The interplay between Grothendieck topoi and logic

MSc. Thesis, Master Mathematical Sciences, Universiteit Utrecht

Jasper Mulder, 3363120

August 26, 2013

Supervisor/First assessor:
Dr. J. van Oosten

Second assessor:
Dr. B. van den Berg

Abstract

Grothendieck topoi are generalizations of the categories of sheaves over topological spaces (familiar from algebraic geometry). The structure of a Grothendieck topos permits an interpretation of a kind of first-order logic called *geometric logic*. We explore the interplay between these two viewpoints, guided by the 1987 article *On Representations of Grothendieck Toposes* by M. Barr and M. Makkai, and the 1977 book *First Order Categorical Logic* by G.E. Reyes and M. Makkai.

Contents

Introduction

This thesis, broadly speaking, intends to provide a multi-faceted introduction to categorical logic. Within this field, primary attention goes out to the so-called Grothendieck topoi (first introduced by Grothendieck et al. as a supplementary tool in algebraic geometry – which is still a primary source of applications for Grothendieck topoi). They provide a generalization of categories of *sheaves over a topological space*, replacing the space by a category endowed with a *Grothendieck topology*; such categories will be called *sites*

In Chapter 1, Grothendieck topoi are defined, and some basic results are proved. We briefly introduce the notion of an *elementary topos*, a generalization of Grothendieck topoi that can intuitively be regarded as an "alternative universe of sets"; these elementary topoi have come to play an important role in categorical logic. A good introduction into the theory of elementary topoi is [MM92].

Next, in Chapter 2, we introduce the logic under consideration for most of the remainder of the work, following the exposition of [MR77]. This logic is called *many-sorted infinitary first-order logic*; in particular, the *geometric fragment* thereof will be extensively used. It is explained how suitable languages with this logic can express properties about categories, and a number of examples of this expressivity are considered. The chapter concludes with a short account of *model theory* in this categorical context.

The material of the first two chapters is brought together in Chapter 3. The notion of an ∞-*pretopos* is introduced; subsequently, it is shown that Grothendieck topoi are ∞-pretopoi, which in particular implies that the logic from Chapter 2 applies to Grothendieck topoi. This gives rise to a first example of how these two interact, in the form of the *Kripke-Joyal semantics* for a Grothendieck topos. This semantics is an important tool in determining if it is possible, and how, to carry over familiar set-theoretical constructions (like for example, the real numbers) to a given Grothendieck topos. The Kripke-Joyal semantics has applications in e.g. axiomatic set theory, see [MM92].

In Chapter 4, we start to consider morphisms between Grothendieck topoi, the so-called *geometric morphisms*. Building on Chapter 2, a detailed derivation of the existence of *classifying topoi* is given; roughly speaking, the classifying topos of a (logical) theory can be regarded as the "categorification" of said theory. Classifying topoi provide a second example of the interplay between logic and Grothendieck topoi.

Finally, in Chapter 5, we introduce a slightly different notion of model compared to that of Chapter 2 which is also applicable to arbitrary sites: *continuous models*. Using results from the preceding chapters, we prove a number of interesting topos-theoretic results about sites and Grothendieck topoi, most of which can be found in [BM87].

All in all, this thesis intends to make the connections between Grothendieck topoi and logic explicit by exhibiting three particular forms of it: Kripke-Joyal semantics, classifying topoi and "continuous model theory". After reading this work, a good deal of the research literature should be accessible.

1 Grothendieck topoi

In this chapter, the main objects of study for the rest of this thesis will be introduced, namely **Grothendieck topoi**. We will assume the reader is familiar with standard category theory, approximately up to the Yoneda Lemma and the notion of adjointness. These subjects may be found in e.g. [Mac71] and [Awo10].

We will write $\operatorname{ob}\mathbf{C}$ and $\operatorname{mor}\mathbf{C}$ for the objects and morphisms of a category $\mathbf{C}$, respectively. Even though these may not be sets, we will find it convenient to use the natural expressions $C \in \operatorname{ob}\mathbf{C}$ and $f \in \operatorname{mor}\mathbf{C}$ to mean "C is an object of $\mathbf{C}$" and "f is a morphism of $\mathbf{C}$", respectively.

1.1 Presheaves

We begin with a series of definitions necessary to state the Yoneda Lemma, which is included for unambiguous reference and completeness.

Definition 1.1. A **presheaf** on a small category $\mathbf{C}$ is a contravariant functor $F : \mathbf{C} \to \mathbf{Set}$, or, equivalently, a covariant functor $F : \mathbf{C}^{\mathrm{op}} \to \mathbf{Set}$. The category $\mathbf{Set}^{\mathbf{C}^{\mathrm{op}}}$ of presheaves (with morphisms natural transformations) is denoted by $\hat{\mathbf{C}}$.

Definition 1.2. For a locally small category $\mathbf{C}$ and an object C of $\mathbf{C}$, the **representable presheaf** $\mathbf{y}_C \in \hat{\mathbf{C}}$ is defined as $\mathbf{y}_C := \mathrm{Hom}_{\mathbf{C}}(-, C)$.

Definition 1.3. The **Yoneda embedding** is the functor $\mathbf{y} : \mathbf{C} \to \hat{\mathbf{C}}$ defined on an object C of $\mathbf{C}$ by $\mathbf{y}_C = \mathrm{Hom}_{\mathbf{C}}(-, C)$, and on a morphism $f : C \to D$ by $\mathbf{y}_f(g) = fg$ for $g : B \to C \in \mathrm{Hom}_{\mathbf{C}}(B, C)$.

An appropriate form of the Yoneda Lemma is stated below; a detailed proof may be found in [Awo10, pp.188–192]. Note that this form differs from the presentation in [Mac71].

Proposition 1.4 (Yoneda Lemma)**.** *Let C be an object of a category $\mathbf{C}$, and let $P \in \hat{\mathbf{C}}$ be a presheaf on $\mathbf{C}$. There is a bijective correspondence:*

$$\begin{aligned} \mathrm{Hom}_{\hat{\mathbf{C}}}(\mathbf{y}_C, P) &\rightleftarrows PC \\ \eta &\longmapsto \eta_C(\mathrm{id}_C) \\ (f \mapsto Pf(p)) &\longleftarrow p \end{aligned}$$

natural in both P and C; this naturality is expressed by the following diagrams:

$$\begin{tikzcd} P \arrow[d, "\phi"'] & \mathrm{Hom}_{\hat{\mathbf{C}}}(\mathbf{y}_C, P) \arrow[r] \arrow[d, "(\eta \mapsto \phi\eta)"'] & PC \arrow[d, "\phi_C"] \\ Q & \mathrm{Hom}_{\hat{\mathbf{C}}}(\mathbf{y}_C, Q) \arrow[r] & QC \end{tikzcd} \qquad \begin{tikzcd} C \arrow[d, "f"'] & \mathrm{Hom}_{\hat{\mathbf{C}}}(\mathbf{y}_C, P) \arrow[r] & PC \\ D & \mathrm{Hom}_{\hat{\mathbf{C}}}(\mathbf{y}_D, P) \arrow[r] \arrow[u, "(\eta \mapsto \eta\mathbf{y}_f)"] & PD \arrow[u, "Pf"'] \end{tikzcd}$$

Presheaves have a very convenient structure when it comes to limits and colimits:

Proposition 1.5. *Let $D : \mathbf{J} \to \hat{\mathbf{C}}$ be a (small) diagram in $\hat{\mathbf{C}}$. Then the limit $\varprojlim_j D_j$ and colimit $\varinjlim_j D_j$ in $\hat{\mathbf{C}}$ exist, and are given on objects by:*

$$\left(\varprojlim_j D_j\right) C = \varprojlim_j (D_j C) \qquad \left(\varinjlim_j D_j\right) C = \varinjlim_j (D_j C)$$

Proof Since **Set** is both complete and cocomplete, this follows from Theorem V.3.1 in [Mac71]. □

Furthermore, there are also exponentials in $\hat{\mathbf{C}}$:

Proposition 1.6. *Let $P, Q \in \hat{\mathbf{C}}$ be presheaves. The exponential P^Q is the presheaf defined by $P^Q(C) := \mathrm{Hom}(\mathbf{y}_C \times Q, P)$ and $P^Q(f) := \eta \mapsto \eta \circ \langle \mathbf{y}_f, \mathrm{id}_Q \rangle$*

Proof Since the exponential arises as a right adjoint to the categorical product, if P^Q were to exist, it would have to satisfy $\mathrm{Hom}(R, P^Q) \cong \mathrm{Hom}(R \times Q, P)$. In particular, with $R = \mathbf{y}_C$, we find that by the Yoneda lemma $P^Q(C) \cong \mathrm{Hom}(\mathbf{y}_C, P^Q) \cong \mathrm{Hom}(\mathbf{y}_C \times Q, P)$. This leads us to *define* $P^Q(C)$ by the rightmost expression. The associated evaluation $e : P^Q \times Q \to P$ is defined by $e_C(\theta, p) := \theta_C(\mathrm{id}_C, p)$. That e, so defined, satisfies the UMP for the exponential (i.e., is the counit of the adjunction) is left to the reader; a proof can also be found in [MM92, §I.6]. $\square$

It is sometimes convenient to be able to restrict attention to representable presheaves only. To this end, the following proposition is useful:

Proposition 1.7. *Every presheaf $P \in \hat{\mathbf{C}}$ is in a canonical fashion the colimit of a small diagram of representable presheaves.*

Proof Fix an arbitrary presheaf P. First, let us define the small index category associated to it. It is called the **category of elements of** P and denoted by $\int_{\mathbf{C}} P$. Its objects are pairs (x, C) such that $x \in PC$; a morphism $f : (p, C) \to (p', C')$ is a morphism $f : C \to C'$ of $\mathbf{C}$ satisfying $Pf(p') = p$. Composition is given by composition in $\mathbf{C}$. Now define a diagram $D : \int_{\mathbf{C}} P \to \hat{\mathbf{C}}$ and cocone η to P by:

$$\begin{tikzcd}
(p,C) \arrow[dd, "f"'] & & \mathbf{y}_C \arrow[dd, "\mathbf{y}_f"] \arrow[dr, "\eta_{(p,C)}"] & \\
 & \overset{D}{\Longrightarrow} & & P \\
(p',C') & & \mathbf{y}_{C'} \arrow[ur, "\eta_{(p',C')}"'] &
\end{tikzcd}$$

where $\eta_{(p,C)}$ is taken to be the morphism $\mathbf{y}_C \to P$ associated to p in the Yoneda Lemma. That the right triangle commutes is a consequence of the naturality in the object C in the Yoneda Lemma.
To show that the cocone so defined actually is the colimit, suppose we have a cocone:

$$\begin{tikzcd}
\mathbf{y}_C \arrow[dd, "\mathbf{y}_f"] \arrow[dr, "\nu_{(p,C)}"] & \\
 & Q \\
\mathbf{y}_{C'} \arrow[ur, "\nu_{(p',C')}"'] &
\end{tikzcd}$$

If $\mu : P \to Q$ is to exist, it must be that $\mu \circ \eta_{(p,C)} = \nu_{(p,C)}$ from the Yoneda lemma; in particular, one needs $\mu_C(\eta_{(p,C)})_C(\mathrm{id}_C) = (\nu_{(p,C)})_C(\mathrm{id}_C)$, i.e. $\mu_C(p) = (\nu_{(p,C)})_C(\mathrm{id}_C)$. That μ, so defined, is a natural transformation follows from the naturality in the presheaf P in the Yoneda Lemma. Moreover, μ is unique, as it is completely determined by the conditions it needs to fulfil; we conclude that indeed P is the colimit of D in $\hat{\mathbf{C}}$. $\square$

1.2 Grothendieck topologies

The notion of a **Grothendieck topology** is a translation of the familiar notion of **topology** to categorical context. Its primary aim is to facilitate the definition of **sheaves** in category-theoretic terms.

Definition 1.8. Let C be an object of a category $\mathbf{C}$. Let us write $m \leq m'$ for two monomorphisms $m : M \rightarrowtail C$ and $m' : M' \rightarrowtail C$ such that there exists a morphism $f : M' \rightarrowtail M$ with $m' = m \circ f$ as in:

$$\begin{tikzcd}
M \arrow[r, tail, "m"] & C \\
M' \arrow[u, tail, "f"] \arrow[ur, tail, "m'"'] &
\end{tikzcd}$$

Since $mf = m'$, it readily follows that f is monic. It is clear that $\leq$ is a preordering on monos to C, and hence induces an equivalence relation $\sim$, by putting $m \sim m'$ iff $m \leq m'$ and $m' \leq m$. A **subobject** of C is an equivalence class under $\sim$ of monos to C. The collection of subobjects of C is denoted $\mathrm{Sub}(C)$. The subobject corresponding to the mono id_C is called the **maximal subobject** of C.

The preordering $\leq$ then becomes an ordering on $\mathrm{Sub}(C)$. Pullback serves as a binary infimum (or "meet") in this ordered set, because the pullback of a mono is again monic, and so is the composition of monos:

$$\begin{tikzcd}
M \arrow[r, tail, "m"] & C \\
M \wedge M' \arrow[u, tail] \arrow[r, tail] & M' \arrow[u, tail, "m'"']
\end{tikzcd}$$

It is readily verified that this is well-behaved on equivalence classes. In the sequel, the distinction between a mono $m : M \rightarrowtail C$ and its equivalence class will often tacitly be ignored.

In categories of presheaves, subobjects take a particularly simple form.

Definition 1.9. For two presheaves $P, Q \in \hat{\mathbf{C}}$, say Q is a **subpresheaf** of P if for each C, we have an inclusion $QC \subseteq PC$.

Proposition 1.10. *Let P be a presheaf. There is a bijective correspondence between subobjects of P and subpresheaves of P.*

Proof Every subpresheaf Q of P obviously induces a subobject (corresponding to the inclusion $Q \hookrightarrow P$). Conversely, let S be a subobject of P, represented by $m : M \rightarrowtail P$. Define $\overline{S}$ by $\overline{S}C = \{p \in PC \mid \exists x \in MC : m_C x = p\} \subseteq PC$, and $\overline{S}f$ to be the restriction of Pf to $\overline{S}(\operatorname{cod} f)$. Then clearly $\overline{S}$ is well-defined and a subpresheaf of P. Because each m_C is an injection, it is a bijection onto its image, which is $\overline{S}C$. It follows that $m : M \rightarrowtail \overline{S}$ is iso, and thus the subobject S induces a subpresheaf. The constructions are clearly inverse to each other. $\square$

As a useful and nontrivial consequence, we have that $\mathrm{Sub}(P)$ is a set, for each presheaf P. We are specifically interested in subpresheaves of representable presheaves:

Definition 1.11. Let C be an object of a category $\mathbf{C}$. A **sieve** on C is a subobject S of $\mathbf{y}_C$ (in $\hat{\mathbf{C}}$). Equivalently, a sieve on C is a collection S of morphisms of $\mathbf{C}$ with codomain C, subject to:

$$f \in S \to fg \in S$$

for every morphism g of $\mathbf{C}$ with $\operatorname{dom} f = \operatorname{cod} g$. This precisely constitutes a subpresheaf of $\mathbf{y}_C$. The maximal subobject t_C of $\mathbf{y}_C$ is called the **maximal sieve** on C.

With the notion of sieve under our belt, we can now state the definition of a Grothendieck topology.

Definition 1.12. A **Grothendieck topology** on a small category $\mathbf{C}$ is a mapping J that assigns to each object C of $\mathbf{C}$ a collection of sieves JC, subject to:

(i) $t_C \in JC$;

(ii) If $S \in JC$ and $h : D \to C$ is a morphism, then $h^*S := \{g \mid \operatorname{cod} g = D, hg \in S\} \in JD$;

(iii) If $S \in JC$ and $h^*R \in JD$ for each $h : D \to C \in S$, then also $R \in JC$.

The conditions (ii) and (iii) are referred to as the **stability** and **transitivity axiom**, respectively. The elements of JC are called **covering sieves** (on C). A collection of morphisms $\{f_i : C_i \to C \mid i \in I\}$ is called **covering**, or **J-covering**, if the generated sieve $\{f_i g \mid \operatorname{cod} g = \operatorname{dom} f_i\}$ is in JC; a morphism that is covering by itself is also called a **singleton cover**.

In specifying a Grothendieck topology, it is often simpler to specify sieves, and topologies, by **generators**. This is possible because both sieves and topologies are defined using certain closure conditions; it is well-known that such closure conditions are preserved under arbitrary intersection. The definitions are as follows:

Definition 1.13. Let $\mathbf{C}$ be a small category and let $C \in \operatorname{ob}\mathbf{C}$. For a collection $F = \{f_i : C_i \to C\}$ of morphisms, the **sieve** (F) **generated by** F is the smallest sieve that contains all f_i. Explicitly, it can be given as: $(F) = \{f_i g \mid \operatorname{cod} g = C_i\}$. Given a topology J on $\mathbf{C}$, we say that F **covers** if $(F) \in JC$.

For a collection $\mathcal{S} = \{\mathcal{S}C : C \in \operatorname{ob}\mathbf{C}\}$ of sieves, the **topology generated by** $\mathcal{S}$ is the smallest topology $J = J(\mathcal{S})$ that makes all sieves in $\mathcal{S}$ covering. We also say that $\mathcal{S}$ **generates** J or that $\mathcal{S}$ is a **basis for** J.

Of course, these two definitions can be combined to define a topology J by specifying certain covering collections (whose generated sieves then in turn generate J).

If $\mathbf{C}$ has pullbacks, the generated topology $J(\mathcal{S})$ can be constructed in three steps: First, add the maximal sieves to $\mathcal{S}$. Next, let $J_0(\mathcal{S})D = \{h^*S \mid h : D \to C, S \in \mathcal{S}C\}$. It is clear that $J_0(\mathcal{S})$ contains $\mathcal{S}$. Lastly, we let:

$$J(\mathcal{S})C = \{R \mid \exists S \in J_0(\mathcal{S})C : \forall h \in S : h^*R \in J_0(\mathcal{S})\}$$

We observe that $J_0(\mathcal{S})$ is contained in $J(\mathcal{S})$. That $J(\mathcal{S})$ is a topology is a formal category-theoretic consequence of its construction; the reader is invited to draw the relevant diagrams. That $J(\mathcal{S})$ is indeed the smallest topology is entirely obvious.

The notation h^*S in the above definitions is justified by the subobject form of sieves. Namely, h^*S fits into the following pullback diagram:

$$\begin{tikzcd} \mathbf{y}_D \arrow[r, "\mathbf{y}_h"] & \mathbf{y}_C \\ h^*S \arrow[u, tail] \arrow[r] & S \arrow[u, tail] \end{tikzcd} \tag{1.1}$$

which at first glance only suggests the notation $\mathbf{y}_h^*S$; however, a common abuse of notation in the literature is to write h instead of the more accurate $\mathbf{y}_h$.

There are some useful properties of covering sieves that are immediate from the definition:

Proposition 1.14. *Let $C \in \operatorname{ob}\mathbf{C}$, and let R, S be sieves on C.*

(i) *If $S \in JC$ is covering and $S \leq R$, then R is also covering.*

(ii) *If $R, S \in JC$ are covering, then so is their meet $R \wedge S$ (in $\operatorname{Sub}(\mathbf{y}_C)$). Thus, covering sieves are stable under meet.*

Proof (i) By the stability axiom, it follows that if $h : D \to C \in S$, then $h^*S = t_D \in JD$. Hence also $h^*R = t_D \in JD$; it thus follows by transitivity that $R \in JC$ – i.e., R is covering.
(ii) It is easily seen that $R \wedge S = R \cap S$ in the set-theoretic perspective on sieves. Then for each $h \in S$:

$$h^*(R \cap S) = h^*R \cap h^*S = h^*R$$

where the last step follows because h^*S is again the maximal sieve; since h^*R is covering (by stability), it follows that $R \wedge S$ is covering as well (by transitivity). $\square$

Definition 1.15. A **site** is a pair $(\mathbf{C}, J)$, where $\mathbf{C}$ is a small category, and J is a Grothendieck topology on $\mathbf{C}$.

The notion of a site is central to our discussion. While it originated as an abstraction of a topological space and led Grothendieck et al. to great successes in algebraic geometry, its applications are nowadays much more diverse than just that. What makes the study of these objects so interesting is that they support suitable forms of **continuity** and, as mentioned, of **sheaf**, in a categorical context. We will focus on sheaves now, and return to the continuity aspect later, in Chapter 4.

1.3 Sheaves

Like in the topological setting, a sheaf on a site is a piece of data that can be reassembled from its, in a sense to be made precise, "local" properties. In this section, we will give the definition of sheaves on a site, which properly generalizes sheaves on topological spaces. For this purpose, let us fix a site $(\mathbf{C}, J)$.

Definition 1.16. Let $P \in \hat{\mathbf{C}}$ be a presheaf, and let S be a sieve on $C \in \operatorname{ob}\mathbf{C}$. A **matching family** $\langle x_f \rangle_{f \in S}$ for S by P is a collection of elements x_f of P such that for $f : D \to C$, $x_f \in PD$, and furthermore, whenever the composition fg is defined:

$$x_{fg} = Pg(x_f)$$

It is common to denote $Pg(x_f)$ as $x_f \cdot g$. Taking a more abstract viewpoint, a matching family may also be described as a natural transformation $S \to P$ in $\hat{\mathbf{C}}$ (given by $f \mapsto x_f$).

An **amalgamation** for $\langle x_f \rangle_{f \in S}$ is an $x \in PC$ such that for all $f \in S : x_f = x \cdot f$. Via Yoneda's lemma, this can be viewed as an extension:

$$\begin{tikzcd} S \arrow[r] \arrow[d, hook] & P \\ \mathbf{y}_C \arrow[ur, "x"'] & \end{tikzcd}$$

Definition 1.17. A presheaf $P \in \hat{\mathbf{C}}$ is said to be **separated** (for J) if for each covering sieve $S \in JC$, each matching family for S by P has *at most* one amalgamation. If additionally every matching family for a covering sieve actually has an amalgamation, then P is called a **sheaf** on $\mathcal{C}$.

It is in general not true that a representable presheaf is a sheaf; whether representable presheaves are sheaves on a site is important information, meriting a dedicated definition:

Definition 1.18. Let $\mathbf{C}$ be a small category. A Grothendieck topology J on $\mathbf{C}$ for which all representable presheaves on $\mathbf{C}$ are sheaves is called **subcanonical**. The **canonical topology** on a category $\mathbf{C}$ is the largest subcanonical topology on $\mathbf{C}$.

We are now can now state the definition of a Grothendieck topos.

Definition 1.19. The **category of sheaves** on $(\mathbf{C}, J)$, denoted $\mathbf{Sh}(\mathbf{C}, J)$, is the full subcategory of $\hat{\mathbf{C}}$ whose objects are sheaves for J (so with morphisms all morphisms between them in $\hat{\mathbf{C}}$). A **Grothendieck topos** is a category that is equivalent to some category of sheaves on a site.

Much of what follows will revolve around Grothendieck topoi, be it as carrying structures (Chapters 2 and 3) or as the main objects of study (Chapters 4 and 5).

1.4 The associated sheaf functor

The discussion in this section largely follows the elegant presentation in [MM92, §III.5].
Since $\mathbf{Sh}(\mathbf{C}, J)$ is a subcategory of $\hat{\mathbf{C}}$, there is a forgetful functor $U : \mathbf{Sh}(\mathbf{C}, J) \to \hat{\mathbf{C}}$. We will now set out to prove that this U has a left adjoint, $\mathbf{a} : \hat{\mathbf{C}} \to \mathbf{Sh}(\mathbf{C}, J)$. Thus, for a presheaf P, $\mathbf{a}P$ will be a sheaf such that for every other sheaf F, $\operatorname{Hom}_{\hat{\mathbf{C}}}(P, UF) \cong \operatorname{Hom}_{\mathbf{Sh}(\mathbf{C})}(\mathbf{a}P, F)$ and this correspondence will be *dinatural* (contravariant in P, covariant in F). The situation can be conveniently illustrated with a diagram in $\hat{\mathbf{C}}$: for each $\eta : P \to F$, there will be a morphism fitting the dashed arrow in:

$$\begin{tikzcd} \mathbf{a}P \arrow[r, dashed] & F \\ P \arrow[u, "\mathbf{a}"] \arrow[ur, "\eta"'] & \end{tikzcd}$$

This diagram suggests that we think of $\mathbf{a}P$ as the "sheaf form" of P. Without further ado, let us continue to construct $\mathbf{a}P$.

When $S \leq R$ for $R, S \in JC$, we say that S **refines** R, or is a **refinement** of R. A consequence of Proposition 1.14(ii) is then that every two covering sieves have a common refinement. Thus every JC is a so-called **filtered poset** under the refinement ordering – see Section A.1 from the Appendix for some results about filtered categories.

Suppose S refines R. If P is a presheaf and $x : R \to P$ is a matching family, then $x \restriction_S : S \to P$ is a matching family for S. Thus the poset structure of JC "translates" to the matching families. More categorically, the assignment:

$$R \mapsto \mathrm{Match}_C(R, P), \qquad (S \leq R) \mapsto (\restriction_S : \mathrm{Match}_C(R, P) \to \mathrm{Match}_C(S, P))$$

where $\mathrm{Match}_C(R, P)$ is the set of matching families for R by P (i.e., the hom-set $\hat{\mathbf{C}}(R, P)$), is contravariantly functorial. In particular, the filtered poset JC is mapped to a cofiltered diagram $(\mathrm{Match}_C(R, P))_{R \in JC}$ in **Set**. Similar to the construction of the associated sheaf in topology, we can now form the colimit $\varinjlim_{R \in JC} \mathrm{Match}_C(R, P)$, which we will call $P^+(C)$. We remark that this colimit can be constructed in the familiar way, via a disjoint sum and an equivalence relation.

Given $h : D \to C$, $R \in JC$ and $x \in \mathrm{Match}_C(R, P)$, x can be precomposed with the morphism $h^*R \to R$ from diagram (1.1) to obtain $h^*x \in \mathrm{Match}_C(h^*R, P)$. It is clear that $h^* : \mathrm{Match}_C(R, P) \to \mathrm{Match}_D(h^*R, P)$ commutes with restriction. Thus it yields a mapping $h^* : P^+(C) \to P^+(D)$ between the colimits. The assignment $h \mapsto h^*$ is contravariant, and hence P^+ is a presheaf, whose effect on morphisms is given by $P^+h = h^*$.

Before turning to the properties that P^+ has, it is good to remark that $P \mapsto P^+$ is a functor from $\hat{\mathbf{C}}$ to itself. Indeed, given $\eta : P \to Q$, there are, for each C, evident induced functors $\mathrm{Match}_C(-, P) \to \mathrm{Match}_C(-, Q)$ given by composition with η. These in turn induce a functor $\eta^+_C : P^+(C) \to Q^+(C)$ for each C. That these η^+_C match together to form a natural transformation $\eta^+ : P^+ \to Q^+$ follows from a diagram chase.

The following three lemmata will establish the required properties of the plus functor.

Lemma 1.20. *For all presheaves P, P^+ is a separated presheaf.*

Proof Suppose that $x, y \in P^+C$ both amalgamate a matching family $S \to P$ for a cover S, i.e. for all $f : D \to C, f \in S$ we have $x \cdot f = y \cdot f$. With moderate abuse of notation, let $x : R \to P$ and $y : R' \to P$ be matching families representing x and y respectively. Then $x \cdot f : f^*R \to P$ being equal to $y \cdot f : f^*R' \to P$ in P^+D means that there exists a covering sieve $T_f \leq f^*R \cap f^*R'$ on D with $(x \cdot f) \cdot g = (y \cdot f) \cdot g$ for all $g \in T_f$. Now the sieve $T = \{fg : f \in S, g \in T_f\}$ satisfies $T \leq R \cap R'$ and $x \cdot t = y \cdot t$ for all $t \in T$. Furthermore, if $f \in S$ then $T_f \leq f^*T$; Proposition 1.14(i) ensures f^*T is covering, and by the transitivity axiom, T is as well. By definition of the equivalence relation on P^+C, we conclude $x = y$; hence P^+ is separated. □

Lemma 1.21. *If P is a separated presheaf, then P^+ is a sheaf.*

Proof Suppose that $m : S \to P^+$ is a matching family for a covering sieve S on C. Thus for all $f \in S$ and g composable with f, it holds that $m_f \cdot g = m_{fg}$ as elements of P^+. Hence if $m_{f,(-)} : R_f \to P$ and $m_{fg,(-)} : R_{fg} \to P$ represent m_f and m_{fg}, respectively, then $m_f \cdot g = m_{fg}$ amounts to the existence of a covering sieve $T_{f,g} \leq g^*(R_f) \cap R_{fg}$ such that for each $t \in T_{f,g} : m_{f,gt} = m_{fg,t}$.
Now define the sieve $T := \{fg : f \in S, g \in R_f\}$ on C, which is covering because for all $f \in S$, $f^*T \geq R_f$ is covering. We purport a well-defined matching family $x : T \to P$ is given by $x_{fg} = m_{f,g}$; if this definition is valid then x amalgamating m is equivalent to m_f and $x \cdot f$ having a common refinement. Now $R_f \leq f^*T$ provides the desired refinement: for each $g \in R_f$, we have $m_f \cdot g = m_{f,g} = x_{fg} = x_f \cdot g$ as desired.
It only remains to verify that x is well-defined. So suppose that $fg = f'g'$. We will show that $m_{f,g}$ and $m_{f',g'}$ are both amalgamations for the covering sieve $T_{f,g} \cap T_{f',g'}$, as is shown by taking an arbitrary $h \in T_{f,g} \cap T_{f',g'}$:

$$\begin{aligned} m_{f,g} \cdot h &= m_{f,gh} && \text{(as } m_f \text{ is matching for } P) \\ &= m_{fg,h} && \text{(since } h \in T_{f,g}) \\ &= m_{f'g',h} = m_{f',g'h} = m_{f',g'} \cdot h \end{aligned}$$

where the last equalities use that $h \in T_{f',g'}$ and $m_{f'}$ is matching for P too. Now since P is separated, we conclude $m_{f,g} = m_{f',g'}$ and the proof is finished. □

Lemma 1.22. *Let P be a presheaf, and let F be a sheaf. Then any morphism $\eta : P \to F$ factors (uniquely) through $P \to P^+$.*

Proof Suppose $x : R \to P$ is an element of P^+C. For $f \in RD$, $x_f \in PD$ corresponds to $\{x_f \cdot g : g \in t_D\}$ under $P \to P^+$. Observe that these form a matching family for R in P^+; moreover, x is an amalgamation for this matching family. Now $\tilde{\eta} : P^+ \to F$, if it is to exist, must preserve amalgamations. Thus $\tilde{\eta}_C(x)$ must be the amalgamation of the $\tilde{\eta}_D(\{x_f \cdot g : g \in t_D\}) = \eta_D(x_f)$. Hence $\tilde{\eta}$ is uniquely dictated by η, and it exists because F is a sheaf. Technically it is still to be verified that $\tilde{\eta}$ is a natural transformation, but this is completely trivial. $\square$

Definition 1.23. Let $(\mathbf{C}, J)$ be a site. The **associated sheaf functor** (also known as the **sheafification functor**) $\mathbf{a} : \hat{\mathbf{C}} \to \mathbf{Sh}(\mathbf{C}, J)$ is defined by $\mathbf{a}P := (P^+)^+$.

The preceding lemmata now show that indeed $\mathbf{a}$ takes presheaves to sheaves, and that it has the desired factorization property. These facts are summarized in the following theorem:

Theorem 1.24. *The sheafification functor* $\mathbf{a}$ *is left adjoint to the forgetful functor* $U : \mathbf{Sh}(\mathbf{C}, J) \to \hat{\mathbf{C}}$. *Moreover,* $\mathbf{a}U$ *is naturally isomorphic to the identity functor on* $\mathbf{Sh}(\mathbf{C}, J)$.

Proof Applying Lemma 1.22 twice shows that we have correspondences:

$$\mathrm{Hom}_{\hat{C}}(P, UF) \cong \mathrm{Hom}_{\hat{C}}(P^+, UF) \cong \mathrm{Hom}_{\hat{C}}((P^+)^+, UF) = \mathrm{Hom}_{\mathbf{Sh}(\mathbf{C},J)}(\mathbf{a}P, F)$$

The required dinaturality is immediate upon drawing the diagrams. This establishes the adjunction $\mathbf{a} \dashv U$. The remaining part of the theorem statement follows from applying Lemma 1.22 to $\mathrm{id}_F : F \to F$ for a sheaf F, which implies that $F \to \mathbf{a}F$ is an isomorphism. $\square$

1.5 Limits, colimits and exponentials of sheaves

The sheafification functor $\mathbf{a}$ enables an explicit description of the behaviour of limits and colimits in $\mathbf{Sh}(\mathbf{C}, J)$.

Proposition 1.25. *The sheafification functor* $\mathbf{a}$ *preserves finite limits. In particular,* $\mathbf{Sh}(\mathbf{C}, J)$ *has all finite limits.*

Proof It will suffice to show that the "plus functor" $P \mapsto P^+$ preserves finite limits. For fixed C and $R \in JC$, the functor $P \mapsto \mathrm{Match}_C(R, P)$ is just the covariant representable functor $\mathrm{Hom}_{\hat{\mathbf{C}}}(R, -)$, which is well-known to preserve finite limits. From the $\mathrm{Match}_C(R, P)$, P^+C is constructed by a cofiltered colimit. In $\mathbf{Set}$, we have from Theorem A.2 that finite limits commute with cofiltered colimits. The pointwise nature of limits and colimits in $\hat{\mathbf{C}}$ means that this property immediately carries over to $\hat{\mathbf{C}}$ as well. Thus all parts comprising the plus functor preserve finite limits; hence, being their composition, so does the plus functor. In conclusion, $\mathbf{a}$ preserves finite limits, which thence exist in $\mathbf{Sh}(\mathbf{C}, J)$. $\square$

Proposition 1.26. $\mathbf{Sh}(\mathbf{C}, J)$ *is cocomplete (i.e., has all (small) colimits).*

Proof Recall from Proposition 1.5 that $\hat{\mathbf{C}}$ is cocomplete. Now $\mathbf{a}$, being a left adjoint, preserves colimits (cf. [Mac71, §V.5]). Since every object of $\mathbf{Sh}(\mathbf{C}, J)$ is (isomorphic to) one of the form $\mathbf{a}P$, this shows $\mathbf{Sh}(\mathbf{C}, J)$ is also cocomplete. $\square$

Proposition 1.27. *Let* P^Q *be a presheaf exponential, and let* P *be a sheaf. Then* P^Q *is also a sheaf.*

Proof This proposition is an abstract consequence of the fact that we have an adjunction $\mathbf{a} \dashv U$ where $U : \mathbf{Sh}(\mathbf{C}, J) \to \hat{\mathbf{C}}$ is an inclusion of a full subcategory and $\mathbf{a}$ preserves finite limits. The general result appears as Proposition A4.3.1 in [Joh02].

Namely, let N be an arbitrary presheaf; then we have bijective correspondences, natural in N:

$$\begin{array}{c}
N \longrightarrow UP^Q \\ \hline
N \times Q \longrightarrow UP \\ \hline
\mathbf{a}N \times \mathbf{a}Q \longrightarrow P \\ \hline
\mathbf{a}U\mathbf{a}N \times \mathbf{a}Q \longrightarrow P \\ \hline
U\mathbf{a}N \times Q \longrightarrow UP \\ \hline
U\mathbf{a}N \longrightarrow UP^Q
\end{array}$$

Here, the third and the penultimate line use that $\mathbf{a}$ preserves products, while the fourth employs that U is an inclusion (so that $\mathbf{a}U$ is naturally isomorphic to the identity). Applying the correspondence between the first and last line to $N = UP^Q$ and the identity on UP^Q, we obtain that UP^Q is the inclusion of $\mathbf{a}(UP^Q)$; i.e. UP^Q is a sheaf. □

Corollary 1.28. $\mathbf{Sh}(\mathbf{C}, J)$ *has exponentials.*

1.6 Subobject classifier

In **Set**, there is the well-known correspondence between $\mathcal{P}(X)$, the power set of a set X, and 2^X, the set of functions from X to the set $2 = \{0, 1\}$, which plays the role of "set of truth-values". It is of paramount importance in set theory, and any attempt at a generalisation of **Set** in categorical language must have an analogue. The theory of topoi (of which Grothendieck topoi are examples) provides such a generalisation by means of the notion of a subobject classifier. We give the definition in full generality, and will then set out to prove that $\mathbf{Sh}(\mathbf{C}, J)$ admits a subobject classifier.

Definition 1.29. Let $\mathbf{C}$ be a category having all finite limits. Denote with 1 its terminal object. A **subobject classifier** for $\mathbf{C}$ is a mono true : $1 \to \Omega$ such that for each subobject $m : M \rightarrowtail X$, there is a unique $\chi_m : X \to \Omega$ making:

$$\begin{tikzcd} M \arrow[r, dashed] \arrow[d, tail, "m"'] & 1 \arrow[d, "\mathrm{true}"] \\ X \arrow[r, "\chi_m"'] & \Omega \end{tikzcd}$$

a pullback diagram.

In this setup, Ω can be thought of as the "object of truth-values", and χ_m as the "characteristic function" of m.

In the familiar situation of **Set**, recall or observe that a subobject is equivalent to a subset. The characteristic function χ_M of $M \subseteq X$ and true : $1 \to 2, \mathrm{true}(*) = 1$ are easily seen to satisfy the definition of a subobject classifier. This observation reassures us that the proposed generalization is indeed a proper generalization. The reader may find it instructive to investigate the subobject classifier for the functor category $\mathbf{Set}^I$ with I a set (i.e., a discrete category).

The subobject classifier Ω gives rise to a bijection $\mathrm{Sub}(C) \cong \mathrm{Hom}(C, \Omega)$, assigning to a subobject its classifying morphism. This bijection is natural, in the following sense:

Proposition 1.30. *For any morphism $f : C \to D$, the following diagram commutes:*

$$\begin{tikzcd} \mathrm{Sub}(D) \arrow[r, "f^*"] \arrow[d, "\cong"'] & \mathrm{Sub}(C) \arrow[d, "\cong"] \\ \mathrm{Hom}(D, \Omega) \arrow[r, "\circ f"'] & \mathrm{Hom}(C, \Omega) \end{tikzcd}$$

Proof Let $X \hookrightarrow D$ be a subobject of D. The result follows from applying the pullback lemma to the following diagram:

$$\begin{tikzcd} f^*X \arrow[r] \arrow[d, hook] & X \arrow[r] \arrow[d, hook] & 1 \arrow[d] \\ C \arrow[r, "f"'] & D \arrow[r] & \Omega \end{tikzcd}.$$

□

In order to describe the subobject classifier of Grothendieck topoi, we need some terminology.

Definition 1.31. Let $(\mathbf{C}, J)$ be a site, and let S be a sieve on an object C. Then S is **closed** if for all $f : D \to C$ it holds that $f^*S \in JD$ iff $f \in S$ (i.e., iff f^*S is maximal).

With the notion of closed comes a notion of closure. The **closure** $\bar{S}$ of S is the smallest closed sieve that contains S. It can also be given explicitly:

Proposition 1.32. *The closure $\bar{S}$ of a sieve S is given by $\bar{S} = \{f : D \to C \mid f^*S \in JD\}$.*

Proof Suppose $f \in \bar{S}$, i.e. $f^*S \in JD$. Since $(fg)^*S = g^*(f^*S)$ for all suitable g it follows by the stability axiom that $fg \in \bar{S}$ as well. Thus $\bar{S}$ is a sieve. Now to show that $\bar{S}$ is closed. Suppose that $h^*\bar{S}$ is a covering sieve. It is to be shown that h^*S is covering as well. For any $k \in h^*\bar{S}$, we have $k^*(h^*S) = (hk)^*S$, and the latter is covering since $hk \in \bar{S}$; hence h^*S is covering, by the transitivity axiom. That $\bar{S}$ is indeed the smallest closed sieve is obvious. □

If S is closed, then it is immediate that f^*S is also closed, for all f. The nature of the pullback operation f^* makes it suitable for defining a contravariant functor, more precisely a presheaf. We thus define a presheaf Ω on $\mathbf{C}$ by $\Omega C := \{S \mid S \text{ is a closed sieve on } C\}$, and $\Omega f(S) = f^*S$.

Proposition 1.33. *The presheaf $\Omega \in \hat{C}$ is a sheaf.*

Proof We begin by showing Ω is separated. Suppose $S \in JC$ and $M, N \in \Omega C$ such that for all $f \in S$, $f^*M = f^*N$. Then for $f : D \to C$, $f \in M \cap S$ means $f^*M = f^*N \in JD$; thus $f \in N$ as N is closed; hence $M \cap S \subseteq N$. Then for $f \in M$, $f^*(M \cap S) = f^*M \cap f^*S \in JD$, and since $M \cap S \subseteq N$, $f^*N \in JD$ by Proposition 1.14. Thus $f \in N$, hence $M \subseteq N$. Interchanging the roles of M and N proves $M = N$.

To prove Ω is a sheaf, let $S \in JC$, and for $f : D \to C$, $f \in S$, let $M_f \in \Omega D$, such that $M_{fg} = g^*M_f$ for all $g : E \to D$. Define $M = \{fg : g \in M_f, f \in S\}$. We will show that $\overline{M}$ amalgamates the M_f. By the easy fact $f^*\overline{M} = \overline{f^*M}$, the condition $f^*\overline{M} = M_f$ will be satisfied as soon as $f^*M = M_f$, since M_f is closed. Of this last equality, the inclusion $M_f \subseteq f^*M$ is obvious. Suppose now that $h \in f^*M$, i.e. $fh = f'g'$ for some $g' \in M_{f'}$. In particular, $M_{fh} = M_{f'g'}$, which by the matching condition implies $h^*M_f = g'^*M_{f'}$. Since $g' \in M_{f'}$, the latter is a maximal sieve; this immediately implies $h \in M_f$, establishing $f^*M \subseteq M_f$. By our earlier argument, we conclude that $f^*\overline{M} = M_f$ for all $f \in S$. Thus Ω is a sheaf. □

Theorem 1.34. *Ω is a subobject classifier for $\mathbf{Sh}(\mathbf{C}, J)$.*

Proof First, let us define the morphism true : $1 \to \Omega$, by true$(C) := t_C$, the maximal sieve on C. So let A be a subsheaf of a sheaf F. Let us define a candidate classifying morphism $\chi_A : F \to \Omega$ by $(\chi_A)_C(x) = \{f : D \to C \mid x \cdot f \in AD\}$. Let us verify that this definition yields a closed sieve: for $g : D \to C$, suppose $g^*((\chi_A)_C(x)) = \{h : D \to E \mid x \cdot (gh) \in AE\} \in JE$. Since $x \cdot (gh) = (x \cdot g) \cdot h$, the only candidate for amalgamation is $x \cdot g$; since A is a sheaf, it follows that $x \cdot g \in AD$, and so $g \in (\chi_A)_C(x)$, making it a closed sieve. For $f : E \to D$ and $g : D \to C$, the computation:

$$\begin{aligned} f \in (\chi_A)_D(x \cdot g) &\iff x \cdot (gf) \in AE \\ &\iff gf \in (\chi_A)_C(x) \\ &\iff f \in g^*((\chi_A)_C(x)) \end{aligned}$$

shows that χ_A is also a natural transformation $F \to \Omega$. That the required diagram

$$\begin{tikzcd} A \arrow[r] \arrow[d] & 1 \arrow[d, "\text{true}"] \\ F \arrow[r, "\chi_A"] & \Omega \end{tikzcd}$$

is a pullback can be verified by considering it as a diagram in $\hat{\mathbf{C}}$, via Proposition 1.25. These are computed pointwise, and it remains to establish that for any C and $x \in FC$, $x \in AC$ iff $(\chi_A)_C(x) = t_C$. This statement is entirely obvious; finally, to establish that χ_A is unique, we can apply the naturality of χ_A and the condition $\mathrm{id}_C \in (\chi_A)_C(x)$ iff $x \in AC$ to conclude $f \in (\chi_A)_C(x)$ iff $x \cdot f \in AD$, for all $f : D \to C$. □

A further investigation of the properties of the subobject classifier in a (Grothendieck) topos can be found in §III.8 and §IV.8 of [MM92]. Most of the results of the preceding sections can be succinctly summarized by introducing the notion of an elementary topos:

Definition 1.35. A category $\mathbf{E}$ is called an **elementary topos**, or simply **topos**, if $\mathbf{E}$ has the following constructs:

(i) finite limits;

(ii) finite colimits;

(iii) exponentials;

(iv) a subobject classifier.

These are not independent conditions (e.g. (ii) is implied by the others), but we will not go into this.

Thus the results from the preceding sections assure us that indeed the notions Grothendieck topos and elementary topos are compatible, since the former are instances of the latter. A result known as Giraud's theorem subsequently characterizes the Grothendieck topoi among the elementary topoi (stronger, among all categories) by a few axioms. As even the statement of Giraud's theorem takes considerable forework, we defer it to Chapter 3, where it appears as Theorem 3.7.

2 Categorical logic

In this chapter, the building blocks for a categorical approach to logic are set out. The main reference used to do this is [MR77], of which Chapters 2 and 3 contain most of the basic material. Later, we will hook into selected material from more advanced chapters of [MR77].

It will turn out that suitably rich categories play a similar role to Boolean algebras for classical propositional logic. In Chapter 4, we will exhibit categories with a universal property reminiscent of that of the Lindenbaum-Tarski algebra in propositional logic.

2.1 Formulae and fragments

This section contains a brief account of many-sorted, infinitary first-order logic (with equality) and some fragments thereof. Readers familiar with these concepts may safely skip it. Because of said brevity, readers unfamiliar with any kind of symbolic logic are advised to consult additional material. Suitable texts for this include [Mar02] and [Poi00].

Definition 2.1. A **signature** L consists of the following disjoint parts:

(i) a collection of objects $s, t, \dots$ called **sorts**;

(ii) a collection of **function symbols** $f, g, \dots$, each accompanied with a finite natural number n, its **arity**, and a sequence of $n+1$ sorts, its **sorting**; these are suggestively denoted $f : s_1 \times \cdots \times s_n \to r$;

(iii) a collection of **relation symbols** $R, S, \dots$, also with an arity n and a sequence of n sorts, its sorting; suggestively, we write $R \subseteq s_1 \times \cdots \times s_n$.

We will take for granted terminology such as "the sort of the ith place of R is s_i"; as a special case, the sort r above may be referred to as the "sort of the value of f". The special cases of 0-ary (*nullary*) functions and relation symbols are referred to as **constants** and **propositional symbols**, respectively.

Before defining what formulae are considered valid in the logical systems ("logics") under discussion, we need some auxiliary symbols. These are as follows:

The following special symbols: $=$ ("identity symbol"); $\neg$ ("negation"); $\to$ ("implication"); $\bigvee$ ("set-indexed disjunction"); $\bigwedge$ ("set-indexed conjunction"), $\exists$ ("existential quantifier"), $\forall$ ("universal quantifier").
Furthermore, for each sort s of the signature, two countably infinite, disjoint collections of **"free" and "bound" variables of sort** s. Both of these are usually denoted $x_s, y_s, \dots$ where the subscript is occasionally dropped.

Beside the above symbols, we use parentheses and commas to ensure that formulae have a unique interpretation. The distinction between free and bound variables is usually dropped in notation; its primary purpose is to make sure that substitution of "terms" for variables is well-behaved on the formal level. Throughout the rest of this section, let us consider a fixed signature L.

Definition 2.2. The **terms** of L are precisely the expressions that can be recursively constructed using the following two rules:

(i) Any free variable x_s of a sort s is a term of sort s;

(ii) For an n-ary function symbol $f : s_1 \times \cdots \times s_n \to r$ and terms t_i of sort s_i, $f(t_1, \dots, t_n)$ is a term of sort r.

Thus terms, like variables, come with a specified sort.

Using terms, we can now define the formulae that comprise the logic for L.

Definition 2.3. The **formulae** of **infinitary first-order logic** $L_{\infty\omega}$ for L are precisely the expressions that can be recursively constructed using the following rules:

(i) For two terms t_1, t_2 of the same sort, $t_1 = t_2$ is a formula;

(ii) For a relation symbol $R \subseteq s_1 \times \cdots \times s_n$ and terms t_i of sort s_i, $R(t_1, \dots, t_n)$ is formula;

(iii) For a formula ϕ, $\neg\phi$ is a formula;

(iv) For formulas ϕ and ψ, $\phi \to \psi$ is a formula.

The rest of the rules need a notion of **occurrence** of a variable. We deem it more insightful to illustrate by example rather than attempt a formal definition: In "$f(x_s, y_r) = x_s$", the variable x_s **occurs** twice, and the variable y_r occurs once; no other variables occur.
To **substitute** a variable x_s with a term t_s (of the same sort!) in a formula ϕ is to replace all occurrences of x_s in ϕ with t_s; the result (which is again a formula) is written $\phi(t_s/x_s)$.

Now, the rest of the rules are as follows:

(v) Let Φ be a *set* of formulas such that the collection of those free variables which occur in at least one $\phi \in \Phi$ is finite. Then $\bigvee \Phi$ is a formula;

(vi) For Φ as above, $\bigwedge \Phi$ is also a formula;

(vii) For ϕ a formula, and x_s, w_s respectively a free and a bound variable of the same sort s, $\exists w_s \phi(w_s/x_s)$ is a formula;

(viii) For ϕ, x_s and w_s as above, $\forall w_s \phi(w_s/x_s)$ is a formula.

We introduce the abbreviations $\top = \bigwedge \varnothing$ and $\bot = \bigvee \varnothing$. To cater for possible ambiguities in reading the resulting formulae (for example, $\neg\phi \to \psi$) parentheses are used as appropriate (resulting in either $(\neg\phi) \to \psi$ or $\neg(\phi \to \psi)$ in our example).

Of particular importance will be the so-called **geometric fragment** (called the **coherent fragment** in [MR77]) $L^g_{\infty\omega}$. It uses only (i), (ii), (v), (vi) and (vii) of the above rules; furthermore, (vi) is restricted to *finite* Φ only. Thus **geometric logic** permits only set-wise disjunction, finite conjunction and existential quantification.

For κ an infinite regular cardinal, we have the fragment $L_{\kappa\omega}$, which restricts (v) and (vi) to sets of cardinality (strictly) less than κ. The **κ-geometric fragment** $L^g_{\kappa\omega}$ is, as one would expect, given by $L^g_{\infty\omega} \cap L_{\kappa\omega}$.

In this definition of formula, we see the rationale behind the terminology "bound" and "free" variable: bound variables are within the *scope* of a quantifier, while free variables are "free of context".

2.2 Structures in categories

We will now define how the signature L can be interpreted in a given category $\mathbf{C}$ with products.

Definition 2.4. Let $\mathbf{C}$ be a category with (binary) products. A **$\mathbf{C}$-valued L-structure** M is a mapping that assigns:

(i) to every sort s of L an object Ms of $\mathbf{C}$;

(ii) to every function symbol $f : s_1 \times \cdots \times s_n \to r$ a morphism $Mf : Ms_1 \times \cdots \times Ms_n \to Mr$ in $\mathbf{C}$;

(iii) to every relation symbol $R \subseteq s_1 \times \cdots \times s_n$ a subobject $MR \hookrightarrow Ms_1 \times \cdots \times Ms_n$ in $\mathbf{C}$.

As a convenience shorthand, we sometimes write $M : L \to \mathbf{C}$ to indicate that M is a $\mathbf{C}$-valued L-structure.

Analogous to standard model theory, the aim is to interpret formulae of $L_{\infty\omega}$ in $\mathbf{C}$ (more formally, in M). Following [MR77, §2.3], we will use Mx for Ms when x is a variable of sort s; $M(\vec{x})$ is defined as $Ms_1 \times \cdots \times Ms_n$ if $\vec{x} = (x_1, \ldots, x_n)$ and x_i has sort s_i. Finally, $t(\vec{x})$ signifies that all (free) variables occurring in t are among $\vec{x}$. The interpretation of terms is such that $M_{\vec{x}}(t)$ is a morphism $M(\vec{x}) \to Ms$.

Definition 2.5. The definition of $M_{\vec{x}}(t)$ proceeds by induction on the complexity of t. Let $t(\vec{x})$ be a term of L of sort s.

(i) If $t = x_i$, then $M_{\vec{x}}(t)$ is the projection $\pi_i : M(\vec{x}) \to M(x_i)$;

(ii) If $t = f(t_1, \ldots, t_m)$ and $M_{\vec{x}}(t_i)$ are known for $i = 1, \ldots, m$, then $M_{\vec{x}}(t)$ is defined as in the following diagram:

$$\begin{tikzcd}
& \prod_i M(s_i) \arrow[r, "Mf"] \arrow[dl, "\pi_i"'] & Ms \\
M(s_i) & & \\
& M(\vec{x}) \arrow[ul, "M_{\vec{x}}(t_i)"] \arrow[uu, "\langle M_{\vec{x}}(t_i)\rangle_i"] \arrow[uur, "M_{\vec{x}}(t)"'] &
\end{tikzcd}$$

Finally, we can turn to the interpretation of $L_{\infty\omega}$ in M. The interpretation of a formula $\phi(\vec{x})$ (using the same convention for $\vec{x}$ as for terms) will be a subobject of $M(\vec{x})$. In order to state the definition of interpretation, there are a number of operations on subobjects that need to be introduced first.

2.3 Interlude: Operations on subobjects

For C an object of a category $\mathbf{C}$, we write $\mathrm{Sub}(C)$ for the set of subobjects of C with the ordering as stated in Definition 1.8. Since $\mathrm{Sub}(C)$ is an ordered set, we have the notions of set-wise **supremum** $\bigvee \mathcal{S}$ and set-wise **infimum** $\bigwedge \mathcal{S}$ for a subset $\mathcal{S}$ of $\mathrm{Sub}(C)$. This notation for suprema and infima attaches a second meaning to the $\bigvee$ and $\bigwedge$ symbols (since they are also used in formulae) in a visually appealing way; it will be clear which is meant, either from the context or by an explicit remark.

Viewing the poset $\mathrm{Sub}(C)$ as a category, we can consider exponentials in it. In this context, one writes X^Y as $Y \to X$, and it is called **Heyting implication**. The terminology "implication" is clear when we expand the product-exponential adjunction in terms of the poset, viz:

$$Z \leq Y \to X \text{ iff } Z \wedge Y \leq X$$

and consider $\leq$ to mean "is logically stronger than". This meaning of $\leq$ will be used again in the description of a categorical model theory, see Definition 2.18 below.

The operations $\to$ and (binary) $\wedge, \vee$ have been studied in great depth in the past, using the notion of a *Heyting algebra*. We give a categorical definition of this concept.

Definition 2.6. A **Heyting algebra** (often abbreviated HA) is a Cartesian closed poset category $\mathbf{H}$ (i.e., $\mathbf{H}$ has finite products and exponentials) that also has finite coproducts.

The concept of a Heyting algebra can be strengthened by adding another operation, $\neg$, called **complementation**. As usual, we write 0 for the initial object (the empty coproduct) and 1 for the terminal object (the empty product).

Definition 2.7. A **Boolean algebra** (BA) is a HA $\mathbf{B}$ endowed with an operation $\neg : \mathbf{B} \to \mathbf{B}$ such that for all objects B, $B \wedge \neg B = 0$ and $B \vee \neg B = 1$.

It is well-known that Boolean algebras provide a sound and complete interpretation for classical propositional (or zeroth order) logic. One can define a generalization of the $\neg$ operation to any HA $\mathbf{H}$, by defining $\neg X := X \to 0$. This corresponds to the largest element of $\mathbf{H}$ such that $X \wedge \neg X = 0$; $\neg X$ is then called the **Heyting complement** of X. Naturally, the "Boolean" complement (if it exists) coincides with the Heyting complement.

The operations on a Heyting algebra (including Heyting complement) can be used to interpret propositional logic. However, because the identity $X \vee \neg X = 1$ need not hold for Heyting $\neg$, this interpretation is complete, but *not sound* for classical propositional logic. It turns out that Heyting algebras form an sound and complete algebraic carrier for so-called **intuitionistic logic**. This logic can be obtained from most of the usual definitions of classical propositional logic by omitting the "Law of Excluded Middle" (stating that $P \vee \neg P$ is always true) or the "Double Negation Elimination" (stating that $\neg\neg P$ entails P) axioms.

We will see in Section 3.1 that for topoi, all $\mathrm{Sub}(C)$ are Heyting algebras, and this gives rise to the statement that "the internal logic of a topos is intuitionistic".

Apart from the operations on HAs and BAs, there are two more operations on subobjects that need to be introduced. These will facilitate the interpretation of formulae involving quantifiers.

For $f : C \to D$ a morphism, there is a "pullback" functor $f^* : \mathrm{Sub}(D) \to \mathrm{Sub}(C)$ defined in the evident way, using that pullbacks preserve monos. The **image** $\exists_f$ and **dual image** $\forall_f$ will turn out to be adjoints of f^*: $\exists_f \dashv f^* \dashv \forall_f$.

Before analyzing these operations in more detail, let us see what happens in **Set** (where the ordering $\leq$ of subobjects is the subset relation $\subseteq$). There we see that f^* is the **inverse image operation** f^{-1}, and we have:

$$X \subseteq f^{-1}(Y) \text{ iff } \{d \in D : \exists c \in X : f(c) = d\} \subseteq Y$$
$$f^{-1}(Y) \subseteq X \text{ iff } Y \subseteq \{d \in D : \forall c \in C : f(c) = d \text{ implies } c \in X\}$$

Now $\exists_f(X)$ will turn out to be the first set on the right-hand side, while $\forall_f(X)$ is the second; the quantifiers occurring in the set specifications justify the notation. With this concrete example in mind, we state:

Definition 2.8. Let $f : C \to D$ be a morphism, and let X be a subobject of C.

The **image** of X under f is the smallest subobject $\exists_f(X)$ of D making
$$\begin{tikzcd} X \arrow[r, hook] \arrow[d] & C \arrow[d, "f"] \\ \exists_f(X) \arrow[r, hook] & D \end{tikzcd}$$
commute.

The **image of** f, $\mathrm{Im}\, f$, is the image $\exists_f(C)$ of C under f.
The **dual image** of X under f, denoted $\forall_f(X)$, is the largest subobject Y of D making
$$\begin{tikzcd} f^*(Y) \arrow[r, hook] \arrow[d] & X \arrow[r, hook] & C \arrow[d, "f"] \\ Y \arrow[rr, hook] & & D \end{tikzcd}$$
commute; i.e. the largest Y such that $f^*(Y) \leq X$ in $\mathrm{Sub}(C)$.

These definitions for $\exists_f$ and $\forall_f$ make them satisfy the asserted adjunctions to the pullback functor f^*:

Proposition 2.9. *With the notation as above, $\exists_f$ and $\forall_f$ are functors $\mathrm{Sub}(C) \to \mathrm{Sub}(D)$, and there are adjunctions $\exists_f \dashv f^* \dashv \forall_f$.*

Proof Working on poset categories, $\exists_f$ and $\forall_f$ are functors precisely when they are order-preserving. This is completely trivial from the definitions. To prove the adjunctions, it suffices to prove the equivalences $X \leq f^*Y \iff \exists_f X \leq Y$ and $f^*Y \leq X \iff Y \leq \forall_f X$. This is because each hom-set consists of zero or one element, and there is a canonical bijection between any two sets of cardinality zero or one.

First of all, we have $f^*Y \hookrightarrow C \xrightarrow{f} D$ factors through Y by definition of f^*Y; hence so does $X \hookrightarrow f^*Y \hookrightarrow C \xrightarrow{f} D$. By definition of $\exists_f X$ we conclude $X \leq f^*Y$ implies $\exists_f X \leq Y$. Conversely, we have a commutative diagram:

$$\begin{tikzcd} X \arrow[r, dashed] \arrow[rr, hook, bend left] \arrow[d] & f^*Y \arrow[r, hook] \arrow[d] & C \arrow[d, "f"] \\ \exists_f X \arrow[r, hook] & Y \arrow[r, hook] & D \end{tikzcd}$$

where the dashed arrow is filled because f^*Y is a pullback. That is, $X \leq f^*Y$. This establishes the adjunction $\exists_f \dashv f^*$.

Now for the adjunction $f^* \dashv \forall_f$, it is to be shown that $f^*Y \preceq X \iff Y \preceq \forall_f X$. The direct implication is immediate from the definition of $\forall_f X$. The reverse implication follows directly from the fact that f^* is order-preserving. $\square$

It is here that we first see merit for discussing logic in categories. Classically, it was a big problem to extend the Boolean algebras, which provide an algebraic context for propositional logic, to first-order logic, precisely because of the quantifiers. By applying the above adjunctions to projection mappings, we will see that this sought "algebraization" of first-order logic can be situated in category theory. This connection was first discovered and published about by Lawvere in the 1960s, and in fact was one of the instigators of **topos theory** as a research field.

2.4 Interpretation of formulae

We are now prepared to give the interpretation of formulae in suitably rich categories. The symbols and conventions from §2.2 are adopted.

Definition 2.10. Again, $M_{\vec{x}}(\phi)$ will be defined by induction on ϕ. The interpretation is considered to be defined only when all the categorical constructions in its specification can be carried out in the underlying category $\mathbf{C}$.

(i) If ϕ is the formula $t_1 = t_2$ with t_1 and t_2 terms of the same sort s, then $M_{\vec{x}}(\phi)$ is (the subobject corresponding to) the following equalizer:

$$\begin{tikzcd} M_{\vec{x}}(t_1 = t_2) \arrow[r, hook] & M(\vec{x}) \arrow[r, shift left, "M_{\vec{x}}(t_1)"] \arrow[r, shift right, "M_{\vec{x}}(t_2)"'] & Ms \end{tikzcd}$$

(ii) If ϕ is the formula $R(t_1, \ldots, t_m)$, then $M_{\vec{x}}(\phi)$ is defined using the following pullback:

$$\begin{tikzcd} M(R) \arrow[r, hook] & \prod_i M(s_i) \\ M_{\vec{x}}(\phi) \arrow[r, hook] \arrow[u] & M(\vec{x}) \arrow[u, "\langle M_{\vec{x}}(t_i) \rangle_i"'] \end{tikzcd}$$

(iii) If ϕ is $\neg\psi$, then $M_{\vec{x}}(\neg\psi) = \neg M_{\vec{x}}(\psi)$ where the $\neg$ on the right denotes Heyting complement in $\mathrm{Sub}(M(\vec{x}))$.

(iv) If ϕ is $\psi \to \chi$, then $M_{\vec{x}}(\psi \to \chi) = M_{\vec{x}}(\psi) \to M_{\vec{x}}(\chi)$ where the $\to$ on the right is Heyting implication.

(v) If ϕ is $\bigvee \Psi$, then $M_{\vec{x}}(\phi) = \bigvee\limits_{\psi \in \Psi} M_{\vec{x}}(\psi)$ where the $\bigvee$ on the right denotes the (set-indexed) supremum operation in $\mathrm{Sub}(M(\vec{x}))$.

(vi) Likewise, if ϕ is $\bigwedge \Psi$, then $M_{\vec{x}}(\phi) = \bigwedge\limits_{\psi \in \Psi} M_{\vec{x}}(\psi)$ where $\bigwedge$ denotes the infimum operation in $\mathrm{Sub}(M(\vec{x}))$.

(vii) If ϕ is $\exists w \psi(w/x)$ then, assuming x is not among $\vec{x}$, $M_{\vec{x}}(\phi) = \exists_\pi(M_{\vec{x},x}(\psi))$ where $\pi : M(\vec{x}) \times Mx \to M(\vec{x})$ is the projection.

(viii) If ϕ is $\forall w \psi(w/x)$ then $M_{\vec{x}}(\phi) = \forall_\pi(M_{\vec{x},x}(\psi))$ where again $\pi : M(\vec{x}) \times Mx \to M(\vec{x})$ is the projection.

All of this lengthy definition serves to define a notion of satisfaction of so-called *sequents*.

Definition 2.11. A **sequent** of $L_{\infty\omega}$ is an expression of the form $\Phi \Rightarrow \Psi$ with Φ and Ψ finite sets of $L_{\infty\omega}$-formulae. If F is a fragment of $L_{\infty\omega}$, an F**-sequent** is one where Φ and Ψ comprise only F-formulae. An $L^g_{\infty\omega}$-sequent is also called a **geometric sequent**.

Definition 2.12. Given a sequent $\Phi \Rightarrow \Psi$, a structure M **satisfies** $\Phi \Rightarrow \Psi$, denoted $M \models \Phi \Rightarrow \Psi$, if $\bigwedge\limits_{\phi \in \Phi} M_{\vec{x}}(\phi) \le \bigvee\limits_{\psi \in \Psi} M_{\vec{x}}(\psi)$ in $\mathrm{Sub}(M(\vec{x}))$, where $\vec{x}$ are the free variables occurring in at least one formula of Φ or Ψ.

If Φ is empty, we will also write $M \models \Psi$ in place of $M \models \varnothing \Rightarrow \Psi$; so $M \models \Psi$ iff $\bigvee\limits_{\psi \in \Psi} M_{\vec{x}}(\psi)$ is the maximal subobject of $M(\vec{x})$.

To get a feeling for the techniques and methods used in working with the above definitions, we discuss some specific examples, phrased in the so-called **canonical language**.

2.5 The canonical language

In this section we will use the preceding definitions to express properties of a category $\mathbf{C}$ in terms of the new language. This is made possible by a suitable choice of signature that has a natural interpretation in $\mathbf{C}$.

Definition 2.13. The **canonical language** $L_{\mathbf{C}}$ of a category $\mathbf{C}$ has as sorts $\mathbf{C}_0$, the objects of $\mathbf{C}$, and for each morphism $f : C \to D$, a corresponding unary function symbol $f : C \to D$. There are no relation symbols.

Intuitively, we can imagine this canonical language $L_{\mathbf{C}}$ as specifying the **underlying graph** of $\mathbf{C}$, that is, forgetting the composition operation of $\mathbf{C}$ (effectively retaining only domain-codomain relationships). Obviously, $\mathbf{C}$ qualifies as a $L_{\mathbf{C}}$-structure, by the "identity interpretation".

As it turns out, we can express many properties of $\mathbf{C}$ using $L_{\mathbf{C}}$, provided $\mathbf{C}$ has enough structure to support the relevant portions of Definition 2.10.

As a first and in some sense most basic example, we consider the "commutative triangle": for $h : X \to Z$, $g : Y \to Z$ and $f : X \to Y$, we claim that, for a variable x of sort X:

$$h = g \circ f \text{ in } \mathbf{C} \quad \text{iff} \quad \mathbf{C} \models hx = gfx$$

According to Definition 2.12, the right-hand side amounts to the identity of subobjects $\mathbf{C}_x(hx = gfx) = \mathrm{id}_X$. Definition 2.10(i) tells us that $\mathbf{C}_x(hx = gfx)$ is the equalizer of $\mathbf{C}_x(hx)$ and $\mathbf{C}_x(gfx)$, which by Definition 2.5 are seen to correspond precisely to $h : X \to Z$ and $g \circ f : X \to Z$.

In conclusion, we need to show that – provided $\mathbf{C}$ has the necessary equalizer – $h = g \circ f$ is equivalent to id_X equalizing h and $g \circ f$. But this is trivially verified.

As a further example, we consider the *graph* of a morphism:

Definition 2.14. Let $\mathbf{C}$ be a category with finite limits. The **graph** of a morphism $f : X \to Y$ is the subobject $\mathbf{C}_{x,y}(fx = y) \hookrightarrow X \times Y$.

Let us expand the definition of $\mathbf{C}_{x,y}(fx = y)$ to get a better feeling for what is going on here. We have the interpretation of terms given by $\mathbf{C}_{x,y}(fx) = X \times Y \xrightarrow{\pi_1} X \xrightarrow{f} Y$ and $\mathbf{C}_{x,y}(y) = X \to Y \overset{\pi_2}{Y}$. An $e : E \to X \times Y$ equalizing these satisfies $\pi_2 e = f\pi_1 e$; thus e factors as $E \xrightarrow{\pi_1 e} X \xrightarrow{\langle \mathrm{id}_X, f\rangle} X \times Y$; by the universal property of the product, it does so uniquely. Hence $\langle \mathrm{id}_X, f\rangle : X \to X \times Y$ is the sought equalizer $\mathbf{C}_{x,y}(fx = y) \hookrightarrow X \times Y$. Thus in **Set**, the notion of graph coincides with what we expect it to be.

Another small and useful result is the following:

Lemma 2.15. *Let $f : X \rightarrowtail Y$ be a monomorphism. Then the subobject of Y associated to it is given by $\mathbf{C}_y(\exists x : fx = y)$.*

Proof By definition 2.10(vii), $\mathbf{C}_y(\exists x : fx = y)$ is $\exists_\pi(\mathbf{C}_{x,y}(fx = y))$ with $\pi : X \times Y \to Y$ the projection. By the characterization of $\mathbf{C}_{x,y}(fx = y)$ as $\langle \mathrm{id}_X, f\rangle$ we see that $\pi_Y \langle \mathrm{id}_X, Y\rangle = f$. Thus f is a subobject through which it factors, and for any other $g : X' \hookrightarrow Y$ it factors through, there is a morphism $h : X \to X'$ with $f = gh$, so that $f \leq g$ as subobjects. The result follows by definition of $\exists_\pi(\mathbf{C}_{x,y}(fx = y))$. $\square$

This result allows us to express arbitrary subobjects $X \hookrightarrow Y$ in the canonical language. We could also have achieved this by introducing relation symbols for them: e.g. a unary symbol $S \subseteq X$ for a subobject S of X, and a binary symbol $S' \subseteq X \times Y$ for a subobject S' of $X \times Y$ (for example, the graph of a morphism). In what follows, it will often be more convenient to refer to such relation symbols rather than to express them using Lemma 2.15. This does not affect the results in any way.

Now we have the following simple proposition:

Proposition 2.16. *Let $S \subseteq X \times Y$ be the graph of some morphism $f : X \to Y$ in $\mathbf{C}$. Then $\mathbf{C}$ models the following sequents:*

$$\begin{aligned} Sxy \wedge Sxy' &\Rightarrow y = y' \\ &\Rightarrow \exists y\, Sxy \end{aligned}$$

Proof Let us begin with the second sequent. By expanding the definitions, we see $\mathbf{C}_x(\exists y\, Sxy)$ is $\exists_{\pi_X}(S)$. But the discussion following the definition of a graph (Definition 2.14) established that $S \rightarrowtail X \times Y \xrightarrow{\pi_X} X = \mathrm{id}_X$. Thus $\mathbf{C}_x(\exists y\, Sxy)$ exists and is the maximal subobject of X. Hence the second sequent is valid.

The whole situation regarding the first sequent can be captured in the following diagram (where Y_1 corresponds to the variable y, and Y_2 to y'):

$$\begin{tikzcd}
\mathbf{C}_{xyy'}(Sxy \wedge Sxy') \arrow[r, hook] \arrow[d, hook] & \mathbf{C}_{xyy'}(Sxy') \arrow[r] \arrow[d, hook] & S = \mathbf{C}_{xy'}(Sxy') \arrow[d, hook] \\
\mathbf{C}_{xyy'}(Sxy) \arrow[r, hook] \arrow[d] & X \times Y_1 \times Y_2 \arrow[r, "\langle \pi_X, \pi_{Y_2} \rangle"] \arrow[d, "\langle \pi_X, \pi_{Y_1} \rangle"'] & X \times Y_2 \arrow[d, shift right, "\pi_{Y_2}"'] \arrow[d, shift left, "f\,\pi_X"] \\
\mathbf{C}_{xy}(Sxy) = S \arrow[r, hook] & X \times Y_1 \arrow[r, shift left, "\pi_{Y_1}"] \arrow[r, shift right, "f\,\pi_X"'] & Y_1 = Y_2
\end{tikzcd}$$

where a word of caution is in order: the bottom-right square does *not* commute. The other three squares are pullbacks, obtained from the interpretations of $\mathbf{C}_{xyy'}(Sxy), \mathbf{C}_{xyy'}(Sxy')$ and $\mathbf{C}_{xyy'}(Sxy \wedge Sxy')$. The bottom row and rightmost column are equalizers due to S being the graph of f.

Now we need to prove that $\mathbf{C}_{xyy'}(Sxy \wedge Sxy') \leq \mathbf{C}_{xyy'}(y = y')$; that is, $i : \mathbf{C}_{xyy'}(Sxy \wedge Sxy') \hookrightarrow X \times Y_1 \times Y_2$ should equalize π_{Y_1} and π_{Y_2}. By chasing the borders of the diagram, we find that $\pi_{Y_1} i = f \pi_X i = \pi_{Y_2} i$, as desired. $\square$

Now what is striking is that the converse also holds: any subobject S for which the two sequents are true, is the graph of a morphism; this morphism is even unique.

Proposition 2.17. *Let $S \subseteq X \times Y$ be a subobject such that $\mathbf{C}$ models the following sequents:*

$$\begin{aligned} Sxy \wedge Sxy' &\Rightarrow y = y' \\ &\Rightarrow \exists y\, Sxy \end{aligned}$$

Then there is a unique morphism $f : X \to Y$ of which S is the graph.

Proof One first shows that $S \xhookrightarrow{\iota_S} X \times Y \xrightarrow{\pi_X} X$ is an isomorphism. Because of the second sequent, it suffices to show that $\pi_X \iota_S$ is a mono, which is the content of [MR77], Proposition 2.4.3. Now it is not hard to show that S is the graph of $\pi_Y \iota_S (\pi_X \iota_S)^{-1} : X \to Y$. $\square$

2.6 Models

Again, let us consider a fixed signature L. It is not practical to discuss L-structures and the sequents they satisfy in any nontrivial depth when they are viewed in isolation. Our category-theoretic setting calls for suitable notions of morphism and associated categories. To this end, we extend the satisfaction relation $\models$ between structures and sequents of Definition 2.12 by introducing the following terminology:

Definition 2.18. A **theory** is a class of sequents T. When all the sequents in T are part of some fragment F (e.g., the geometric fragment $L^g_{\infty\omega}$), T is called an F**-theory**; in the special case $F = L^g_{\infty\omega}$ we speak of a **geometric theory**.

Let $\mathbf{C}$ be a category. A $\mathbf{C}$-valued L-structure M is called a **model** of an F-theory T, denoted $M : T \to \mathbf{C}$, if $M \models \Phi \Rightarrow \Psi$ for all sequents $\Phi \Rightarrow \Psi$ in T. In particular, it is implicit that the underlying category $\mathbf{C}$ of M has all the required structure for interpreting the fragment F.

This notion of model extends the one familiar from classical model theory, which corresponds to an $L_{\omega\omega}$-structure in $\mathbf{Set}$. Next, we introduce some categories centered around these models. First, we define a notion of homomorphism between L-structures in the same category $\mathbf{C}$.

Definition 2.19. Let $M, M' : L \to \mathbf{C}$ be $\mathbf{C}$-valued L-structures. A **homomorphism of L-structures** $H : M \to M'$ is a collection of morphism $Hs : Ms \to M's$, one for each sort s of L, such that H is compatible with the interpretation of function and relation symbols, i.e.:

$$\begin{tikzcd} MR \arrow[r, hook] \arrow[d, dashed] & Ms_1 \times \cdots \times Ms_n \arrow[d, "Hs_1 \times \cdots \times Hs_n"] \\ M'R \arrow[r, hook] & M's_1 \times \cdots \times M's_n \end{tikzcd} \qquad \begin{tikzcd} Ms_1 \times \cdots \times Ms_n \arrow[r, "Mf"] \arrow[d, "Hs_1 \times \cdots \times Hs_n"'] & Mt \arrow[d, "Ht"] \\ M's_1 \times \cdots \times M's_n \arrow[r, "M'f"'] & M't \end{tikzcd}$$

commute for every relation symbol $R \subseteq s_1 \times \cdots \times s_n$ and function symbol $f : s_1 \times \cdots \times s_n \to t$ of L.

With the obvious composition and identity morphisms, this yields a category $\mathrm{Str}(L, \mathbf{C})$ of L-structures in $\mathbf{C}$. We will concern ourselves with full subcategories of this category:

Definition 2.20. The category $\underline{\mathrm{Mod}}(T, \mathbf{C})$ of **T-models in $\mathbf{C}$** is the full subcategory of $\mathrm{Str}(L, \mathbf{C})$ whose objects are T-models.

Now while any functor $F : \mathbf{C} \to \mathbf{D}$ that preserves finite limits induces a functor $\mathrm{Str}(L, F) : \mathrm{Str}(L, \mathbf{C}) \to \mathrm{Str}(L, \mathbf{D})$, we cannot conclude that this functor restricts to one $\underline{\mathrm{Mod}}(T, \mathbf{C}) \to \underline{\mathrm{Mod}}(T, \mathbf{D})$: it is not guaranteed that F will preserve validity of sequents. We describe the conditions required on F for the various logic fragments F that we will consider in Section 4.2 and following.

Definition 2.21. Let κ be an infinite regular cardinal. A functor $F : \mathbf{C} \to \mathbf{D}$ is called **κ-geometrical** if it preserves finite limits, images of subobjects, and suprema of sets of subobjects of cardinality less than κ. F is called **geometrical** if it preserves finite limits, images, and *all* small suprema.

We prefer to use "κ-geometrical" in place of the term "κ-logical", which is used in e.g. [MR77], because these functors will turn out to preserve only *geometric* logic. As a first result in this direction, let us prove that the κ-geometrical functors preserve models of $L^g_{\kappa\omega}$-theories:

Proposition 2.22. *Let $F : \mathbf{C} \to \mathbf{D}$ be a functor. Let κ be an infinite regular cardinal, or ∞, and let T be an $L^g_{\kappa\omega}$-theory. Then the functor $\mathrm{Str}(L, F) : \mathrm{Str}(L, \mathbf{C}) \to \mathrm{Str}(L, \mathbf{D})$ restricts to a functor $\underline{\mathrm{Mod}}(T, F) : \underline{\mathrm{Mod}}(T, \mathbf{C}) \to \underline{\mathrm{Mod}}(T, \mathbf{D})$ of T-models.*

Proof We will start by showing that $(FM)_{\vec{x}}(\phi) = F(M_{\vec{x}}(\phi))$ for κ-geometric ϕ, where FM is short for $\mathrm{Str}(L, F)(M)$.

First of all, for any sort s, $(FM)s = F(Ms)$. Since F preserves finite limits, it follows that $(FM)(\vec{x}) = \prod_i F(Mx_i)$. It is equally immediate that $(FM)f = F(Mf)$ and $(FM)R = F(MR)$ for function symbols f and relation symbols R. This establishes the assertion above that F induces a functor $\mathrm{Str}(L, F) : \mathrm{Str}(L, \mathbf{C}) \to \mathrm{Str}(L, \mathbf{D})$, and gives an explicit (albeit rather tautologous) description of $\mathrm{Str}(L, F)$.

It can be verified by inspection that all of the formation rules for κ-geometric logic involve only finite limits, colimits of cardinality at most κ, and images of subobjects; these are all preserved by F, which establishes the desired equality $(FM)_{\vec{x}}(\phi) = F(M_{\vec{x}}(\phi))$ for every κ-geometric ϕ.

Finally, since F preserves finite limits, it is order-preserving on subobject posets. Hence, as F preserves finite limits and colimits, it follows that $\bigwedge_{\phi \in \Phi} M_{\vec{x}}(\phi) \leq \bigvee_{\psi \in \Psi} M_{\vec{x}}(\psi)$ implies $\bigwedge_{\phi \in \Phi} FM_{\vec{x}}(\phi) \leq \bigvee_{\psi \in \Psi} FM_{\vec{x}}(\psi)$. Thus if $M \models \Phi \Rightarrow \Psi$, also $FM \models \Phi \Rightarrow \Psi$; this concludes the proof. □

We will return to models in Chapter 4 and, to a lesser extent, in Chapter 5. But first, we will connect the material of this chapter with that of Chapter 1.

3 Grothendieck topoi as structures

In this section, we will connect the two main topics we have discussed so far, namely Grothendieck topoi and categorical logic. To this end, we will first introduce and justify the concept of an ∞-pretopos as a carrier for geometric logic.

A preliminary remark about the terminology: "∞-pretopos" nowadays has an unfortunate similarity to the more recent concept of "∞-topos" as introduced by Joyal and Lurie; this should be considered a quirk in mathematical nomenclature rather than an indication of an intimate relationship between the two.

3.1 ∞-pretopoi

Definition 3.1. An **∞-pretopos** is a category having:

(i) finite limits;

(ii) stable suprema of *sets* of subobjects;

(iii) stable images of subobjects;

(iv) quotients of equivalence relations;

(v) disjoint sums of *sets* of objects.

The notions occurring in clauses (iv) and (v) will be introduced below. A category with (i)-(iii) is called **∞-geometrical**, while for a regular infinite cardinal κ, a category is **κ-geometrical** if it satisfies (ii) for sets of cardinality smaller than κ.

The notion of an ∞-pretopos was designed to describe those categories that admit a suitably well-behaved interpretation of geometric logic. The first three conditions are more or less what one would expect, because they are required to interpret certain parts of the geometric logic. Namely, (i) is required to interpret atomic formulae and finite conjunctions; (ii) is required to interpret disjunctions; and (iii) is required to interpret existential quantification.

One now probably wonders why the final two points, and the "stability" in points (ii) and (iii) are included. To begin to understand this, we will first have to define what all this terminology means in the first place.

Definition 3.2. Let $\{X_i : i \in I\}$ be a set of subobjects of some C. Then the supremum $\bigvee_i X_i$ is said to be **stable (under pullbacks)**, or **universal**, if for each $h : C' \to C$, we have the following identity of subobjects of C':

$$\bigvee_{i \in I} h^* X_i = h^* \left(\bigvee_{i \in I} X_i \right)$$

Similarly, for a subobject X of C, the image $\exists_f(X)$ of X is called **stable** if for each pullback diagram:

$$\begin{tikzcd} C' \arrow[r, "h"] \arrow[d, "f'"'] & C \arrow[d, "f"] \\ D' \arrow[r, "h'"'] & D \end{tikzcd}$$

we have the equality of subobjects $h'^*(\exists_f(X)) = \exists_{f'}(h^* X)$. This *stability of images* is often referred to in the literature as the **Beck-Chevalley condition**.

The reason for considering these stability conditions is apparent from the following results:

Proposition 3.3. *Let $X, X_i\,(i \in I)$ and Y be subobjects of C. Let $f : C \to D$ be an arbitrary morphism.*

(i) *Suppose that the supremum $\bigvee_i X_i$ is stable. Then $Y \wedge \bigvee_i X_i = \bigvee_i (Y \wedge X_i)$.*

(ii) *Suppose that the image $\exists_f(X)$ is stable. Then $\exists_f(f^* Y \wedge X) = Y \wedge \exists_f(X)$.*

Proof (i) $Y \wedge \bigvee_i X_i$ is defined as the pullback of $\bigvee_i X_i \hookrightarrow C$ along $Y \hookrightarrow C$. Applying the stability condition to this pullback immediately yields the result.
(ii) We employ the condition from Definition 3.2. For this, we use the following pullback diagram associated to the construction of f^*Y. Concretely, this means that in the notation of Definition 3.2, $f' : f^*Y \to Y$ is the map induced from f; h is $f^*Y \hookrightarrow C$; and h' is $Y \hookrightarrow D$. Interpreting the pullbacks resulting from the stability condition as meets in $\mathrm{Sub}(C)$, we obtain the desired formula. $\square$

Corollary 3.4. *In an ∞-pretopos the following sequents are true for all formulae $\psi, \phi_i\ (i \in I), \chi(\vec{x})$ and $\eta(\vec{x}, y)$ (here $\Leftrightarrow$ is used as an obvious abbreviation):*

$$\psi \wedge \bigvee_{i \in I} \phi_i \Leftrightarrow \bigvee_{i \in I} (\psi \wedge \phi_i) \qquad\qquad \exists y(\chi(\vec{x}) \wedge \eta(\vec{x}, y)) \Leftrightarrow \chi(\vec{x}) \wedge \exists y\, \eta(\vec{x}, y)$$

Definition 3.5. A subobject $R \hookrightarrow C \times C$ is called an **equivalence relation** if $\mathbf{C}$ satisfies the following self-explanatory sequents of its canonical language:

$$\Rightarrow Rcc \qquad\qquad Rcc' \Rightarrow Rc'c \qquad\qquad Rcc' \wedge Rc'c'' \Rightarrow Rcc''$$

A morphism $q : C \to D$ is called a **quotient** of R if the following sequents are satisfied:

$$\Rightarrow \exists c(qc = d) \qquad\qquad Rcc' \Leftrightarrow qc = qc'$$

For objects $\{C_i : i \in I\}$ of $\mathbf{C}$, their **disjoint sum** $\coprod_{i \in I} C_i$ is an object D with injections $\iota_i : C_i \to D$ satisfying the following sequents, for $i, i' \in I$ and $i \neq i'$:

$$f_i c_i = f_i c_i' \Rightarrow c_i = c_i' \qquad\qquad C_i(d) \wedge C_{i'}(d) \Rightarrow \bot \qquad\qquad d = d \Rightarrow \bigvee_{i \in I} C_i(d)$$

What ∞-pretopoi have to do with Grothendieck topoi is expressed in Theorem 3.7 below. For its statement, we need the following definition:

Definition 3.6. Let $\mathbf{C}$ be a category; a set $\mathcal{S} \subseteq \mathrm{ob}\,\mathbf{C}$ of objects of $\mathbf{C}$ is called a **set of generators** if for each $C \in \mathrm{ob}\,\mathbf{C}$, the collection $\{s : S \to C \mid S \in \mathcal{S}\}$ is an **epimorphic family**, i.e., for $f, g : C \to D$, if $fs = gs$ for all such s, then $f = g$.

Theorem 3.7 (Giraud, [MR77, Proposition 3.4.8])**.** *A category is a Grothendieck topos iff it is an ∞-pretopos and has a set of generators.*

Proof We will only show the "only if" direction: that a Grothendieck topos is an ∞-pretopos with a set of generators; the converse is not within the scope of this thesis. So let $\mathbf{Sh}(\mathbf{C}, J)$ be a Grothendieck topos. We remark that $\mathbf{Set}$ is an ∞-pretopos (only the stability of images is not immediate, but still straightforward). Now we first show that $\hat{\mathbf{C}}$ is an ∞-pretopos.

It clearly has finite limits. For the supremum of subobjects, we easily obtain $(\bigvee_i Q_i)C = \bigcup_i Q_i C$ for a set $(Q_i)_{i \in I}$ of subpresheaves of some presheaf P. This concludes the general case by Proposition 1.10; stability of these suprema is immediate. The image of a subobject $i_Q : Q \hookrightarrow P$ under $\eta : P \to P'$ is given by $(\exists_\eta Q)C = \eta_C(QC)$. That this defines a subpresheaf of P' follows from the following computation, for $f : C \to D$:

$$P'f((\exists_\eta Q)D) = P'f(\eta_D(QD)) = \eta_C Pf(QD) = \eta_C QC = (\exists_\eta Q)C$$

That it is indeed an image is obvious. The stability of this image (i.e. the Beck-Chevalley condition) boils down to the equality of two subpresheaves, and hence the equality of the respective sets at each C. Thus we infer stability from the fact that $\mathbf{Set}$ is an ∞-pretopos. It is straightforward but a bit tedious to verify that "R is an equivalence relation on P" in the sense of Definition 3.5 means that R is a subpresheaf of $P \times P$ such that for each C, RC is an equivalence relation on PC. Now define $(P/R)C = PC/RC$, and $q : P \to P/R$ in the natural way; then q is the desired quotient of R. Lastly, the disjoint sum of presheaves is seen to be defined pointwise. In conclusion, $\hat{\mathbf{C}}$ is an ∞-pretopos.

Now let us turn to $\mathbf{Sh}(\mathbf{C}, J)$. Its finite limits, suprema of subobjects, images, quotients, and disjoint sums are all obtained by means of the sheafification functor $\mathbf{a}$ and the corresponding constructions on $\hat{\mathbf{C}}$.

One applies $\mathbf{a}$ to the construction obtained for $\hat{\mathbf{C}}$, and the conclusions follow after suitable applications of the preservation properties (for colimits and finite limits) of $\mathbf{a}$ and Lemma 1.22. The required stability conditions are inherited from their counterparts in $\hat{\mathbf{C}}$.

It remains to prove that $\mathbf{Sh}(\mathbf{C}, J)$ has a set of generators. We claim that $\{\mathbf{ay}_C : C \in \operatorname{ob}\mathbf{C}\}$ is such a set. For suppose that $f, g : F \to G$ are morphisms of sheaves, such that $fh = gh$ for all $h : \mathbf{ay}_C \to F$. It suffices to show that $f = g$ in $\hat{\mathbf{C}}$; to this end, let $k : \mathbf{y}_D \to F$ be arbitrary, with $\mathbf{y}_D$ a representable presheaf. Now k factors through $\mathbf{a}E$ by the universal property of $\mathbf{ay}_D$, so that $fk = gk$. Letting k range over the morphisms that make F a colimit of representable presheaves (see Proposition 1.7), it follows from the uniqueness clause of the universal property of the colimit that $f = g$, as desired. $\square$

For later reference, we give a description of the canonical topology on an ∞-pretopos (and hence, on a Grothendieck topos). First, some terminology.

Definition 3.8. Let $\mathbf{C}$ be a category with finite limits. A collection of morphisms $S_0 = \{f_i : C_i \to C\}$ with common codomain C is called an **effective epimorphic family** if every matching family $x : S \to \mathbf{y}_D$ for a representable presheaf $\mathbf{y}_D$ on the sieve S generated by S_0 has a unique amalgamation $x \in \mathbf{y}_D(C) = \operatorname{Hom}(C, D)$. Equivalently, if for each $D \in \operatorname{ob}\mathbf{C}$, the diagram:

$$\operatorname{Hom}(C, D) \xrightarrow{(f_i)_i} \prod_i \operatorname{Hom}(C_i, D) \underset{f_j^*(f_i)}{\overset{f_i^*(f_j)}{\rightrightarrows}} \prod_{i,j} \operatorname{Hom}(C_i \underset{C}{\times} C_j, D)$$

is an equalizer in $\mathbf{Set}$. (Replacing the $\operatorname{Hom}(-, D) = \mathbf{y}_D$ with a presheaf P yields a diagrammatic condition, both necessary and sufficient, for P to be a sheaf.)

The family S_0 is called **stable effective epimorphic** if for every $f : D \to C$, the collection f^*S_0 of pullbacks is effective epimorphic.

The stable effective epimorphic families are seen to comprise the canonical topology. In ∞-pretopoi, we have the following result, which is very useful for both theoretical purposes and applications.

Proposition 3.9 ([MR77], Proposition 3.4.11)**.** *Let* $\mathbf{E}$ *be an* ∞*-pretopos. Then every epimorphic family* $\{f_i : C_i \to C\}$ *(i.e., any family such that for* $f, g : C \to D$*,* $ff_i = gf_i$ *for all* i *implies* $f = g$*) is stable effective epimorphic.*

3.2 Kripke-Joyal semantics

Next, we discuss a different presentation of the semantics from Chapter 2, called the *Kripke-Joyal semantics*. One of its main virtues is establishing a link between set theory and topoi: the presence of a subobject classifier will allow for attaching in every topos a meaning to a familiar set-theoretic expression like $\{(e, f) \mid \phi(e, f)\}$. Our presentation differs from the one in [MM92, VI.§§5-6] so as to make it clear that it is a restatement of the semantics discussed in the preceding chapter, rather than a new take at the material.

Let $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$ be a Grothendieck topos with subobject classifier Ω. We augment the interpretation of the canonical language $L_{\mathbf{E}}$ of $\mathbf{E}$ from Definition 2.10 by introducing a new relation symbol $\in$ (for "elementhood"):

Definition 3.10. If t is a term of sort S and t' is a term of sort Ω^S, then $t \in t'$ is a formula.

The interpretation $\mathbf{E}_{\vec{x}}(t \in t')$ of $t \in t'$ in $\mathbf{E}$ is given by the subobject classified by $\operatorname{ev}\langle t, t'\rangle$, i.e. that makes the following diagram a pullback:

$$\begin{tikzcd}
\mathbf{E}_{\vec{x}}(t \in t') \arrow[rr] \arrow[d, hook] & & 1 \arrow[d, tail, "\text{true}"] \\
\mathbf{E}(\vec{x}) \arrow[r, "\langle \mathbf{E}_{\vec{x}}(t), \mathbf{E}_{\vec{x}}(t')\rangle"'] & S \times \Omega^S \arrow[r, "\text{ev}"'] & \Omega
\end{tikzcd}$$

Having expanded the language with this extra relation symbol, we now define the central notion of Kripke-Joyal semantics, the *forcing* relation.

Definition 3.11. Let C be an object of $\mathbf{C}$, let $\phi(\vec{e})$ be a formula with free variables $\vec{e}$ and let $\vec{\xi} \in \mathbf{E}(\vec{e})(C)$. Then one says C **forces** $\phi(\vec{\xi})$, denoted:

$$C \Vdash \phi(\vec{\xi})$$

if $\vec{\xi} \in \mathbf{E}_{\vec{e}}(\phi)(C)$. In particular, this interpretation makes sense if $\vec{e}$ is the empty sequence. In that case, $C \Vdash \phi$ if and only if $\mathbf{E} \models \phi$, in the sense of Definition 2.12.

In order to give a full and explicit characterization of the Kripke-Joyal semantics, we need an explicit description of the various operations on subobjects in a Grothendieck topos.

Lemma 3.12. *Let E be a sheaf in a Grothendieck topos $\mathbf{Sh}(\mathbf{C}, J)$, and let A,B, and $(A_i)_{i\in I}$ be subobjects of E.*

(i) *The meet $A \wedge B$ is the sheaf defined by $(A \wedge B)C := AC \cap BC$;*

(ii) *The join $\bigvee_i A_i$ is the sheaf defined by $\xi \in (\bigvee_i A_i)C \iff \{f : D \to C \mid \xi \cdot f \in \bigcup_i A_i D\} \in JC$;*

For a morphism $\eta : E \to E'$, we have the following description of the image operation $\exists_\eta$:

(iii) *The image $\exists_\eta A$ is the sheaf defined by $\xi' \in (\exists_\eta A)C \iff \{f : D \to C \mid \exists a \in AD : \eta_D a = \xi' \cdot f\} \in JC$.*

Finally, in Boolean Grothendieck topoi, we have the following descriptions of Heyting implication, negation, and dual image operations:

(iv) *The Heyting implication $A \to B$ is the sheaf defined by $\xi \in (A \to B)C \iff e \cdot f \in AD$ implies $\xi \cdot f \in BD$ for all $f : D \to C$.*

(v) *The negation $\neg A$ is the sheaf defined by $\xi \in (\neg A)C \iff \xi \cdot f \in AD$ for some $f : D \to C$ implies that $\varnothing \in JD$, i.e. that the empty sieve covers D.*

(vi) *The dual image $\forall_\eta A$ is the sheaf defined by $\xi' \in (\forall_\eta A)C \iff$ for all $f : D \to C$ and $d \in ED$ such that $\xi' \cdot f = \eta_D d$, $d \in AD$.*

Proof Elegant proofs of all these statements may be found in [MM92, §III.8]. □

Using these descriptions, we obtain the following theorem, giving an explicit description of the Kripke-Joyal semantics:

Theorem 3.13. *Let $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$ be a Grothendieck topos. Let ϕ, ψ and $\phi_i, i \in I$ be formulae of the language of $\mathbf{E}$, in free variables $\vec{e} = (e_1, \ldots e_n)$ of types $E_1, \ldots E_n$, and let $\vec{\xi} = (\xi_1, \ldots, \xi_n) \in \prod_{k=1}^n E_k C$. Then:*

(i) *$C \Vdash \phi(\vec{\xi}) \wedge \psi(\vec{\xi})$ iff $C \Vdash \phi(\vec{\xi})$ and $C \Vdash \psi(\vec{\xi})$;*

(ii) *$C \Vdash \bigvee_i \phi_i(\vec{\xi})$ iff there is a covering family $\{f_j : C_j \to C\}$ such that for all j, there exists an $i \in I$ such that $C_j \Vdash \phi_i(\vec{\xi} \cdot f_j)$;*

Suppose now that $\phi = \phi(\vec{e}, e)$ has an additional free variable e of type E. Then:

(iii) *$C \Vdash \exists e \phi(\vec{\xi}, e)$ iff there is a covering family $\{f_j : C_j \to C\}$ such that for all j, there is an $\varepsilon_j \in EC_j$ such that $C_j \Vdash \phi(\vec{\xi} \cdot f_j, \varepsilon_j)$;*

In a Boolean Grothendieck topos, the Kripke-Joyal semantics extend to full infinitary logic by the rules:

(iv) *$C \Vdash \phi(\vec{\xi}) \to \psi(\vec{\xi})$ iff for all $f : D \to C$, $D \Vdash \phi(\vec{\xi} f)$ implies $D \Vdash \psi(\vec{\xi} \cdot f)$.*

(v) *$C \Vdash \neg\phi(\vec{\xi})$ iff for all $f : D \to C$, $D \Vdash \phi(\vec{\xi} \cdot f)$ implies $\varnothing \in JD$.*

(vi) *$C \Vdash \forall e \phi(\vec{\xi}, e)$ iff for all $f : D \to C$ and $\varepsilon \in ED$, $D \Vdash \phi(\vec{\xi} \cdot f, \varepsilon)$.*

The proof of this theorem is an entirely obvious writing exercise in view of the preceding lemma.

The Kripke-Joyal semantics also give a perspective as to why the disjoint sums and quotients of equivalence relations we have imposed on ∞-pretopoi are useful. Namely, many important constructions in set theory, but also for example analysis (construction of $\mathbb{R}$) and algebra (tensor and exterior spaces) rely in one way or another on the availability of these disjoint sums and quotients in the category **Set**. Their presence in ∞-pretopoi enhances the expressivity of these "alternative universes for set theory" and hence their usefulness in the study of axiomatic independence and alternative approaches to set theory and logic.

4 Geometric morphisms and classifying topoi

In the previous chapters, we have only regarded Grothendieck topoi as stand-alone objects, discussing their properties and several methods to analyze these properties. In the spirit of category theory we will now discuss a suitable category of topoi and so-called *geometric morphisms*. A connection with the preceding material is found in *classifying topoi*, which relate models of theories (formulated in the framework of Chapter 2) and certain Hom-sets of geometric morphisms.

4.1 Geometric morphisms and continuous functors

The definition of a geometric morphism is founded on the important example of a topological space X and sheaves on the associated category of open sets and inclusions $\mathcal{O}(X)$. For a continuous mapping $f : X \to Y$ we have a functor $f_* : \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$ defined by $(f_*F)U := F(f^{-1}U)$, and a left adjoint $f^* : \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ of f_* which is defined by pullback of étale bundles over Y (which are equivalent to sheaves on Y). Moreover, f^* preserves finite limits. An extensive discussion can be found in Chapter II of [MM92].

These adjoint functors can be thought of as an "approximation" of the underlying continuous function $f : X \to Y$. In this light, we introduce the definition of a morphism between arbitrary topoi:

Definition 4.1. Let $\mathbf{E}$ and $\mathbf{E}'$ be topoi. A **geometric morphism** $f : \mathbf{E} \to \mathbf{E}'$ comprises adjoint functors $f^* \dashv f_*$:

$$f_* : \mathbf{E} \rightleftarrows \mathbf{E}' : f^*$$

such that f^* preserves finite limits. Here f_* is called the **direct image part**, and f^* the **inverse image part**, of f.

The identity geometric morphism and the general fact that the composition of two adjoint pairs of functors yields another pair of adjoint functors (see [Mac71, Theorem IV.8.1]) provide us with a category **Topos** of topoi and geometric morphisms.

As it turns out, the hom-sets $\mathrm{Hom}_{\mathbf{Topos}}(\mathbf{F}, \mathbf{E})$ can also be made into categories themselves: a morphism $\alpha : f \to g$ between two geometric morphisms $f, g : \mathbf{F} \to \mathbf{E}$ is a natural transformation $\alpha : f^* \to g^*$ between their inverse image parts. There is no asymmetry here: the Yoneda lemma ensures a correspondence $\mathrm{Nat}(f^*, g^*) \cong \mathrm{Nat}(g_*, f_*), \alpha \sim \beta$, determined by the commutative diagram:

$$\begin{tikzcd}
\mathbf{F}(f^*E, F) \arrow[r, "\cong"] & \mathbf{E}(E, f_*F) \\
\mathbf{F}(g^*E, F) \arrow[r, "\cong"] \arrow[u, "\mathbf{F}(\alpha_E, F)"] & \mathbf{E}(E, g_*F) \arrow[u, "\mathbf{E}(E, \beta_F)"']
\end{tikzcd}$$

where the horizontal equivalences result from the adjunctions defining the geometric morphisms f and g.

The resulting category of geometric morphisms and natural transformations will be denoted by $\underline{\mathrm{Hom}}(\mathbf{F}, \mathbf{E})$.

The notion of geometric morphism was abstracted from the case of a continuous function. We can define functors between sites that are analogous to continuous functions:

Definition 4.2. Let $(\mathbf{C}, J)$ and $(\mathbf{D}, K)$ be sites, and let $F : \mathbf{C} \to \mathbf{D}$ be a functor. Then F is said to be **continuous** if:

(i) F preserves finite limits;

(ii) F takes J-covers to K-covers, i.e., for any J-covering collection $\{f_i : C_i \to C\}$ of morphisms (see Definition 1.12), $\{Ff_i : FC_i \to FC\}$ is K-covering.

The category of sites and continuous functors between them is denoted **Site**, and we also refer to a continuous functor as a **morphism of sites**.
The category of continuous functors $\mathbf{C} \to \mathbf{D}$ is denoted with $\mathrm{Con}(\mathbf{C}, \mathbf{D})$.

An important example of a continuous functor concerns a site and its associated Grothendieck topos.

Proposition 4.3. *Let $(\mathbf{C}, J)$ be a site, with associated Grothendieck topos $\mathbf{Sh}(\mathbf{C}, J)$. Let $\mathbf{Sh}(\mathbf{C}, J)$ be endowed with the canonical topology. Then the composition $\epsilon = \mathbf{ay} : \mathbf{C} \to \mathbf{Sh}(\mathbf{C}, J)$ of the Yoneda embedding $\mathbf{y} : \mathbf{C} \to \hat{\mathbf{C}}$ and the sheafification functor $\mathbf{a} : \hat{\mathbf{C}} \to \mathbf{Sh}(\mathbf{C}, J)$, is a continuous functor.*

Proof Since $\mathbf{y}$ and $\mathbf{a}$ both preserve finite limits, so does ϵ; now suppose that $\{C_i \to C\} \in JC$ is a cover. By Proposition 3.9, we only need to verify that $\{\epsilon C_i \to \epsilon C\}$ is an epimorphic family; we will only need that the stability in "stable effective epimorphic family" is automatic. By the adjunction $\mathbf{a} \dashv U$, we have, for a fixed sheaf F, canonical isomorphisms $\mathrm{Hom}(\mathbf{y}C, UF) \cong \mathrm{Hom}(\epsilon C, F)$. This gives rise to the following diagram (suppressing occurrences of U):

$$\begin{tikzcd}
\mathrm{Hom}(\mathbf{y}C, F) \arrow[r] \arrow[d, "\cong"] & \prod \mathrm{Hom}(\mathbf{y}C_i, F) \arrow[r, shift left] \arrow[r, shift right] \arrow[d, "\cong"] & \mathrm{Hom}(\mathbf{y}(C_i \times_C C_j), F) \arrow[d, "\cong"] \\
\mathrm{Hom}(\epsilon C, F) \arrow[r] & \prod \mathrm{Hom}(\epsilon C_i, F) \arrow[r, shift left] \arrow[r, shift right] & \mathrm{Hom}(\epsilon(C_i \times_C C_j), F)
\end{tikzcd}$$

It follows that the bottom row is an equalizer in **Set** iff the top row is. But the top row being an equalizer just expresses (through the Yoneda lemma and the remark in Definition 3.8) that F is a sheaf. Since F was an arbitrary sheaf, it follows that $\{\epsilon C_i \to \epsilon C\}$ is an effective epimorphic family, as desired. □

We will refer to ϵ as the **canonical continuous functor** on $(\mathbf{C}, J)$. We already encountered it briefly in Chapter 3, which allows for a short proof of the following lemma:

Lemma 4.4. *Let $(\mathbf{C}, J)$ be a site, and let $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$. Then every $E \in \mathrm{ob}\,\mathbf{E}$ has a cover by objects of the form ϵC.*

Proof This is a direct consequence of Proposition 3.9 and the fact that the ϵC form a generating set for $\mathbf{E}$ (as established in the proof of Theorem 3.7). □

The subcategory of $\mathbf{Sh}(\mathbf{C}, J)$ determined by ϵ intuitively should look a lot like the original site $(\mathbf{C}, J)$. The following nontrivial result in this direction (whose proof is beyond our current scope in both length and methods) occurs as Proposition 1.3.3 in [MR77]:

Proposition 4.5. *Suppose that $\{\epsilon f_i : \epsilon C_i \to \epsilon C\}$ is a cover in $\mathbf{Sh}(\mathbf{C}, J)$. Then $\{f_i : C_i \to C\}$ is a cover in $(\mathbf{C}, J)$.*

A conceptually important result is the following, which indicates a "completeness" property of Grothendieck topoi. Namely, that every sheaf over a Grothendieck topos is representable.

Lemma 4.6 ([MR77, Lemma 1.3.14])**.** *Let $\mathbf{E}$ be a Grothendieck topos, and let $\mathbf{Sh}(\mathbf{E})$ be the category of sheaves over $\mathbf{E}$. Then $\epsilon : \mathbf{E} \to \mathbf{Sh}(\mathbf{E})$ is an equivalence of categories.*

The analogy between continuous functors and continuous functions that motivated the definition of the former follows through even further: every continuous functor induces a suitable geometric morphism.

Theorem 4.7 ([MR77, Theorem 1.3.10])**.** *Let $f : \mathbf{C} \to \mathbf{D}$ be a continuous functor. Then there exists a geometric morphism $\mathbf{Sh}(\mathbf{D}, K) \to \mathbf{Sh}(\mathbf{C}, J)$, also denoted f, whose left adjoint $f^* : \mathbf{Sh}(\mathbf{C}, J) \to \mathbf{Sh}(\mathbf{D}, K)$ makes the following diagram commute:*

$$\begin{tikzcd}
\mathbf{C} \arrow[r, "f"] \arrow[d, "\epsilon"'] & \mathbf{D} \arrow[d, "\epsilon"] \\
\mathbf{Sh}(\mathbf{C}, J) \arrow[r, "f^*"'] & \mathbf{Sh}(\mathbf{D}, K)
\end{tikzcd}$$

Corollary 4.8. *Let $(\mathbf{C}, J)$ be a site. For each Grothendieck topos $\mathbf{E}$, $\epsilon : (\mathbf{C}, J) \to \mathbf{Sh}(\mathbf{C}, J)$ induces an equivalence of categories $\mathrm{Con}(\mathbf{C}, \mathbf{E}) \cong \mathrm{Con}(\mathbf{Sh}(\mathbf{C}, J), \mathbf{E})$.*

In the next chapter, we will come back to continuous functors; for now, we focus on the notion of the **classifying topos** of a (geometric) theory T: a topos that can intuitively be regarded as the "categorification" of T.

4.2 Classifying topoi

Consider a fixed geometric theory T. Then for each Grothendieck topos $\mathbf{E}$ we can consider the category $\underline{\mathrm{Mod}}(T, \mathbf{E})$ as defined in Section 2.6.

Definition 4.9. A **classifying topos** for T is a topos $\mathbf{S}(T)$ such that for all Grothendieck topoi $\mathbf{E}$, there is an equivalence of categories:

$$\underline{\mathrm{Hom}}(\mathbf{E}, \mathbf{S}(T)) \cong \underline{\mathrm{Mod}}(T, \mathbf{E})$$

which is natural in $\mathbf{E}$.

Alternatively, we can take the point of view of the **universal model of** T, i.e. the model $U = U(T) \in \underline{\mathrm{Mod}}(T, \mathbf{S}(T))$ corresponding to the identity natural transformation on $\mathbf{S}(T)$. In this setup, the pair $(\mathbf{S}(T), U)$ is a **classifying topos** iff for each $\mathbf{E}$, the obvious functor:

$$\underline{\mathrm{Hom}}(\mathbf{E}, \mathbf{S}(T)) \to \underline{\mathrm{Mod}}(T, \mathbf{E}), \quad f \mapsto f^*(U)$$

is an equivalence of categories.

As is usual with these definitions, it is clear that $\mathbf{S}(T)$ is unique up to isomorphism. However, the surprising fact is:

Theorem 4.10. *For any geometric theory T, there is a Grothendieck topos $\mathbf{S}(T)$ that is a classifying topos for T.*

Thus for any theory T, there is a "canonical context" in which to study T, which is very useful for practical purposes. The construction of $\mathbf{S}(T)$ proceeds in several steps, and will arise as only one of several results involving "classifying categories".

4.3 Classifying categories for κ-geometrical theories

Let κ be an infinite regular cardinal, or ∞. We recall the notions of κ-geometrical category and functor from Definitions 3.1 and 2.21:

Definition 4.11. A category $\mathbf{C}$ is called **κ-geometrical** if it has:

(i) finite limits;

(ii) stable suprema of sets of subobjects of cardinality less than κ;

(iii) stable images of subobjects.

A functor $F : \mathbf{C} \to \mathbf{D}$ is called **κ-geometrical** if it preserves finite limits, images of subobjects, and suprema of sets of subobjects of cardinality less than κ. It is understood that all sets have cardinality less than ∞.

The category of κ-geometrical categories and likewise functors is denoted by **κ-Geo**.

Each κ-geometrical category $\mathbf{C}$ has an associated theory $T_{\mathbf{C}}$ in the language $L^g_{\kappa\omega}$, where $L = L_{\mathbf{C}}$ is the canonical language for $\mathbf{C}$.

Definition 4.12. Let $\mathbf{C}$ be a κ-geometrical category. The theory $T_{\mathbf{C}}$ **associated to $\mathbf{C}$** consists of the following axioms (where we identify objects and morphisms of $\mathbf{C}$ with the corresponding symbols in $L_{\mathbf{C}}$)

(i) For each object C of $\mathbf{C}$, the axiom: $c = \mathrm{id}_C\, c$ (an obvious shorthand for the formal version $\varnothing \Rightarrow \{c = \mathrm{id}_C\, c\}$);

(ii) For each commutative triangle $h = gf$, the axiom: $hc = gfc$;

(iii) For each terminal object T of $\mathbf{C}$, the axioms: $t = t'$ and $\exists t(t = t)$;

(iv) For each product $C = A \times B$ of $\mathbf{C}$ with projections π_A and π_B, the axioms:

$$\begin{aligned} \pi_A c = \pi_A c' \wedge \pi_B c = \pi_B c' &\Rightarrow c = c' \\ &\Rightarrow \exists c(\pi_A c = a \wedge \pi_B c = b) \end{aligned}$$

(v) For each equalizer $B \xrightarrow{e} C \overset{f}{\underset{g}{\rightrightarrows}} D$, the axioms:

$$\begin{aligned} eb = eb' &\Rightarrow b = b' \\ &\Rightarrow feb = geb \\ fc = gc &\Rightarrow \exists b(eb = c) \end{aligned}$$

(vi) For each identity of subobjects $Y = \bigvee_i X_i$ of C with $i \in I$ and $|I| < \kappa$, the axiom: $\bigvee_i X_i(c) \Leftrightarrow Y(c)$;

(vii) For each extremal epimorphism $f : C \to D$, the axiom: $\exists c(fc = d)$.

That $\mathbf{C}$ is actually a model of $T_{\mathbf{C}}$ is Theorem 2.4.5 of [MR77]. Using this result, we can now derive:

Theorem 4.13. *For each κ-geometrical category $\mathbf{D}$, there is a bijection, natural in $\mathbf{D}$:*

$$\kappa\text{-}\mathbf{Geo}(\mathbf{C}, \mathbf{D}) \to \underline{\mathrm{Mod}}(T_{\mathbf{C}}, \mathbf{D}), \quad F \mapsto \underline{\mathrm{Mod}}(T_{\mathbf{C}}, F)([\mathbf{C}])$$

where $[\mathbf{C}]$ denotes $\mathbf{C}$ considered as the canonical model of $T_{\mathbf{C}}$, and $\underline{\mathrm{Mod}}(T_{\mathbf{C}}, F)$ is as in Proposition 2.22.

Proof By Proposition 2.22, $\underline{\mathrm{Mod}}(T_{\mathbf{C}}, F)([\mathbf{C}])$ is a model for each κ-geometrical functor $F \in \kappa\text{-}\mathbf{Geo}(\mathbf{C}, \mathbf{D})$. Conversely, given a model $M : T_{\mathbf{C}} \to \mathbf{D}$, we define $F_M : \mathbf{C} \to \mathbf{D}$ by $F_M C = MC$ and $F_M f = Mf$. Then by clause (ii) of 4.12, it is ensured that F_M is a functor. By Theorem 2.4.5 of [MR77], we know that each of the stated sets of axioms expresses precisely what it purports to. For example, let us derive from clause (v) that Me is the equalizer of Mf and Mg.

The first axiom expresses that Me is monic: $M_{b,b'}(eb = eb')$ is the equalizer of

$$MB \times MB \overset{Me\pi_2}{\underset{Me\pi_1}{\rightrightarrows}} MC$$

and $M_{b,b'}(b = b')$ is the equalizer of the projections $\pi_1, \pi_2 : MB \times MB \to MB$. Now suppose $Me\alpha = Me\beta$ for some α and β. This is equivalent to:

$$Me\pi_2\langle\alpha, \beta\rangle = Me\pi_1\langle\alpha, \beta\rangle$$

so that $\langle\alpha, \beta\rangle$ factors through $M_{b,b'}(eb = eb')$. By the axiom, this subobject is smaller than $M_{b,b'}(b = b')$, so that $\langle\alpha, \beta\rangle$ also factors through this latter subobject. But this means $\alpha = \beta$, i.e. Me is monic.

As established in Section 2.5, the second axiom translates to $M(fe) = M(ge)$, which means Me equalizes Mf and Mg due to clause (ii) of 4.12. From the material in referenced section, we also obtain that $M_c(fc = gc)$ is the equalizer of Mf and Mg. Since by Lemma 2.15, $M_c(\exists b(eb = c))$ is the subobject classified by the mono $Me : Mb \rightarrowtail Mc$, the second and third axiom combine to the statement that Me is the equalizer of Mf and Mg, as asserted.

Hence, due to the axioms of $T_{\mathbf{C}}$, we can establish that F_M preserves equalizers, and indeed that it is κ-geometrical. The two operations described clearly establish the desired bijection. □

This theorem gives us a connection between categories and theories. We will now set out to a result in the converse direction, associating a category $\mathbf{C}_T$ to a theory T. The upshot is to obtain the following result:

Theorem 4.14 ([MR77, Theorem 8.3.1])**.** *Let T be an $L^g_{\kappa\omega}$-theory. Then there exist a κ-geometrical category $\mathbf{C}_T$ and a model M_0 of T in $\mathbf{C}_T$ such that for each κ-geometrical category $\mathbf{D}$ and model $M \in \underline{\mathrm{Mod}}(T, \mathbf{D})$, there is a κ-geometrical functor $F : \mathbf{C}_T \to \mathbf{D}$ such that $M = FM_0$.*

Equivalently: For each κ-geometrical category $\mathbf{D}$, there is an equivalence of categories:

$$\underline{\mathrm{Mod}}(T, \mathbf{D}) \cong \kappa\text{-}\mathbf{Geo}(\mathbf{C}_T, \mathbf{D})$$

from which M_0 can be recovered as the model corresponding to the identity functor on $\mathbf{C}_T$.

Definition 4.15. Let L be a signature. A formula ϕ of $L^g_{\infty\omega}$ is called **κ-simple** if it is of the form:

$$\phi = \bigvee_{i \in I} \psi_i$$

where $|I| < \kappa$, each ψ_i is of the form $\exists x_1 \cdots \exists x_n(\chi_1 \wedge \cdots \wedge \chi_m)$, and each χ is atomic.

The κ-simple formulae are like the Conjunctive and Disjunctive Normal Forms of propositional calculus: every formula ϕ of $L^g_{\kappa\omega}$ is "equivalent" to a κ-simple formula ψ, in the sense of provable equivalence in a proof system for $L^g_{\infty\omega}$. For completeness and ease of reference, we include the main result about this proof system:

Theorem 4.16 (Cf. Corollary 5.2.3 of [MR77])**.** *There is a proof system $\vdash$ for $L^g_{\infty\omega}$ that is both sound and complete for models as introduced in Definition 2.18. More specifically, for each κ (possibly ∞), for each κ-geometric theory T and κ-geometric sequent $\Phi \Rightarrow \Psi$, $T \vdash \Phi \Rightarrow \Psi$ holds if and only if $M \models \Phi \Rightarrow \Psi$ for each model M of T in a κ-geometric category $\mathbf{C}$.*

The κ-simple formula ψ which is equivalent to ϕ, i.e. satisfies $\phi \vdash \psi$ and $\psi \vdash \phi$, can be explicitly constructed from ϕ, and we will henceforth denote the ψ arising from that construction by $[\phi]$. For details on this construction and proof system, we refer to Proposition 8.3.2 of [MR77] and the results quoted there. These formulae $[\phi]$ will become the objects of $\mathbf{C}_T$, and we denote with $\mathcal{O}_\kappa$ the set of all $[\phi]$.

The proof system recurs in the definition of what will become the morphisms of $\mathbf{C}_T$:

Definition 4.17. Let $\phi(\vec{x}), \psi(\vec{y}) \in \mathcal{O}_\kappa$. A **premorphism** $\mu : \phi \to \psi$ is a formula $\mu(\vec{x}', \vec{y}')$ (where e.g. the sorts of $\vec{x}$ and $\vec{x}'$ match), such that:

$$\begin{aligned}
&T \vdash \mu(\vec{x}', \vec{y}') \Rightarrow \phi(\vec{x}') \wedge \psi(\vec{y}') \\
&T \vdash \mu(\vec{x}', \vec{y}'), \mu(\vec{x}', \vec{y}'') \Rightarrow \vec{y}' = \vec{y}'' \\
&T \vdash \phi(\vec{x}') \Rightarrow \exists \vec{y}' \mu(\vec{x}', \vec{y}')
\end{aligned}$$

Two premorphisms $\mu, \mu' : \phi \to \psi$ are **equivalent**, denoted $\mu \sim \mu'$, if T proves their equivalence, i.e. $T \vdash \mu(\vec{x}', \vec{y}') \Leftrightarrow \mu'(\vec{x}', \vec{y}')$.

Definition 4.18. The category $\mathbf{C}_T$ has as objects the set $\mathcal{O}_\kappa$, and a morphism $\phi \to \psi$ is an equivalence class $\mu/{\sim}$ of premorphisms $\mu : \phi \to \psi$. The composition of $\mu : \phi \to \psi$ and $\nu : \psi \to \chi$ is given by $\lambda(\vec{x}', \vec{z}') = \exists \vec{y}'(\mu(\vec{x}', \vec{y}') \wedge \nu(\vec{y}', \vec{z}'))$.

We define a structure M_0 in $\mathbf{C}_T$ by putting $M_0 X = [x = x]$ for a sort X (with x a variable of sort X), $M_0 R = [R(\vec{x})]$ for a relation symbol $R \subseteq \prod_i X_i$, and $M_0 f = (f(\vec{x}) = y)/{\sim}$ for a function symbol $f : \prod_i X_i \to Y$.

Lemma 4.19. *$\mathbf{C}_T$ is a category and has all finite limits.*

Proof That $\lambda : \phi \to \chi$ is a premorphism, and that $[\lambda]$ is well-defined given $[\mu]$ and $[\nu]$, and the other axioms for a category may be verified by reasoning in an arbitrary T-model by Theorem 4.16, which is straightforward but tedious. Similarly, one may prove that indeed the expression $f(\vec{x}) = y$, used to define $M_0 f$, constitutes a premorphism.

Now, to verify that $\mathbf{C}_T$ has finite limits. For products, we propose $\phi(\vec{x}) \times \psi(\vec{y}) = \phi(\vec{x}) \wedge \psi(\vec{y})$, with premorphisms $\pi_1 : \phi \wedge \psi \to \phi$ given by:

$$\pi_1(\vec{x}, \vec{y}; \vec{x}') = \phi(\vec{x}) \wedge \psi(\vec{y}) \wedge \vec{x} = \vec{x}'$$

and π_2 defined analogously. It is immediate that π_1 and π_2 define premorphisms. Now to verify the universal property of the product. Given $\mu : \chi(\vec{z}) \to \phi(\vec{x})$ and $\nu : \chi(\vec{z}) \to \psi(\vec{y})$, it is natural to consider $\mu \wedge \nu$ as a premorphism $\chi \to \phi \wedge \psi$. We need to establish its uniqueness, so suppose that $\theta : \chi \to \phi \wedge \psi$ satisfies $\pi_1 \theta = \mu$ and $\pi_2 \theta = \nu$. The formula corresponding to $\pi_1 \theta(\vec{z}, \vec{x}')$ is:

$$\exists \vec{x}, \vec{y}(\theta(\vec{z}; \vec{x}, \vec{y}) \wedge \phi(\vec{x}) \wedge \psi(\vec{y}) \wedge \vec{x} = \vec{x}')$$

which is seen to be equivalent to $\exists\vec{y}(\theta(\vec{z};\vec{x}',\vec{y}) \wedge \phi(\vec{x}') \wedge \psi(\vec{y}))$. By the second provability axiom for a premorphism, this is furthermore equivalent to $\exists\vec{y}\theta(\vec{z};\vec{x}',\vec{y})$. The hypothesis $\pi_1\theta = \mu$ now implies that this is T-provably equivalent to $\mu(\vec{z},\vec{x})$. Similarly, $\exists\vec{x}'\theta(\vec{z};\vec{x}',\vec{y})$ is T-provably equivalent to $\nu(\vec{z},\vec{y})$. Since θ is a premorphism, there can be at most one $\vec{x}',\vec{y}$ such that $\theta(\vec{z};\vec{x}',\vec{y})$, and the above equivalences translate to $\mu(\vec{z},\vec{x}) \wedge \nu(\vec{z},\vec{x})$ being provably equivalent to $\theta(\vec{z};\vec{x},\vec{y})$, as desired.

Now for equalizers. We claim that the equalizer of $\mu,\nu : \phi \to \psi$ is given by $\tau : \chi \to \phi$, with $\chi(\vec{x})$ the formula $\exists\vec{y}(f(\vec{x},\vec{y}) \wedge g(\vec{x},\vec{y}))$, while $\tau(\vec{x},\vec{x}')$ is defined as $\chi(\vec{x}) \wedge (\vec{x} = \vec{x}')$. The verification is analogous to the above case of products. Finally, it is obvious how to turn the tautology $\top = \bigvee\varnothing$ into a terminal object. $\square$

Lemma 4.20. *Every formula of $L^g_{\kappa\omega}$ can be interpreted in $\mathbf{C}_T$ by means of M_0. Moreover, M_0 is a model of T in $\mathbf{C}_T$ such that $M_0 \models \Phi \Rightarrow \Psi$ iff $T \vdash \Phi \Rightarrow \Psi$.*

Proof The interpretation of M_0 is of a rather tautologous nature. For example, the equalizer that interprets $t_1 = t_2$ is, by the preceding lemma, given by $[\exists y(\tilde{t}_1(\vec{x},y) \wedge \tilde{t}_2(\vec{x},y))]$ where $\tilde{t}_i$ is the premorphism $M_0 t_i : M_0\vec{X} \to M_0 Y$; this is the same as the object $[t_1(\vec{x}) = t_2(\vec{x})]$. Similarly we can deal with the other atomic case $R(t_1,\dots t_n)$ to establish that $(M_0)_{\vec{x}}(R(t_1,\dots,t_n))$ is $[R(t_1,\dots,t_n)]$ (with the context $\vec{x}$ appropriately extended). We will not bore the reader with the details of $\bigvee_{i\in I}[\phi_i] = \left[\bigvee_{i\in I}\phi_i\right]$ or the analogous case for conjunction. The existential quantifier case, however, is more interesting.

Given $\phi(\vec{x},y)$, we have by definition of interpretation and the induction hypothesis that $(M_0)_{\vec{x}}(\exists y\phi(\vec{x},y)) = \exists_\pi(M_0)_{\vec{x},y}(\phi(\vec{x},y)) = \exists_\pi[\phi(\vec{x},y)]$ with π the suitable projection $M_0(\vec{x},y) \to M_0(\vec{x})$. We expect that $\exists_\pi[\phi(\vec{x},y)] = [\exists y\phi(\vec{x},y)]$. The associated premorphism $\mu : \phi \to \exists y\phi$ is given by $\mu(\vec{x},y,\vec{x}') = \phi(\vec{x},y) \wedge \vec{x} = \vec{x}'$. Now let us verify that it is indeed the image.

Suppose that there were

Now if $T \vdash \Phi \Rightarrow \Psi$, then we see that there is a premorphism $[\bigwedge\Phi] \to [\bigvee\Psi]$ given by $\bigwedge\Phi(\vec{x}) \wedge \vec{x} = \vec{x}'$. This premorphism clearly commutes with the respective inclusions in $[\vec{x} = \vec{x}]$, so that indeed $M_0 \models \Phi \to \Psi$. Conversely, the existence of a premorphism $\mu : \phi \to \psi$ implies that $T \vdash \phi \Rightarrow \psi$; this is an easy consequence of the definition of premorphism. $\square$

Lemma 4.21. *The category $\mathbf{C}_T$ is κ-geometrical.*

Proof The constructions in the preceding lemma essentially say that $L^g_{\kappa\omega}$ is the canonical language for $\mathbf{C}_T$, and that M_0 is the associated identity interpretation of $L^g_{\kappa\omega}$ in $\mathbf{C}_T$. Combining that M_0 only models sequents that are consequences of T with Definition 4.12, only three things remain. Namely, first of all, to identify the candidates for images and sups, and secondly to verify that these candidate fulfil the requirements; finally, to check that they are stable.

As to the first, we purport that for $\mu : \phi(\vec{x}) \to \psi(\vec{y})$, we can define the image $\exists_\mu\phi(\vec{y})$ by $\mu' : \phi(\vec{x}) \to \exists_\mu\phi(\vec{y}'), \mu'(\vec{x},\vec{y}') = \mu(\vec{x},\vec{y}')$ and $\nu : \exists_\mu\phi(\vec{y}') \to \psi(\vec{y}), \nu(\vec{y}',\vec{y}) = \exists\vec{x}\mu(\vec{x},\vec{y}) \wedge \vec{y} = \vec{y}'$. The case for suprema is even simpler. A reasoning as in the image case of the above lemma establishes that these definitions yield images and suprema.

Finally, the stability (as defined in Definition 3.2) can be expressed using certain sequents in $L^g_{\kappa\omega}$ (using (iii) and (iv) of Definition 4.12 to express the pullbacks) and hence is a consequence of Theorem 4.16. $\square$

Together, the preceding three lemmata prove Theorem 4.14.

4.4 The κ-pretopos completion of a κ-geometrical category

The next step in constructing classifying topoi is a general construction of a κ-pretopos $\mathbf{E}$ from a κ-geometrical category $\mathbf{C}$, which will have the following universal property.

Theorem 4.22. *Let $\mathbf{C}$ be a small, κ-geometrical category. Then there exists a κ-pretopos $\mathbf{E}$ such that for each κ-pretopos $\mathbf{F}$, there is an equivalence of categories:*

$$\kappa\text{-}\mathbf{Geo}(\mathbf{C},\mathbf{F}) \cong \kappa\text{-}\mathbf{Geo}(\mathbf{E},\mathbf{F})$$

which, as before, can also be expressed using a κ-geometrical functor $I : \mathbf{C} \to \mathbf{E}$.

Moreover, if $\kappa < \infty$, $\mathbf{E}$ is small; if $\kappa = \infty$, $\mathbf{E}$ is a Grothendieck topos.

Lemma 4.23. *Let $\mathbf{C}$ be a small, κ-geometrical category. Then there exist a κ-geometrical category $\mathbf{C}^{\coprod}$ and a κ-geometrical functor $I : \mathbf{C} \to \mathbf{C}^{\coprod}$ such that:*

(i) *$\mathbf{C}^{\coprod}$ has disjoint sums of fewer than κ objects;*

(ii) *For any κ-geometrical category $\mathbf{D}$ and κ-geometrical functor $F : \mathbf{C} \to \mathbf{D}$ such that $\mathbf{D}$ satisfies (i), there is a unique κ-geometrical functor $F' : \mathbf{C}^{\coprod} \to \mathbf{D}$ such that $F = F'I$;*

(iii) *I "reflects the subobject lattices": If $IA \leq IB$ for subobjects A, B of C, then $A \leq B$;*

(iv) *I is "full on subobjects": If $Z \leq IB$, then $Z \cong IA$ for some subobject A;*

(v) *Every object C of $\mathbf{C}^{\coprod}$ is isomorphic to $\coprod_{j \in J} IC_j$ for objects $C_j \in \operatorname{ob} \mathbf{C}$, with $j \in J$ and $|J| < \kappa$.*

Proof Let $T_{\mathbf{C}}$ be the theory of $\mathbf{C}$, and let $L_{\mathbf{C}}$ be the signature of $T_{\mathbf{C}}$. We extend $L_{\mathbf{C}}$ by adding for each subset $J \subseteq \operatorname{ob} \mathbf{C}$ with $|J| < \kappa$ a sort C_J with intended interpretation $\coprod_{j \in J} C_j$; also, for each $j \in J$, a function symbol $i_j : C_j \to C_J$. In this extended signature, we consider the theory T which consists of $T_{\mathbf{C}}$ together with, for each J, the axioms from Definition 3.5 expressing that C_J is indeed the disjoint sum $\coprod_{j \in J} C_j$ with injections i_j. Now we define $\mathbf{C}^{\coprod}$ to be the classifying category of T. The functor I arises from the universal property of $\mathbf{C}$ as a model of $T_{\mathbf{C}}$.

It is obvious that $\mathbf{C}^{\coprod}$, so defined, satisfies (i) and (ii); for (ii), observe that any model in $\mathbf{D}$ of $T_{\mathbf{C}}$ which satisfies (i) has an extension to a model of T. For the proofs of the remaining statements, we refer the reader to [MR77], Theorem 8.4.2. $\square$

Lemma 4.24. *Let $\mathbf{C}$ be a small, κ-geometrical category. Then there exists a κ-geometrical catgory $\mathbf{C}^{\sim}$ and a κ-geometrical functor $I^{\sim} : \mathbf{C} \to \mathbf{C}^{\sim}$ such that:*

(i) *$\mathbf{C}^{\sim}$ has quotients by equivalence relations;*

(ii) *For any κ-geometrical category $\mathbf{D}$ and κ-geometrical functor $F : \mathbf{C} \to \mathbf{D}$ such that $\mathbf{D}$ satisfies (i), there is a unique κ-geometrical functor $F' : \mathbf{C}^{\sim} \to \mathbf{D}$ such that $F = F'I^{\sim}$;*

(iii) *I reflects subobject lattices and is full on subobjects;*

(iv) *If $\mathbf{C}$ has disjoint sums of cardinality less than κ, then $\mathbf{C}^{\sim}$ is a κ-pretopos.*

Proof As in the other lemma, we extend the theory $T_{\mathbf{C}}$. This time, we introduce sorts $Q(C, R)$, one for each equivalence relation $R \hookrightarrow C \times C$ in $\mathbf{C}$, as well as function symbols $q_R : C \to Q(C, R)$. We add to $T_{\mathbf{C}}$ suitable forms of the axioms from Definition 3.5, expressing that q_R is a quotient of R; denote the resulting theory by T. Again, we define $\mathbf{C}^{\sim}$ to be the classifying category for T.

The first two items are again immediate from the relevant universal properties; the latter two require a more involved proof and we refer the reader to [MR77] once again; the relevant result is Theorem 8.4.3. $\square$

The preceding two lemmata together are seen to imply Theorem 4.22. Combining theorem 4.22 with 4.14, we obtain:

Theorem 4.25. *Let T be an $L^g_{\kappa\omega}$-theory. Then there exists a κ-pretopos $\mathbf{E}_T$ such that for each κ-pretopos $\mathbf{F}$, there is an equivalence of categories:*

$$\underline{\mathrm{Mod}}(T, \mathbf{F}) \cong \kappa\text{-}\mathbf{Geo}(\mathbf{E}_T, \mathbf{F})$$

We need the following lemma takes care of the last gap before we can conclude the proof of Theorem 4.10.

Lemma 4.26. *Let $\mathbf{E}$ and $\mathbf{F}$ be Grothendieck topoi, and let $F : \mathbf{E} \to \mathbf{F}$ be an ∞-geometrical morphism. Then F is continuous with respect to the canonical topologies on $\mathbf{E}$ and $\mathbf{F}$.*

Proof By Proposition 3.9, it suffices to show that F preserves epimorphic families. So let $\{f_i : E_i \to E\}$ be epimorphic; because each f_i factors through $\operatorname{Im} f_i \hookrightarrow E$, and F preserves images, we may assume without loss of generality that the f_i are monos. But then the condition that $\{f_i\}$ be epimorphic is just the statement that $\bigvee_i E_i = E$. Since F preserves suprema as well, we conclude that F is continuous. $\square$

Finally, now, we can conclude the existence of a classifying topos as in Theorem 4.10.

Proof of Theorem 4.10 By Theorem 4.25, there is a natural correspondence:

$$\frac{T \longrightarrow \mathbf{E}}{\mathbf{E}_T \longrightarrow \mathbf{E}}$$

By Lemma 4.26, the ∞-geometrical functor $f : \mathbf{E}_T \to \mathbf{E}$ is continuous with respect to the canonical topologies on these Grothendieck topoi. Theorem 4.7 then establishes a geometric morphism $g_T : \mathbf{Sh}(\mathbf{E}) \to \mathbf{Sh}(\mathbf{E}_T)$. But by Lemma 4.6, $\epsilon : \mathbf{E} \to \mathbf{Sh}(\mathbf{E})$ is an equivalence for all Grothendieck topoi $\mathbf{E}$; it follows that $\mathbf{Sh}(\mathbf{E})$ is also a classifying pretopos for T, and that g_T fits the universal property. It is now straightforward to verify that $T \iff g_T$ is the sought correspondence. $\square$

5 Presentations and representations of Grothendieck topoi

In this section we use some of the methods developed to establish a number of non-trivial results in the theory of Grothendieck topoi. We prove a number of results regarding existence of particular underlying sites of certain classes of Grothendieck topoi (which may be thought of "presentations" of GTs), as well as results about embedding Grothendieck topoi in richer ones ("representations" of GTs) such as categories of presheaves.

5.1 Size restrictions on sites and Grothendieck topoi

The following definition and proposition are needed to establish Proposition 5.6.

Definition 5.1. Let μ be an infinite regular cardinal, and let $\mathbf{E}$ be a Grothendieck topos. An object E of $\mathbf{E}$ is called **μ-compact** if every covering family of E has a refinement comprising fewer than μ morphisms. The object E is called **μ-coherent** if it is μ-compact and for any pullback diagram:

$$\begin{array}{ccc} E_1 & \longrightarrow & E_2 \\ \downarrow & & \downarrow \\ E_3 & \longrightarrow & E \end{array}$$

in which E_2 and E_3 are μ-compact, E_1 is also μ-compact.
The full subcategory of $\mathbf{E}$ whose objects are the μ-coherent objects is denoted $\mathrm{Coh}_\mu(\mathbf{E})$.

The notion of μ-compact is clearly reminiscent of the analogous topological concept of compactness. In fact, as we have seen before, it is a strict generalization: For take the "sheaf of subsets" $\mathcal{P}(X)$ on a topological space (X, τ), that is, on the site $(\tau, \subseteq)$ with the obvious topology. Then each subset X' of X determines a sheaf by $X'(U) = X' \cap U$; it is not hard to see that this sheaf X' is an $\aleph_0$-compact object of $\mathcal{P}(X)$ iff X' is compact in the usual sense.

For suitable sites $(\mathbf{C}, J)$, the category $\mathrm{Coh}_\mu(\mathbf{Sh}(\mathbf{C}, J))$ has a factorization property, as expressed in the following proposition. Moreover, it is a μ-pretopos.

Proposition 5.2 (cf. [MR77, Theorem 9.2.2])**.** *Suppose that $(\mathbf{C}, J)$ is a site whose topology J is generated by covering families of cardinality less than μ, and denote $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$. Then $\mathrm{Coh}_\mu(\mathbf{E})$ is a μ-pretopos. Moreover, the canonical continuous functor $\epsilon : \mathbf{C} \to \mathbf{E}$ factors through $\mathrm{Coh}_\mu(\mathbf{E})$.*

Definition 5.3. Let μ be a cardinal. A site $(\mathbf{C}, J)$ is said to have **size at most μ** if $\mathbf{C}$ has cardinality at most μ, and if there is a collection $J_0 \subseteq J$ of cardinality at most μ that generates J.

If $(\mathbf{C}, J)$ has size at most μ, then $\mathbf{Sh}(\mathbf{C}, J)$ is said to be **μ-presentable**. In the case $\mu = \aleph_0$, we also speak of **separable** sites and Grothendieck topoi.

Definition 5.4. The site $(\mathbf{C}, J)$ is called **regular** if the following conditions hold:

(i) Every cover is **extremal**: If $\{f_i : C_i \to C\}$ is covering, and A is a subobject of C such that each f_i factors through the inclusion $A \hookrightarrow C$, then $A = C$;

(ii) Every morphism $f : D \to C$ factors as $f = iq$, with i a monomorphism, and q a singleton cover.

It is an immediate consequence of this definition that a singleton cover in a regular site is precisely the same as an extremal epimorphism. Also, clause (ii) implies that every morphism $f : D \to C$ has an image $\mathrm{Im}\, f \hookrightarrow C$ in the sense of Definition 2.8. A simple verification establishes that a Grothendieck topos, with the canonical topology, is a regular site. There is a simple characterization of the topology of a regular site in terms of these extremal epis and covers by monos.

Lemma 5.5. *Let $(\mathbf{C}, J)$ be a regular site. Then the collection of singleton covers (extremal epis) and the collection of covering families of monos together generate the topology J.*

Proof Suppose that $F = \{f_i : C_i \to C\}$ is a J-cover. Since $(\mathbf{C}, J)$ is regular, we decompose every f_i as $j_i q_i$, with $j_i : A_i \to C$ mono, and $q_i : C_i \to A_i$ a singleton cover. Let $S = \{j_i : A_i \to C\}$; we will show that S is covering. It is immediate upon drawing the diagram that $f_i^* j_i = \mathrm{id}_{C_i}$. Thus for each $f_i \in F$,

$f_i^*(S)$ contains id_{C_i}, hence is covering. By the transitivity axiom, S is covering for C. Conversely, if S is covering, then $q_i \in j_i^*(F)$ for each i, whence by transitivity F is covering. Thus the arbitrary cover F may be generated by singleton covers and covers by monomorphisms. □

The site of definition of a μ-presentable Grothendieck topos may be taken to be regular:

Proposition 5.6. *Let* $\mathbf{E}$ *be a* μ*-presentable Grothendieck topos. Then there is a regular site of definition* $\mathbf{C}$ *for* $\mathbf{E}$ *that has size at most* μ.

Proof Since $\mathbf{E}$ is μ-presentable, let $(\mathbf{C}_0, J)$ be a defining site for $\mathbf{E}$ of size at most μ. We construct $\mathbf{C}$ inductively as a subcategory of Coh_μ. Let $\mathbf{C}_1 = \epsilon\mathbf{C}_0$. Now inductively, for n odd, we define $\mathbf{C}_{n+1}$ by adding for each $f : C_n \to C'_n$ an extremal epi/mono factorization $f = iq$, $q : C_n \to D_{n+1}, i : D_{n+1} \to C'_n$ with $D_{n+1} \in \mathrm{ob}\,\mathrm{Coh}_\mu(\mathbf{E})$ (these factorizations exist by virtue of Proposition 5.2). For n even, we take $\mathbf{C}_{n+1}$ to be the finite-limit closure of $\mathbf{C}_n$ in $\mathrm{Coh}_\mu(\mathbf{E})$. Finally, we define $\mathbf{C} = \bigcup_n \mathbf{C}_n$.

It is now clear that; $\epsilon : \mathbf{C}_0 \to \mathrm{Coh}_\mu(\mathbf{E})$ factors through $\mathbf{C}$; that $\mathbf{C}$ has finite limits; that each morphism of $\mathbf{C}$ has an extremal epi/mono factorization. The topology K on $\mathbf{C}$ we take to be generated by the ϵ-images of a basis J_0 for the topology J on $\mathbf{C}_0$, together with the extremal epis; it is clear that $(\mathbf{C}, K)$ is of size at most μ. Then we have continuous functors $\mathbf{C}_0 \to \mathbf{C} \to \mathbf{E}$, which hence induce morphisms $\mathbf{E} \to \mathbf{Sh}(\mathbf{C}, K) \to \mathbf{Sh}(\mathbf{E})$. But by Lemma 4.6, $\mathbf{E} \simeq \mathbf{Sh}(\mathbf{E})$; hence $\mathbf{Sh}(\mathbf{C}, K) \simeq \mathbf{E}$, so $(\mathbf{C}, K)$ is a defining site for $\mathbf{E}$.

Lastly, it remains to verify that $(\mathbf{C}, K)$ is regular; since we know that there are extremal epi/mono factorizations, it suffices to show that any cover is extremal. By continuity, any K-cover is also a cover in $\mathbf{E}$; because the inclusion is faithful, it follows that K-covers are extremal. □

We will make extensive use of this proposition, together with Lemma 5.5. The combination of these two makes reasoning about arbitrary covers in the underlying site much simpler (and, of course, we can obtain stronger results). Hence, most of what follows will deal with μ-presentable topoi.

5.2 Prime-generated and atomic sites

The prime-generated and atomic sites form another cornerstone of the results we will obtain, especially in combination with regularity. Their definition employs the notions of *prime* and *atom*. In this section, we fix a site $(\mathbf{C}, J)$ on a category $\mathbf{C}$ with finite limits.

Definition 5.7. Let C be an object of $\mathbf{C}$. Then:

(i) C is called **empty** if the empty sieve covers C;

(ii) C is called an **atom** if it is non-empty, and any morphism $f : D \to C$ such that D is not empty is a cover of C;

(iii) C is called a **prime** if every cover $\{f_i : C_i \to C\}$ has a refinement by a singleton cover $f : D \to C$,

The site $(\mathbf{C}, J)$ is called **atomic**, respectively **prime-generated**, if every object C has a cover $\{C_i \to C\}$ with each C_i an atom, respectively a prime. More generally, a collection $\mathcal{C}$ of objects of $\mathbf{C}$ is called **dense** if each $C \in \mathrm{ob}\,C$ has a cover $\{C_i \to C\}$ with each $C_i \in \mathcal{C}$.
The notions of **atomic** and **prime-generated** Grothendieck topoi are obtained by considering a GT as a site, with its canonical topology.

It is obvious that every atom is a prime. Also, covering morphisms take atoms to atoms, and primes to primes:

Lemma 5.8. *Let* $f : D \to C$ *be a singleton cover. If* D *is an atom, respectively a prime, then so is* C.

Proof Suppose D is an atom, and suppose that $g : C' \to C$ is a morphism with C' non-empty. By stability, $g^*f : D \times_C C' \to C'$ is covering, and if $D \times_C C'$ were empty, then so would C' be, by transitivity. Then $f^*g : D \times_C C' \to D$ is covering, since D is an atom and $D \times_C C'$ is non-empty. Again by transitivity, we conclude that g is covering, and C is an atom.
Now suppose D is prime, and that $\{g_i : C_i \to C\}$ is a cover for C. By stability, $\{f^*g_i : D \times_C C_i \to D\}$ covers D; suppose $f^*g_i : D \times_C C_i \to D$ is the singleton refinement. Then $f(f^*g_i) : D \times_C C_i \to C$ covers C, and since $f(f^*g_i) = g_i(g_i^*f)$ is contained in the original cover, it follows that C is also prime. □

If $(\mathbf{C}, J)$ is regular, the canonical continuous functor $\epsilon : (\mathbf{C}, J) \to \mathbf{Sh}(\mathbf{C}, J)$ is well-behaved with respect to primes and atoms, as is made precise by the following results. In an effort to make the statements less convoluted, we will use parentheses to reduce inessential case distinctions between "prime" and "atom".

Lemma 5.9. *Let $\epsilon : (\mathbf{C}, J) \to \mathbf{Sh}(\mathbf{C}, J)$ be the canonical continuous functor. Let C be an object of $\mathbf{C}$. Then C is a prime (an atom) iff ϵC is a prime (an atom).*

Proof Suppose C is a prime, and let $\{E_i \to \epsilon C\}$ be a cover of ϵC. Let $\{\epsilon C_{ij} \to E_i\}$ cover E_i; then by transitivity, $\{\epsilon C_{ij} \to \epsilon C\}$ is also a cover. By Proposition 4.5, we infer that $\{C_{ij} \to C\}$ covers C; since C is prime, find a singleton cover $C_{ij} \to C$; then by continuity of ϵ, $\epsilon C_{ij} \to \epsilon C$ covers C, and C is a prime.

Conversely assume ϵC is a prime, and $\{C_i \to C\}$ covers C. Since ϵC is a prime, the cover $\{\epsilon C_i \to \epsilon C\}$ has a refinement to a singleton cover $\epsilon C_i \to \epsilon C$; again by Proposition 4.5, this means $C_i \to C$ is a singleton cover of C, as desired.

For atoms, Proposition 4.5 can be employed in a similar, even easier way. $\square$

Proposition 5.10. *Let $(\mathbf{C}, J)$ be regular. Then $(\mathbf{C}, J)$ is prime-generated (atomic), if and only if $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$ is prime-generated (atomic).*

Moreover, if this is the case, the primes (atoms) of $\mathbf{E}$ are precisely those E such that there exist a prime (atom) C of $\mathbf{C}$ and a singleton cover $\epsilon C \to E$.

Proof Suppose $\mathbf{E}$ is prime-generated (atomic). Let C be an object of $\mathbf{C}$. By Lemma 5.5, we find a cover $\{E_i \rightarrowtail \epsilon C\}$ of monomorphisms, with each E_i a prime (an atom). By Lemma 4.4, there exist covers $\{\epsilon C_{ij} \to E_i\}$ for all i. Since E_i is a prime, find a C_{ij} such that $\epsilon C_{ij} \to E_i$ is covering; denote it with C_i.

Since ϵ is full (which is just a combination of the Yoneda Lemma and Theorem 1.24) the composite $\epsilon C_i \to E_i \rightarrowtail \epsilon C$ is of the form ϵg_i with $g_i : C_i \to C$. Because $\mathbf{C}$ is regular, we can split g_i as $m_i q_i$ with $q_i : C_i \to C'_i$ a singleton cover and $m_i : C'_i \rightarrowtail C$ a mono. This yields another factorization of ϵg_i, as $(\epsilon m_i)(\epsilon q_i)$. A third factorization is given by $\epsilon C_i \to \mathrm{Im}\,\epsilon g_i \rightarrowtail \epsilon C$. These three factorizations fit into the following diagram:

$$\begin{tikzcd}
 & & \epsilon C'_i \arrow[dr, tail] & \\
\epsilon C_i \arrow[r] \arrow[urr, bend left, two heads] \arrow[drr, bend right, two heads] & \mathrm{Im}\,\epsilon g_i \arrow[ur, tail] \arrow[rr, tail] \arrow[dr] & & \epsilon C \\
 & & E_i \arrow[ur, tail] &
\end{tikzcd}$$

and because both of the curved arrows are extremal epis, it follows that $\epsilon C'_i \cong \mathrm{Im}\,\epsilon g_i \cong E_i$. Since E_i is a prime (an atom), we infer by Lemma 5.9 that C'_i is a prime (an atom); finally, by Lemma 5.8, C_i is a prime (an atom) as well, and we conclude that $\{C_i \to C\}$ is a cover of C by primes (atoms), so that $(\mathbf{C}, J)$ is prime-generated (atomic).

Conversely, suppose $(\mathbf{C}, J)$ is prime-generated (atomic). Then by Lemma 4.4, each object E of $\mathbf{E}$ has a cover $\{\epsilon C_i \to E\}$ by ϵC_is; since $(\mathbf{C}, J)$ is prime-generated (atomic), each of these C_i has a cover $\{C_{ij} \to C_i\}$ by primes (atoms); it follows that the composites $\{\epsilon C_{ij} \to C_i \to E\}$ constitute a cover of E by primes (atoms). Moreover, if E is a prime (an atom), this cover has a refinement by a singleton cover $C_{ij} \to E$. This concludes the proof. $\square$

Combining this result with Proposition 5.6, we obtain:

Corollary 5.11. *Let $\mathbf{E}$ be a μ-presentable prime-generated (atomic) Grothendieck topos. Then $\mathbf{E}$ has a site of definition $\mathbf{C}$ of size at most μ, which is regular and prime-generated (atomic).*

This concludes our account of the technical results that the remainder of this chapter builds upon. Now, we present a slightly different take at the models studied in earlier chapters.

5.3 Continuous models

The concept of join in subobject lattices can be altered to be compatible with sites, more precisely, with Grothendieck topologies. This change (which by Proposition 5.13 does not affect Grothendieck topoi) allows us to define "models" on sites.

Definition 5.12. Let $(\mathbf{C}, J)$ be a site. For subobjects $A_0, (A_i)_{i\in I}$ of some object A, A_0 is the ***J*-join**, or more generally **topological join**, of the A_i, denoted $A_0 = \bigvee_{i\in I}^{(J)} A_i$, if the family $\{A_i \hookrightarrow A_0\}$ is a cover of A_0.

Without further information on the topology J, we cannot prove much about this new operator $\bigvee^{(J)}$, not even its uniqueness. However, for the subcanonical topologies, we can say something:

Proposition 5.13. *Let $(\mathbf{C}, J)$ be a site. Suppose that J is subcanonical. Then every J-join is also a join of subobjects.*

For a Grothendieck topos $\mathbf{E} = \mathbf{Sh}(\mathbf{C}, J)$ with the canonical topology, the converse holds, i.e. every join of subobjects is also a J-join.

Proof Let $\bigvee_{i\in I}^{(J)} A_i = A_0$ for subobjects of an object C of $\mathbf{C}$. Then since J is subcanonical, for each D and morphisms $f_i : A_i \to D$ such that $f_i = f_j \iota_{ji}$ for each $\iota_{ji} : A_i \to A_j$ (that is, each matching family for $\{A_i \hookrightarrow A_0 : i \in I\}$ by $\mathbf{y}_D$), there is a unique amalgamation $f : A_0 \to D$ such that $f_i = f\iota_{0i}$. It follows that A_0 satisfies the universal property for $\varprojlim_{i\in I} A_i$, and it is clear that this implies $A_0 = \bigvee_{i\in I} A_i$.

Suppose $\bigvee_{i\in I} A_i = A_0$ for subobjects of $E \in \operatorname{ob}\mathbf{E}$. Suppose that $\eta_i : A_i \to F$ form a matching family for $\mathbf{y}_F$; we need to show that there is a unique $\eta : A_0 \to F$ that amalgamates these η_i. So let $C \in \operatorname{ob}\mathbf{C}$ and $\xi \in A_0C$. We need to determine $\vartheta \in FC$ uniquely on the basis of the η_i. By Lemma 3.12(ii), we find a covering sieve $\{f : D \to C\} \in JC$ such that $\xi \cdot f \in \bigcup_i A_iD$ for each f. Since the η_i are matching, each $\xi \cdot f$ determines a unique $\vartheta_f = (\eta_i)_D(\xi \cdot f) \in FD$. Since F is a sheaf and $\{f : D \to C\}$ covers, these ϑ_f amalgamate to a unique $\vartheta \in FC$. The definition $\eta_C(\xi) = \vartheta$ now obviously defines a natural transformation $A_0 \to F$, as desired. Since F was arbitrary, it follows that the $A_i \hookrightarrow A_0$ together form an effective epimorphic family; by Proposition 3.9, this family is also stable. In conclusion, $\bigvee_{i\in I}^{(J)} A_i = A_0$, as desired. □

From here on, we will primarily deal with the standard model-theoretic situation, which comes down to considering functors to **Set** instead of arbitrary Grothendieck topoi. We also introduce a new kind of model which takes the topology of a site into account.

Definition 5.14. Let $(\mathbf{C}, J)$ be a site. A continuous functor $M : \mathbf{C} \to \mathbf{Set}$ will be called a **continuous model** of $\mathbf{C}$. The category of continuous models of $(\mathbf{C}, J)$ is denoted $\mathbf{cMod}(\mathbf{C}, J)$.

A **continuous premodel** of $\mathbf{C}$ is a functor $M : \mathbf{C} \to \mathbf{Set}$ that preserves finite limits and topological joins (considering **Set** with its canonical topology).

Hereafter we will often drop the "continuous" prefix and refer to models and premodels instead, when this causes no ambiguity. We start with investigating models of two specific kinds.

Lemma 5.15. *Let $(\mathbf{C}, J)$ be a site, and let J_0 be a basis for the topology. Let $D : \mathbf{I} \to \mathbf{C}$ be a filtered diagram, and let $M : \mathbf{C} \to \mathbf{Set}$ be the functor represented by D, i.e.:*

$$MC = \varinjlim_{I\in\mathbf{I}} \operatorname{Hom}_{\mathbf{C}}(DI, C)$$

Then M is a model iff for every J_0-sieve $\{C_i \to C\}$ and $f : DI \to C$, there exists an i such that for some $I' \to I \in \operatorname{mor}\mathbf{I}$ and $DI' \to C_i$, the following square commutes:

$$\begin{tikzcd} DI' \arrow[r] \arrow[d] & DI \arrow[d, "f"] \\ C_i \arrow[r] & C \end{tikzcd}$$

Proof By Theorem A.4, M preserves finite limits. Thus it is to be shown that the condition mentioned is equivalent to M being continuous.

Suppose first that M is a model, and that $\{C_i \to C\}$ is a J_0-sieve. Then $\{MC_i \to MC\}$ is a surjective family; so for $[g] \in MC$ (where square brackets indicate the equivalence class, which arises from the colimit) there exists $f_i : C_i \to C$ and $h : DI' \to C_i$ such that $[g] = [f_i h] = Mf_i[h]$. The definition of M as a colimit now immediately implies the existence of a morphism $I' \to I$ with the sought properties.

For the converse, note that the steps in the above are all reversible, so that we find for J_0-sieves $\{C_i \to C\}$ that $\{MC_i \to MC\}$ is a surjective family; that is, M preserves J_0-joins. Since J_0 is a basis for J, we conclude that M is a model. □

Lemma 5.16. *Let $(\mathbf{C}, J)$ be a regular, prime-generated site. A continuous premodel $M : \mathbf{C} \to \mathbf{Set}$ of $\mathbf{C}$ is a model iff for every covering epi $f : D \to C$ between primes and $c \in MC$, there is a $d \in MD$ such that $Mf(d) = c$.*

Proof By Lemma 5.5, the topology J is generated by covers of monos (that is, J-joins) and singleton covers; by inspecting the proof and recalling Lemma 5.8, we see that only singleton covers between primes are used.

Since M is already a continuous premodel, it preserves topological joins, and therefore, the covers of monos. It remains to verify that M preserves the singleton covers as well. In **Set**, a singleton is a cover iff it is a surjective mapping. But this is precisely what the condition on M expresses. Hence M is a model. □

An important piece of the model-theoretic view is the notion of a type.

Definition 5.17. Let $M : \mathbf{C} \to \mathbf{Set}$ be a premodel, and let $c \in MC$. The **type** of c, denoted $\mathbf{t}^M_C(c)$ or simply $\mathbf{t}(c)$, is the set of subobjects A of C such that $c \in MA$. The fact that M preserves products allows for a natural extension of the concept to tuples, for example $\mathbf{t}(c, c')$.

This definition is to be understood in the context of the interpretation of formulae as subobjects from Chapter 2 and also of the Kripke-Joyal semantics from Chapter 3; it follows that this definition is the natural analogue of the classical model-theoretical definition $\mathbf{t}(c) = \{\phi(x) \mid M \models \phi(c)\}$.

It is clear that $\mathbf{t}(c)$ is upper closed. Since M preserves finite limits, we conclude that $M(A \wedge B) = MA \cap MB$ for subobjects A, B; therefore, if $A, B \in \mathbf{t}(c)$ then so is $A \wedge B$. That is, $\mathbf{t}(c)$ is a filter in $\mathrm{Sub}(C)$.

Moreover, if M is a model, then the fact that M preserves J-joins implies that $\mathbf{t}(c)$ is **J-prime**, i.e. $\bigvee^{(J)}_{i\in I} A_i \in \mathbf{t}(c)$ implies that for some i, $A_i \in \mathbf{t}(c)$.

In the special case of a separable site $(\mathbf{C}, J)$ and $C \in \mathrm{ob}\,\mathbf{C}$, any J-prime filter in $\mathrm{Sub}(C)$ arises as a type.

Proposition 5.18. *Let $(\mathbf{C}, J)$ be a regular separable site, C be an object of $\mathbf{C}$ and suppose $\mathbf{t}$ is a J-prime filter in $\mathrm{Sub}(C)$. Then there exist a countable continuous model M and $c \in MC$ such that $\mathbf{t} = \mathbf{t}^M_C(c)$.*

Proof Let J_0 be a countable basis for J, and let $(A_k)_k$ be an enumeration of $\mathbf{t}$. There are countably many triples $(m, D \to C, \{C_i \to C\})$ with $m \in \mathbb{N}$, $D \to C \in \mathrm{mor}\,\mathbf{C}$, and $\{C_i \to C\}$ a J_0-cover of C, say $\tau_1, \tau_2, \ldots$. If we concatenate the initial segments of this sequence of lengths $1, 2, \ldots$, we get a sequence in which each τ_k occurs infinitely often; hence, without loss of generality $(\tau_k)_k$ contains all triples infinitely often. With this setup, we will now construct a filtered diagram, and use Lemma 5.15 to establish that the functor it represents is a model. This model will have $\mathbf{t}$ as a type.

We construct a filtered diagram of type ω^{op}, that is, a sequence $D_0 \leftarrow D_1 \leftarrow \ldots$.
Let $D_0 \to C$ be $\mathrm{id}_C : C \to C$. We will ensure that $\mathrm{Im}\, D_n \hookrightarrow C$ is in $\mathbf{t}$. For the inductive step, we distinguish between even and odd n:

- $n = 2k$ is even: Consider $\tau_k = (m, D \to C, \{C_i \to C\})$. If $m > n$ or $D \neq D_m$, put $D_{n+1} = D_n$ and let $D_{n+1} \to D_n$ be id_{D_n}.

If on the other hand $m \le n$ and $D = D_m$, then we have a composite mapping $f_n : D_n \to D_m \to D$. For each $C_i \to C$ of the cover in τ_k, let $D_{n,i} \to D_n$ be the pullback $f_n^*(C_i \to C)$. Then $\{D_{n,i} \to D_n\}$ covers D_n, and by regularity of $\mathbf{C}$, we have the J-join $\operatorname{Im} D_n = \bigvee_{i\in I}^{(J)} \operatorname{Im} D_{n,i}$ of subobjects of C. Because $\mathbf{t}$ is J-prime, we find a $D_{n,i}$ such that $\operatorname{Im} D_{n,i} \in \mathbf{t}$. We define $D_{n+1} \to D_n$ as $D_{n,i} \to D_n$.

- $n = 2k + 1$ is odd: Consider A_k, the kth subobject in $\mathbf{t}$, and define $D_{n+1} = D_n \times S_k$. Then $\operatorname{Im} D_{n+1} = \operatorname{Im} D_n \wedge S_k \in \mathbf{t}$ as $\mathbf{t}$ is a filter. The morphism $D_{n+1} \to D_n$ is taken to be the projection.

Define M to be the functor represented by D, as in Lemma 5.15; we will show that it satisfies the conditions mentioned there. For, given a cover $S = \{C_i \to C\}$, an $m \in \mathbb{N}$ and a morphism $f_m : D_m \to C$, we find an $n = 2k \ge m$ such that $\tau_k = (m, f_m, S)$. Then $D_{n+1} \to D_n$ was constructed precisely to satisfy the condition of Lemma 5.15. That is, M is a model of $\mathbf{C}$.

By construction, we have that id_C (or rather, its equivalence class $[\mathrm{id}_C]$) is in MC. We will show that $\mathbf{t}_C^M(\mathrm{id}_C) = \mathbf{t}$. Suppose that $A \in \operatorname{Sub} C$, and that $[\mathrm{id}_C] \in MA \subseteq MC$; then there must be a D_n such that $D_n \to D_0 = C$ factors through A, i.e. $\operatorname{Im} D_n \le A$. Since $\operatorname{Im} D_n \in \mathbf{t}$ and $\mathbf{t}$ is a filter, it follows that $A \in \mathbf{t}$.

Conversely, suppose that $A \in \mathbf{t}$. Then $A = A_k$ for some k; by the construction of $(D_n)_n$, $\operatorname{Im} D_{2k+1} \le A$, and so $D_n \to D_0$ factors through A, i.e. $A \in \mathbf{t}_C^M(\mathrm{id}_C)$. □

Recall that a filter F is called **principal** if there is an $f \in F$ such that $f' \in F$ iff $f \le f'$.

Definition 5.19. Let M be a premodel, and let $c \in MC$. Then c is called **principal** if $\mathbf{t}_C^M(c)$ is a principal filter. In the special case that $\mathbf{t}_C^M(c)$ contains only the maximal subobject id_C, c is called **generic** for C.

We say M is **principal** if for all objects C of $\mathbf{C}$, each $c \in MC$ is principal.

It is clear that if $c \in MC$ is principal, and its filter is generated by $A \hookrightarrow C$, then c is generic for A.

Lemma 5.20. *Let M, M' be premodels, and suppose that $(c, d) \in MC \times MD$, $(c', d') \in M'C \times M'D$ are such that $\mathbf{t}(c, d) \subseteq \mathbf{t}(c', d')$. Then $\mathbf{t}(c) \subseteq \mathbf{t}(c')$.*

If M is principal, the converse holds: If $\mathbf{t}(c) \subseteq \mathbf{t}(c')$ then for all $d \in MD$ there is a $d' \in M'D$ such that $\mathbf{t}(c, d) \subseteq \mathbf{t}(c', d')$.

Proof For any $A \hookrightarrow C$, if $c \in MA$, then $A \times D \in \mathbf{t}(c, d)$. Hence $A \times D \in \mathbf{t}(c', d')$, and it follows that $c' \in M'A$; hence $\mathbf{t}(c) \subseteq \mathbf{t}(d)$.

For the converse, let $A \hookrightarrow C$ be the generator of $\mathbf{t}(c)$, and $i : S \hookrightarrow C \times D$ the generator of $\mathbf{t}(c, d)$; then $\pi i : S \hookrightarrow C \times D \to C$ is in $\mathbf{t}(c)$; moreover, since $S \le A \times D \hookrightarrow C \times D$, it follows that $\pi i S \le A$ and since A generates $\mathbf{t}(c)$, $\pi i : S \to A$ is an isomorphism. In particular, it is a singleton cover, and so $M'S \to M'A$ is surjective. Now let $s' \in M'S$ be such that $M\pi i(s') = a'$; then $s' = (c', d')$ for some $d' \in M'D$. Since $S \in \mathbf{t}(c', d')$, it follows immediately that $\mathbf{t}(c, d) \subseteq \mathbf{t}(c', d')$. □

We can recognize models of regular sites that are representing a diagram by properties of those diagrams.

Proposition 5.21. *Let $(\mathbf{C}, J)$ be a regular site. Let $D : \mathbf{I} \to \mathbf{C}$ be a filtered diagram, and let $M : \mathbf{C} \to$ **Set** be the corresponding functor. Then:*

(i) *If DI is prime for each I, M is a premodel;*

(ii) *If additionally $D\iota : DI' \to DI$ is a singleton cover for every $\iota : I' \to I$ in $\mathbf{I}$, M is principal;*

(iii) *If (i) holds, and for each $k : DI \to C$ and singleton cover $f : B \to C$, there are $\iota : I' \to I$ in $\mathbf{I}$ and $l : DI' \to B$ such that:*

$$\begin{array}{ccc} DI' & \xrightarrow{D\iota} & DI \\ {\scriptstyle l}\downarrow & & \downarrow{\scriptstyle k} \\ B & \xrightarrow[f]{} & C \end{array}$$

commutes, M is a model.

Proof (i) Since M preserves finite limits, it suffices to prove that it preserves topological joins. Now if $f : DI \to A$ and $A = \bigvee_i^{(J)} A_i$, then by the stability axiom for topologies, $DI = \bigvee_i^{(J)} f^* A_i$; since DI is prime, there is a singleton cover $f^* A_i \to DI$. But since $A_i \rightarrowtail A$ is mono, $f^* A_i \simeq DI$. Hence, $DI \to A$ factors through an A_i; thus any $f \in \mathrm{Hom}(DI, A)$ corresponds to some $f_i \in \mathrm{Hom}(DI, A_i)$, and passing to the colimit we find that $\{MA_i \to MA\}$ is a surjective family, as desired.
(ii) Let $c \in MC$, and let $f : DI \to C$ represent c. Let $f : DI \xrightarrow{e} A \hookrightarrow C$ be the singleton cover/mono factorization of f. Then clearly $c \in MA$. Conversely if $c \in MA'$ for a subobject $i_{A'} : A' \hookrightarrow C$, then there must exist some DI', $g : DI' \to A'$ and $\iota : I' \to I$ such that $fD\iota = i_{A'}g$; since $D\iota$ is a singleton cover, $i_A(eD\iota)$ is a singleton cover/mono factorization of $i_{A'}g$. It follows that $A \leq A'$; hence $\mathbf{t}(c)$ is generated by A, and since c and C were arbitrary, M is principal.
(iii) This is a combination of part (i), Lemma 5.15 and Lemma 5.5. $\square$

5.4 Models for $\aleph_1$-presented sites

The aim of this section is to prove the following result.

Theorem 5.22. *Let $(\mathbf{C}, J)$ be a regular prime-generated site of size at most $\aleph_1$. If the empty sieve is not a J-cover for the terminal object 1, then $(\mathbf{C}, J)$ has a model.*

Note that the condition that the empty sieve does not cover 1 is very weak; indeed, if $\varnothing$ does cover 1, then by stability, it covers *every* object of $\mathbf{C}$. This is of course a rather pathological situation.

Definition 5.23. Let $M : \mathbf{C} \to \mathbf{Set}$ preserve finite limits. A set $\{(C_i, c_i \in MC_i) : i \in I\}$ **generates** M if for every proper subfunctor M', there is an i such that $c_i \notin M'C_i$.

For a cardinal μ, we call M **μ-generated** if there is a generating set for M of cardinality at most μ. M is called **μ-presented** if additionally for each $c \in MC$, $\mathbf{t}_C^M(c)$ is generated (as a filter) by a set of cardinality at most μ.

There are several technical results about μ-generated and -presented functors that we need to state first.

Lemma 5.24. *Let $\mathbf{C}$ be a category with finite limits and images, and let $M : \mathbf{C} \to \mathbf{Set}$ preserve finite limits. Then:*

(i) *If $M = \mathrm{Hom}(C, -)$ for some $C \in \mathrm{ob}\,\mathbf{C}$, then M is $\aleph_0$-presented;*

(ii) *If $M = \varinjlim_i M_i$ for μ-generated (μ-presented) M_i, then M is also μ-generated (μ-presented);*

(iii) *If M is μ-presented, then for each dense collection $\mathcal{C}$ of objects of $\mathbf{C}$, there exists a diagram $(C_i)_i$ with objects in $\mathcal{C}$ such that $M = \varinjlim_i \mathrm{Hom}(C_i, -)$;*

(iv) *Moreover, if M is $\aleph_0$-presented, then there is a subdiagram $(C_n)_n$ of $(C_i)_i$ of type ω such that $M = \varinjlim_n \mathrm{Hom}(C_n, -)$.*

Proof Only the proof of (iii) is not straightforward; we refer the reader to Proposition 1.4 of [BM87]. $\square$

Lemma 5.25. *Let M be an $\aleph_0$-presented premodel, and let $f : D \to C$ be a singleton cover. For any natural transformation $\eta : \mathrm{Hom}(C, -) \to M$, there exists an $\aleph_0$-presented premodel M' and a mono $i : M \rightarrowtail M'$ such that η can be extended (along i) to a natural transformation $\eta' : \mathrm{Hom}(D, -) \to M'$:*

$$\begin{tikzcd} \mathrm{Hom}(C,-) \arrow[r, tail, "\mathrm{Hom}(f,-)"] \arrow[d, "\eta"] & \mathrm{Hom}(D,-) \arrow[d, "\eta'"] \\ M \arrow[r, tail, "i"'] & M' \end{tikzcd}$$

Proof Since M is a premodel, it is a model for the site $(\mathbf{C}, J')$ where J' is generated by the covers by monos of J. Due to Lemma 5.24, we have that $M = \varinjlim_n \mathrm{Hom}(P_n, -)$ for a diagram $(P_n)_n$ of type ω in $\mathbf{C}$, whose objects are primes. Hence, $\eta : \mathrm{Hom}(C, -) \to M$ is seen to be determined by some $p : P_n \to C$ which satisfies $[p] = \eta_C(\mathrm{id}_C)$.

Since $f : D \to C$ is a singleton cover, so is $p^*f : P_n \times_C D \to P_n$. Because the topology J is prime-generated, we find a prime subobject P'_n of $P_n \times_C D$ such that $P'_n \hookrightarrow P_n \times_C D \to P_n$ is a singleton cover. Iterating the construction with $P'_n \to P_n$, we build a chain

$$D \leftarrow P'_n \leftarrow P'_{n+1} \leftarrow \dots$$

such that $P'_{n+k} \to P_{n+k}$ is a singleton cover. Now we define M' as the functor represented by the diagram:

$$P_0 \leftarrow \dots \leftarrow P_n \leftarrow P'_n \leftarrow P'_{n+1} \leftarrow \dots$$

Then η' is induced by the morphism $P'_n = P_n \times_C D \to D$; since the $P'_{n+k} \to P_{n+k}$ are singleton covers, it follows that they induce a mono $i : M \rightarrowtail M'$. By Lemma 5.24, M' is $\aleph_0$-presented; it readily follows from applying Lemma 5.15 to the topology J' that M' is again a premodel. $\square$

We are now ready to prove Theorem 5.22.

Proof of Theorem 5.22 Denote with $M \subseteq M'$ that M is a subfunctor of M'; we will construct a $\subseteq$-chain $(M_\alpha)_{\alpha<\omega_1}$ of length ω_1 of $\aleph_0$-presentable premodels. The colimit $M_{\omega_1} = \varinjlim_\alpha M_\alpha$ will be the sought model. Effectively, we only seek a method to ensure that M_{ω_1} takes every singleton cover to a surjection. We will ensure this by a technical construction, effectively adding the necessary elements one by one.

For an $\aleph_0$-presentable premodel M, let ΦM be the set of pairs $(f : D \to C, c \in MC)$ with f a singleton cover. Since $\mathbf{C}$ is of size at most $\aleph_1$, it is clear that $|\Phi M| \leq \aleph_1$. Also, $M \subseteq M'$ is seen to imply $\Phi M \subseteq \Phi M'$. As a further ingredient, let $\omega_1 = \bigcup_{\alpha<\omega_1} X_\alpha$, with $|X_\alpha| = \aleph_1$ for all α.

In addition to the M_α, we will construct a sequence of $g_\alpha : \bigcup_{\beta\leq\alpha} X_\alpha \to \Phi M_\alpha$, such that g_α extends g_β for $\alpha > \beta$, and such that the preimage of each $(f, c) \in \Phi M_\alpha$ has cardinality $\aleph_1$.

Now, finally, we start the induction. Let M_0 be the functor defined on objects by:

$$M_0 C = \begin{cases} \varnothing & \text{if the empty sieve covers } C \\ \{*\} & \text{otherwise;} \end{cases}$$

while M_0 is already uniquely determined on morphisms. Then M_0 is an $\aleph_0$-presentable premodel, and ΦM_0 is seen to be the collection of singleton covers between non-empty objects of $(\mathbf{C}, J)$. Note that ΦM_0 is non-empty since the empty sieve does not cover 1. Let $g_0 : X_0 \to \Phi M_0$ be any suitable mapping.

For the successor step, suppose M_α and g_α are given.
If $\alpha \notin \operatorname{dom} g_\alpha$, let $M_{\alpha+1} = M_\alpha$ and extend g_α by any mapping $X_{\alpha+1} \to \Phi M_\alpha$.
If on the other hand $\alpha \in \operatorname{dom} g_\alpha$, let $g_\alpha(\alpha) = (f, c)$ for some $f : D \to C$ and $c \in M_\alpha C$. Then c induces a natural transformation $\operatorname{Hom}(C, -) \to M_\alpha$ by $(h : C \to C') \mapsto M_\alpha h(a)$. By Lemma 5.25, we find an extension $M_\alpha \subseteq M_{\alpha+1}$ (this extension will make Lemma 5.16 apply to the colimit). Define $g_{\alpha+1}$ by extending g_α with a suitable bijection $X_{\alpha+1} \to \Phi M_{\alpha+1} \setminus \Phi M_\alpha$.

For the limit step, we let $M_\alpha = \varinjlim_{\beta<\alpha} M_\beta = \bigcup_{\beta<\alpha} M_\beta$, and g_α is uniquely determined by the condition that it extends all g_β. It is obvious that M_α is again an $\aleph_0$-presented premodel.

Now let $M_{\omega_1} = \varinjlim_{\alpha<\omega_1} M_\alpha$. It is clear that M_{ω_1} is a premodel. We will verify the condition of Lemma 5.16. So suppose that $f : D \to C$ is a singleton cover and $c \in M_{\omega_1} C$. Then $c \in M_\alpha C$ for some $\alpha < \omega_1$, so $(f, c) \in \Phi M_\alpha$. Since $g_\alpha^{-1}(f, c) \subseteq \bigcup_{\beta\leq\alpha} X_\alpha$ is uncountable, there is some $\alpha' > \alpha$ such that $g_\alpha(\alpha') = (f, c)$. By construction of the sequence $(g_\alpha)_\alpha$, $g_{\alpha'}(\alpha') = (f, c)$. From the construction of the inclusion $M_{\alpha'} \subseteq M_{\alpha'+1}$ we obtain a commutative diagram as in Lemma 5.25. Now dually to the Yoneda Lemma, the natural transformation $\eta' : \operatorname{Hom}(D, -) \to M_{\alpha'+1}$ is determined by some $d \in M_{\alpha'+1} D$. Chasing $\mathrm{id}_C \in \operatorname{Hom}(C, C)$ through the diagram tells us that $c = M_{\alpha'+1} f(d)$; since c was arbitrary and $M_{\alpha'+1} f(d) = M_{\omega_1} f(d)$, we conclude that M_{ω_1} is a model by Lemma 5.16. $\square$

Now, we wonder: why does this proof not follow through for higher cardinalities? This is because we made explicit use of Lemma 5.24(iv) in Lemma 5.25. Subsequently we relied on the latter Lemma in the proof of the theorem. A strengthening of this approach fails, because when we attempt to generalize Lemma 5.24(iv) to $\aleph_\alpha$, constructing a subdiagram of type ω_α requires that the underlying filtered category $\mathbf{I}$ admits cocones for diagrams of type ω.

5.5 Representation theorems for Grothendieck topoi

In this last section, we use the results we have established, suitably augmented by accounts of *sufficient collections* of models and *prime models*, to give a necessary and sufficient condition for a Grothendieck topos $\mathbf{E}$ to be embeddable in a presheaf category $\hat{\mathbf{C}}$ in a subobject-preserving way (Theorem 5.38).

Definition 5.26. Let $\mathbf{E}$ be a Grothendieck topos, and let $\mathcal{M}$ be a collection of models of $\mathbf{E}$. Then $\mathcal{M}$ is called **sufficient** if the induced continuous functor $\mathbf{E} \to \mathbf{Set}^{\mathcal{M}}$ reflects isomorphisms. Equivalently, $\mathcal{M}$ is sufficient iff for every *non*-cover $\{C_i \to C\}$ in $\mathbf{E}$, there is a model $M \in \mathcal{M}$ such that $\{MC_i \to MC\}$ is not covering either.

For an arbitrary site $(\mathbf{C}, J)$, the two conditions are inequivalent; the definition of sufficiency is extended to this general case by the latter condition (which is stronger). Finally, a category $\mathbf{M}$ of models of $\mathbf{E}$ or $(\mathbf{C}, J)$ is **sufficient** if $\operatorname{ob} \mathbf{M}$ is.

Theorem 5.27. *Let $\mathbf{E}$ be an $\aleph_1$-presentable, prime-generated Grothendieck topos. Then the collection of models of $\mathbf{E}$ is sufficient.*

Proof By Corollary 5.11, there exists a regular prime-generated site $(\mathbf{C}, J)$ of size at most $\aleph_1$ for $\mathbf{E}$. It suffices to establish, for each non-cover $\{C_i \to C\}$, a model M for $(\mathbf{C}, J)$ such that $\{MC_i \to MC\}$ is also a non-cover (i.e., is not surjective).

It readily follows from Proposition A.6 from the Appendix (which discusses a site on the slice category $\mathbf{C}/C$) that we may reduce to the case $C = 1$. In this simpler case, the task at hand is reduced to establishing a model M such that each MC_i is the empty set. To force this, we add to the topology J the empty sieve on each C_i. The resulting topology J' is seen to be defined by the following condition:

$$\{C'_i \to C'\} \in J'C \iff \{C'_i \to C'\} \cup \{C' \times C_i \to C'\} \in JC$$

It follows that each C_i is empty for J', and that $\{C_i \to 1\}$ is not a J'-cover. Now a J-prime P is either J'-empty or a J'-prime; in particular, $(\mathbf{C}, J')$ is still prime-generated. Hence Theorem 5.22 applies, and $(\mathbf{C}, J')$ has a model M; then M is automatically also a model of $(\mathbf{C}, J)$ and hence of $\mathbf{E}$, in which $\{MC_i \to M1\}$ is not surjective (for, all MC_i are empty). $\square$

Definition 5.28. Let M be a continuous model of $(\mathbf{C}, J)$. Then M is called **prime** if for each model M', object C of $\mathbf{C}$ and $c \in MC, c' \in M'C$ such that $\mathbf{t}(c) \subseteq \mathbf{t}(c')$, there is a natural transformation $\alpha : M \to M'$ such that $\alpha_C(c) = c'$.

Theorem 5.29. *Let $(\mathbf{C}, J)$ be a separable site and M be a countable principal model of $\mathbf{C}$. Then M is prime.*

Proof Suppose M' is a model, and that $c \in MC, c' \in M'C$ are such that $\mathbf{t}(c) \subseteq \mathbf{t}(c')$. Enumerate the elements of M as $(c_n)_n$ such that $c_1 = c$ (this can be done as $\operatorname{ob} \mathbf{C}$ as well as each MC is countable). By iterated application of Lemma 5.20, we find a sequence $(c'_n)_n$ of elements of M' such that for each n, $\mathbf{t}(c_1, \ldots, c_n) \subseteq \mathbf{t}(c'_1, \ldots, c'_n)$.

We propose a natural transformation $\alpha : M \to M'$ by $\alpha(c_n) = c'_n$; if now $Mf(c_n) = c_m$ for $f : C_n \to C_m \in \operatorname{mor} \mathbf{C}$, then by a suitable variant of Lemma 5.20 we find $\mathbf{t}(c_n, c_m) \subseteq \mathbf{t}(c'_n, c'_m)$; now observe that the graph of f (see Definition 2.14) is in $\mathbf{t}(c_n, c_m)$, hence also in $\mathbf{t}(c'_n, c'_m)$. But the latter implies that $M'f(c'_n) = c'_m$. Hence $M'f\alpha = \alpha Mf$, i.e. α is a natural transformation, and M is a prime model. $\square$

It is worthwhile to note the similarity of the preceding theorem with the following classical model-theoretic result:

Proposition 5.30 ([Blu12, Lemma E2.3.11])**.** *Every countable atomic model is prime.*

The two are basically the same theorem in different contexts; for, in classical model theory, a structure is **atomic** if each realized type is **isolated**, which means that it consists of the (provable) consequences of a single formula ϕ. The similarity with a principal continuous model is now obvious if we bring to mind the intuition of subobjects corresponding to formulae.

Corollary 5.31. *For any separable, prime-generated site, the collection of principal prime models is sufficient. In particular, this holds for Grothendieck topoi.*

Next, we focus on the special case of categories of presheaves, and establish some conditions under Grothendieck topoi embed in these. Fix a category of presheaves $\hat{\mathbf{C}}$ with $\mathbf{C}$ small and finitely complete.

Lemma 5.32. *For each object C of $\mathbf{C}$, the functor $\mathrm{ev}_C : \hat{\mathbf{C}} \to \mathbf{Set}$, defined by $\mathrm{ev}_C(F) = FC$ and $\mathrm{ev}_C(\eta : F \to G) = \eta_C : FC \to GC$, is a principal model of $\hat{\mathbf{C}}$.*

Proof By applying Corollary 4.8 to $\mathbf{E} = \mathbf{Set}$, we find that $\mathbf{cMod}(\hat{\mathbf{C}})$ and $\mathbf{cMod}(\mathbf{C}, J)$ (where a sieve $S = \{C_i \to C\}$ is a J-cover iff $\mathrm{id}_C \in S$) are equivalent. Since J is evidently subcanonical, it remains to verify that $\mathbf{C} \to \mathbf{Set}, D \mapsto \mathrm{ev}_C\, \mathbf{y}_D = \mathrm{Hom}(C, D)$ is a model. But this is immediate from the definition of J. Hence ev_C is a model.

To see that ev_C is principal, suppose that $f \in \mathrm{ev}_C\, F = FC$. By Proposition 1.10, the type $\mathbf{t}(f)$ is given by the set of presheaves $\mathcal{G} = \{G \in \hat{\mathbf{C}} : f \in GC, \forall D \in \mathrm{ob}\, \mathbf{C} : GD \subseteq FD\}$. But the intersection of presheaves is again a presheaf, and it follows that the presheaf $\bigcap \mathcal{G}$ (defined pointwise) generates $\mathbf{t}(c)$. Hence, ev_C is principal. □

Proposition 5.33. *The evaluation functors $\mathrm{ev}_C : \hat{\mathbf{C}} \to \mathbf{Set}$ are prime models.*

Proof To show that ev_C is prime, fix $F \in \mathrm{ob}\, \hat{\mathbf{C}}$ and $f \in \mathrm{ev}_C\, F = FC$, a model $M : \hat{\mathbf{C}} \to \mathbf{Set}$ and an $m \in MF$, such that $\mathbf{t}(f) \subseteq \mathbf{t}(m)$. We need to find a natural transformation $\alpha : \mathrm{ev}_C \to M$ such that $\alpha_F(f) = m$.

Since ev_C is principal, we may replace F by the subobject $F' \hookrightarrow F$ generating f, so that $\mathbf{t}(f)$ contains only id_F; then $\mathbf{t}(f) \subseteq \mathbf{t}(m)$ is automatic, and it suffices to show that we can define α.

Because $\mathbf{t}(f) = \{\mathrm{id}_F\}$, we have by the Yoneda lemma that $f : \mathbf{y}_C \to F$ is an extremal epi, i.e. a singleton cover. Hence $M\mathbf{y}_C \to MF$ is surjective, and we find an $y \in M\mathbf{y}_C$ such that $Mf(y) = m$. Now we define $\alpha : \mathrm{ev}_C \to M$ by $\alpha_G(g) = Mg(y)$, where we identified $g \in GC$ with $g : \mathbf{y}_C \to G$; it is clear that $\alpha_F(f) = m$, and it only remains to show that α is indeed a natural transformation. So let $\eta : G \to G' \in \mathrm{mor}\, \hat{\mathbf{C}}$; we need to show that:

$$\begin{tikzcd} GC \arrow[r, "\eta_C"] \arrow[d, "\alpha_G"'] & G'C \arrow[d, "\alpha_{G'}"] \\ MG \arrow[r, "M\eta"'] & MG' \end{tikzcd}$$

commutes. So let $g \in GC$. Then:

$$M\eta(\alpha_G(g)) = M\eta(Mg(y)) = M(\eta g)(y) = M(\eta_C(g))(y) = \alpha_{G'}(\eta_C(g))$$

where we used the correspondence of the Yoneda Lemma to conclude $\eta g = \eta_C(g)$. Hence $\alpha : \mathrm{ev}_C \to M$ is a natural transformation, proving that ev_C is a prime model. □

Corollary 5.34. *The collection of principal prime models of $\hat{\mathbf{C}}$ is sufficient.*

This Corollary constitutes one of the building blocks for the theorems we aim for. A last technical result is:

Lemma 5.35. *Let $(\mathbf{C}, J)$ be a regular site. If the collection $\mathcal{P}$ of principal models of $(\mathbf{C}, J)$ is sufficient, then $(\mathbf{C}, J)$ is prime-generated.*

Proof Recall that $\mathcal{P}$ is sufficient iff for every non-cover $\{C_i \to C\}$, there is a model $P \in \mathcal{P}$ such that $\{PC_i \to PC\}$ is not covering either.

Let $P \in \mathcal{P}$ and $c \in PC$. Since P is principal, there is a generator C_c^P for $\mathbf{t}_C^P(c)$. If $\{f_i : C_i \to C_c^P\}$ is a J-cover, then $\{Pf_i : PC_i \to PC_c^P\}$ is surjective, hence for some i we have $c \in \mathrm{Im}\, Pf_i$. That is, $\mathrm{Im}\, f_i \in \mathbf{t}(c)$, and since C_c^P generates $\mathbf{t}(c)$, it follows that $\mathrm{Im}\, f_i \simeq C_c^P$. This means $f_i : C_i \to C_c^P$ is a singleton cover, and we conclude that C_c^P is prime.

Having established each C_c^P is prime, it now suffices to show that $\{C_c^P \to C : P \in \mathcal{P}, c \in PC\}$ is covering. Since $\mathcal{P}$ is sufficient, we need only show that for each P', $\{P'C_c^P \to P'C\}$ is surjective. This is clear because $c \in P'C_c^{P'}$ for each $c \in P'C$. Hence $\{C_c^P \to C\}$ is a J-cover, and $(\mathbf{C}, J)$ is prime-generated.
□

Lastly, we need adequate terminology to phrase the theorems. The following is a more refined version of Definition 5.28, basically restricting the notion of primeness to a subcategory of $\mathbf{cMod}(\mathbf{C}, J)$.

Definition 5.36. Let $(\mathbf{C}, J)$ be a site. Let $\mathbf{M}$ be a subcategory of $\mathbf{cMod}(\mathbf{C}, J)$, and let M be an object of $\mathbf{M}$. Then M is called **prime relative to M** if for each model M' in $\mathbf{M}$, object C of $\mathbf{C}$ and $c \in MC, c' \in M'C$ such that $\mathbf{t}(c) \subseteq \mathbf{t}(c')$, there is a natural transformation $\alpha : M \to M'$ in $\mathbf{M}$ such that $\alpha_C(c) = c'$.

Definition 5.37. A functor $F : \mathbf{C} \to \mathbf{D}$ is called **powerful** if for each $C \in \operatorname{ob} \mathbf{C}$, F induces a bijection $\operatorname{Sub}(C) \cong \operatorname{Sub}(FC)$. F is called an **embedding** if it is full and faithful.

At last, we are fully set up to state, and prove, the main theorem.

Theorem 5.38 ([BM87, Theorem 3.22])**.** *Let $\mathbf{E}$ be a Grothendieck topos. Then $\mathbf{E}$ has a continuous powerful embedding into a category of presheaves iff there is a sufficient category $\mathbf{M}$ of principal models which are all prime relative to $\mathbf{M}$.*

More specifically, if $\mathbf{M}$ satisfies the mentioned conditions, then the induced functor $\mathrm{ev} : \mathbf{E} \to \mathbf{Set}^{\mathbf{M}}$ (given by $E \mapsto \mathrm{ev}_E$) is a powerful embedding.

Proof Suppose we have a continuous powerful embedding $I : \mathbf{E} \to \hat{\mathbf{C}}$. By Corollary 5.34, $\hat{\mathbf{C}}$ has a sufficient collection of principal prime models $\mathcal{M}$. Since I is powerful, it preserves suprema subobjects, i.e. $I\left(\bigvee_i E_i\right) = \bigvee_i IE_i$. By Proposition 5.13 and faithfulness of I, this means that $\{E_i \to E\}$ is a cover in $\mathbf{E}$ iff $\{IE_i \to IE\}$ is in $\hat{\mathbf{C}}$; since $\mathcal{M}$ is sufficient for $\hat{\mathbf{C}}$, it follows that $\mathcal{M}I = \{MI : M \in \mathcal{M}\}$ is sufficient for $\mathbf{E}$. Moreover, $\mathcal{M}I$ consists of principal models, which by Corollary 5.34 are prime relative to the full subcategory $\mathbf{M}$ of $\mathbf{cMod}(\mathbf{E})$ whose objects are $\mathcal{M}I$.

Suppose now that $\mathbf{M}$ is a sufficient category of principal models which are prime relative to $\mathbf{M}$; by Corollary 5.34, the functors ev_M, $M \in \operatorname{ob} \mathbf{M}$ are sufficient for $\mathbf{Set}^{\mathbf{M}}$. The functor $\mathrm{ev} : \mathbf{E} \to \mathbf{Set}^{\mathbf{M}}$ preserves finite limits, as each $M \in \operatorname{ob} \mathbf{M}$ does. Similarly, since the ev_M are sufficient for $\mathbf{Set}^{\mathbf{M}}$, we have for a cover $\{f_i : E_i \to E\}$, that $\{\mathrm{ev}_{f_i} : \mathrm{ev}_{E_i} \to \mathrm{ev}_E\}$ covers iff $\{\mathrm{ev}_M(\mathrm{ev}_{f_i}) : \mathrm{ev}_M(\mathrm{ev}_{E_i}) \to \mathrm{ev}_M(\mathrm{ev}_E)\}$ is surjective for each M. But this is just $\{Mf_i : ME_i \to ME\}$, which is surjective because each M is continuous. Hence $\mathrm{ev} : \mathbf{E} \to \mathbf{Set}^{\mathbf{M}}$ is continuous.

Since $\mathbf{M}$ is sufficient, ev reflects isomorphisms; as ev also preserves equalizers, we conclude it is full and faithful. In particular, we have an injection $\mathrm{ev} : \operatorname{Sub}(E) \rightarrowtail \operatorname{Sub}(\mathrm{ev}_E)$. It remains to show that this induced operation on subobjects is surjective. To this end, let $F \hookrightarrow \mathrm{ev}_E$. By Proposition 1.10, F is a presheaf $\mathbf{M} \to \mathbf{Set}$ such that $FM \subseteq ME$ for all M.

Let $E_0 = \bigvee\{E' \in \operatorname{Sub}(E) \mid \forall M : ME' \subseteq FM\}$; then clearly $\mathrm{ev}_{E_0} \leq F$, and it remains to show that $F \leq \mathrm{ev}_{E_0}$. For arbitrary M, let $f \in FM \subseteq ME$. Since M is principal, we find an $E' \in \operatorname{Sub}(E)$ such that $\mathbf{t}_E^M(f)$ is generated by E'; then f is generic for ME'. Since M is prime in $\mathbf{M}$, for each model M' and $f' \in M'E'$, there is an $\alpha : M \to M'$ such that $\alpha_{E'}(f) = f'$; in particular, $\alpha_E(f) = \mathrm{ev}_E(\alpha)(f) = f'$. Since F is a subpresheaf of ev_E, the fact that $f \in FM$ implies that $F\alpha(f) = f' \in FM'$. Hence $M'E' \subseteq FM'$ for all $M' \in \operatorname{ob} \mathbf{M}$. Hence $E' \leq E_0$, and we conclude that for all M and $f \in FM$ we have $f \in ME_0$; that is, $F \leq \mathrm{ev}_{E_0}$. This concludes the proof that ev is powerful, and hence the proof of the theorem. $\square$

Corollary 5.39. *If $\mathbf{E}$ has a sufficient collection of principal prime models, then $\mathbf{E}$ has a continuous powerful embedding into a category of presheaves.*

In the special case that $\mathbf{E}$ is separable and prime-generated, it has such an embedding into $\hat{\mathbf{C}}$ for a countable category $\mathbf{C}$.

Our last result is a slight restatement of Theorem 5.38 which can be somewhat easier to work with, as we do not have to keep track of, nor have to identify, a special category of models. It basically arises from using the "old" Definition 5.28 in place of the more "refined" version Definition 5.36. This simplicity goes at the expense of losing the equivalence and the powerfulness of ev.

Theorem 5.40. *Suppose that $\mathbf{E}$ has a sufficient collection of principal prime models. Then $\mathrm{ev} : \mathbf{E} \to \mathbf{Set}^{\mathbf{cMod}(\mathbf{E})}$ is a continuous embedding.*

Proof That ev is continuous and faithful is clear, in view of the proof of Theorem 5.38. It thus only remains to show that ev is full. So let $\eta : \mathrm{ev}_E \to \mathrm{ev}_{E'}$ be a morphism of $\mathbf{Set}^{\mathbf{cMod(E)}}$. We have an obvious restriction $\mathbf{Set}^{\mathbf{cMod(E)}} \to \mathbf{Set}^{\mathbf{M}}$, where $\mathbf{M}$ is the full subcategory of principal prime models. Observe that Theorem 5.38 applies to ev' : $\mathbf{E} \to \mathbf{Set}^{\mathbf{M}}$.

Let $R \subseteq \mathrm{ev}(E \times E')$ be the graph of η as defined in Chapter 2, Definition 2.14. Since ev' is powerful, we find a $R' \subseteq E \times E'$ such that $\mathrm{ev'}(R')(M) = MR$ for all principal prime models M.

Since R is the graph of a morphism the sequents (phrased in the canonical language) of Proposition 2.16 are valid in $\mathbf{Set}^{\mathbf{M}}$. As ev' is powerful, this implies that in fact R' also must satisfy these sequents. By Proposition 2.17, the converse of Proposition 2.16, it follows that R' is the graph of some $f : E \to E'$. We immediately conclude that for each $M \in \mathrm{ob}\,\mathbf{M}$, $\mathrm{ev}(f)M = Mf = \eta_M$.

As a last step, we need to ensure that $\eta_M = \eta'_M$ for all $M \in \mathrm{ob}\,\mathbf{M}$ implies that $\eta = \eta'$. It suffices to show that for every $N \in \mathbf{cMod}(\mathbf{E})$ and $e \in NE$, there exist an $M \in \mathrm{ob}\,\mathbf{M}$, a morphism $h : M \to N$ and an $e' \in ME$ such that $e = h_E(e')$. By Lemma 5.35, $\mathbf{E}$ is prime-generated. More specifically, an inspection of the proof yields that for every prime P of $\mathbf{E}$ there is a model $M \in \mathrm{ob}\,\mathbf{M}$ and a $p \in MP$ such that p is generic for P. In conclusion, since $\mathbf{E}$ is prime-generated, we first reduce to the case of E a prime. Now we find M and $e' \in ME$ such that e' is generic for E. By primeness of M and the now vacuous condition that $\mathbf{t}(e') \subseteq \mathbf{t}(e)$, we conclude that there exists a morphism $h : M \to N$ with the sought properties. Therefore, each η is fully determined by the η_M for $M \in \mathrm{ob}\,\mathbf{M}$, and we conclude $\eta = \eta'$, as desired; hence ev is full. □

A nice touch in the above proof is the use of Propositions 2.16 and 2.17; results obtained by logical methods (viz the categorization of graphs of morphisms using sequents) are used to prove nontrivial results in a topological, or rather topos-theoretical, context. For, despite the model-theoretical terminology employed, the result obtained is a typical topos-theoretic one.

Discussion

During the course of this work, several connections between (geometric) logic and Grothendieck topoi have been exhibited. It is clear that the connection between the two is deep – so deep, that this work is effectively not but an introduction, and an invitation to explore it further. The questions provoked, and the corners cut due to space and time limitations (which e.g. shows in the scarcity of examples), together could easily multiply this document in volume, and effectively turn it into a book, while still retaining the same sense of "incompleteness" and loose ends that its current form may portray.

Let me mention some of the questions that have cropped up with me during the writing of this work. One of the notions that caught my eye very early (but which, in spite of my enthusiasm, I didn't fully understand until the late stages of my writing) was the definition of *types* in [BM87] (see Definition 5.17). It is a very clear and accurate translation of the corresponding concept in classical model theory; it would be very interesting to see if it can be meaningfully extended beyond the context of continuous models to arbitrary geometrical or other logical functors, allowing the codomain **Set** to vary over e.g. Grothendieck topoi or more generally, any κ-geometrical category.

Another interesting research topic would be to attempt to apply the continuous-model techniques from Chapter 5 to the classifying topoi from Chapter 4. It seems to me that – due to the very explicit construction of classifying topoi – it should be possible to obtain nontrivial results about, say, geometric theories via this route, which would be yet another great confirmation of the intimate relationship between (geometric) logic and Grothendieck topoi.

Acknowledgements

My gratitude goes out to Jaap van Oosten, who kindly agreed to supervise this thesis, and provided many of the references that have played an important role in writing it. Most importantly, on a number of occasions, the appointments and deadlines we set provided the necessary discipline to persevere and keep delving in. His supervision and constructive criticism have ensured both a substantial increase in quality and a timely completion of my thesis.

Further, I would like to thank all of my friends, family and acquaintances, both near and far, for their continued support and understanding – particularly during some of the tougher times, when all progress seemed to have come to a grinding halt. I couldn't have done it without you.

A Appendix

This Appendix contains very brief accounts of some material that crops up in the main sections of the thesis on only one or two occasions. It is presented here so as to sustain flow and coherence of the main work.

A.1 Filtered categories and -colimits

Definition A.1. A category $\mathbf{I}$ is **cofiltered** if:

(i) $\mathbf{I}$ is non-empty;

(ii) for all $i, i' \in \operatorname{ob} \mathbf{I}$, there is a $j \in \operatorname{ob} \mathbf{I}$ such that there are morphisms $i \to j$ and $i' \to j$;

(iii) for all parallel morphisms $f, f' : i \to i'$ there is a $g : i' \to j$ such that $gf = gf'$.

The prefix "co" stems from duality: $\mathbf{I}$ is **filtered** if $\mathbf{I}^{\mathrm{op}}$ is cofiltered.
When I is (co-)filtered, a functor $D : \mathbf{I} \to \mathbf{C}$ is called a **(co-)filtered diagram** (in $\mathbf{C}$); similarly, the colimit $\varinjlim_i D_i$ is called **(co-)filtered** if D is.

It is not hard to see that a category $\mathbf{I}$ is cofiltered iff it admits a cocone for every finite diagram $\mathbf{I}' \to \mathbf{I}$.

A word of caution is in order: in the literature, "filtered" and "cofiltered" are often interchanged (for example, in [Mac71]), so one needs to pay attention as to which convention is in use. In our nomenclature, a filter on a poset P is precisely a filtered subcategory of P when considered as a category; this is considered by the author to be the most intuitive of the two commonly used options.

Theorem A.2 (cf. [Mac71], Theorem IX.2.1)**.** *In* **Set***, finite limits commute with (small) cofiltered colimits. That is, for any functor $D : \mathbf{I} \times \mathbf{J} \to \mathbf{Set}$ with $\mathbf{I}$ small and cofiltered and $\mathbf{J}$ finite, there is an isomorphism:*

$$\kappa : \varinjlim_i \varprojlim_j D(i,j) \to \varprojlim_j \varinjlim_i D(i,j)$$

Definition A.3. Let $\mathbf{C}$ be a small category, and let $D : \mathbf{I} \to \mathbf{C}$ be a diagram. The functor $M : \mathbf{C} \to \mathbf{Set}$ **represented by** D is defined by:

$$MC = \varinjlim_{I \in \mathbf{I}} \operatorname{Hom}(DI, C)$$

on objects, while on morphisms it is induced by the representable functors $\operatorname{Hom}(DI, -)$.

The significance of represented functors in our setting is apparent from the following result (the "if" part of which readily follows from the preceding Theorem A.2):

Theorem A.4 (cf. [AKT01], Theorem 2.2)**.** *Let $\mathbf{C}$ be a category with finite limits. A functor $M : \mathbf{C} \to \mathbf{Set}$ preserves finite limits iff it can be represented by a filtered diagram.*

A.2 The Grothendieck topology induced on a slice category

Definition A.5. Let $(\mathbf{C}, J)$ be a site, and let $C \in \operatorname{ob} \mathbf{C}$. The **induced topology** J/C on the slice category $\mathbf{C}/C$ is defined by declaring that $\{f_i : d_i \to d\}$ is a J/C-cover iff $\{f_i : D_i \to D\}$ (where e.g. $D_i = \operatorname{dom} d_i$) is a J-cover.

The following results are all straightforward:

Proposition A.6. *Let $(\mathbf{C}, J), C$, and J/C be as above. Then:*

(i) *J/C is a topology;*

(ii) *If $(\mathbf{C}, J)$ is regular, so is $(\mathbf{C}/C, J/C)$;*

(iii) *If $(\mathbf{C}, J)$ is (sub)canonical, so is $(\mathbf{C}/C, J/C)$;*

(iv) $\operatorname{Sub}_{\mathbf{C}}(C) \cong \operatorname{Sub}_{\mathbf{C}/C}(\mathrm{id}_C)$*;*

(v) *$F : \mathbf{C} \to \mathbf{C}/C$ defined by $FD = \pi_1 : C \times D \to C$, $F(f : D \to D') = \mathrm{id}_C \times f : C \times D \to C \times D'$ is continuous;*

(vi) *If D is a prime in $(\mathbf{C}, J)$, then FD is a prime in $(\mathbf{C}/C, J/C)$. Hence if $(\mathbf{C}, J)$ is prime-generated, so is $(\mathbf{C}/C, J/C)$.*

References

[AKT01] J. Adámek, V. Koubek, and V. Trnková. How Large are Left Exact Functors? *Th. Appl. Categories*, VIII(13):377–390, 2001.

[Awo10] S. Awodey. *Category Theory*, volume 52 of *Oxford Logic Guides*. Oxford University Press, 2010.

[Blu12] A. Blumensath. *Logic, Algebra and Geometry.* 2012. Digital version available at http://www.mathematik.tu-darmstadt.de/ blumensath/.

[BM87] M. Barr and M. Makkai. On Representations of Grothendieck Toposes. *Can. J. Math.*, XXXIX(1):168–221, 1987.

[Joh02] P.T. Johnstone. *Sketches of an Elephant – A Topos Theory Compendium.* Oxford University Press, 2002.

[Mac71] S. Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics.* Springer, 1971.

[Mar02] D. Marker. *Model Theory: An Introduction*, volume 217 of *Graduate Texts in Mathematics.* Springer, 2002.

[MM92] S. Mac Lane and I. Moerdijk. *Sheaves in Geometry and Logic.* Universitext. Springer, 1992.

[MR77] M. Makkai and G.E. Reyes. *First Order Categorical Logic*, volume 611 of *Lecture Notes in Mathematics.* Springer, 1977.

[Poi00] B. Poizat. *A Course in Model Theory.* Universitext. Springer, 2000.