We define the functor $\mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E}) \rightarrow \mathbb{T}\text{-mod}(\mathcal{E})$ by sending a flat $J_{\mathbb{T}}$ -continuous functor F to $F(M_{\mathbb{T}})$, which we must check is a model of \mathbb{T} in \mathcal{E} . Indeed, if we have an axiom $\forall \vec{x} [\varphi(\vec{x}) \Rightarrow \psi(\vec{x})]$ in \mathbb{T} , then we have a subobject

$$[\varphi(\vec{x}) \wedge \vec{y} = \vec{x}] : \llbracket \vec{x}.\varphi(\vec{x}) \rrbracket \rightarrow \llbracket \vec{y}.\psi(\vec{y}) \rrbracket \quad (80)$$

in $\mathcal{C}_{\mathbb{T}}$, and since F preserves monomorphisms we have

$$\{\vec{x} \mid \varphi(\vec{x})\}^{F(M_{\mathbb{T}})} \leq \{\vec{x} \mid \psi(\vec{x})\}^{F(M_{\mathbb{T}})}, \quad (81)$$

in $\mathbf{Sub}_{\mathcal{E}}(A_1 \times \cdots \times A_n)$. So the axiom is satisfied in $F(M_{\mathbb{T}})$.

For the reverse functor $\mathbb{T}\text{-mod}(\mathcal{E}) \rightarrow \mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$, given a model N of \mathbb{T} in a Grothendieck topos \mathcal{E} , we have a functor F_N sending $\llbracket \vec{x}.\varphi(\vec{x}) \rrbracket$ to $\{\vec{x} \mid \varphi(\vec{x})\}^N$ and sending a morphism $[\theta]$ to the composite

$$\{\vec{x} \mid \varphi(\vec{x})\} \rightarrow \{(\vec{x}, \vec{y}) \mid \theta(\vec{x}, \vec{y})\} \rightarrow \{\vec{y} \mid \psi(\vec{y})\}, \quad (82)$$

noting that this composite is independent of our choice of representative θ because any other such choice is \mathbb{T} -provably equivalent. It follows quickly from the definition of finite limits in $\mathcal{C}_{\mathbb{T}}$ that this functor indeed preserves finite limits, and thus is flat by Lemma 2.23.

Moreover, if S is a $J_{\mathbb{T}}$ -covering sieve then by soundness,

$$F(M_{\mathbb{T}}) \models \forall \vec{y} \bigvee_{[\theta] \in S} \exists \vec{x} \theta(\vec{x}, \vec{y}); \quad (83)$$

but this is precisely the statement (internal to the logic of $F(M_{\mathbb{T}})$) that $F(S)$ is an epimorphic family in \mathcal{E} . So we see that F is $J_{\mathbb{T}}$ -continuous.

Furthermore, if $h : N \rightarrow N'$ is a homomorphism in $\mathbb{T}\text{-mod}(\mathcal{E})$, then h induces morphisms $\{\vec{x} \mid \varphi(\vec{x})\}^N \rightarrow \{\vec{x} \mid \varphi(\vec{x})\}^{N'}$ for each formula φ , so this indeed defines a functor $\mathbb{T}\text{-mod}(\mathcal{E}) \rightarrow \mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$.

Finally, we must check that these functors form an equivalence. Indeed, certainly $F_N(M_{\mathbb{T}}) \cong N$ for any \mathbb{T} -model N in \mathcal{E} . For the converse, suppose $F : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{E}$ is a flat $J_{\mathbb{T}}$ -continuous functor, and set $N = F(M_{\mathbb{T}})$. Then by induction over the structure of formulae, for each formula φ we have a natural isomorphism $F(\llbracket \vec{x}.\varphi(\vec{x}) \rrbracket) \cong \{\vec{x} \mid \varphi(\vec{x})\}^N$, which form a natural isomorphism $F \cong F_N$, naturally in F . So the functors indeed form an equivalence. \square

Now we can apply Diaconescu's equivalence (Corollary 2.26) to deduce

$$\mathbf{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})) \simeq \mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E}) \simeq \mathbb{T}\text{-mod}(\mathcal{E}), \quad (84)$$

and thus we have:

Theorem 4.9: For any geometric theory \mathbb{T} , $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})$ is a classifying topos for \mathbb{T} .

Despite now having an explicit description of a classifying topos for \mathbb{T} , the syntactic nature of our construction means that we do not understand the structure of this topos particularly well. In the next section, we will give more helpful descriptions of classifying toposes for certain theories - for this to be useful, we first show that:

Proposition 4.10: For a geometric theory \mathbb{T} , the classifying topos $\mathbf{Set}[\mathbb{T}]$ is unique up to equivalence.

Proof: Suppose \mathcal{F} and \mathcal{G} are classifying toposes for \mathbb{T} ; that is, they contain universal models M and N (respectively) for \mathbb{T} . Then by the universal property of classifying toposes, there are

unique geometric morphisms whose inverse images map M to (an isomorphic copy of) N and vice versa. Then the composite of these geometric morphisms maps M to an isomorphic copy of itself, and thus (again by the universal property) is naturally isomorphic to the identity. \square

4.3 Examples of Classifying Toposes

As remarked in the last section, our construction of the classifying topos for a geometric theory has a syntactic flavour; we seek more concrete descriptions of the classifying toposes for familiar theories.

Example 4.11: Let \mathbb{O} be the “theory of objects”; that is, the theory with a single sort and no axioms. Its models in any category are simply the objects of that category, so $\mathbb{O}\text{-mod}(\mathcal{E}) \cong \mathcal{E}$. Then the universal property of $\mathbf{Set}[\mathbb{O}]$ becomes (for any Grothendieck topos \mathcal{E})

$$\mathcal{E} \simeq \mathbf{Geom}(\mathcal{E}, \mathbf{Set}[\mathbb{O}]). \quad (85)$$

We claim that $[\mathbf{Finset}, \mathbf{Set}]$ classifies \mathbb{O} . To show this we will proceed similarly to the proof of Proposition 4.8 and show that

$$\mathcal{E} \simeq \mathbf{Flat}(\mathbf{Finset}^{\text{op}}, \mathcal{E}), \quad (86)$$

where $\mathbf{Flat}(\mathbf{Finset}^{\text{op}}, \mathcal{E})$ consists of the finite-limit-preserving functors by Lemma 2.23 (since $\mathbf{Finset}^{\text{op}}$ has finite limits). In one direction, we have a functor $F \mapsto F(\{*\})$ which evaluates a given flat functor at the one-point set $\{*\}$. Conversely, we note that any finite set is a finite coproduct of $\{*\}$ with itself (and thus a finite product in $\mathbf{Finset}^{\text{op}}$). So given an object C of \mathcal{E} , we have a flat functor $F_C : \mathbf{Finset}^{\text{op}} \rightarrow \mathcal{E}$ sending $S \mapsto \prod_{s \in S} C$.

To check these form an equivalence, we first note that given an object C of \mathcal{E} , $F_C(\{*\}) \cong C$. Conversely, given a flat functor G , since G commutes with products in $\mathbf{Finset}^{\text{op}}$ (i.e. coproducts in \mathbf{Finset}), we have an isomorphism $F_{G(\{*\})}(S) = \prod_{s \in S} G(\{*\}) \cong G(S)$, naturally in S . So $F_{G(\{*\})} \cong G$, and we have the required equivalence.

Applying Diaconescu’s theorem (Theorem 2.22), we see that $[\mathbf{Finset}, \mathbf{Set}]$ classifies \mathbb{O} .

In fact, the example worked through here is a special case of a more general phenomenon; when we restrict to *universal Horn theories*, that is, theories whose axioms are all of the form

$$\forall \vec{x} [\varphi_1(\vec{x}) \wedge \dots \wedge \varphi_n(\vec{x}) \Rightarrow \psi(\vec{x})] \quad (87)$$

where φ_i and ψ are atomic formulae, we have a useful semantic characterisation of the classifying topos:

Theorem 4.12: Let \mathbb{T} be a universal Horn theory. Then \mathbb{T} is classified by $[\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set}), \mathbf{Set}]$, where $\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set})$ is the category of *finitely presented* models of \mathbb{T} in \mathbf{Set} and homomorphisms between them.

Proof (Sketch): A proof of this fact can be found in [10, Theorem 1]; the key idea is to construct a simplified version $\mathcal{C}'_{\mathbb{T}}$ of the syntactic category, generated only by conjunctions of atomic formulae (rather than the wider class of geometric formulae as in Definition 4.6). It then suffices to show that:

- (i) This simplified syntactic category $\mathcal{C}'_{\mathbb{T}}$ is equivalent to $\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set})^{\text{op}}$;
- (ii) $[\mathcal{C}'_{\mathbb{T}}{}^{\text{op}}, \mathbf{Set}]$ classifies \mathbb{T} .

To prove the first claim, we construct a map $\mathcal{C}'_{\mathbb{T}} \rightarrow \text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set})^{\text{op}}$, sending $[\vec{x}.\varphi(\vec{x})]$ to the model $\langle \vec{x} \mid \varphi(\vec{x}) \rangle$ of \mathbb{T} whose elements are equivalence classes of terms $t(\vec{x})$ under the relation

$$s(\vec{x}) \sim t(\vec{x}) \quad \text{iff} \quad \mathbb{T} \vdash \forall \vec{x} [\varphi(\vec{x}) \Rightarrow s(\vec{x}) = t(\vec{x})]; \quad (88)$$

that is, the model “freely generated” by the formal symbols x_1, \dots, x_n subject to the relation φ . We check that this is indeed a model of \mathbb{T} ; it is not hard to show that this map extends to a full and faithful functor $\mathcal{C}'_{\mathbb{T}} \rightarrow \mathbb{T}\text{-mod}(\mathbf{Set})^{\text{op}}$, whose image is precisely the finitely presented models of \mathbb{T} . Thus $\mathcal{C}'_{\mathbb{T}}$ is equivalent to $\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set})^{\text{op}}$.

The second claim follows in much the same way as Proposition 4.8. \square

In particular, following through the equivalences we see that the universal model $U_{\mathbb{T}}$ of a universal Horn theory \mathbb{T} in $[\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set}), \mathbf{Set}]$ is induced by the inclusion $\text{f.p.}\mathbb{T}\text{-mod}(\mathbf{Set}) \hookrightarrow \mathbb{T}\text{-mod}(\mathbf{Set})$, and thus for each sort A , $A^{U_{\mathbb{T}}}$ is the functor which sends a finitely presented model M to the set A^M (and similarly for function and relation symbols).

We now have a nice description of the classifying topos for familiar theories:

Example 4.13: The theory of groups is classified by $[\mathbf{Grp}_{\text{fp}}, \mathbf{Set}]$, where \mathbf{Grp}_{fp} is the category of finitely presented groups and homomorphisms between them. Likewise the theory of (commutative) rings is classified by $[\mathbf{Rng}_{\text{fp}}, \mathbf{Set}]$; the universal model of each of these theories is given by the forgetful functor sending a finitely presented group or ring to its underlying set.

Let’s explore one more example of a geometric theory with applications to algebraic geometry, whose classifying topos we want to better understand: the theory of *local rings*. There are, of course, a few different formulations of this theory which are classically equivalent to one another: for our purposes, we will take the condition that a ring R is a local ring iff, for all $x \in R$, either x or $1 - x$ is a unit. That is, we will take LocRing to be the theory of rings together with an axiom

$$\forall x[(\exists y. x \cdot y = 1) \vee (\exists y. (1 - x) \cdot y = 1)]. \quad (89)$$

Since we can view $\text{LocRing}\text{-mod}(\mathcal{E})$ as a full subcategory of $\text{Ring}\text{-mod}(\mathcal{E})$, we naturally should expect the classifying topos $\mathbf{Set}[\text{LocRing}]$ to be some full subcategory of the classifying topos for rings $\mathbf{Set}[\text{Ring}] = [\mathbf{Rng}_{\text{fp}}, \mathbf{Set}]$. In other words, we seek a suitable Grothendieck topology J on $\mathbf{Rng}_{\text{fp}}^{\text{op}}$ such that the J -continuous flat functors correspond to models of LocRing in Grothendieck toposes.

Note that \mathbf{Rng}^{op} is equivalent to the category \mathbf{Aff} of affine schemes - we will thus refer to the objects of $\mathbf{Rng}_{\text{fp}}^{\text{op}}$ corresponding to a ring R by $\text{Spec}(R)$. We will define our topology in terms of the Zariski topology of the corresponding schemes.

Definition 4.14: A finite family $\{f_i : \text{Spec}(A_i) \rightarrow \text{Spec}(R) \mid i \in I\}$ on $\text{Spec}(R)$ is said to be a *Zariski-covering* when the following conditions hold [11]:

- (i) Each A_i is a localisation $A_i = R[r_i^{-1}]$ of a single element $r_i \in R$;
- (ii) Each f_i is the canonical inclusion $\text{Spec}(R[r_i^{-1}]) \hookrightarrow \text{Spec}(R)$;
- (iii) The open subsets $\text{Spec}(R[r_i^{-1}])$ cover $\text{Spec}(R)$:

$$\bigcup_{i \in I} \text{Spec}(R[r_i^{-1}]) = \text{Spec}(R) \quad (90)$$

(equivalently, there are elements $s_i \in R$ such that $\sum_{i \in I} r_i s_i = 1$).

The *Zariski topos* is the category $\mathcal{Z} = \mathbf{Sh}(\mathbf{Rng}_{\text{fp}}^{\text{op}}, J)$, where J consists of the sieves which contain a Zariski-covering.

To show that this is indeed the classifying topos $\mathbf{Set}[\text{LocRing}]$, we will prove:

Proposition 4.15: Let \mathcal{E} be a Grothendieck topos, R a ring object in \mathcal{E} , and $\varphi_R : \mathbf{Rng}_{\text{fp}}^{\text{op}} \rightarrow \mathcal{E}$ the flat functor it corresponds to under the equivalence $\text{Ring}\text{-mod}(\mathcal{E}) \simeq \mathbf{Flat}(\mathbf{Rng}_{\text{fp}}^{\text{op}}, \mathcal{E})$. The following are equivalent:

- (i) R is a model of LocRing in \mathcal{E} ;

(ii) φ_R sends (the opposite of) the morphisms

$$\begin{aligned} \mathbb{Z}[X] &\longrightarrow \mathbb{Z}[X, Y]/(XY - 1), & [= \mathbb{Z}[X][X^{-1}]] \\ \mathbb{Z}[X] &\longrightarrow \mathbb{Z}[X, Y]/(XY - Y + 1) & [= \mathbb{Z}[X][(1 - X)^{-1}]] \end{aligned} \quad (91)$$

in \mathbf{Rng}_{fp} to an epimorphic family in \mathcal{E} ;

(iii) φ_R sends Zariski-coverings to epimorphic families in \mathcal{E} .

Proof: (i) \Leftrightarrow (ii): For a given finitely presented ring $A = \mathbb{Z}[X_1, \dots, X_n]/(P_1, \dots, P_k)$, $\varphi_R(\text{Spec}(A))$ is given by the equaliser

$$\varphi_R(\text{Spec}(A)) \rightrightarrows R^n \begin{array}{c} \xrightarrow{(P_1, \dots, P_k)} \\ \xrightarrow{(0, \dots, 0)} \end{array} R^k, \quad (92)$$

where the action of the polynomials P_i is given by the interpretation of the corresponding term in the language of rings. So φ_R sends the first of the maps of (91) to the subobject $\{x \mid \exists y \ x \cdot y = 1\} \rightrightarrows R$, and the second map is sent to $\{x \mid \exists y \ (1 - x) \cdot y = 1\} \rightrightarrows R$. Then (ii) holds iff the union of these subobjects is R itself, that is, iff the axiom (89) is satisfied by R . But this is precisely the condition that R is a model of LocRing .

(ii) \Leftrightarrow (iii): The reverse implication (iii) \Rightarrow (ii) holds since (the opposites of) the two homomorphisms of (91) form a Zariski-covering of $\text{Spec}(\mathbb{Z}[X])$, as they are the localisations of $\mathbb{Z}[X]$ at X and $1 - X$.

For the converse, assume (ii) holds. Let A be a finitely presented ring and $a_1, \dots, a_n \in A$ such that $\sum_{i=1}^n a_i = 1$. We'll consider the case $n = 2$; the general case is an induction on n and can be found in [1, Lemma VIII.6.2]. We have pullback diagrams in $\mathbf{Rng}_{\text{fp}}^{\text{op}}$:

$$\begin{array}{ccccc} \text{Spec}(\mathbb{Z}[X][X^{-1}]) & \rightrightarrows & \text{Spec}(\mathbb{Z}) & \longleftarrow & \text{Spec}(\mathbb{Z}[X][(1 - X)^{-1}]) \\ \uparrow & & \uparrow & & \uparrow \\ \text{Spec}(A[a^{-1}]) & \rightrightarrows & \text{Spec}(A) & \longleftarrow & \text{Spec}(A[(1 - a)^{-1}]). \end{array} \quad (93)$$

Since φ_R is flat, these are sent to a corresponding pullback diagram in \mathcal{E} ; but since the image of the top row under φ_R is sent to an epimorphic family by (ii), so too is their pullback, the image of the bottom row. So (iii) holds. \square

The equivalence (i) \Leftrightarrow (iii) of the above proposition tells us that the J -continuous flat functors φ_R are precisely those which correspond to local rings R in a given Grothendieck topos \mathcal{E} . In particular, we can conclude:

Theorem 4.16: The Zariski topos \mathcal{Z} classifies the theory of local rings.

Remark 4.17: The topos \mathcal{Z} can be viewed as a special case of a more general scheme-theoretic construction: Given a scheme S , there is a category $\mathbf{Aff}_{\text{fp}}/S$ of affine schemes $\text{Spec}(R)$ (with R finitely presented) with a morphism $\text{Spec}(R) \rightarrow S$.

We can form a *Zariski topology* on $\mathbf{Aff}_{\text{fp}}/S$ whose covering sieves are those, as in Definition 4.14, which contain finite families of localisations of R whose spectra cover $\text{Spec}(R)$. The topos of sheaves $\text{Zar}(S)$ on this site is known as the *big Zariski topos* associated to S , and if $S = \text{Spec}(A)$ then it can be shown [1, Ch. VIII.6] that this topos classifies the theory of local A -algebras. Our Zariski topos \mathcal{Z} is thus $\text{Zar}(\text{Spec}(\mathbb{Z}))$.

As in Section 3.3, the internal language of $\text{Zar}(S)$ and its Kripke-Joyal semantics can be used to translate intuitionistically-valid statements into *scheme-theoretic* theorems. This idea is systematically exploited in [4, Part III], and the universal property of $\text{Zar}(S)$ as a classifying topos features prominently.

Conclusion

Although the general theory of classifying toposes is well-understood, their manifold applications leave open a world of possibilities for this rich theory to be put to work. The interplay between the intuitionistic logic of a topos and the complexity of scheme-theoretic Algebraic Geometry, and the ability to move between these two viewpoints, suggests that a fuller constructive account of Algebraic Geometry should be within reach. In particular, the ongoing work of Cherubini *et al.* [12] in developing such a foundation for Algebraic Geometry within the internal logic of the Zariski topos could open up new avenues of proof in the field. Many questions remain unanswered in this developing research – open problems include whether the statement ‘ X is affine’ is equivalent in this account to its double negation [13]. Future work could explore these and other open problems.

Bibliography

- [1] S. MacLane and I. Moerdijk, *Sheaves in Geometry and Logic*. Springer Science & Business Media, 2012.
- [2] O. Caramello, *Theories, Sites, Toposes: Relating and Studying Mathematical Theories through topos-theoretic 'bridges'*. Oxford University Press, 2018.
- [3] P. T. Johnstone, *Sketches of an Elephant: A Topos Theory Compendium*. Oxford University Press, 2002.
- [4] I. Blechschmidt, “Using the internal language of toposes in algebraic geometry.” 2021. doi: 10.48550/arxiv.2111.03685.
- [5] M. Shulman, “Comparing material and structural set theories,” *Annals of Pure and Applied Logic*, vol. 170, pp. 465–504, 2019, doi: 10.1016/j.apal.2018.11.002.
- [6] F. Borceux, *Handbook of Categorical Algebra: Volume 3, Categories of Sheaves*. Cambridge University Press, 1994.
- [7] R. Vakil, “The Rising Sea: Foundations Of Algebraic Geometry.” Accessed: Mar. 07, 2024. [Online]. Available: <https://math.stanford.edu/~vakil/216blog/>
- [8] A. S. Troelstra and D. van Dalen, *Constructivism in Mathematics*, vol. 1. Elsevier, 2014.
- [9] M. H. Sørensen and P. Urzyczyn, *Lectures on the Curry-Howard Isomorphism*. Elsevier, 2006.
- [10] A. Blass and A. Scedrov, “Classifying Topoi and Finite Forcing,” *Journal of Pure and Applied Algebra*, vol. 28, pp. 111–140, 1983, doi: 10.1016/0022-4049(83)90085-3.
- [11] nLab authors, “Zariski site.” [Online]. Available: <https://ncatlab.org/nlab/show/Zariski+site>
- [12] F. Cherubini, T. Coquand, and M. Hutzler, “A Foundation for Synthetic Algebraic Geometry.” 2023. doi: 10.48550/arXiv.2307.00073.
- [13] F. Cherubini, “Synthetic Algebraic Geometry in the Zariski-Topos.” [Online]. Available: <https://github.com/felixwellen/synthetic-zariski>