Universiteit Utrecht

A Spatial Model of the Lambda Calculus

by Jeroen Goudsmit
under the supervision of prof. dr. I. Moerdijk
second examiner: dr. J. van Oosten

Acknowledgements

This thesis has been written as a master's thesis at the University of Utrecht, in order to complete the master's program 'Mathematical Sciences'. Several people helped me in writing this, and I would like to take this opportunity to thank them.

My supervisor prof. dr. Ieke Moerdijk helped me these months, guiding me through the literature, for which I am very thankful. My understanding of the spatial cover greatly benefited from our meetings. I very much appreciate the time we spent working on this thesis, I believe to have learned a great deal due to his guidance. Furthermore I would like to thank my second examiner dr. Jaap van Oosten, for steering me in the right direction when I came across a difficult design choice. My understanding of λ-calculus profited from the conversations I had with prof. dr Vincent van Oostrom, for which am much obliged. Several times I worked together with Sori Lee, his company boosted my performance quite a bit. I am especially grateful for his presence during my final rehearsal.

On a more personal note I would like to thank Barbara Harzevoort. She proof-read my work and pointed out anything in need of refinement. The quality of this work greatly improved due to her suggestions. When the time came to prepare a talk, she was there to rehearse it with me. She provided me with the attentive audience I needed to condense my ideas into the appropriate time span. Finally, I would like to thank my parents for supporting me throughout my time as a student.

Contents

Acknowledgements iii

Introduction 1

1 Grothendieck Toposes 5
- 1.1 Sheaves . . . 5
- 1.2 Étale Spaces . . . 12
  - 1.2.1 Sheaves and Étale Spaces . . . 16
  - 1.2.2 Base Change . . . 19
  - 1.2.3 Locally Connected Étale Spaces . . . 23
  - 1.2.4 Points . . . 25
- 1.3 Grothendieck Topology . . . 26
- 1.4 Geometric Morphisms . . . 29

2 Spatial Cover 31
- 2.1 Enumeration Space . . . 32
- 2.2 The Space of Points . . . 33
- 2.3 A Beck–Chevalley Condition . . . 37

3 λ-Calculus 43
- 3.1 Cartesian Closed Categories . . . 44
- 3.2 The λ-Calculus . . . 51
- 3.3 Models of λ-Calculus . . . 58

Bibliography 63

Introduction

In this master's thesis we will construct a *complete* and *functionally complete model* of the *λ-calculus*. In the following it will become clear what most of these terms mean, and how this result relates to earlier work. For reasons of brevity, we skip over a great deal of details in these few pages, but rest assured that all will be made precise in the following three chapters.

The λ-calculus is, in its most basic form, a theory of functions. In order for this thesis to be self-contained, we cover it with a fair amount of detail in Chapter 3. For the moment, think of λ-calculus as the formal system of equations on terms generated by the following inference system. We say that M is a *λ-term* of *type* ϕ within the *context* Γ whenever $\Gamma \vdash M : \phi$ can be deduced from the inference system below. We remain purposely vague on what types and context are precisely, this will all be specified to great detail in Section 3.2. It is important to see however that the types ϕ mentioned within this system are elements of the closure of some set **V** (the set of basic types) under the type-forming operation $(-) \Rightarrow (-)$.

$$\frac{}{\Gamma, x : \phi \vdash x : \phi}\ (\text{ID})$$

$$\frac{\Gamma, x : \phi \vdash M : \psi}{\Gamma \vdash \lambda x^{\phi}.M : \phi \Rightarrow \psi}\ (\Rightarrow\text{I}) \qquad \frac{\Gamma \vdash M : \phi \Rightarrow \psi \qquad \Delta \vdash N : \phi}{\Gamma, \Delta \vdash M\ N : \psi}\ (\Rightarrow\text{E})$$

Terms of this system include for example

(i) $\lambda x^{\phi}.x^{\phi}$ of type $\phi \Rightarrow \phi$;
(ii) $\lambda x^{\phi}.\lambda y^{\psi}.x^{\phi}$ of type $\phi \Rightarrow \psi \Rightarrow \phi$;
(iii) $\lambda f^{\phi \Rightarrow \psi \Rightarrow \chi}.\lambda g^{\phi \Rightarrow \psi}.\lambda x^{\phi}.(f\ x)\ (g\ x)$ of type $(\phi \Rightarrow \psi \Rightarrow \chi) \Rightarrow (\phi \Rightarrow \psi) \Rightarrow \phi \Rightarrow \chi$,

all within the empty context. The term (i) can be thought of as the identity function, (ii) is the 'constant-function function' and (iii) 'distributes' its third argument to its first and second argument. The types of the two latter terms are precisely the axioms of Hilbert-style zeroth-order logic.[1]

Of course, the above intuitive description makes no sense at all until we introduce equations on λ-terms. The reason that $\lambda x^{\phi}.x^{\phi}$ resembles the identity function is that when 'applied' to a λ-term of type ϕ, it 'yields' this term. The two basic equalities are concerned with extensionality and reduction, and given by

$$\lambda x^{\phi}.M\ x \equiv M \qquad\qquad (\lambda x^{\phi}.M)\ N \equiv M\,[x \mapsto N]\,,$$

where $M\,[x \mapsto N]$ is to be read as 'the term resulting from replacing all occurrences of x within M with N'. It thus follows from the second equality that

$$(\lambda x^{\phi}.x^{\phi})\ M \equiv M$$

holds for any term M of type ϕ within some context.

[1] The calculus concerned with these three combinators was studies by Schönfinkel, 1924, who called these terms respectively I (Identiteitsfunktion), C (Konstantfunktion) and S (Verschmelzungsfunktion). Nowadays these are called the SKI-combinators, their relation with λ-calculus and logic is discussed in Sørensen and Urzyczyn (2006, Chapter 5).

Suppose one assigns to each element p of **V** a set $[\![p]\!]$. This mapping can be extended inductively so as to send the type $\phi \Rightarrow \psi$ to the set $\mathbf{Set}\left([\![\phi]\!],\ [\![\psi]\!]\right)$.[2] A context, being a finite set of variables, can thus be sent to some product $\prod_{x\,:\,\gamma\in\Gamma}[\![\gamma]\!]$. One can interpret terms M of type ϕ within a context Γ as set-theoretic functions from the interpretation of their context to the interpretation of their type. The mapping is inductive along the deduction used to construct them with the inference system given above, and defined below.

$$\frac{}{[\![\Gamma, x\,:\,\phi \vdash x^\phi\,:\,\phi]\!] : [\![\Gamma, x\,:\,\phi]\!] \to [\![\phi]\!], \quad p \mapsto p_x}\ (\textsc{id})$$

$$\frac{\Gamma, x\,:\,\phi \vdash M\,:\,\psi}{[\![\Gamma \vdash \lambda x^\phi.M\,:\,\phi \Rightarrow \psi]\!] : [\![\Gamma]\!] \to \mathbf{Set}\left([\![\phi]\!],\ [\![\psi]\!]\right), \quad p \mapsto \big(x \mapsto [\![\Gamma, x\,:\,\phi \vdash M\,:\,\psi]\!](p, x)\big)}\ (\Rightarrow\text{I})$$

$$\frac{\Gamma \vdash M\,:\,\phi \Rightarrow \psi \qquad \Delta \vdash N\,:\,\phi}{[\![\Gamma, \Delta \vdash M\ N\,:\,\psi]\!] : [\![\Gamma, \Delta]\!] \to [\![\psi]\!], \quad (p, q) \mapsto [\![\Gamma \vdash M\,:\,\phi \Rightarrow \psi]\!](p)\big([\![\Delta \vdash N\,:\,\phi]\!](q)\big)}\ (\Rightarrow\text{E})$$

This *interpretation* provides a *model* for the simply typed λ-calculus. Indeed, one can interpret an equality $M \equiv N$ of λ-terms M and N of the same type ϕ within the same context Γ as the actual equality of the functions

$$[\![\Gamma \vdash M\,:\,\phi]\!] = [\![\Gamma \vdash N\,:\,\phi]\!].$$

We call this interpretation a model because every equation which holds in the theory (of equations between λ-terms within a given context), is interpreted as a valid equation. The interpretation given here is a very concrete version of the interpretation described in Section 3.3. It is easy to see that the λ-terms (i) and (ii) are mapped to the functions below, when interpreted within the empty context.

$$* \mapsto \big(x \mapsto x\big) \qquad\qquad * \mapsto \big(x \mapsto (y \mapsto x)\big)$$

It has been proven by Plotkin (1973) that the collection of models in **Set** where each basic type is interpreted by a finite set is complete for this λ-calculus.[3] A collection of models is said to be *complete* when an expression holds in the theory if the interpretation of said expression holds in all models. This means that if the interpretation of an equality holds in all models with 'finite basic types', then it must hold in the λ-calculus as well. This does *not* mean that if an equality holds in *some* model in **Set** with finite interpretations of basic types, it also holds in the theory. Let us demonstrate this little fact.

Suppose that the symbol $\mathbb{N}$ is contained within the set of basic types **V**. This means that we have the type $\mathbf{N} := (\mathbb{N} \Rightarrow \mathbb{N}) \Rightarrow \mathbb{N} \Rightarrow \mathbb{N}$. For each natural number n we have the a term within the empty context of this type **N**. These are the so-called *Curch numerals*, defined as

$$\overline{n} := \lambda s^{\mathbb{N}\Rightarrow\mathbb{N}}.\lambda x^{\mathbb{N}}.s^n\ x,$$

where $s^n\ x$ is to be read as the n-fold application of s to x. The equation below on the left is an actual syntactic equality, the equality on the right follows from the imposed equations.

$$\overline{0} = \lambda s^{\mathbb{N}\Rightarrow\mathbb{N}}.\lambda x^{\mathbb{N}}.x \qquad\qquad \overline{n+1} \equiv \lambda s^{\mathbb{N}\Rightarrow\mathbb{N}}.\lambda x^{\mathbb{N}}.\overline{n}\ s\ (s\ x)$$

When we interpret the type $\mathbb{N}$ as the set $\mathbb{N}$, then it is clear that $\overline{n} \equiv \overline{m}$ holds precisely if $n = m$. Note that we obtain this result quite easily through our model, whereas it would be a lot harder to prove purely syntacticly. From this we deduce that there are infinitely many 'really' distinct terms of type **N**. So if we were to interpret the type $\mathbb{N}$ as a finite set, it would follow that **N** is interpreted as a finite set. Now suppose that this model is *complete*, that is, every equality which holds in the model holds in the theory as well. This would force a non-existent equality in the

[2] We will often use the notation $\mathcal{C}\left(X,\ Y\right)$ to denote the set of arrows between the objects X and Y in the category $\mathcal{C}$. In this particular instance, we refer to the set of set-theoretic functions between $[\![\phi]\!]$ and $[\![\psi]\!]$. This is sometimes denoted $[\![\psi]\!]^{[\![\phi]\!]}$ but due to ideologic and cosmetic reasons we do not adopt this latter notation at all.

[3] See Selinger (2007, Section 10) for a small description of this theorem. The following example is motivated by his notes.

theory of equations governing this λ-calculus, a contradiction. Friedman (1975) proved the existence of a complete model of the λ-calculus as defined above with only one basic type.

Let us now take a small step back. We can construct a model for the λ-calculus given any set **V** of basic types, and the entirety of such models where the basic types are interpreted by finite sets is complete. If we restrict **V** to be a one-point set, then we can find a single model which is complete. Although these models are nice, they might seem superfluous as the problem of equality within these λ-calculi is decidable. Indeed, one can give a *direction* to the equalities given here and thusly form a *rewrite system*. It has been proven that two λ-terms are equal precisely if their normal forms with respect to this rewrite system coincide, see for instance Sørensen and Urzyczyn (2006).[4] The rewrite system has some very nice properties. It is for instance bounded, that is to say, one can find an *a priori* upper bound on the amount of reductions needed to reach the normal form. Unfortunately, this upper bound itself can not be bounded by any elementary function, in particular it can not be bounded by $\exp_r(n)$ in the size n of a term.[5] Even more unfortunate is the fact that this approach does not work when we add additional equalities to the λ-calculus.

Indeed, it was shown by Scott (1980, Section 3) that the untyped λ-calculus can be modeled as a typed λ-calculus, by adding appropriate equations and constants. It is a well-known fact that the problem of equality for the untyped λ-calculus is undecidable, in particular, the rewrite system given above is no longer normalizing. This means that the approach above does not generalize to λ-calculus with additional equalities.

One might wonder whether it is possible to give each such λ-calculus with additional equalities and constants (or *basic terms*) a model in **Set** which is complete, akin to the result of Friedman. The answer is a resounding 'no', as we will now demonstrate. Consider the λ-calculus Λ_{ref} which has one basic type **D**, two basic terms $i : (\mathbf{D} \Rightarrow \mathbf{D}) \Rightarrow \mathbf{D}$ and $r : \mathbf{D} \Rightarrow (\mathbf{D} \Rightarrow \mathbf{D})$ and comes equipped with the additional equation

$$\vdash \lambda x^{\mathbf{D}\Rightarrow\mathbf{D}}.r\ (i\ x) \equiv \lambda x^{\mathbf{D}\Rightarrow\mathbf{D}}.x$$

It is clear that any model in **Set** of this λ-calculus must interpret $x^{\mathbf{D}} \vdash r\ x : \mathbf{D} \Rightarrow \mathbf{D}$ as a surjection $[\![\mathbf{D}]\!] \to [\![\mathbf{D} \Rightarrow \mathbf{D}]\!]$. For cardinality reasons, this entails that $[\![\mathbf{D}]\!]$ is a one-point set. This trivializes the model, and every equation becomes provable. Yet non-trivial models of this exist, as discussed by Scott (1980, Section 3) and described in depth in Lambek and Scott (1988, Chapter 1.15 – 1.18). In particular there is a model where the equation

$$\vdash \lambda x^{\mathbf{D}}.i\ (r\ x) \equiv \lambda x^{\mathbf{D}}.x$$

does not hold. Consequently, no (standard) model of Λ_{ref} in **Set** can be complete. We thus abandon the search for complete and standard semantics for λ-calculus in **Set**, and seek better suited categories.

The category **Set** shares some structure with a great deal of other categories. For instance, it has a chosen categorical binary product in the form of the cartesian product of sets. It also has a chosen one-point set $\{\emptyset\}$, and the set of functions between the sets X and Y is itself a set. We thus have functors $(-) \wedge (-) : \mathbf{Set} \times \mathbf{Set} \to \mathbf{Set}$ and $(-) \Rightarrow (-) : \mathbf{Set}^{\text{op}} \times \mathbf{Set} \to \mathbf{Set}$ which respectively assign to a pair of sets X and Y the set $X \times Y$ and the set $\mathbf{Set}(X, Y)$. There is an obvious isomorphism

$$\mathbf{Set}(X \times Y, Z) \cong \mathbf{Set}(X, \mathbf{Set}(Y, Z)),$$

natural in X and Z. This means that the functor $(-) \wedge Y$ is left-adjoint to the functor $Y \Rightarrow (-)$, we call a right adjoint of this product functor an *exponent*. A category which has a terminal object, binary products and exponents is usually called *cartesian closed*. We will cover such categories in Section 3.1, where we capture this notion by means of an inference system and equations between deductions therein.

It has been proven that to each λ-calculus Λ, with additional equations and basic terms, there exists a cartesian closed category $\mathcal{C}$ whose objects are the types of Λ and whose arrows are formed from terms. This syntactic category is covered in great deal by Lambek and Scott (1988), and we treat it in Section 3.2. One can form a model of Λ, not in the category **Set** but in its syntactic category $\mathcal{C}$, generalizing the notion of a model we gave above. Indeed, to

[4]This fact also holds for the basic λ-calculus extended with 'surjective pairing', we cover a small extension of this in Chapter 3. Lambek and Scott (1988) provide a proof extended to this setting, which in essence is an adaptation of a proof by de Vrijer (1987). A very readable version of the latter can be found in Terese (2003).

[5]The function $\exp_r(n)$ is defined inductively by $\exp_0(n) = n$ and $\exp_r(n) = 2^{\exp_r(n)}$.

interpret λ-terms we merely used the cartesian closed structure present in **Set**, so we can easily adapt the notion of an interpretation to any cartesian closed category. We formalize this in Section 3.3.

We can thus find complete semantics of any λ-calculus in the collection of cartesian closed categories. By the Yoneda embedding we know that each (small) category $\mathcal{C}$ is isomorphic to the categories of functors $\mathcal{C}^{\text{op}} \to \mathbf{Set}$, denoted here $\mathbf{PSh}(\mathcal{C})$. It is easy to prove, as we do in Proposition 3.1.1, that $\mathbf{PSh}(\mathcal{C})$ is a cartesian closed category. Moreover, the Yoneda-embedding $\mathbf{y} : \mathcal{C} \to \mathbf{PSh}(\mathcal{C})$ preserves the cartesian closed structure, as we demonstrate in Proposition 3.1.2. It thus suffices to look at the collection of all categories of the form $\mathbf{PSh}(\mathcal{C})$ to find complete semantics for any λ-calculus.

This collection was refined by Mitchell and Moggi (1991), who provide a complete model to each λ-calculus in a category of the form $\mathbf{PSh}(P)$ with P a partially ordered set. There is however a significant difference between their notion of a model and ours. An interpretation in **Set** as defined above makes sure that the *type* $X \Rightarrow Y$ is mapped to the *set* $X \Rightarrow Y = \mathbf{Set}\,(X, Y)$. This was implicitly understood to be generalized to the cartesian closed setting in the previous two paragraphs. Yet Mitchell and Moggi do not demand this, so they give what Awodey (2000) calls a *non-standard model*.

Our aim is to provide a *complete* model to each λ-calculus in the category of sheaves over a certain topological space $\mathbf{X}$. As each sheaf on $\mathbf{X}$ is in particular a presheaf on the partially ordered set made up out of the opens of $\mathbf{X}$, this is again a refinement. In order to reach this goal we follow the reasoning of Awodey (2000), and at a crucial point we employ the *spatial cover* constructed by Butz and Moerdijk (1999). Additionally, our model is *functionally complete* as well, which means that every arrow whose domain is terminal and whose codomain is the interpretation of a type arises as the interpretation of a λ-term, unique up to equality.

The spatial cover of Butz and Moerdijk is the main ingredient in restricting from presheaves over a small category to sheaves over a topological space, so we prove this result in Chapter 2. In order to form this proof we need a fair amount of theory about sheaves on a topological space. We provide an introduction to sheaves in Chapter 1, and prove almost all necessary results from scratch. As a consequence, this text is largely self-contained, at the cost of spending quite a lot of space on results that might be considered basic.

Chapter 1

Grothendieck Toposes

The purpose of this chapter is threefold. First and foremost it is meant to introduce the reader to topos theory by covering some of the basic material in this field. Secondly, it serves as a reference, providing definitions, notation and results which will be referred to in the remainder of this thesis. And finally, it provides the author the opportunity to internalize the basic material through explanation. All the material covered in this chapter is 'common' knowledge, which makes it difficult to provide the proper references. A great deal of the material here is taken from Mac Lane and Moerdijk (1991) and Johnstone (2002a,b).

In order to introduce the notion of a topos, we first consider the motivating example of sheaves on a topological space in Section 1.1. All notions covered there will be generalized to the setting of the Grothendieck toposes in Section 1.3. Many sheaves arise as sections of covering (or the more general étale) spaces, we treat these in Section 1.2. This notion is of special interest to us, as we will use it often in Chapter 2. Finally, we cover morphisms between (Grothendieck) toposes in Section 1.4. The material in Subsection 1.2.3 concerns an essential geometric morphism between sheaves on a locally connected and connected topological space and **Set**, this will be crucial in the following chapter. Points of toposes play a pivotal role in the next chapter, we treat those extensively in Subsection 1.2.4.

Section 1.1 Sheaves

Consider some topological space X. For any open subset $U \subseteq X$, one can consider the set of continuous maps $\mathbf{Top}(U, \mathbb{R})$, that is to say, the set of arrows between U and $\mathbb{R}$ in the category of topological spaces. An inclusion $V \subseteq U \subseteq X$ gives rise to a restriction $\mathbf{Top}(U, \mathbb{R}) \to \mathbf{Top}(V, \mathbb{R})$, which simply restricts a continuous map $f : U \to \mathbb{R}$ to the 'smaller' map $f \restriction V : V \to \mathbb{R}$. This assignment is *functorial*, in the sense that $f \restriction U$ equals f and that the restriction $(f \restriction V) \restriction W$ equals the restriction $f \restriction W$ for any chain of inclusions $W \subseteq V \subseteq U$ of open subsets of X.

This assignment satisfies a certain 'collation' or 'amalgamation' property. Suppose we have an open cover $\mathcal{V}$ of U and a continuous map $f_V : V \to \mathbb{R}$ for each V in the cover $\mathcal{V}$, subject to the condition that these maps coincide on intersections. More formally, the restrictions $f_V \restriction (V \cap W)$ and $f_W \restriction (V \cap W)$ ought to be equal for any pair of opens V and W in the cover $\mathcal{V}$. From elementary topology we know that to such a collection of maps there is a unique continuous map $f : U \to \mathbb{R}$ with the property that the restrictions $f \restriction V$ equal f_V for all $V \in \mathcal{V}$. An assignment similar to $U \mapsto \mathbf{Top}(U, \mathbb{R}) : \mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}$, in the sense that it satisfies the properties described above, is called a *sheaf* and will be defined more formally below.

There are many different ways to define the notion of a sheaf, the following is 'economic' in the sense that it requires little auxiliary definitions. In Proposition 1.1.1 we give an alternative definition which generalizes well to the setting in Section 1.3. The definition below is taken from Mac Lane and Moerdijk (1991), in mainstream mathematical texts one more often come across the definition as given by Johnstone (1997, Section 5.1), as for instance in Forster (1981, Definiton 6.1).

Definition 1.1.1 (Sheaf). A *sheaf* on a topological space X is a functor $F : \mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}$ such that for any open cover $\mathcal{V}$ of any open $U \subseteq X$ the following diagram is an equalizer.

$$F(U) \xrightarrow{e} \prod_{V \in \mathcal{V}} F(V) \overset{p}{\underset{q}{\rightrightarrows}} \prod_{V,W \in \mathcal{V}} F(V \cap W) \tag{1.1}$$

In the above, the maps e, p and q send elements to their sensible restrictions.[1] We often write $f \restriction V$ for $F(V \subseteq U)(f)$ where $f \in F(U)$, as if actually restricting a function to a smaller domain of definition.

Remark 1.1.1. Note that many incarnations of this definition, such as the definition given by Hartshorne (1977, Section 2.1), have a different definition of presheaf. We consider a presheaf to be nothing more than a presheaf in categorical context, that is, a contravariant functor from $\mathrm{Opens}(X)$ to **Set**. Other texts might additionally demand that a presheaf sends the empty open to the empty set, trivial group, trivial ring or what have you. One can immediately deduce from the diagram in (1.1) that as the empty set has an empty cover, $F(\emptyset)$ must be terminal in the relevant category whenever F is a sheaf. So this small difference in definition poses no real problem, for we are mostly concerned with sheaves.

Because (1.1) has to be an equalizer, we know that to each set S with a map $f : S \to \prod_{V \in \mathcal{V}} F(V)$ such that $ps = qs$, there must be a unique map $g : S \to F(U)$ such that $eg = f$. Such a map f is simply an assignment which sends $s \in S$ to the $\mathcal{V}$-indexed tuple $(s_V)_{V \in \mathcal{V}}$, such that this tuple behaves nicely with respect to restrictions in the sense that $s_V \restriction V \cap W = s_W \restriction V \cap W$ for any pair of opens V and W in the cover $\mathcal{V}$. The map g assigns to s an element of $F(U)$, such that $s \restriction V = s_V$ for all $V \in \mathcal{V}$. Below a small list of examples, the first few to indicate some common sheaves, followed by some more concrete examples. The final few examples provide some motivation for Section 1.3. Only Example 1.1.2 is of immediate relevance to the following.

Example 1.1.1 (Sheaf of Continuous Functions). Let us first revisit our motivating example. To this end, consider some topological space X. Now take as above the functor $F : \mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}$ as given by $F(U) := \mathbf{Top}\,(U, \mathbb{R})$, the restriction of $\mathbf{Top}\,((-), \mathbb{R}) : \mathbf{Top}^{\mathrm{op}} \to \mathbf{Set}$ to the open subsets of X. This is a sheaf, as follows directly from the introduction and the paragraph above.

Example 1.1.2 (Representable Sheaf). Consider the representable presheaf

$$\mathbf{y}U := \mathrm{Opens(X)}\,((-), U) : \mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}.$$

Note that $\mathbf{y}U(V)$ is simply the one-point set for any $V \subseteq U$. This makes the diagram in Definition 1.1.1 an equalizer, because all objects involved are isomorphic to $\mathbf{1}$.

Example 1.1.3 (Sheaf of Holomorphic Functions). Let $f : X \to Y$ be a continuous map between Riemann surfaces. Now consider the charts $\phi_U : X \supseteq U \to \mathbb{C}$ and $\psi_V : Y \supseteq V \to \mathbb{C}$ belonging to the atlases of X and Y respectively. We say that the map f is *holomorphic* exactly when the composite

$$\mathbb{C} \supseteq \left(\phi_U \restriction f^{-1}(V)\right)^{-1}(U) \xrightarrow{\phi_U^{-1} \restriction f^{-1}(V)} X \xrightarrow{f} V \xrightarrow{\psi_V} \mathbb{C}$$

is holomorphic in the usual complex analytic sense, see Lang (1999) for a definition. It is obvious that for functions on $\mathbb{C}$ the two notions of holomorphicity coincide. Denote the set of holomorphic maps between the Riemann surfaces X and Y by $\mathcal{O}(X, Y)$.

Using the above notation we can construct the presheaf $\mathcal{O}((-), \mathbb{C}) : \mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}$ which maps opens to the set of holomorphic functions to $\mathbb{C}$ starting in that set. Using similar reasoning as in Example 1.1.1 one can prove this presheaf to be a sheaf. Note that $\mathcal{O}(U, \mathbb{C})$ is a ring for each $U \subset X$, using the obvious operations of point-wise addition and multiplication.

Example 1.1.4 (Sheaf of Meromorphic Functions). Similar to the definition of holomorphic functions in the above example, one can define the presheaf $\mathcal{M}((-), \mathbb{C})$ of *meromorphic* maps on a Riemann surface X. As a reminder, a complex analytical function $\mathbb{C} \supseteq U \to \mathbb{C}$ is said to be meromorphic precisely when it is holomorphic but for finitely many points, which are poles of this function. The ring structure on $\mathcal{M}(U, \mathbb{C})$ is more than just a ring, it is in fact a field. Indeed, any function $f : U \to \mathbb{C}$ has as a multiplicative inverse the map $x \in U \mapsto 1/f(x)$, which is meromorphic when f is. This operation interchanges roots and zeroes. Note that each element of $\mathcal{M}(U, \mathbb{C})$ can be expressed as a quotient of elements of $\mathcal{O}(U, \mathbb{C})$.

[1] Formally, $e(x) = \left(x \restriction V\right)_{V \in \mathcal{V}}$, $p(\mathbf{x}) = \left(x_V \restriction V \cap W\right)_{V,W \in \mathcal{V}}$ and $q(\mathbf{x}) = \left(x_W \restriction V \cap W\right)_{V,W \in \mathcal{V}}$.

Example 1.1.5 (Projective Plane). As a more concrete instance of 1.1.4 consider the projective line over the complex numbers, denoted by $\mathbb{P}^1(\mathbb{C})$. An important basic fact about meromorphic functions is that there is a natural isomorphism between the sheaves $\mathcal{M}((-), \mathbb{C})$ and $\mathcal{O}((-), \mathbb{P}^1(\mathbb{C}))$ on any Riemann surface X. One can think of the map from left to right as sending a function $f : X \to \mathbb{C}$ to the function $X \to \mathbb{P}^1(\mathbb{C})$ which is exactly the same on the affine part, but sends poles of f to the point at infinity.

Example 1.1.6 (Periodic Functions). Consider the real line $\mathbb{R}$ and the integers $\mathbb{Z}$. For each non-zero $\lambda \in \mathbb{R}$ we can consider the space $X := \mathbb{R}/2\lambda\pi\mathbb{Z}$, where one has to consider the quotient as in the category of topological spaces. Think of this space as a circle of radius λ. Take for instance $\lambda = 1$ and consider the map $\sin : \mathbb{R} \to \mathbb{R}$, this gives rise to a map $\sin : X \to X$. More generally, all trigonometric functions are singly periodic, and thus give rise to a continuous map on X. This shows us that, amongst others, $\sin \upharpoonright U$ is an element of $\mathbf{Top}(U, \mathbb{R}/2\pi\mathbb{Z})$ for any $U \subseteq \mathbb{R}/\mathbb{Z}$.

Example 1.1.7 (Doubly Periodic Functions). Take a lattice (in the geometric, not order-theoretic sense) in the complex plane $\mathbb{C}$, given as $\Lambda = \omega_1 \cdot \mathbb{Z} + \omega_2 \cdot \mathbb{Z}$. When we take the (topological) quotient, this gives rise to a torus $\mathbb{C}/\Lambda$. One can endow this space with the structure of a Riemann surface, and consider the sheaf of meromorphic functions on $\mathbb{C}/\Lambda$. These are precisely the doubly periodic functions, amongst which $\wp_\Lambda$, Weierstraß's elliptic function, as defined below.

$$\wp_\Lambda : \mathbb{C} \to \mathbb{C}, \qquad z \mapsto \frac{1}{z^2} + \sum_{0 \neq \omega \in \Lambda} \left(\frac{1}{(z-\omega)^2} + \frac{1}{\omega^2} \right)$$

To prove that $\wp_\Lambda$ is well-defined, meromorphic and doubly-period is a standard result, which can be found in textbooks on modular forms or elliptic curves, see for instance Diamond and Shurman (2005) or Silverman (1994). The same texts provide a proof of the fact that $\wp'_\Lambda$ is doubly-periodic and meromorphic as well.

Example 1.1.8 (Algebraic Geometry). Let k be an algebraically closed field, for instance $\mathbb{C}$. Now consider the space $X = k^n$ for some $n \in \mathbb{N}$ as endowed with the Zariski topology. This topology is generated by the basic closed sets $\mathrm{Z}(I)$ for I an ideal in $k[x_1, \ldots, x_n]$, where $\mathrm{Z}(I)$ is the joint zero-set of all polynomials in the ideal I. Given a polynomial $f \in k[x_1, \ldots, x_n]$, it is easy to see that f restricted to the zero-set of the ideal I is fully determined by $f + I \in k[x_1, \ldots, x_n]/I$. In a sense, I specifies a relation between the *coordinates* of the space k^n, any function restricted to the zero-set of I need only be defined up to this relation.

Due to the Hilbert Nullstellensatz we know of a correspondence between closed subsets of X and radical ideals of the ring $k[x_1, \ldots, x_n]$. This allows us to define $\mathrm{A}(Y)$, the *coordinate ring* of a closed subset of X, as the quotient of the ring $k[x_1, \ldots, x_n]$ by the radical ideal associated to Y. In the case that Y is *irreducible*, that is to say, has no non-trivial decomposition into closed subsets, we know that the ideal associated to Y is a prime ideal. This entails that $\mathrm{A}(Y)$ is a domain. In particular, $\mathrm{A}(X)$ equals the domain $k[x_1, \ldots, x_n]$. When $Y \subseteq X$ is closed and irreducible, we call Y a *variety*.

We call a function $k^n \to k$ regular if it is given by a polynomial, note that this polynomial is unique as k is algebraically closed. A *regular* function on a variety in X to k is simply the restriction of a regular function on all of k^n. Hence regular functions on the variety $Y \subseteq X$ correspond exactly with elements $\mathrm{A}(Y)$.

The above definition of regular functions can be extended to Zariski-open sets $U \subseteq Y$ of a variety $Y \subseteq k^n$. Define a function $f : U \to k$ to be *regular* at P exactly when there exists an open neighbourhood $V \subseteq U$ of P where $f = \frac{g}{h}$ for $g, h \in \mathrm{A}(Y)$ and of course $\mathrm{Z}(h) \cap U = \emptyset$. A function $U \to k$ is now said to be *regular* when it is regular on every point of its domain. It can be proven that on all of Y, the functions that are regular are exactly those that were regular in the definition above.

With the above definition one can define a k-algebra of regular functions $\mathcal{O}_X(U)$ for each open U in X. It is an easy matter to prove that $\mathcal{O}_X((-))$ defines a sheaf on X, due to the inherently local nature of a regular function. Recall Remark 1.1.1 and see that $\mathcal{O}_X(\emptyset)$ now must be the trivial k-algebra, which makes perfect sense, as the empty function is the only function on $\emptyset$ and it is regular on all of its domain. For more details on all of this, see for instance Hartshorne (1977).

Example 1.1.9 (Regular Functions on a Circle). Consider the unit circle X in the complex plane $\mathbb{C}^2$. It is easy to see that it can be defined by the zero-set of the ideal generated by the irreducible polynomial $f := x^2 + y^2 - 1$.

The north pole $P = (0, 1)$ yields a closed subset,[2] and we can thus define stereographic projection on the open $U = X - \{P\}$ as the quotient $f = \frac{x}{y-1}$. It is clear that $f \in \mathcal{O}_X(U)$.

Now consider the points $\pm i$ on the complex line $\mathbb{C}$ and let V be $\mathbb{C} - \{\pm i\}$. On this open we can define the regular functions $g = \frac{2t}{t^2+1}$ and $h = \frac{t^2-1}{t^2+1}$. It is not hard to see that both $\langle g, h\rangle \circ f = \mathrm{id}_U$ and $f \circ \langle g, h\rangle = \mathrm{id}_V$ hold. When one sees the above regular functions as functions on $\mathbb{R}$, then it is clear that the functions are everywhere defined, thus providing some sort of isomorphism between the unit circle without a point and the real line.

Example 1.1.10 (Elliptic Curves Algebraically). We can now put Example 1.1.7 in the more algebraic setting of Example 1.1.8. First consider the field of meromorphic functions of $\mathcal{C}/\Lambda$, and recall that it both contains $\wp$ and $\wp'$. It is a basic result that the field of meromorphic functions on $\mathcal{C}/\Lambda$ satisfies an equation of the form $\wp_\Lambda'^2 = 4\wp_\Lambda^3 - g_2(\Lambda)\wp_\Lambda - g_3(\Lambda)$, where $g_2(\Lambda)$ and $g_3(\Lambda)$ are normalized Eistein series related to the lattice Λ.

We can thus make a mapping

$$\mathcal{C}/\Lambda \to \mathcal{C}, \quad z \mapsto \left(\wp_\Lambda'(z), \wp_\Lambda(z)\right)$$

and note that every element in the image is contained in the Zariski-closed $X = \mathrm{Z}\left(y^2 - g_2(\Lambda)x - g_3(\Lambda)\right)$. Now consider the sheaf $\mathcal{O}_X((-))$ of regular functions on X, which turns out to be an algebraic curve of genus 1. It is a consequence of the above and the Riemann–Roch theorem that the field of regular functions on all of X is isomorphic to the field of meromorphic functions on $\mathcal{C}/\Lambda$. For more details on this see Silverman (1994, Chapter 4).

Example 1.1.11 (Maximal Ideal Spectrum). Consider a variety X in k^n and a variety Y in k^m for some algebraically closed field k. A function $f = (f_1, \ldots, f_m) : X \to Y$ is said to be *regular* when all $f_i : X \to k$ are. One can prove a one-to-one correspondence between k-algebra homomorphisms $\mathcal{O}_X(X) \cong \mathrm{A}(X) \to \mathrm{A}(Y)$ and regular maps $Y \to X$. Moreover, realize that the maximal ideals in X are in one-to-one correspondence with the points of X. Now as $\mathrm{A}(X)$ is a finitely generated and reduced k-algebra, the above suggests working with this finitely generated (reduced) k-algebra instead of X.

Let A be a finitely generated k-algebra, then one can define Specm A as the collection of maximal ideals in A. Given a maximal ideal $\mathfrak{m}$ of A, it can be proven that $A/\mathfrak{m} \cong k$, so there are maps $\rho_x : A \to k, x \mapsto x \mod \mathfrak{m}$. An element $a \in A$ gives rise to a map $f_a = x \mapsto \rho_a(x) : A \to k$, and we can define $\mathrm{Z}(a) := \ker f_a$. Equipping Specm A with the topology generated by the closed sets $\mathrm{Z}(a)$ for $a \in A$, we see that any k-algebra homomorphism $\phi : A \to B$ gives rise to a continuous map $\mathrm{Specm}\, B \to \mathrm{Specm}\, A$ defined by $\mathfrak{m} \mapsto \phi^{-1}(\mathfrak{m})$. When $A = \mathrm{A}(X)$, it can be proven that the spaces Specm A and X are homeomorphic. Through this correspondence, taking $f \mod \mathfrak{m}_p$ with $f \in \mathrm{A}(X)$ and $\mathfrak{m}_p$ the maximal ideal of $\mathrm{A}(X)$ corresponding to the point $p \in X$, is the same thing as evaluating f at the point p.

Consider the map

$$A \to \mathbf{Top}(\mathrm{Specm}\, A,\, k), \quad a \mapsto f_a = \left(\mathfrak{m} \mapsto a \mod \mathfrak{m}\right)$$

and realize that this is a map of k-algebras. Let a be a nilradical of A, which means that $a^l = 0$ for some l. This means that a is contained within all prime ideals, and hence contained within all maximal ideals. As a consequence, $a \mod \mathfrak{m} = 0$ for all maximal ideals, so f_a is the constant zero map. One can prove the converse to be true as well, so when A is reduced, the above mapping is injective. In particular, when A is the coordinate ring of the variety X, the above embeds $\mathrm{A}(X)$ into the k-valued functions on $\mathrm{Specm}\,\mathrm{A}(X) \cong X$.

Think for a moment about A as if it were the coordinate ring R of X. An element f in the field of fractions of A (function field of X) is said to be *regular* at a point $\mathfrak{m}$ of A (maximal ideal of R) exactly when $f = \frac{g}{h}$ for $g, h \in A$ and $g \notin \mathfrak{m}$. Take care to note that this only makes sense if A is reduced, for otherwise the field of fractions does not even exist. Using this, it makes sense to define a sheaf of regular functions on A by setting $\mathcal{O}_{\mathrm{Specm}\, A}(U)$ as the k-algebra of regular functions on U in this new sense. In the case that A in fact is the coordinate ring of X, both meanings of regular coincide.

It is also not hard to see that $\mathcal{O}_X(X - \{p\}) = \mathfrak{m}_p^{-1}\mathrm{A}(X)$, that is to say, the regular functions of X on X but a point are the coordinates, localized at the maximal ideal associated to p. In the setting of maximal ideal spectra, the similar fact $\mathcal{O}_{\mathrm{Specm}\, A}(\mathrm{Specm} - \{\mathfrak{m}\}) = \mathfrak{m}^{-1}A$ holds. This leads us to define the sheaf of regular functions on a (not necessarily reduced) k-algebra A as follows. On a basic open $U = \mathrm{Specm}\, A - \mathrm{Z}(a)$ with $a \notin \sqrt{(0)}$, define

[2] Using a dimension argument one can even prove that all closed subsets of X are finite unions of points.

$\mathcal{O}_{\mathrm{Specm}\,A}(U)$ as the localization of A at the multiplicative system $\{1, a, a^2, \ldots\}$. Then extend this appropriately to unions of these basic opens.

Example 1.1.12 (A More Sober Approach). Consider some algebraically closed field k, and realize that k itself is irreducible when endowed with the Zariski topology. Every point λ in k is closed, as it can be defined as the zero-set of the polynomial $f = x - \lambda$. It thus follows that k is not a *sober space*,[3] for k is not equal to a single point.

Given an arbitrary ring R define Spec R to be the *prime ideal spectrum* of R, which as a set consists of all prime ideals of R. Endow Spec R with the topology generated by the basic closed sets $\mathrm{Z}(f) := \{\mathfrak{p} \in \mathrm{Spec}\, R \mid f \in \mathfrak{p}\}$ for all $f \in R$. Now if R is a finitely generated k-algebra, we can consider the space Specm R and the space Spec R. It is clear that the former is a subset of the latter, and as a closed subset $\mathrm{Z}(a) = \ker f_a = \ker(\mathfrak{m} \mapsto a \mod \mathfrak{m})$ is the intersection of $\mathrm{Z}(a)$ with Specm A we can see it as a subspace as well. Similar to the situation of the maximal ideal spectrum, one can define the sheaf of regular functions on Spec R at a basic open $U_a = \mathrm{Spec}\, R - \mathrm{Z}(a)$ by $\mathcal{O}_{\mathrm{Spec}\,R}(U_a) = \{1, a, a^2, \ldots\}^{-1} R$.

Now consider the space Spec k, and let $\mathfrak{p} = (0)$ be the zero-ideal, which is prime as k is a domain. Suppose that Y is a closed set containing $\mathfrak{p}$, which means that it is the finite union of some basic closed sets. In particular, it contains $\mathrm{Z}(f)$ for some $f \in k$. It must follow that $f \in (0)$, so $f = 0$. This proves that $Y = \mathrm{Spec}\, k$, hence the closure of the single (and unique) point $\mathfrak{p}$ is the entire space k. We thus know that every irreducible closed set within k is the closure of a unique point, so Spec k is a sober space. It can be proven that Spec R is sober for every ring R.

As each sheaf is a functor $\mathrm{Opens}(X)^{\mathrm{op}} \to \mathbf{Set}$, it is in particular a presheaf on X. For convenience, define $\mathbf{PSh}(X)$ to be the *category of presheaves on X*. That is to say, the category whose objects are contravariant functors from $\mathrm{Opens}(X)$ to **Set** and whose arrows are natural transformations.

Definition 1.1.2 (Topos of Sheaves). Given a topological space X, define $\mathbf{Sh}(X)$ to be the *category of sheaves on X*, whose objects are sheaves on X and whose arrows are natural transformations between these sheaves.

The category $\mathbf{Sh}(X)$ is a *topos*, which in particular means that it has small limits, small colimits and a subobject classifier. Moreover, the inclusion functor which maps sheaves to presheaves preserves limits. Proofs for these statements can be found in many standard texts, see for instance Mac Lane and Moerdijk (1991, Chapter 2). We now wish to formulate an alternative definition of a sheaf. To this end we re-capture the notion of an open cover. Given a collection $\mathcal{U}$ of opens in X one can form the presheaf $S_{\mathcal{U}}$, which sends V to the one-point set $\mathbf{1}$ if $V \subseteq U$ for some $U \in \mathcal{U}$ and to $\emptyset$ otherwise. As a consequence, $S_{\mathcal{U}}$ is a subfunctor of $\mathbf{y} \bigcup \mathcal{U} = \mathrm{Opens(X)}\left((-), \bigcup \mathcal{U}\right)$.

Definition 1.1.3 (Sieve). Given a topological space X, a *sieve* on the open U in X is a subfunctor of $\mathbf{y}U = \mathrm{Opens(X)}\left((-), U\right)$. To a sieve S we associate the set $\{V \in \mathrm{Opens}(X) \mid (V \subseteq U) \in S(V)\}$. A sieve is said to be *covering* exactly when the associated set covers U.

Remark 1.1.2 (Sieve as Sheaf). A sieve on U is said to be *principal* when $S = \mathbf{y}V$ for some open $V \subseteq U$. Note that a sieve is principal precisely if it is a sheaf. The one direction follows immediately from Example 1.1.2.

To prove the other, assume S is a sheaf. Consider the set covered by S, say $V := \bigcup S$ and the inclusion $i : S \to \mathbf{y}V$ of S into the maximal sieve on V. We will show that this map is in fact an isomorphism. It suffices to show that it is a surjection on components, as it already is an injection there. To show this, we need to show that for any open $W \subseteq V$ the inclusion $W \subseteq V$ is in the image of i. Take such an open W and construct the following open cover of W.

$$\mathcal{W} := \{T \cap W \mid T \text{ open and } (T \cap W \subseteq U) \in S(T \cap W)\}$$

For each $T \in \mathcal{W}$ we have $f_T := T \subseteq U \in S(T)$. This gives us a map from $\mathbf{1}$ to $\prod_{T \in \mathcal{W}} S(T)$, satisfying the here vacuous requirements of equalizing p and q. Consequently, this gives us an element of $S(W)$, which was what we set out to prove.

In summary, representable presheaves $\mathbf{y}U$ for U open in X are sheaves, for they are exactly the principal sieves.

[3] A topological space X is said to be sober when every irreducible closed subset is the closure of a unique point.

In Definition 1.1.1 we see that a map η starting in **1** which equalizes both p and q gives rise to the unique element in $f \in F(U)$ satisfying $\eta(*) = e(f)$. Think of such an element as the 'collation' of the components of η into one element of $F(U)$. We now define a notion quite similar to a map equalizing p and q, which we use in Proposition 1.1.1 to characterize sheaves in a different way.

Definition 1.1.4 (Matching Family). A *matching family* for a presheaf P on a sieve S is a map $S \to P$.

An element $f \in P(U)$ for a presheaf P gives rise to a matching family for any sieve S in U. Indeed, simply define $\eta : S \to P$ on components as

$$\eta_V = \big(V \subseteq U \mapsto f \restriction V\big) \text{ if } (V \subseteq U) \in S(V), \text{ otherwise } \emptyset$$

This now clearly is a natural transformation, which sends f to its respective restrictions. Realize that if S is principal, as in Remark 1.1.2, then we have $\mathbf{PSh}(\mathrm{X})\,(\mathbf{y} \bigcup S,\ P) \cong P(\bigcup S)$. This means that matching families of principal sieves are always of the above form. This directly applies to the following proposition, where we note that $\mathbf{PSh}(\mathrm{X})\,(\mathbf{y}U,\ P)$ is in fact isomorphic to $P(U)$ via the Yoneda-lemma. Consequently, the proposition states that every matching family can be uniquely *amalgamated* to an element of $P(U)$.

Proposition 1.1.1 (Characterization of Sheaves). *A presheaf P on a topological space X is a sheaf if and only if the induces map below is an isomorphism for every covering sieve S on any open U*

$$\mathbf{PSh}(\mathrm{X})\,(\mathbf{y}U,\ P) \xrightarrow{\mathbf{PSh}(\mathrm{X})(i_S : S \to \mathbf{y}U,\ P)} \mathbf{PSh}(\mathrm{X})\,(S,\ P)\,. \tag{1.2}$$

Proof. Fix an arbitrary presheaf P on a topological space X and an open set U. Let $\mathcal{V}$ be an arbitrary open cover of U and let S be the associated sieve. Note that we touch upon every possible sieve S in this manner. Consider the following diagram, which is used in the definition of a sheaf above.

$$E \xrightarrow{d} \prod_{V \in \mathcal{V}} F(V) \underset{q}{\overset{p}{\rightrightarrows}} \prod_{V,W \in \mathcal{V}} F(V \cap W)$$

Using the canonical construction of equalizers in **Set**, we see E to be the set of tuples $\mathbf{x} \in \prod_{V \in \mathcal{V}} F(V)$ such that $x_V \restriction V \cap W = x_W \restriction W \cap V$ for all opens V and W in the cover $\mathcal{V}$. The map d in this setting is simply the inclusion map.

We will now identify such a tuple $\mathbf{x}$ with a matching family $\eta^{\mathbf{x}} : S \to P$. We distinguish between $S(V)$ being empty or non-empty, or rather, on whether $V \subseteq U \in S$. Realize that $S(V) \neq \emptyset$ exactly if $V \subseteq W$ for some $W \in \mathcal{V}$. For such a V define $\eta^{\mathbf{X}}_V$ to map the inhabitant of $S(V)$ to $x_W \restriction V$. Note that this is well-defined, for if $V \subseteq W'$ then

$$x_W \restriction V = x_W \restriction W \cap V = x_V \restriction W \cap V = x_V \restriction W' \cap V = x_{W'} \restriction W' \cap V = x_{W'} \restriction V.$$

Now to treat the case when $S(V) = \emptyset$, take $\eta^{\mathbf{X}} : S \to P$ to be the unique map starting in the empty set. The map d now sends η to $\big(\eta_V(*)\big)_{V \in \mathcal{V}}$, and starts in $\mathbf{Sh}(\mathrm{X})\,(S,\ P)$.

Now consider the diagram below, where i denotes $\mathbf{Sh}(\mathrm{X})\,(i_S,\ P)$. See that the square on the left commutes.

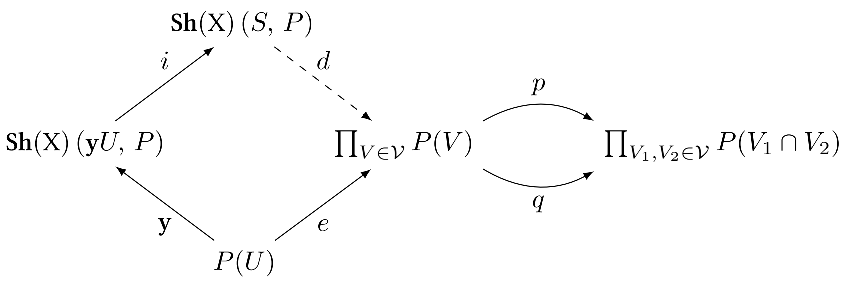

Suppose that P is a sheaf. This makes the lower-right of the diagram an equalizer. Now as $\mathbf{Sh}(\mathrm{X})(S, P)$ is also an equalizer, we know these $P(U)$ to be isomorphic to $\mathbf{Sh}(\mathrm{X})(S, P)$. By the Yoneda-embedding, the former is isomorphic to $\mathbf{Sh}(\mathrm{X})(\mathbf{y}U, P)$, proving the desired. Conversely, if (1.2) is an isomorphism, then the commuting square on the top makes e an equalizer as desired. ☐

Remark 1.1.3 (Invariant under Isomorphism). Suppose that P and Q are isomorphic presheaves. This induces the following isomorphisms.

$$\mathbf{PSh}(\mathrm{X})(\mathbf{y}U, P) \cong \mathbf{PSh}(\mathrm{X})(\mathbf{y}U, Q) \qquad \mathbf{PSh}(\mathrm{X})(S, P) \cong \mathbf{PSh}(\mathrm{X})(S, Q)$$

We thus see that an isomorphism as in (1.2) for P gives rise to such an isomorphism for Q. Consequently, P is a sheaf exactly if Q is, making the notion of sheaf invariant under isomorphism.

Via the techniques as displayed in the above proofs we can identify a collection $\{f_V \in P(V)\}_{V\in\mathcal{V}}$ for an open cover $\mathcal{V}$ of U in X satisfying $f_V \upharpoonright W \cap V = f_W \upharpoonright W \cap V$ for any pair of opens V and W in the cover $\mathcal{V}$ with a matching family $S_{\mathcal{V}} \to P$ for the covering sieve $S_{\mathcal{V}}$ of U. It thus suffices to show that each such *compatible family* has a unique amalgamation, in order to prove P to be a sheaf.

The terminal object in the category $\mathbf{Sh}(X)$ is the representable sheaf $\mathbf{y}X = \mathrm{Opens}(\mathrm{X})((-), X)$. Indeed, any sheaf F on X can be mapped to $\mathbf{y}X$ on components by using the unique map, proving the existence of a unique map $F \to \mathbf{y}X$. Maps into $\mathbf{y}X$ carry information about the space X. Below we show that subobjects of $\mathbf{y}X$ corresponds to opens in X. To see this, first observe that any subobject can be represented by a subsheaf. Pick a representative $m : F \to X$ and note that when considered as a map of presheaves, it is a component-wise injection.[4] This makes $F(U)$ isomorphic to a subset of $\mathbf{y}X(U)$, and as naturality is vacuous we get a subsheaf of $\mathbf{y}X$ isomorphic to F.

Proposition 1.1.2 (Sheaves Reflect Opens). *Let X be a topological space. There is an isomorphism*

$$\mathrm{Opens}(X) \cong \mathrm{Sub}_{\mathbf{Sh}(X)}(\mathbf{1}),$$

in other words, subsheaves of the terminal object are opens of X.

Proof. We need to define a functor $\phi : \mathrm{Opens}(X) \to \mathrm{Sub}_{\mathbf{Sh}(X)}(\mathbf{1})$ and a functor ψ in the opposite direction, such that these are mutually inverse. Define ϕ using the Yoneda-embedding in $\mathrm{Opens}(X)$, mapping an open U to the subsheaf $\mathrm{Opens}(\mathrm{X})((-), U)$. Given a subobject $S \subseteq \mathbf{y}X$, represented as subsheaf, we have in fact a sieve and we can map this to the open $\bigcup S$. Recall that a sieve S is principal — which would entail that $S = \phi(\bigcup S)$ — precisely if it is a sheaf. This clearly makes the functors mutually inverse. ☐

We now define the notion of a stalk. When considering the sheaf of continuous maps on a topological space, a stalk contains information about the behavior of functions near a certain point.

Definition 1.1.5 (Stalk). Given a topological space X, a point $p \in X$ and a sheaf F on X, we define the *stalk* of F at p to be

$$F_p := \operatorname*{colim}_{p\in U\in\mathrm{Opens}(X)} F(U).$$

Example 1.1.13 (Local Behaviour of Holomorphic Functions). Let X be a Riemann surface and $p \in X$ a point. Recall the sheaf of holomorphic functions $F = \mathcal{O}((-), \mathbb{C})$ on X, and consider now F_p. It consists of pairs $\langle U, f : U \to \mathbb{C}\rangle$ where $p \in U \subseteq X$ is an open neighbourhood of p and f is a holomorphic function defined on U, up to common refinement. To be a bit more precise, pairs $\langle U, f\rangle$ and $\langle V, g\rangle$ are considered equal when $f \upharpoonright W = g \upharpoonright W$ for some open neighbourhood W of p contained within both U and V. It is easy to see that this is a ring, and that $f \in \mathcal{O}(U, \mathbb{C})$ represents an invertible element exactly if $f(p) \neq 0$. As a consequence, the sum of two non-invertible elements of F_p is again non-invertible, proving F_p to be a local ring.

Example 1.1.14 (Local Behaviour of Regular Functions). Take X to be a variety in k^n for some algebraically closed field k, endowed with the sheaf $\mathcal{O}_X((-))$ of regular functions. Similar to the above, those regular functions which evaluate to zero on a point $p \in X$ are exactly the representatives of the non-invertible elements in $\mathcal{O}_X((-))$. It is easy to see that they form an ideal, precisely the kernel of evaluation at p. This is the maximal ideal $\mathfrak{m}_p$, the ideal corresponding to the point p. It thus it not hard to derive that $\mathcal{O}_X((-))_p$ in fact equals the localization of $\mathrm{A}(X)$ at the maximal ideal $\mathfrak{m}_p$.

[4] This follows from the earlier remark that the inclusion of sheaves into presheaves preserves limits, and a monomorphism can be expressed as a pullback.

SECTION 1.2 Étale Spaces

In this section we study another way to specify sheaves over a topological space X. The notion of a sheaf was first formulated in the manner we describe here, originally by Leray and in the following years by Cartan (1950-1951), Serre (1955) and Grothendieck (1955). The latter two transported sheaves from algebraic topology and analysis to algebraic geometry, where they gained a lot of popularity.

We will see a sheaf on a topological space X as some special kind of fibre space, or space over X. First we formally define the type of space over X which we will later show to be equivalent to a sheaf on X, and prove some elementary properties. In Subsection 1.2.1 we actually prove this acclaimed equivalence, which we will use throughout this text.

In this topological setting we can more easily show what type of a functor a continuous map $f : X \to Y$ induces between $\mathbf{Sh}(Y)$ and $\mathbf{Sh}(X)$. We will construct an inverse image map $f^* : \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ and a direct image map $f_* : \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$ in Subsection 1.2.2. We prove that the map f^* is left adjoint to f_* and show that it preserves finite limits. This situation is the motivating example for geometric morphisms as covered in Section 1.4. The remainder of this section is spent on discussing two edge-cases of this. In Subsection 1.2.3 we study the unique map to the one-point space, and we will show that the related inverse image map has a left adjoint when X is locally connected. Information about these maps will be very important to Section 2.3. A point of the space X is essentially the same thing as a map $\mathbf{1} \to X$. We study such maps in Subsection 1.2.4, for these are crucial for the construction discussed in Section 2.2.

Definition 1.2.1 (Local Homeomorphism). A local homeomorphism is a continuous map $p : E \to X$ between topological spaces such that for each point $e \in E$ there is an open neighbourhood $e \in U \subseteq E$ which has an open p-image homeomorphic to U via p.[5]

Consider the category $\mathbf{Top}/X$ for some topological space X. Its objects are simply continuous maps $p : E \to X$, or spaces over X. We call such an object an étale space when p is a local homeomorphism, think of E as a space 'spread out' over X. There is a category $\mathbf{Etale}(X)$, which is the full subcategory of $\mathbf{Top}/X$ whose objects are étale spaces.

Example 1.2.1. Every homeomorphism is a local homeomorphism, in particular $\mathrm{id}_X : X \to X$ is an étale space. Note that every object $p : E \to X$ of $\mathbf{Etale}(X)$ has a unique morphism to id_X, being the map p itself. This shows that id_X is the terminal object in the category $\mathbf{Etale}(X)$.

Given a local homeomorphism $E \to X$, many topological properties hold for X precisely when they hold for E. Amongst these properties are locally connected, locally path-connected, locally compact and first-countable. We prove the first of these, as it will be of importance later on.

Lemma 1.2.1. *Let $p : E \to X$ be an étale space. If X is locally connected, then so is E.*

Proof. Let $e \in E$ be an arbitrary point and let W be an open neighbourhood of e. There is an open neighbourhood U of e such that $p(U)$ is open and $p \restriction U$ is a homeomorphism, fix such an open U. We now see that $p(U \cap W)$ is open, and hence contains a connected open neighbourhood of $p(e)$, say V. As $p \restriction U$ is a homeomorphism, we know $e \in p^{-1}(V)$ to be open and connected, proving the desired. □

Remark 1.2.1 (Local Behaviour of Sections). On a small neighbourhood, a section $\sigma : U \to X$ of an étale space $p : E \to X$ looks like 'p^{-1}'. More precisely, for any point $x \in U$ there is an open $x \in V \subseteq U$ such that $p \restriction V$ is a homeomorphism, and as such, $(p \restriction V)^{-1}$ is a section of p as well. We know that $p \circ \sigma = \mathrm{id}_U$, so by composing with the inclusion map we obtain $(p \restriction V) \circ (\sigma \restriction p(V)) = \mathrm{id}_{p(V)}$. From this we readily derive that locally, the section is the inverse to p, as follows from

$$(p \restriction V)^{-1} \circ (p \restriction V) \circ (\sigma \restriction V) = \sigma \restriction V = (p \restriction V)^{-1}$$

[5]For a local homeomorphism $p : E \to X$, we often take some open U around a point e of E such that $p(U)$ is open and homeomorphic to U via p. For reasons of brevity we often say that $p \restriction U$ is a homeomorphism, in the understanding that $p \restriction U$ denotes the map p restricted to U, with the codomain restricted to the open $p(U)$.

There are different ways to specify a local homeomorphism $E \to X$. In Definition 1.2.1 we mention points of the space E, this is not necessary. The following lemma proves several notions of local homeomorphism to be equivalent.

Lemma 1.2.2 (Notions of Local Homeomorphism). *Let $p : E \to X$ be a continuous map. The following are equivalent:*

(i) p is a local homeomorphism;

(ii) there exists an open cover $\mathcal{U}$ of E such that $p(U)$ is open and $p \restriction U$ is a homeomorphism for all $U \in \mathcal{U}$;

(iii) both p and the diagonal $\Delta : E \to E \times_X E$ are open maps.

Proof. First suppose that (i) holds. Assign to each point $e \in E$ the collection $\mathcal{U}_e$ of all open neighbourhoods U of e such that $p(U)$ is open and $p \restriction U$ is a homeomorphism. It is easy to see that $\mathcal{U} := \bigcup_{e \in E} \mathcal{U}_e$ is an open cover of E which satisfies the properties needed to prove (ii). The converse is equally easy.

Now suppose that (ii) holds. To prove that p and Δ are open, fix some open V in E. Note that for each U in the open cover $\mathcal{U}$ we know $V \cap U$ to be open in U, and as $p \restriction U$ is a homeomorphism we know $p(U \cap V)$ to be open as well. With this data we can compute

$$p(V) = p\left(\bigcup_{U \in \mathcal{U}} V \cap U\right) = \bigcup_{U \in \mathcal{U}} p(U \cap V),$$

which makes it clear that $p(V)$ is open.

To prove that Δ is open, we need to show that $\Delta(V)$ is open as well. We are done when we can show that $\Delta(U \cap V)$ equals the intersection of $\pi_1^{-1}(U \cap V)$ and $\pi_2^{-1}(U \cap V)$. It is easy to see that

$$\pi_1^{-1}(U \cap V) \cap \pi_2^{-1}(U \cap V) = \{(x, y) \mid x \in U \cap V,\ y \in U,\ p(x) = p(y)\}$$

holds by construction. Now as p is a homeomorphism on U, the statement $p(x) = p(y)$ implies $x = y$. As a consequence, the above equals $(U \cap V) \times (U \cap V) = \Delta(U \cap V)$.

Finally, suppose (iii) holds. We wish to provide an open cover $\mathcal{U}$ of E such that $p \restriction U$ is injective.[6] It is easy to see that if $p \restriction U$ is injective then $p \restriction U$ is in fact a homeomorphism, as p was assumed to be open. We now let $\mathcal{U}$ be the collection of opens U in E such that $U \times U \cap E \times_X E$ is contained within the diagonal $D := \Delta(E)$. Compute

$$D = D \cap \bigcup_{U, V \text{ open in } E} \Big((U \times V) \cap \big(E \times_X E\big)\Big) = \bigcup_{\substack{U, V \text{ open in } E \\ (U \times V) \cap (E \times_X E) \subseteq D}} U \times V = \bigcup_{U \in \mathcal{U}} U \times U,$$

so as $\pi_1(D) = E$ it follows that $\mathcal{U}$ is indeed an open cover of E. It is easy to see that $p \restriction U$ is injective for any U in the cover. Indeed, if $p(x) = p(y)$ holds for $x, y \in U \in \mathcal{U}$, then $(x, y) \in U \times U \subseteq D$, so $x = y$ must follow. This proves (ii) to hold. ☐

Remark 1.2.2. Let $p : E \to X$ be a local homeomorphism. For each subset $U \subseteq E$ it holds that U is open precisely if $p(U)$ is. The one direction follows directly from (iii) of the above lemma. To prove the other, suppose $p(U)$ is open. Now $U \subseteq p^{-1}p(U)$, so by the above we know of an open cover $\mathcal{V}$ of $p^{-1}p(U)$ such that p restricts to homeomorphisms to opens on elements of this cover. Take $V \in \mathcal{V}$ and note that $U \cap V$ is mapped homeomorphically to $p(U \cap V) = p(U) \cap p(V)$, which is open. Hence $U \cap V$ is open for all $V \in \mathcal{V}$, proving U to be open.

With this lemma we can more easily prove properties of étale spaces. Below we prove that maps between étale spaces must be local homeomorphisms.

Lemma 1.2.3. *A map between étale spaces is itself a local homeomorphism.*

[6] This is always surjective, as we by convention defined $p \restriction U$ to be the map which is equal to p, but with U as domain and $p(U)$ as codomain.

Proof. The situation is as in the following diagram, where both p and q are local homeomorphisms.

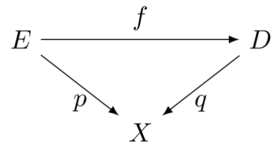

By Lemma 1.2.2 we know of a cover $\mathcal{U}_E$ of E such that $p(U)$ is open in X and $p \restriction U$ is a homeomorphism for any open $U \in \mathcal{U}_E$. On the other hand, we also know of a cover $\mathcal{U}_D$ of D such $q(U)$ is open in X and $q \restriction U$ is a homeomorphism for any $U \in \mathcal{U}_E$. Now define

$$\mathcal{U} := \Big\{ U \cap p^{-1}(q(V)) \mid U \in \mathcal{U}_E,\ V \in \mathcal{U}_D \Big\}.$$

See that as $p(U)$ is open in X, it follows that $q^{-1}(p(U)) \cap V$ is open in D and thus in V. This makes its q-image open in X, which in turn makes its p-preimage open in E. This clearly is a cover of E, and for any $W \in \mathcal{U}$ we see that $\big(q \restriction q^{-1}(p(W))\big) \circ \big(f \restriction W\big) = p \restriction W$, which entails that $f = \big(q \restriction q^{-1}(p(W))\big)^{-1} \circ \big(p \restriction W\big)$. From here, the desired is clear. □

Now that we have some machinery, we can look at examples of étale spaces. One probably encountered Example 1.2.2 in a basic course on topology, where it is used to compute the fundamental group of the circle. Elliptic curves are a central object of study in mainstream mathematics, so Example 1.2.3 can be found in many incarnations throughout basic literature. We point out in particular Diamond and Shurman (2005), where the relation between the tori, elliptic curves and moduli spaces is explicated. We also used the spaces as given in these examples to give examples of sheaves in the previous section.

Example 1.2.2 (Spiral Cover). Consider the real line $\mathbb{R}$ lying above the unit circle. More precisely, consider the quotient map $p : \mathbb{R} \to \mathbb{R}/\mathbb{Z}$. It is easy to see that each interval of size less than one in $\mathbb{R}$ is mapped homeomorphically to an arc on the unit circle, thus providing $\mathbb{R}$ with the cover needed to make p a local homeomorphism.

Example 1.2.3 (Elliptic Curve). Let Λ be a lattice in $\mathbb{C}$. The quotient map $p : \mathbb{C} \to \mathbb{C}/\Lambda$ is a local homeomorphism. This can easily be proven by considering translations of a fixed fundamental domain of Λ in the complex plane $\mathbb{C}$, on the interior of which p clearly is homeomorphic.

Example 1.2.4 (Covering Map). Let X be a connected manifold. There exists a connected and simply connected manifold $\tilde{X}$ and a covering map $p : \tilde{X} \to X$, which in is the universal covering space of X. For a proof of this, see for instance Forster (1981). In particular, this universal covering space is an étale space. This because a covering map $p : E \to X$ is a surjective, continuous map where each point in X has an open neighbourhood whose pre-image decomposes into disjoint opens, each of which homeomorphic to U via p. It is easy to see that this entails p to be a local homeomorphism.

Example 1.2.5 (**Set** as Category of Étale Spaces). Consider for a moment the category **Etale**(**1**), and note that objects simply are arrows $! : S \to \mathbf{1}$ in **Top**. Realize that ! can only be a local homeomorphism if S is endowed with the discrete topology. From here it is clear that **Etale**(**1**) $\cong$ **Set**, so étale spaces are a generalization of sets.

In order to formulate the following two examples, we need the small fact proven below. In Example 1.2.6 we construct a functor sending sets to étale spaces, this functor is constructed by pullback in the category **Top**. We then construct an étale space out of a presheaf in Example 1.2.7, setting the first step to prove that sheaves and étale spaces are basically 'the same thing'.

Lemma 1.2.4 (Étale Stable under Pullback). *Let $p : E \to Y$ be an étale space and $f : X \to Y$ be a continuous map. Now the pullback of p along f (in the category of topological spaces) is an étale space as well.*

Proof. Let $(x, e) \in f^*E \subseteq E \times E$ be arbitrary, and recall that $f(x) = p(e)$. Take some open neighbourhood U around e such that $p(U)$ is open and $p \restriction U$ a homeomorphism. Now $f^{-1}(p(U)) \times U$ is an open neighbourhood of (x, e), and there is a map $g : f^{-1}(p(U)) \to f^{-1}(p(U)) \times U$ defined by $g(x) = (x, p^{-1}(f(x)))$ which is clearly continuous and mutually inverse to f^*p. This proves the desired. □

Remark 1.2.3. Suppose $p : U \to Y$ is an isomorphism and $f : X \to Y$ a map, both in the category **Top**. We can consider the pullback of p along f, and remark that $f^{-1}(p(U))$ is this pullback, with the obvious inclusions. This in particular entails that if $U \subseteq Y$ is an open such that the étale map $p : E \to Y$ restricts to an isomorphism $p \restriction p^{-1}(U)$, then the pullback of this isomorphism along a continuous map f is simply $f^{-1}(U)$.

Example 1.2.6 (Functor between **Etale**$(\mathbf{1})$ and **Etale**(X)). Recall that **Etale**$(\mathbf{1}) \cong$ **Set**, as demonstrated in Example 1.2.5. Now consider the pullback diagram,

$$\begin{array}{ccc} X \times_{\mathbf{1}} S & \longrightarrow & S \\ {\scriptstyle \pi_X}\big\downarrow & & \big\downarrow{\scriptstyle !} \\ X & \xrightarrow{\ !\ } & \mathbf{1} \end{array}$$

which simply expresses the product of X and S with the usual projections. By Lemma 1.2.4 we now know $\pi_X : X \times S \to X$ to be an étale space as well. This yields a functor $\pi^* : \mathbf{Etale}(\mathbf{1}) \to \mathbf{Etale}(X)$, defining the map on arrows in the obvious manner:

$$T \xrightarrow{f} S \qquad \mapsto \qquad \begin{array}{ccc} X \times T & \xrightarrow{\langle \pi_X, f \rangle} & X \times S \\ & {\scriptstyle \pi_X}\searrow \quad \swarrow{\scriptstyle \pi_X} & \\ & X & \end{array}$$

Example 1.2.7 (Display Space of a Presheaf). Let $P : \mathrm{Opens}(X)^{\mathrm{op}} \to$ **Set** be a presheaf on a topological space X. One can construct the display space

$$\Lambda P := \coprod_{x \in X} P_x$$

This set comes equipped with a canonical projection $\pi : \Lambda P \to X$ which sends elements of the component at x to x. Given an open U of X and an element $s \in P(U)$ we have a map $\rho_s : U \to \Lambda P$ which sends each x to the stalk of s at x, in symbols, $s_x \in P_x$. We now endow ΛP with the topology generated by the basic open sets $\rho_s(U)$ for each open U in X and $s \in P(U)$.

This topology makes π is continuous, for $\pi^{-1}(U) = \bigcup_{s \in P(U)} \rho_s(U)$ holds by the construction of ΛP. For each open U of X we know that the open $\rho_s(U) =: V$ of ΛP is mapped to U. Moreover, both $(\pi \restriction V) \circ \rho_s = \mathrm{id}_U$ and $\rho_s \circ (\pi \restriction V) = \mathrm{id}_V$ hold. We thus know π to be a local homeomorphism by Lemma 1.2.2.

Consider opens U and V in X, and pick some point x in the intersection. If there are $t \in P(U)$ and $s \in P(V)$ such that $\rho_t(x) = \rho_s(x)$, then this means we have some open neighbourhood W of y within $U \cap V$ such that $s \restriction W = t \restriction W$. As such, the set of all points x within $U \cap V$ such that $\rho_s(x)$ equals $\rho_s(y)$ is open. From this it is clear that the pre-image of $\rho_t(U)$ for some open U in X under the map $\rho_s : U \to \Lambda P$ is itself open, proving continuity of the maps ρ_s. By the above paragraphs we see that these maps are homeomorphisms.

Given a map of presheaves $f : P \to Q$ we define a map $\Lambda f : \Lambda P \to \Lambda Q$ by sending the representative $\langle s, U \rangle$ of a stalk at x on P to the representative $\langle f_U(s), U \rangle$ of a stalk at x on Q. Due to the naturality of f, this is well-defined. This map also clearly respects the projections. Now as the pre-image of the basic open $\rho_s(U)$ under the map Λf is simply the pre-image of $\rho_{f_U(s)}$ (U), we see the map Λf to be continuous by the reasoning of the previous paragraph.

Summarizing the above, we know of a functor $\Lambda = \Lambda((-)) : \mathbf{PSh}(X) \to \mathbf{Top}/X$ which maps presheaves over X to étale spaces over X. We could indeed have chosen the codomain of Λ to be **Etale**(X), but defining the map like this gives a stronger result in Lemma 1.2.6.

There also is a map in the other direction. In fact, when we take a mere continuous map $p : E \to X$ we can construct $\Gamma((-), p) : \mathrm{Opens}(X)^{\mathrm{op}} \to$ **Set**, the *(pre)sheaf of cross-sections* over X. Define it by the following object and arrow maps.

$$U \in \mathrm{Opens}(x) \mapsto \mathbf{Top}/\mathrm{X}\left(\begin{array}{c} U \\ \cup\!\!\mid\ \big\downarrow \\ X \end{array} , \begin{array}{c} E \\ {\scriptstyle p}\big\downarrow \\ X \end{array} \right) \qquad V \subseteq U \mapsto \mathbf{Top}/\mathrm{X}\left(\begin{array}{ccc} V & \xrightarrow{\ \subseteq\ } & U \\ & {\scriptstyle \subseteq}\searrow \quad \swarrow{\scriptstyle \supseteq} & \\ & X & \end{array} , \begin{array}{c} E \\ {\scriptstyle p}\big\downarrow \\ X \end{array} \right)$$

Explicating the constructions above a bit, the object map sends an open U of X to the set of continuous maps $f : U \to E$ such that $pf = \mathrm{id}_U$. An inclusion $V \subseteq U$ sends a continuous map $f : U \to E$ to the actual restriction $f \upharpoonright V : V \to E$, it is easy to see that the necessary diagram commutes.

When given a map $f : p \to q$ of spaces $p : E \to X$ and $g : D \to X$ over X, we can construct a map of presheaves $\Gamma((-), f) : \Gamma((-), p) \to \Gamma((-), q)$ by assigning to $g : U \to X$ the map $g \circ f : U \to D$. The diagram

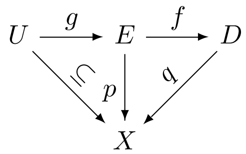

commutes clearly, so we in fact have a functor $\Gamma = \Gamma((-), (-)) : \mathbf{Top}/X \to \mathbf{PSh}(X)$ which maps spaces over X to presheaves on X. As in Example 1.2.7, we can show that anything in the image of Γ is something stronger than merely something in its codomain.

Lemma 1.2.5 (Sheaf of Cross Sections). *For any space $p : E \to X$ over X, the presheaf $\Gamma((-), p : E \to X)$ is a sheaf.*

Proof. Let $\mathcal{V}$ be an open cover of U open in X, and consider a compatible system $\{f_V \in \Gamma(V, p)\}_{V \in \mathcal{V}}$. One can now simply construct $f := \bigcup_{V \in \mathcal{V}} f_V$, which is clearly continuous and still a section of p. This f is the unique amalgamation of the above compatible system, proving the desired. $\square$

With this construction we can see ρ_s for $s \in P(U)$ in a new light. By letting this vary along all values of $P(U)$, we see that $\rho_{(-)}$ gives rise to a natural transformation $\eta : \mathrm{id}_{\mathbf{PSh}(X)} \to \Gamma\Lambda$. This natural transformation is defined as

$$(\eta_P)_U : s \in P(U) \mapsto (x \in X \mapsto \rho_s(x) \in P_x).$$

Given a space $p : E \to X$ over X we can consider the space $\Lambda\Gamma(p : E \to X)$. Realize that elements of $\Gamma(U \subseteq X, p)$ are actual sections of p, so $\Lambda\Gamma(p : E \to X)$ consists of elements of the form $\rho_s(x)$ where $s : U \to E$ is a section of p. As a consequence, one can construct the map

$$\rho_s(x) \in \Lambda\Gamma\left((-), E \xrightarrow{p} X\right) \mapsto s(x) \in E.$$

This map is well-defined, for if $\rho_s(x) = \rho_t(x)$, then $s(x) = t(x)$ holds in particular. Moreover, it gives rise to a natural transformation $\epsilon : \Lambda\Gamma \to \mathrm{id}_{\mathbf{Top}/X}$ by varying over the spaces over X. The following two lemma's occur in many standard texts, a nice formulation is given by Johnstone (1982) in Section 5.1.

Lemma 1.2.6. *The functor Λ is left adjoint to Γ.*

Proof. It is easy to verify that the natural transformations as above constitute an adjunction. $\square$

SUBSECTION 1.2.1 Sheaves and Étale Spaces

We can recognize étale spaces and sheaves as being those spaces over X and presheaves on X which make the corresponding natural transformations into isomorphisms, as formalized in the lemma below. For both statements, one direction is immediate. Indeed, if η_P is an isomorphism, then P is isomorphic to $\Gamma\Lambda P$. By Lemma 1.2.5 we know this to be a sheaf, so P is a sheaf too because the notion of sheaf is invariant under isomorphism as demonstrated in Remark 1.1.3. A similar argument works for this direction in (ii).

Lemma 1.2.7 (Recognizing Sheaves and Étale Spaces). *Given a presheaf P and a space $p : E \to X$, both over X, the following hold:*

(i) P is a sheaf if and only if $\eta_P : P \to \Gamma\Lambda(P)$ is an isomorphism;
(ii) $p : E \to X$ is an étale space if and only if $\epsilon_p : \Lambda\Gamma(p) \to p$ is an isomorphism.

Proof. Suppose that P is a sheaf. We are done when we can show $(\eta_P)_U$ to be both injective and surjective for each open U in X, so fix such an open. Let t and s be of $P(U)$ and suppose that $(\eta_P)_U(s) = \rho_s = \rho_t = (\eta_P)_U(t)$. This means that for all $x \in U$ there is some open V_x of X such that $s \restriction V_x = t \restriction V_x$. Clearly, these V_x cover X, so as P is a sheaf we know this compatible system to have only one solution, forcing $s = t$.

To prove surjectivity, let $h : U \to \Lambda\, P$ be a section. Assign to each $x \in U$ the U_x and $s^x \in P(U_x)$ such that $h(x) = (s^x)_x$. Define $V_x := h^{-1}(\rho_{s^x}(U)) \cap U$ and note that V_x is an open in U as it is the pre-image of a basic open of $\Lambda\, P$ under a continuous map. We thus have a cover $\{V_x\}_{x\in U}$ of U and an element $f_x := s^x \restriction V_x \in P(V_x)$. On all restrictions we know ρ_{f_x} to agree with h for $(h \restriction V_x)(x) = h(x) = (s^x)_x = (f^x)_x$ holds by construction. By the reasoning of the previous paragraph, this proves that $(f_x)_{x\in U}$ is a compatible system. This gives us a unique $f \in P(U)$ such that $f \restriction V_x = f_x = s_x \restriction V_x$. Consequently, we know that $(\eta_P)_U(f)(x) = \rho_f(x) = (s^x)_x = h(x)$, proving (i).

In order to prove (ii), we construct an inverse map θ. Given $e \in E$, we know of an open neighbourhood $e \in U \subseteq E$ such that $p(U)$ is open and $p \restriction U$ is a homeomorphism. This gives us a section $s := (p \restriction U)^{-1} : p(U) \to U \subseteq E$. Now define $\theta(e) = s_{p(e)}$, and note that it is well-defined and continuous. That $\epsilon_p\theta$ is the identity map is clear. From Remark 1.2.1 it is also immediate that $\theta\epsilon_p$ is the identity map. ☐

Remark 1.2.4 (Sheafification). Consider the inclusion functor $i : \mathbf{Sh}(X) \to \mathbf{PSh}(X)$ which maps a sheaf F to the presheaf F; the act of forgetting that F is a sheaf. In the opposite direction we have the sheafification or associated sheaf functor $\mathbf{a} := \Gamma\Lambda : \mathbf{PSh}(X) \to \mathbf{Sh}(X)$, which sends a presheaf to its 'best approximation' by a sheaf. There is an isomorphism

$$\mathbf{Sh}(\mathrm{X})\,(\mathbf{a}P,\ F) \cong \mathbf{PSh}(\mathrm{X})\,(P,\ i(F))\,,$$

that is to say, sheafification is left adjoint to the inclusion of sheaves into presheaves. In the one direction we can send a map $\tau : \mathbf{a}P \to F$ to the composite below on the left. Conversely, given a map $\gamma : P \to i(F) = F$ we send this to the composite below on the right.

$$P \xrightarrow{\eta_P} \mathbf{a}P \xrightarrow{\tau} F \qquad\qquad \mathbf{a}P \xrightarrow{\mathbf{a}\gamma} \mathbf{a}F \xrightarrow{\eta_F^{-1}} F.$$

To see that these maps are mutually inverse, consider the diagrams below.

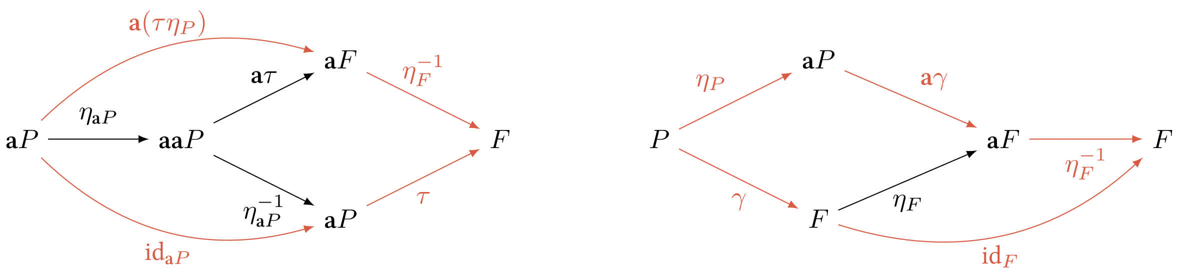

We first inspect the left diagram. Here the left square follows from naturality of η^{-1}, and this natural transformation exists due to Lemma 1.2.7. Note that the arrow $\eta_{\mathbf{a}P}^{-1}$ exists because $\mathbf{a}P$ is a sheaf, as proven in Lemma 1.2.5. The square in the right diagram is naturality of η. We have the desired isomorphism, and this makes $\mathbf{a}$ left adjoint to i, which in particular proves that it preserves colimits.

Theorem 1.2.1 (Sheaves and Étale Spaces). *The categories* $\mathbf{Sh}(X)$ *and* $\mathbf{Etale}(X)$ *are equivalent.*

Proof. This is a consequence of the adjunction from Lemma 1.2.6, the isomorphisms guaranteed by Lemma 1.2.7 and the fact that the functors involved restrict from $\mathbf{Top}/X$ to $\mathbf{Etale}(X)$ and $\mathbf{PSh}(X)$ to $\mathbf{Sh}(X)$. For a formal proof, see Mac Lane and Moerdijk (1991, p. 90). ☐

Remark 1.2.5. A map of sheaves $h : F \to G$ is, due to Theorem 1.2.1 the same thing as a continuous map $\Lambda\, h : \Lambda\, F \to \Lambda\, G$. Note that a map into $\Lambda\, G$ is continuous if and only if its right-composite with $\rho_s : U \to \Lambda\, G$ is continuous, where $s \in G(U)$ and U open in X. This allows us to further reformulate h into a map of sets $h' : \Lambda\, F \to \Lambda\, G$, such that the maps $h' \circ \rho_s$ are continuous. Finally, deconstructing the definition of $\Lambda\, F$ and $\Lambda\, G$, we see that $\Lambda\, h$ is simply a family of functions $h_x : F_x \to G_x$, where F_x and G_x respectively denote the stalks of F and G at x.

The space $\Lambda\, P$ is not always Hausdorff, as illustrated in Example 1.2.8. However, under some mild conditions it is in fact Hausdorff. These mild conditions are satisfied by sheaves on a lot of spaces of common interest to mainstream mathematicians. This lemma is not relevant to the remainder of this text.

Lemma 1.2.8 (Hausdorff Display Space, Forster (1981, Theorem 6.10)). *Let X be a locally connected Hausdorff space and P a presheaf on X. Assume that for each open connected set $U \subseteq X$ and each pair of sections $f, g \in P(U)$ we know that if there is some $x \in U$ such that $\rho_f(x) = \rho_g(x)$, then $f = g$. Now $\Lambda\, X$ is Hausdorff as well.*

Proof. Let $s, t \in \Lambda\, X$ be unequal. First suppose $\pi(x) = z = \pi(y)$. Choose representatives $\langle f, U\rangle$ and $\langle g, V\rangle$ for s_z and t_z respectively. Pick some connected open $z \in C \subseteq U \cap V$ and realize that $\rho_f(C)$ and $\rho_g(C)$ are open neighbourhoods for s and t respectively. Suppose these are not disjoint, then we can pick some u in their intersection. This means that $\rho_f(c) = u = \rho_g(c)$ for some $c \in C$. By assumption, this entails $f = g$, which as assumed not to be true, a contradiction.

We are thus left with the case that $\pi(x) \neq \pi(y)$. Simply pick disjoint neighbourhoods U and V of $\pi(x)$ and $\pi(y)$ respectively, which we know to exist by the assumption that X is Hausdorff. Now $x \in \pi^{-1}(U)$ and $y \in \pi^{-1}(V)$ are the desired distinct neighbourhoods. □

Example 1.2.8 (Not Hausdorff). Let F be the sheaf of real-valued functions on the real line. Consider the constant zero map 0 and the map f which maps x to either 0 if $x \leq 0$ and x otherwise. It is easy to see that $0, f \in F(\mathbb{R})$ and $0_0 \neq f_0$. Yet for each $x < 0$ we have $0_x = f_x$, so there is no open neighbourhood around 0 discerning 0 and f.

Example 1.2.9 (Riemann Surfaces). Given a Riemann surface, the sheaves of holomorphic and meromorphic functions on it, as described in Example 1.1.4 and 1.1.3, satisfy Lemma 1.2.8 and thus yield a Hausdorff display space.

Instead of constructing the display space of a presheaf in order to make a left adjoint to Γ, we might also have taken a more technical approach. Consider the functor $A : \mathrm{Opens}(X) \to \mathbf{Top}/X$ which sends the open set U to the étale space $U \subseteq X$ and the inclusion $V \subseteq U$ to the obvious map of étale spaces. This gives rise to the map $R = \mathbf{Top}/\mathrm{X}\,(A(-),\,(-))$ as defined below on objects.

$$R : \mathbf{Top}/X \to \mathbf{PSh}(X), \qquad \begin{array}{c} Y \\ f\big\downarrow \\ X \end{array} \mapsto \left(U \in \mathrm{Opens}(X) \mapsto \mathbf{Top}/\mathrm{X}\left(A(U),\ \begin{array}{c} Y \\ f\big\downarrow \\ X \end{array} \right) = \mathbf{Top}/\mathrm{X}\left(\begin{array}{c} U \\ \cup\!\!\mid\big\downarrow \\ X \end{array},\ \begin{array}{c} Y \\ f\big\downarrow \\ X \end{array} \right) \right)$$

Note that this map R is exactly the same thing as Γ, formulated slightly differently. The following general lemma provides this map with a left adjoint $L : \mathbf{PSh}(X) \to \mathbf{Top}/X$. Due to the uniqueness up to isomorphism of left adjoints, this map must be isomorphic to Λ.[7] The lemma uses the category as defined below.

Definition 1.2.2 (Category of Elements). Given a presheaf $A : C^{\mathrm{op}} \to \mathbf{Set}$ on $\mathcal{C}$, define its *category of elements*

$$\int P := \Big(\mathbf{1} \downarrow P \Big).$$

The objects of this category are pairs $\langle C, p\rangle$ with $p \in P(C)$, its arrows $f : \langle X, x\rangle \to \langle Y, y\rangle$ are maps $f : X \to Y$ such that $P(f)(x) = y$.

Lemma 1.2.9. *Let $A : \mathcal{C} \to \mathcal{D}$ be a functor from a small category to a cocomplete category. Consider the functors*

$$\begin{array}{rcllcl} \mathcal{D}\,(A(-),\,(-)) = R & : & \mathcal{D} \to \mathbf{PSh}(\mathcal{C}), & D & \mapsto & \big(C \mapsto \mathcal{D}\,(A(C),\,D)\,\big) \\ L & : & \mathbf{PSh}(\mathcal{C}) \to D, & P & \mapsto & \mathrm{colim}\left(\int P \xrightarrow{\pi} \mathcal{C} \xrightarrow{A} \mathcal{D}\right). \end{array}$$

Now L is left adjoint to R.

[7] See Mac Lane (1997, Corollary 4.1.1) for a proof of this uniqueness.

Proof. This lemma can easily be proven when factoring the desired isomorphism through a convenient intermediate category as below.

$$\mathcal{D}\left(L(P),\, D\right) \cong \mathcal{D}^{\int \mathrm{P}}\left(A\pi,\, \Delta\, D\right) \cong \mathbf{PSh}(\mathcal{C})\left(P,\, R(D)\right)$$

The left isomorphism is immediate from the definition of colimits, as the middle Hom-set is the set of cocones from $A\pi$ to D. We now construct the isomorphism on the right, thus concluding the proof.

Given a natural transformation $\tau : P \to R(D) = \mathcal{D}\left(A((-)),\, D\right)$ we can assign the natural transformation $\gamma : A\pi \to \Delta\, D$, by defining it on components to be

$$\gamma_{\langle C,p\rangle} := \tau_C(p) : A\pi(\langle C,p\rangle) = A(C) \to D = \Delta\, D(\langle C,p\rangle).$$

Naturality of τ ensures that for $y \in P(Y)$ and $f : X \to Y$ in $\mathcal{C}$ we have $\tau_Y(y)A(f) = \tau_X(P(f)(y))$. This entails that for an arrow $f : \langle X,x\rangle \to \langle Y,y\rangle$ the following diagram commutes.

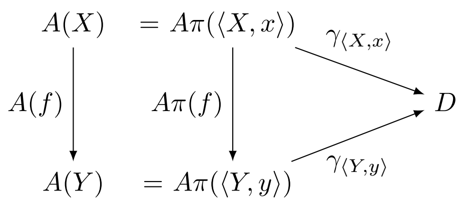

In the other direction, one can map $\gamma : A\pi \to \Delta\, D$ to the natural transformation $P \to R(D)$ defined on components by

$$\tau_C = p \in P(C) \mapsto \alpha_{\langle C,p\rangle} \in D,$$

naturality here follows similarly. It is clear that these maps are mutually inverse, thus proving the desired. □

SUBSECTION 1.2.2 Base Change

Given a map $f : X \to Y$ of topological spaces we can create a functor $f_* : \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$ which sends a sheaf to its '*direct image* under f'. This functor is defined by sending a sheaf F on X to the presheaf $f_*(F)$, defined to map opens U to $F(f^{-1}(U))$ and inclusions $V \subseteq U$ to the obvious restriction. One can easily verify that $f_*(F)$ in fact is a sheaf. A map of sheaves $h : F \to G$ is sent to the map $f_*(h)$, defined on components U by $h_{f^{-1}(U)}$. From this description it is clear that f_* in fact is a functor.

One can also for a map in the other direction. By Lemma 1.2.4 we know that such a map $f : X \to Y$ gives rise to a map $f^* : \mathbf{Etale}(Y) \to \mathbf{Etale}(X)$, which sends an étale space $p : E \to X$ to its *inverse image* under f, that is to say, to its pullback along f. We know that the categories of sheaves and étale spaces are equivalent, so this map gives rise to the composite

$$\mathbf{Sh}(Y) \xrightarrow{\Lambda} \mathbf{Etale}(Y) \xrightarrow{f^*} \mathbf{Etale}(X) \xrightarrow{\Gamma} \mathbf{Sh}(X),$$

which we also give the name f^* for convenience. The maps f^* and f_* are intimately related, Theorem 1.2.4 will show the former to be left adjoint to the latter. Furthermore, f^* preserves finite limits as we will prove in Theorem 1.2.2. To prove this theorem, it is useful to realize that $f^* : \mathbf{Etale}(Y) \to \mathbf{Etale}(X)$ is simply a restriction of the functor $\mathbf{Top}/Y \to \mathbf{Top}/X$ which sends a space $p : E \to X$ over X to its pullback along f. This is a specific instance of performing the pullback along an arrow $f : X \to Y$, which yields a functor $\mathcal{C}/Y \to \mathcal{C}/X$ for any category $\mathcal{C}$ with pullbacks. From the following general lemma we derive a left-adjoint to taking pullbacks.

Lemma 1.2.10. *Let $\mathcal{C}$ be a category with finite products and X an object in $\mathcal{C}$. The forgetful functor $\Sigma : \mathcal{C}/X \to \mathcal{C}$ is left adjoint to the pullback functor $!^* : \mathcal{C} \to \mathcal{C}/X$ of the unique map $! : X \to \mathbf{1}$.*

Proof. We need to show the existence of a natural isomorphism

$$\mathcal{C}\left(Y,\, Z\right) = \mathcal{C}\left(\Sigma(p : Y \to X),\, Z\right) \cong \mathcal{C}/\mathrm{X}\left(p : Y \to X,\, !^*(Z)\right) = \mathcal{C}\left(p : Y \to X,\, \pi_X : Z \times Y \to X\right).$$

The right-hand side equality is immediate from the definition of pullbacks, similar to the reasoning as we saw in Example 1.2.6. On the left-hand side, equality follows from the definition of Σ.

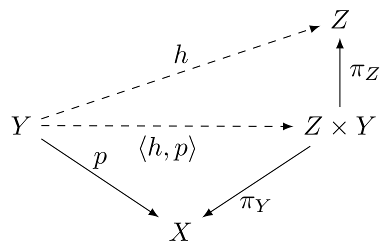

Consider the diagram to the left, which illustrates the desired (natural) isomorphism. There is a unique correspondence between arrows $Y \to X$, and arrows $Y \to Z \times Y$ which make the lower triangle commute. One can readily derive this from the universal property of the product. This provides the middle isomorphism.

□

Corollary 1.2.1. *Let $\mathcal{C}$ be a category which has pullbacks and let $f : X \to Y$ be an arrow in $\mathcal{C}$. The functor f^* : $\mathcal{C}/Y \to \mathcal{C}/X$ of pulling back along f has a left adjoint $\Sigma_f : \mathcal{C}/X \to \mathcal{C}/Y$.*

Proof. The arrow f is an object in the category $\mathcal{C}/Y$, and as such it yields a new slice category $(\mathcal{C}/Y)/f$. Objects in this category are maps $a : \big(d : A \to Y\big) \to \big(f : X \to Y\big)$ in $\mathcal{C}/Y$ such that $f \circ a = d$. Note that this object is wholly determined by the map $a : A \to Y$. A map between two such objects $a : A \to Y$ and $b : A \to Y$ is simply a map $h : A \to B$ such that $h \circ a = b$. This makes it clear that there is an isomorphism of categories $\phi_f : (\mathcal{C}/Y)/(f) \to \mathcal{C}/X$.

Now note that $\mathcal{C}/Y$ has finite products. Consider the unique map $! : \big(f : X \to Y\big) \to (\mathrm{id}_Y : Y \to Y)$ in $\mathcal{C}/Y$, and note that the product there of the objects $f : X \to Y$ and $h : Z \to Y$ is simply the pullback of h along f in $\mathcal{C}$. The previous lemma tells us that there is a left adjoint $\Sigma : \mathcal{C}/Y \to (\mathcal{C}/Y)/f$ to the functor $!^* : (\mathcal{C}/Y)/f \to \mathcal{C}/Y$. Applying the isomorphism of the previous paragraph we see that $\Sigma_f := \phi_f \Sigma$ is left adjoint to $!^* \phi_f^{-1} = f^*$. □

We now know that taking the pullback along a map preserves finite limits in the category of topological spaces over X, because $f_* : \mathbf{Top}/Y \to \mathbf{Top}/X$ has a left adjoint due to the above lemma. In order to show that $f_* : \mathbf{Etale}(Y) \to \mathbf{Etale}(X)$ preserves finite limits it now suffices to show that taking a finite limit of objects in $\mathbf{Etale}(Z)$ yields the same object as taking that finite limit within $\mathbf{Top}/Z$ for any topological space Z. Indeed, it then follows that the image under $f^* : \mathbf{Etale}(Y) \to \mathbf{Etale}(X)$ of $\lim D$ for some finite diagram $D : \mathcal{D} \to \mathbf{Etale}(Y)$ is canonically isomorphic to the limit of D composed with $f^* : \mathbf{Top}/Y \to \mathbf{Top}/X$.

It is clear that the terminal objects of these categories coincide. The product of two étale spaces over X is their product in $\mathbf{Top}/X$, which we know to be étale by Lemma 1.2.4. To complete the proof, we now need to show that the equalizer of two étale spaces in again étale. This can be proven in a (more lax and intuitive) topological manner, or a bit more dryly, using only already proven categorical properties.

Lemma 1.2.11. *Let $p : A \to X$ and $q : B \to X$ be étale spaces over X and let $g, h : p \to q$ be a parallel pair of arrows. Now the equalizer $e : E \to A$ of g and h is étale.*

Proof via Topology. When E is canonically constructed it is simply a subspace of A, so it suffices to prove that it is open. Note that the maps f and g give rise to a map $\langle f, g\rangle : A \to B \times_X B$ into the pullback of q along itself. Now look at the diagonal map $\Delta : B \to B \times_X B$ and realize that by Lemma 1.2.2 we know it to be open. See that $E = \langle f, g\rangle^{-1} \Delta(B)$, proving the desired. □

Proof via Universal Properties. We claim that E is in fact the pullback of $\Delta : B \to B \times_X B$ along $\langle g, h\rangle : A \to B \times_X B$, which proves the desired by Lemma 1.2.4. In this description, Δ is the diagonal map as mentioned in Lemma 1.2.2 and $\langle g, h\rangle$ is the unique map induced by the pullback as depicted below to the left.

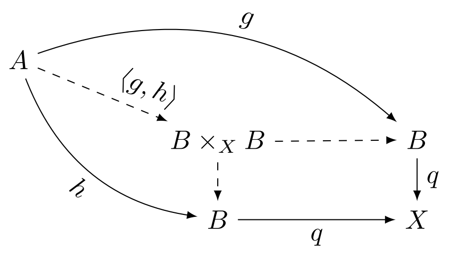

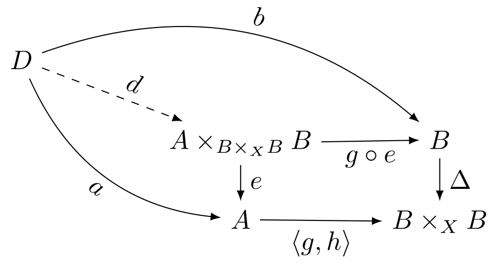

It is not hard to see that E fits in the pullback diagram on the right. To prove this, let consider D with arrows as indicated. Now as $\langle b, b\rangle = \Delta b = \langle g, h\rangle\, a = \langle ga, ha\rangle$ holds, we know $ga = ha$. Using the universal property of E, we obtain a map $d : D \to E$ unique such that $e \circ d' = e$, proving the desired. □

Combining the above, we have proven the following theorem. Indeed, due to the above it is clear that the inverse image functor $f^* : \mathbf{Etale}(Y) \to \mathbf{Etale}(X)$ preserves finite limits, so by composition with the equivalence of Theorem 1.2.1 the result is immediate.

Theorem 1.2.2. *For any map $f : X \to Y$, the inverse image functor $f^* : \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ preserves finite limits.*

This would have been considerably easier to prove if we knew $f : X \to Y$ to be a local homeomorphism. Indeed, one can easily define a functor $f_! : \mathbf{Etale}(X) \to \mathbf{Etale}(Y)$ which assigns to $p : E \to X$ the étale space $f \circ p : E \to Y$ and acts as the identity on arrows. This is of course only possible if the composite $f \circ p$ is always an étale space, which is the case whenever f is étale.

Theorem 1.2.3. *For a local homeomorphism $f : X \to Y$, the functor f^* is right adjoint to the functor $f_!$. This gives a natural isomorphism as below.*

$$\mathbf{Etale}(\mathrm{Y})\left(f_!\left(\begin{matrix}E\\ p\big\downarrow\\ X\end{matrix}\right),\ \begin{matrix}D\\ q\big\downarrow\\ X\end{matrix}\right) \cong \mathbf{Etale}(\mathrm{X})\left(\begin{matrix}E\\ p\big\downarrow\\ X\end{matrix},\ f^*\left(\begin{matrix}D\\ q\big\downarrow\\ X\end{matrix}\right)\right)$$

Proof. The result is immediate from the diagram below. Indeed, due to the universal property of the pullback there is a correspondence between the upper indicated arrow which makes the triangle commute, and the lower indicated arrow. This provides us with the desired natural isomorphism.

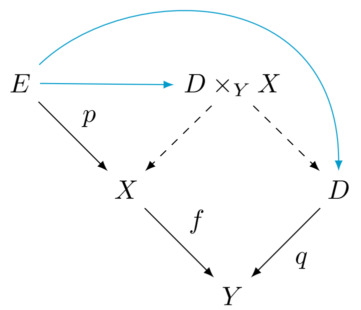

□

In the situation of the above theorem one knows $f^* : \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ to preserve all limits, for it has a left adjoint. In general it is the case that the direct image functor f^* preserves all colimits, for it has a right adjoint as we will show below. To this end we first take a closer look at f^*G for a sheaf G over Y, and $f : X \to Y$ a map of topological spaces. By construction, we know f^* to be the sheaf of sections of $f^*\Lambda G$. For any open U in X we now know that an element of f^*G is a section $t : U \to f^*\Lambda G$ such that $\pi \circ s$ is simply the inclusion of U into X.

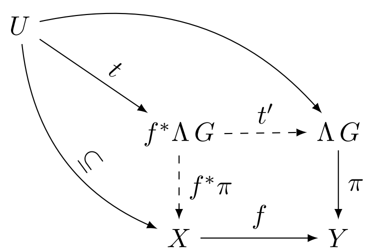

In the above diagram we see that there is a correspondence between maps $t : U \to f^*\Lambda\, G$ which make the left-most diagram commute, and maps $t' : U \to \Lambda\, G$ such that $\pi \circ t' = f|U$. This makes sections of $f^*\Lambda\, G$ simply lifts of f along the projection π.

Recall that during our discussion of the display space in Example 1.2.7 we saw maps $\rho_s : Y \supseteq U \to \Lambda\, G$ for U open in Y and $s \in G(U)$. An important property of such maps was that $\pi \circ \rho_s$ is simply the inclusion of U into Y. Hence by pre-composing with f we obtain the composite

$$f^{-1}(U) \xrightarrow{f} U \xrightarrow{\rho_s} \Lambda\, G,$$

which one now can readily see to be a lift of f along π. By the above we now know this to give rise to a section $t_s : f^{-1}(U) \to f^*\Lambda\, G$, which we know to have an open image due to Remark 1.2.2. The images of all such sections t_s for $s \in G(U)$ and U open in Y yield an (open) cover of $\Lambda\, G$. As a consequence, any map (of sets) $h : f^*\Lambda\, G \to H$ is continuous if and only if the composite $h \circ t_s$ is for all such sections t_s.

It will be useful to have an explicit description of points in $f^*\Lambda\, G$. Using the canonical method of constructing pullbacks one can easily see that points are pairs $\langle x, a\rangle$, where a is a germ at $f(x)$. Of course, a can be represented by some section $\langle s, V\rangle$ with $s \in G(V)$, and we will not hesitate to do so.

Theorem 1.2.4. *For any map $f : X \to Y$, the inverse image functor f^* is left adjoint to the direct image functor f_*. That is to say, there exists a natural isomorphism*

$$\mathbf{Sh}(\mathrm{X})\,(f^*(G),\ F) \cong \mathbf{Sh}(\mathrm{Y})\,(F,\ f_*(F))\,. \tag{1.3}$$

Proof. We claim that there exists a natural isomorphism as below

$$\mathbf{Etale}(\mathrm{X})\,(f^*\Lambda\, G,\ \Lambda\, F) \cong \mathbf{Sh}(\mathrm{Y})\,(G,\ f_*\Gamma\Lambda\, F)\,. \tag{1.4}$$

Assuming that this is in fact the case, we are done due to the following chain of natural isomorphisms.

$$\mathbf{Sh}(\mathrm{X})\,(f^*G,\ F) \cong \mathbf{Etale}(\mathrm{X})\,(f^*\Lambda\, G,\ \Lambda\, F) \cong \mathbf{Sh}(\mathrm{Y})\,(G,\ f_*\Gamma\Lambda\, F) \cong \mathbf{Sh}(\mathrm{Y})\,(G,\ f_*F)$$

The left-hand isomorphism is simply the equivalence between sheaves and étale spaces as proven in Theorem 1.2.1. The final isomorphism is post-composition with the inverse of this adjunctions's unit, which we know to be an isomorphism due to Lemma 1.2.7. We see that this indeed proves the desired, so we are left to prove that (1.4) is in fact a natural isomorphism.

First realize that a map of étale spaces $h : f^*\Lambda\, G \to \Lambda\, F$ is simply a map of sets over X, such that the pre-composition with all sections t_s is continuous as discussed in the paragraph above this proof. This allows us to define the mapping

$$\phi :\ f^*\Lambda\, G \xrightarrow{h} \Lambda\, F \ \mapsto\ \left(s \in G(U) \mapsto f^{-1}(U) \xrightarrow{h \circ t_s} \Lambda\, F\right)_{U \subseteq Y} : G(U) \to f_*\Gamma\Lambda\, F(U) = \Gamma\left(\Lambda\, F, f^{-1}(U)\right).$$

This mapping sends a map h to the natural transformation defined on components by the map which sends sections s to the composite $h \circ t_s$.

In the other direction we construct a map

$$\psi :\quad G \xrightarrow{\tau} f_*\Gamma\Lambda\, F \quad\mapsto\quad \Big(\langle x, \langle s, U\rangle\rangle \in f^*\Lambda\, G \mapsto \tau_U(s)(x)\Big) : f^*\Lambda\, G \to \Lambda\, F,$$

which can easily be checked to be well-defined. Given $s \in G(U)$ we can now check that for all $x \in X$ we have

$$(\psi(\tau) \circ t_s)(x) = \psi(\tau)(\langle x, \langle s, U\rangle\rangle) = \tau_U(s)(x),$$

which makes $\psi(\tau) \circ t_s = \tau_U(s)$ a continuous map. This makes it a bona fide map in the opposite direction of (1.4).

We now need to check that these maps are mutually inverse. Take a map $h : f^*\Lambda\, G \to \Lambda\, F$ and consider $\psi\phi(h)$, we can compute that for any point $\langle x, \langle s, U\rangle\rangle \in f^*\Lambda\, G$ we get

$$\psi\phi(h)(\langle x, \langle s, U\rangle\rangle) = \phi(h)(s)(x) = (h \circ t_s)(x) = h(x, \langle s, U\rangle).$$

In the other direction, given a map $\tau : G \to f_* \Gamma \Lambda\, F$, an open $U \subseteq Y$ and $s \in G(U)$ we see

$$\phi\psi(\tau)(s) = \psi(\tau) \circ t_s = \tau_U(s).$$

This proves that these mapping indeed are mutually inverse, proving that (1.4) is an isomorphism. We neglect to verify that this is indeed natural. ☐

SUBSECTION 1.2.3 Locally Connected Étale Spaces

Recall $\mathbf{Etale}(\mathbf{1})$ as discussed in Example 1.2.5. There is an obvious (and unique) map $\pi : X \to \mathbf{1}$ for any topological space X. This gives rise to a functor $\pi^* : \mathbf{Set} \cong \mathbf{Sh}(\mathbf{1}) \to \mathbf{Sh}(X)$. One can wonder whether this map has a left adjoint. We already know by Theorem 1.2.3 that this is true in case π is étale. Unfortunately, this is only so when X is endowed with the discrete topology, as mentioned in Example 1.2.5. Here we will construct a left adjoint to $\pi^* : \mathbf{Set} \to \mathbf{Sh}(X)$, assuming that X is locally connected.

Let us first take a closer look at $\pi_0(E)$, the set of connected components of an étale space $p : E \to X$. We can define a relation $R \subseteq E \times E$ by setting $a \; R \; b$ precisely if there is an open $U \subseteq E$ such that U is connected and $a, b \in U$. This relation is always transitive and symmetric, but for reflexivity one needs E to be locally connected. This must be the case when X is locally connected, due to Lemma 1.2.1. In this case R is an equivalence relation, and it is clear that cosets of R are precisely the connected components. Topologize R as subspace of the product $E \times E$. Consequently, the map $q : E \to \pi_0(E)$ which sends elements to the connected component they are in is the coequalizer of the left- and right projection of R to E.

We can now define the functor $\pi_! : \mathbf{Etale}(X) \to \mathbf{Set}$, which sends étale spaces to their connected components and arrows to the induced arrows on the connected components as arising from the universal property. Concretely, this means that a map $f : E \to D$ is mapped to $\pi_!(f)$, which sends the connected component U to the connected component containing $f(U)$. This is well-defined by virtue of f being continuous.

Theorem 1.2.5. *Let X be a locally connected topological space. There is an adjunction*

$$\mathbf{Set}\left(\pi_!\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array}\right), S\right) \cong \mathbf{Etale}(\mathrm{X})\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array}, \pi^*(S)\right).$$

Proof. Given a map $h : E \to \pi^*(S)$ we know that $\pi_S h(x) = \pi_S h(y)$ holds when $x \; R \; y$. Indeed, take $x, y \in U$ open and connected and see that $\pi_S h(U) = \{h(x)\}$ by continuity of h and the fact that S is endowed with the discrete topology. As such, h gives rise to a map $\pi_0(E) \to S$ via the universal property of the coequalizer $q : E \to \pi_0(E)$. This direction, and the converse, can easily be read out of the following diagram.

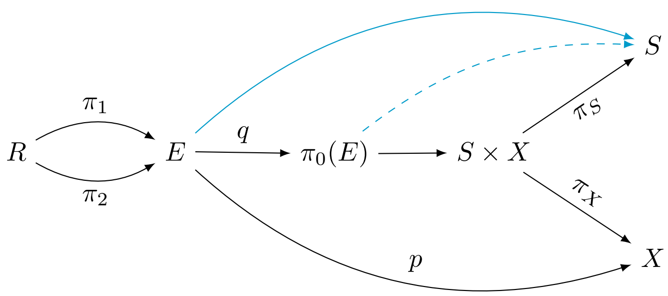

Given a map of sets $g : \pi_0(E) \to S$, we obtain the top indicated arrow by pre-composition with q. This gives rise to the desired map $\langle g \circ q, p\rangle : E \to S \times X$ of étale spaces. In the opposite direction, the discussion above teaches us that a map $h : E \to S \times X$ gives rise to a map $\overline{\pi_S \circ h} : \pi_0(E) \to S$. These operations are mutually inverse by virtue of their construction, as can be easily checked. ☐

By composing the adjunction derived above with the equivalence of categories between **Sh**(X) and **Etale**(X) we see that the obvious functors $\pi_! : \mathbf{Sh}(X) \to \mathbf{Set}$ and $\pi^* : \mathbf{Set} \to \mathbf{Sh}(X)$ are such that the former is left adjoint to the latter.

When a topological space is connected and locally connected, we have extra information on the composite $\pi_!\pi^*$. The following proposition leads to this result. We will not prove this proposition in full, see Lemma C1.5.7 of Johnstone (2002b) for more details.[8] Neither do we fully prove the corollary in which we derive the property we will use later on. The only missing ingredient is a lemma concerning general category theory, there is no point in explicating it here. We do prove Lemma 1.2.12, at it can be solved with a topological argument which is fairly straightforward given the discussion above, and we require the existence of the isomorphism described therein later on in Theorem 3.1.1.

Proposition 1.2.1. *Let X be a topological space, now X is connected if and only if π^* is full and faithful.*

Corollary 1.2.2. *If X is a connected and locally connected topological space, then $\epsilon : \pi_!\pi^* \to$ **Set** is an isomorphism.*

Proof. As we know $\pi_!$ to be left-adjoint to π^* by Theorem 1.2.5, and π^* is full and faithful due to Proposition 1.2.1, the desired follows from Theorem 4.3.1 of Mac Lane (1997). □

Lemma 1.2.12. *Let X be a connected and locally connected topological space. Now $\epsilon : \pi^!\pi_* \to \mathrm{id}_{\mathbf{Set}}$ is an isomorphism, and the composite given in the diagram below is an isomorphism, natural both in the étale space $p : E \to X$ and the (small) set S. Moreover, π^* is full and faithful.*

$$\pi_!\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array} \times \pi^*(S)\right) \xrightarrow{\left\langle \pi_!(\pi^{\mathrm{l}}), \pi_!(\pi^{\mathrm{r}})\right\rangle} \pi_!\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array}\right) \times \pi_!\pi^*(S) \xrightarrow{\mathrm{id}_{\pi_!(p)} \wedge \epsilon} \pi_!\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array}\right) \times S$$

Proof. Let S be any small set, and realize that $\pi^!\pi_*(S)$ is given as the set of connected components of the étale space $\pi_X : S \times X \to X$. The map $\epsilon : \pi^!\pi_*(S) \to S$ is the image of the identity on this space under the isomorphism described in Theorem 1.2.5. This map simply sends (the class of) $\langle s, x\rangle$ to s. This map is clearly surjective, and injectivity is also easy to derive. Indeed, if $\langle s, x\rangle$ and $\langle t, y\rangle$ yield the same ϵ-image, then $s = t$ follows, so as X is connected we know both to lie within the same connected open $\{s\} \times X$. We can thus remark that two representatives $\langle s, x\rangle$ and $\langle t, y\rangle$ represent the same connected component exactly if $t = s$.

Now consider the space $p \times \pi^*(S)$, which we know to be given as the pullback of $\pi_X : S \times X \to X$ along p, with the sensible map down to X. The space can be canonically described as the set

$$\langle\langle s, x\rangle, e\rangle \mid \langle s, x\rangle \in S \times X,\ e \in E\ p(e) = x\}$$

endowed with the subspace topology inherited from $(S \times X) \times E$. As a consequence, the map $\left\langle \pi_!(\pi^{\mathrm{l}}), \pi_!(\pi^{\mathrm{r}})\right\rangle$ can be described on representatives as

$$\langle\langle s, x\rangle, e\rangle \in (S \times X) \times E \mapsto \langle\langle s, x\rangle, e\rangle \in \pi_!(p) \times \pi_!\pi^*(S).$$

This map is well-defined, for when two representatives lie within the same connected open, this open is projected down to connected opens by π^{l} and π^{r}. It is surjective due to the above remark. To prove it to be injective, take $\langle\langle s, x\rangle, e\rangle$ and $\langle\langle s', x'\rangle, e'\rangle$ and assume that both have the same image. This entails that $s = s'$ and e and e' lie within some connected open U of E. It is now easy to see that $(\{s\} \times X) \times U$ is connected and open in the domain, so the two representatives represent the same class. This proves the composite to be an isomorphism as desired.

The final statement is an immediate consequence of the first statement and the cited lemma in Corollary 1.2.2. □

[8] This proof refers back to Johnstone (2002b, Lemma C1.3.15), a lemma on locales. We prove the case for topological spaces in Proposition 1.1.2.

A point x of a topological space X is virtually the same as the map $\mathbf{1} \to X$ which sends the sole inhabitant of $\mathbf{1}$ to the point x. As such we make no syntactic difference between the two, and just write x to also mean this map $\mathbf{1} \to X$. It is clear that this map is always continuous, so we know of functors $x^* : \mathbf{Sh}(X) \to \mathbf{Sh}(\mathbf{1}) \cong \mathbf{Set}$ and $x_* : \mathbf{Set} \cong \mathbf{Sh}(\mathbf{1}) \to \mathbf{Sh}(X)$. From now on we write x^* and x_* to actually mean the composition of the above functors with the isomorphism between **Set** and $\mathbf{Sh}(\mathbf{1})$. By Theorem 1.2.4 we know the former functor to be left adjoint to the latter. Moreover, the functor x^* is *left exact*, that is to say, it preserves finite limits.

Let us turn to the situation of étale spaces, where a point $x \in X$ gives rise to $x^* : \mathbf{Etale}(X) \to \mathbf{Set}$. One readily sees that $x^*(p : E \to X) = p^{-1}(\{x\})$. Indeed, given any map $f : A \to E$ with $p \circ f = x$, it is clear that the image of f is contained in $p^{-1}(\{x\})$, and this restriction is the unique map $h : A \to p^{-1}(\{x\})$ such that $(\subseteq) \circ h = f$.

We now step through the equivalence between étale spaces and sheaves, and look at what $x^* : \mathbf{Sh}(X) \to \mathbf{Set}$ does to a sheaf. To this end, fix a sheaf $F : \mathrm{Opens}(X)^{\mathrm{op}}$ and compute

$$x^*(F) = x^*(\pi : \Lambda F \to X) = \left(\pi : \coprod_{y \in X} \operatorname*{colim}_{y \in U \in \mathrm{Opens}(X)} F(U) \to X\right)^{-1}(x) = \operatorname*{colim}_{x \in U \in \mathrm{Opens}(X)} F(U).$$

This means that $x^*(F)$ is equal to the *stalks of F at x*. The following lemma shows us what stalks of sections of étale spaces are.

Lemma 1.2.13. *For every point $x \in X$, there is a natural isomorphism $x^*\Gamma \cong x^*$.*

Proof. Let $p : E \to X$ be an étale space over X. It is easy to compute $x^*(p)$, which is simply the pre-image of x under p. We also know that $x^*\Gamma\left((-), p\right)$ consists of the stalk of $\Gamma\left((-), p\right)$ at x. This allows us to explicitly define the map

$$\phi_p : x^*\Gamma\left((-), \begin{array}{c} E \\ p\big\downarrow \\ X \end{array}\right) \to x^*\left(\begin{array}{c} E \\ p\big\downarrow \\ X \end{array}\right), \qquad \sigma : U \to X \mapsto \sigma(x) \in p^{-1}(\{x\}),$$

which is clearly well-defined. To construct a mapping in the other direction, take a point $e \in p^{-1}(\{x\})$ and realize that there exists an open neighbourhood U of x such that $p \restriction U$ is a homeomorphism. This gives us a section $(p \restriction U)^{-1} : U \to X$. It is not hard to compute that ϕ is in fact natural in p.

It is clear that the point associated to the section associated to a point, is the point we started with. To prove the other direction, fix a section $\sigma : U \to X$ and an open neighbourhood V in X of $\sigma(x)$ such that $p \restriction V$ is a homeomorphism to an open. We need to show that σ and $(p \restriction V)^{-1}$ agree on some open, but this follows directly from the local behavior of sections as mentioned in Remark 1.2.1. ☐

Finally, let us have a closer look at $x_* : \mathbf{Set} \to \mathbf{Sh}(X)$. Recall that we identify sheaves on $\mathbf{1}$ with sets simply by seeing the set S as the sheaf of sections S on the étale space $! : S \to \mathbf{1}$. Now as $\mathrm{Opens}(X)$ consists of $\{*\}$ and $\emptyset$, we know everything there is to know about this functor when we realize that $S(\emptyset) = \emptyset$ and $S(\{*\}) = S$. The former is of course a consequence of Remark 1.1.1. Using this information we see that $x^*(S)$ is the functor that sends opens U to the set $S(x^{-1}(U))$. This means precisely that $x^*(S)(U)$ equals S if $x \in U$ and $\emptyset$ otherwise, so $x^*(S)$ is the *skyscraper sheaf*.

It is an immediate consequence of Theorem 1.2.3 that the stalk functor is left adjoint to the skyscraper functor. Each point gives rise to a stalk functor. Below we will show that the totality of these stalk functors preserves a lot of information.

Theorem 1.2.6. *A map of sheaves $h : F \to G$ on X is an (monomorphism) epimorphism exactly if $x^*(h)$ is for all points x of X.*

Proof. Recall Remark 1.2.5, which states that a map of sheaves $f : G \to H$ is fully determined by a family of maps $f_x : F_x \to G_x$ with a continuity requirement. Retracing this argument we indeed see that the equivalence between sheaves and étale spaces teaches us that f is fully determined by Λf. Now inspect $x^*(f)$ for some $x \in X$ and see that $\langle x^*(f)\rangle_{x\in X} = \Lambda f$.

Fix a map of sheaves $h : F \to G$. Suppose that $x^*(h)$ is epi (a surjection) for all points $x \in X$. Let $f, g : G \to H$ be such that $fh = gh$, and derive that

$$x^*(f)x^*(h) = x^*(fh) = x^*(gh) = x^*(g)x^*(h),$$

so as $x^*(h)$ is epi this entails that $x^*(f) = x^*(g)$ for all $x \in X$. By the above paragraph we now deduce that $f = g$. The argument for h mono works in exactly the same way.

The converse follows from the general theory we already built up. Because x^* has a right-adjoint by Theorem 1.2.3, we know that x^* preserves colimits, so in particular it preserves epis. On the other hand we know x^* to preserve finite limits by Theorem 1.2.2, forcing it to preserve monos. From this the desired is immediate. □

Corollary 1.2.3. *A map of sheaves $h : F \to G$ on X is an isomorphism precisely if $x^*(h)$ is an isomorphism for all points of X.*

Proof. As a map of sheaves is an isomorphism exactly when it is an isomorphism point-wise, this is immediate from the fact that isomorphisms in **Set** are exactly those maps which are both injective and surjective. □

SECTION 1.3 Grothendieck Topology

Recall the setting of Example 1.1.11, which discussed the sheaf of regular functions on the maximal ideal spectrum of a finitely generated k-algebra A for some algebraically closed field k. The open sets in Specm A are given as (unions of) hyperplane complements $\mathrm{Z}(a)$ for $a \in A$. As a consequence, the collection of (basic) opens $\{U_\iota := \text{Specm } A - \mathrm{Z}(a_\iota)\}_{\iota=1,\ldots,m}$ covers Specm A when the ideal generated by $a_1, \ldots, a_m$ is the entirety of A.

Take for example as A the coordinate ring of $\mathbb{C}$, where the field k is of course taken to be $\mathbb{C}$. Basic opens of $\mathbb{C}$ as endowed with the Zariski topology are complements of finitely many points. A basic open U is thus given by the zero-set of an ideal generated by $a_I = \prod_{p\in I}(x - p)$ with I some finite subset of $\mathbb{C}$. Given a collection $\{\mathbb{C} - I = \mathrm{Z}(a_I) \mid I \in \mathcal{I}\}$ of basic open sets, we know that it is a cover of $\mathbb{C}$ precisely when $\bigcap \mathcal{I}$ is empty, which is equivalent to the statement that the ideal generated by $\{f_I \mid I \in \mathcal{I}\}$ is the whole ring $\mathrm{A}(X)$.

A hypersurface complement $U_a = \text{Specm } A - \mathrm{Z}(a)$ can be seen as a k-algebra homomorphism

$$A \to \{1, a, a^2, \ldots\}^{-1}A = A[\frac{1}{a}],$$

and note that the co-domain of this map is the trivial k-algebra when a is a nilpotent element. Generalizing k into an arbitrary (commutative) ring R one can consider the (small) category R-Alg of finitely generated R-algebras. Its objects are R-algebras of the form $R[x_1, \ldots, x_n]/I$ for some $n \in \mathbb{N}$ and ideal $I \subseteq R[x_1, \ldots, x_n]$, and arrows are R-algebra homomorphisms. Intuitively, the R-algebra A is covered by the 'hypersurface complements' $A[\frac{1}{a_\iota}]$ when, as before, the ideal generated by a_ι for $\iota = 1, \ldots, m$ contains 1. A *Grothendieck topology* allows one to formally specify this idea of covering on an arbitrary small category $\mathcal{C}$, and this allows for one to define what it means to be a sheaf on $\mathcal{C}$.

We first specify a generalization of the notion of sieve as defined in Definition 1.1.3, replacing the category Opens(X) with the more general (small) category $\mathcal{C}$.

Definition 1.3.1 (Sieve). A *sieve* on an object C of $\mathcal{C}$ is a subfunctor of $\mathbf{y}C = \mathcal{C}((-), C)$.

Equivalently, a sieve is a set S of morphisms of $\mathcal{C}$ into C such that if $f : X \to C$ is in S, then for any arrow $g : Y \to X$ of $\mathcal{C}$ we have $fg \in S$. One might think of S as being a right-ideal, generalizing that a sieve is downwards closed as in Section 1.1. Given a sieve S on C, we can form the pullback of $S \subseteq \mathbf{y}C$ along $\mathcal{C}((-), h)$ for any arrow $h : D \to C$. This gives rise to a sieve

$$h^*(S) = \{g : X_g \to D \mid \mathcal{C}((-), h)(g) = hg \in S(X_g)\}.$$

Definition 1.3.2 (Grothendieck Topology). A *Grothendieck topology* on $\mathcal{C}$ is a function $\mathbf{J}$ which assigns to objects a set of sieves over that object. Any sieve $S \in J(C)$ is said to *cover* C. The assignment is subject to the following laws:

MAXIMALITY the maximal sieve $\mathbf{y}C$ covers C;

STABILITY if S covers C, then for any $h : D \to C$, the pullback of S along h covers D;

TRANSITIVITY if R is a sieve on C and S covers C, then R covers C when the pullback of R along every arrow $h : D \to C$ in S covers D.

The pair $\langle \mathcal{C}, \mathbf{J} \rangle$ is called a *site*.

Example 1.3.1 (Topological Space). Consider a topological space X. Now let $\mathcal{C}$ be the category of opens of X. We define the Grothendieck topology $\mathbf{J}$ by asigning to every open subset U of X the set of sieves $S \subseteq \mathbf{y}U$ that are covering in the sense of Definition 1.1.3. This means that the sieves which *cover* U in the new sense are exactly those sieves whose associated open covers actually covers U.

This assignment clearly satisfies maximality, as all opens below U certainly cover U. To prove that stability holds consider an inclusion $V \subseteq U$ and a sieve S which covers U, and let $\mathcal{V}$ be its associated open cover of U. Note that the pullback $R := \left(V \subseteq U\right)^*(S)$ is taken point-wise, and for the object W it is given by either the singleton set containing the inclusion $W \subseteq V$ in the case that $W \subseteq V$ and $(W \subseteq V \subseteq U) \in S$, and empty otherwise. Performing several computations, we derive that S indeed is a covering sieve.

To show transitivity, let both R and S be sieves on U, with S assumed to be covering. Write $\mathcal{V}$ for the open cover associated to S. Assume that the pullback along each inclusion $V \subseteq U$ of R is covering. This makes the set of opens $\mathcal{V}_V$ associated to $(V \subseteq U)^*(R)$ an open cover of V. From this it is easy to compute that $\mathcal{V}$ indeed covers U, proving $\mathbf{J}$ to be a Grothendieck topology.

Example 1.3.2 (Trivial Topology). For any small category $\mathcal{C}$ we can form the *trivial Grothendieck topology*, where $\mathbf{J}(C)$ consists only of the maximal sieve $\mathbf{y}C$. This clearly satisfies maximality. Stability follows as the pullback of $\mathbf{y}C$ along any arrow $D \to C$ is simply all of $\mathbf{y}D$. Transitivity is easy, for the pullback of R along id_C is forced to equal $\mathbf{y}C$, which in turn equals R again, making it covering.

A *matching family* in this setting is exactly the same as in Definition 1.1.4. We are now ready to define what a sheaf is in this new setting. The definition below is a straighforward generalization of the characterization obtained in Proposition 1.1.1.

Definition 1.3.3 (Sheaf). Let $\langle \mathcal{C}, \mathbf{J} \rangle$ be a site. A presheaf P on $\mathcal{C}$ is a sheaf exactly when for each object C and each covering sieve S for $\mathbf{J}$ the inclusion map $S \xrightarrow{i_S} \mathbf{y}C$ gives rise to an isomorphism

$$\mathbf{PSh}(\mathcal{C})(\mathbf{y}C, P) \xrightarrow{\mathbf{PSh}(\mathcal{C})(i_S, P)} \mathbf{PSh}(\mathcal{C})(S, P). \tag{1.5}$$

The category of all sheaves on a certain site $\langle \mathcal{C}, \mathbf{J} \rangle$ is denoted $\mathbf{Sh}(\mathcal{C}, \mathbf{J})$, this category is a subcategory of $\mathbf{PSh}(\mathcal{C})$. A category equivalent to the category of sheaves on a site is called a *Grothendieck topos*.

Example 1.3.3 (Presheaves as Sheaves). Recall the trivial topology from Example 1.3.2 and note that any presheaf is a sheaf, as the induced map in (1.5) is simply the identity map.

Example 1.3.4 (Sheaves on Topological Spaces). Let X be a topological space and P be a presheaf on X as in the first section of this chapter. Recall the definition of a Grothendieck topology on the category Opens(X) of Example 1.3.1, and realize that P is a sheaf in the sense of Grothendieck topologies exactly when it is a sheaf in the sense of Definition 1.1.1.

In Example 1.3.3 we saw that all presheaves are sheaves when the category is endowed with the trivial topology, so in particular, all representable presheaves are sheaves. Recall that when one considers sheaves over a topological space, all representable functors are sheaves, as proven in Remark 1.1.2. We now set out to prove that all sheaves are in a natural way related to representable presheaves. First, we will show that each presheaf is the colimit of representable presheaves. We then posit the existence of an adjunction between $\mathbf{PSh}(\mathcal{C})$ and $\mathbf{Sh}(\mathcal{C},\mathbf{J})$, as a generalization of the sheafification functor displayed in Remark 1.2.4. The desired result is then but an easy corollary.

Proposition 1.3.1. *Let $\mathcal{C}$ be a category. Each object of $\mathbf{PSh}(\mathcal{C})$ is canonically isomorphic to a colimit of representable objects $\mathbf{y}C = \mathcal{C}((-), C)$.*

Proof. Apply Lemma 1.2.9 with $A = \mathbf{y} : \mathcal{C} \to \mathbf{PSh}(\mathcal{C})$, which is applicable as $\mathbf{PSh}(\mathcal{C})$ is clearly cocomplete. We now see that $R(P)(C) = \mathbf{PSh}(\mathcal{C})(\mathbf{y}C, P)$, by the Yoneda lemma this is isomorphic to $P(C)$. This gives us a natural isomorphism $R \cong \mathrm{id}_{\mathbf{PSh}(\mathcal{C})}$. By the uniqueness up to isomorphism of left adjoints we know that L is (canonically) naturally isomorphic to the identity functor, providing the desired isomorphism

$$P \cong \operatorname{colim}\left(\int P \xrightarrow{\pi} \mathcal{C} \xrightarrow{\mathbf{y}} \mathbf{PSh}(\mathcal{C})\right). \qquad \square$$

We now simply state the following theorem, for a proof see for instance Mac Lane and Moerdijk (1991, Theorem 3.5.1).

Theorem 1.3.1 (Sheafification). *To each site $\langle \mathcal{C}, \mathbf{J} \rangle$ there exists an associated sheaf functor $\mathbf{a} : \mathbf{PSh}(\mathcal{C}) \to \mathbf{Sh}(\mathcal{C},\mathbf{J})$ left-adjoint to the inclusion functor $i : \mathbf{Sh}(\mathcal{C},\mathbf{J}) \to \mathbf{PSh}(\mathcal{C})$.*

Corollary 1.3.1 (Sheaves as Colimits of Representables). *Each object of $\mathbf{Sh}(\mathcal{C},\mathbf{J})$ is a colimit of representables.*

Proof. Given a sheaf F we can consider $i(F)$ and note that by Proposition 1.3.1 we know F to be a colimit of representables in the category $\mathbf{PSh}(\mathcal{C})$. Using the sheafification functor of Theorem 1.3.1 the following is immediate, because $\mathbf{a}$ must preserve colimits as it has a right adjoint.

$$F \cong \operatorname{colim}\left(\int P \xrightarrow{\pi} \mathcal{C} \xrightarrow{\mathbf{y}} \mathbf{PSh}(\mathcal{C}) \xrightarrow{\mathbf{a}} \mathbf{Sh}(\mathcal{C},\mathbf{J})\right) \qquad \square$$

Every Grothendieck topos has a set of objects such that the arrows starting in these objects can 'distinguish' other arrows in some sense. To make this more precise we introduce the two following notions, and prove that there is a generating set to each Grothendieck topos.

Definition 1.3.4 (Jointly Epimorphic). A collection $\mathcal{F}$ of arrows into an object X is said to be jointly epimorphic when for any pair of arrows $f, g : X \to Y$ the equality $f = g$ holds precisely when $fh = gh$ holds for all $h \in \mathcal{F}$.

Definition 1.3.5 (Generating Collection). A collection $\mathcal{G}$ of objects of a category $\mathcal{C}$ is said to generate (or separate) $\mathcal{C}$ if for each object the collection of arrows into that object starting at an object in $\mathcal{G}$ is jointly epimorphic. If $\mathcal{G} = \{G\}$, then G is said to be a generating object of $\mathcal{C}$.

Lemma 1.3.1. *Consider the Grothendieck topos $\mathcal{E} := \mathbf{Sh}(\mathcal{C},\mathbf{J})$. Now the collection $\mathcal{G} := \{\mathbf{ay}C \mid C \text{ object of } \mathcal{C}\}$ generates $\mathcal{E}$.*

Proof. Suppose that $f, g : F \to G$ is a parallel pair of arrows between sheaves in $\mathcal{E}$. Now assume that $fh = gh$ for all h starting in an object of $\mathcal{G}$. We know that $F = \operatorname{colim} \mathbf{ay}C_\iota$ where the ι are objects of some diagram $\mathcal{I}$. Consider the canonical maps $i_\iota : \mathbf{ay}C_\iota \to F$ and realize that this collection gives rise to natural transformations $\alpha_\iota := f i_\iota$ and $\beta_\iota := g i_\iota$, and see that $\alpha_\iota = \beta_\iota$ holds by assumption. These natural transformations give rise to maps $\operatorname{colim} \mathbf{ay}C_\iota \to G$. We thus compute

$$f \operatorname{colim} i_\iota = \operatorname{colim} \alpha_\iota = \operatorname{colim} \beta_\iota = g \operatorname{colim} i_\iota,$$

and as $\operatorname{colim} i_\iota$ is the identity map the desired follows immediately. $\square$

Remark 1.3.1. Suppose that $\mathcal{F}$ is a collection of arrows into X which is jointly epimorphic in a category which has all colimits. Then one can consider

$$\coprod_{f\in\mathcal{F}} \operatorname{dom} f \xrightarrow{\coprod_{f\in\mathcal{F}} f} X,$$

and write $e := \coprod_{f\in\mathcal{F}} f$ for convenience. For each $(h : C \to X) \in \mathcal{F}$ there is the canonical map $i_h : C \to X$. Consider a parallel pair $f, g : X \to Y$ and suppose that $fe = ge$. Then it follows that $fei_h = gei_h$, but due to the construction of e this entails that $fh = gh$. By assumption we have that $f = g$, proving e to be an epimorphism. Conversely, if this e is an epimorphism then the family $\mathcal{F}$ is jointly epimorphic. As a corollary of the above lemma, this discussion entails that every Grothendieck topos has a generating object.

SECTION 1.4 Geometric Morphisms

Given a map of spaces $f : X \to Y$ we know of the inverse image functor $f^* : \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ and the direct image functor $f_* : \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$. From Theorem 1.2.4 we learn that the inverse image functor is left adjoint to the direct image functor, and Theorem 1.2.2 proves that f^* is *left exact*, that is to say, preserves finite limits. We generalize this to generic toposes in the following definition.

Definition 1.4.1 (Geometric Morphism). A map of toposes $f : \mathcal{E} \to \mathcal{D}$ is given by a pair of functors $f^* : \mathcal{D} \to \mathcal{E}$ and $f_* : \mathcal{E} \to \mathcal{D}$, where f^* is said to be the *inverse image* and f_* the *direct image* of f, such that f^* is left exact and left adjoint to f_*.

From the above it is clear that each continuous map between topological spaces gives rise to a geometric morphism. As a more concrete example, think of the map $X \to \mathbf{1}$ and the geometric morphism $\pi : \mathbf{Sh}(X) \to \mathbf{Set}$ it induces. Indeed, the inverse image $\pi^* : \mathbf{Set} \to \mathbf{Sh}(X)$ sends a set S to the sheaf of sections of the étale space $S \times X$. This functor is left exact and left adjoint to $\pi_* : \mathbf{Sh}(X) \to \mathbf{Set}$. We furthermore proved in Subsection 1.2.3 that if X is locally connected, then π_* has as a left adjoint the functor $\pi_! : \mathbf{Sh}(X) \to \mathbf{Set}$, which sends a sheaf to the connected components of its display space.

Recall the geometric morphism induced by a point, as covered in Subsection 1.2.4. Given a point $x \in X$, this gives rise to a geometric morphism $x : \mathbf{Set} \to \mathbf{Sh}(X)$. The corresponding inverse image sends a sheaf to its stalks over x, and the direct image sends a set to the skyscraper sheaf over x. This case of points of a topological space gives rise to the following generalization.

Definition 1.4.2 (Point). A *point* of a topos $\mathcal{E}$ is a geometric morphism $\mathbf{Set} \to \mathcal{E}$.

In Subsection 1.2.4 we proved Corollary 1.2.3 which states, in the terms of the following definition, that the inverse image functors induced by the points of a topological space X yield a set of jointly conservative functors from $\mathbf{Sh}(X)$ to $\mathbf{Set}$.

Definition 1.4.3 (Jointly Conservative). Given categories $\mathcal{C}$ and $\mathcal{D}$ a collection K of functors $\mathcal{C} \to \mathcal{D}$ is said to be *jointly conservative* from $\mathcal{C}$ to $\mathcal{D}$ when for any arrow $f : X \to Y$ of $\mathcal{C}$ it holds that f is an isomorphism precisely if $F(f)$ is for all $F \in K$.

A topos which has a collection of jointly conservative functors from this topos to $\mathbf{Set}$ is said to have *enough points*. So Corollary 1.2.3 can be summarized as saying that $\mathbf{Sh}(X)$ has enough points.

Definition 1.4.4 (Surjective Geometric Morphism). A geometric morphism $f : \mathcal{C} \to \mathcal{D}$ is said to be *surjective* when its inverse image part f^* is faithful. The morphism f is an *embedding* when its inverse image f^* is both full and faithful.

The following lemma comes from Lemma 7.4.3 of Mac Lane and Moerdijk (1991). In particular, this states that an embedding reflects isomorphisms.

Lemma 1.4.1 (Notions of Surjectivity). *For a geometric morphism $f : \mathcal{C} \to \mathcal{D}$ the following are equivalent:*

(i) f is surjective

(ii) for each object $D \in \mathcal{D}$, f^* induces an injective homomorphism of subobject lattices $\mathrm{Sub}_{\mathcal{D}}(E) \to \mathrm{Sub}_{\mathcal{C}}(f^*E)$.

(iii) f^* reflects isomorphisms

Proof. First suppose that (i) holds. Let $m : X \to D$ be a (representative of a) subobject of $D \in \mathcal{D}$. Consider the pullback diagram specifying its characteristic map χ_X.

$$\begin{array}{ccc} X & \xrightarrow{\;!\;} & 1 \\ {\scriptstyle m}\downarrow & & \downarrow{\scriptstyle \mathrm{true}} \\ D & \xrightarrow{\;\chi_X\;} & \Omega \end{array} \qquad\qquad \begin{array}{ccc} f^*X & \xrightarrow{\;!\;} & f^*1 \\ {\scriptstyle f^*m}\downarrow & & \downarrow{\scriptstyle f^*\mathrm{true}} \\ f^*D & \xrightarrow{\;f^*\chi_X\;} & f^*\Omega \end{array}$$

Now as f^* is left exact, the right-hand diagram is a pullback as well. Let $n : Y \to D$ also be a subobject of D, and suppose that there exists an isomorphism $v : f^*X \to f^*Y$ such that $f^*m \circ v = f^*n$. This is exactly the statement that f^*m and f^*n represent the same subobject. Note that by the diagrams above we derive $\chi_{f^*X} \circ f^*m = \mathrm{true} \circ !_{f^*X}$, so it certainly holds that

$$\chi_{f^*X} \circ f^*n = \chi_{f^*X} \circ f^*m \circ v = \mathrm{true} \circ !_{f^*X} \circ v = \mathrm{true} \circ !_{f^*Y}.$$

One can perform a similar construction with the roles of X and Y reversed. Using the faithfulness of f^* we now derive that n and m represent the same subobject, proving (ii).

Suppose (ii) holds. Let $v : X \to Y$ be an arrow in $\mathcal{D}$ such that f^*v is an isomorphism. Consider the factorization of v and its inverse image under f, which preserves the factorization as f is both left- and right exact.

$$X \xrightarrow{\;m\;} \mathrm{im}\, f \xrightarrow{\;e\;} Y \qquad\qquad f^*X \xrightarrow{\;f^*m\;} f^*\mathrm{im}\, f \xrightarrow{\;f^*e\;} f^*Y$$

We know f^*v to be an isomorphism, so $f^*\mathrm{im}\, f = f^*Y$. From (ii) it now follows that $\mathrm{im}\, f = Y$, so v is an epimorphism. To prove that v is an isomorphism, we now prove it to be mono as well. Consider the pullback which defines the kernel pair $k_1, k_2 : K \to X$ and the factorization of the diagonal arrow $X \to K$, both of which are preserved under f^*.

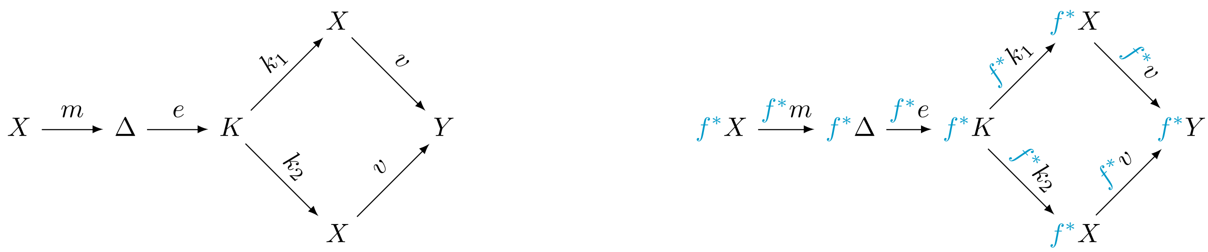

Now as f^*v is iso, we know it to be mono, which forces $f^*k_1 = \mathrm{id}_K = f^*k_2$. As a consequence we see that $f^*me = \mathrm{id}_X$, forcing $f^*\Delta = f^*K$. Using again (ii), we derive $K = \Delta$. This yields $e = \mathrm{id}_K$ which makes $v = m$ a monomorphism, proving (iii).

Now suppose (iii) holds. Let $v, w : X \to Y$ be arrows in $\mathcal{D}$ such that $f^*v = f^*w$. Consider their equalizer on the left, and note that as f^* is left-exact, this yields the equalizer on the right.

$$E \xrightarrow{\;e\;} X \underset{w}{\overset{v}{\rightrightarrows}} Y \qquad\qquad f^*E \xrightarrow{\;f^*e\;} f^*X \underset{f^*w}{\overset{f^*v}{\rightrightarrows}} f^*Y$$

We know that $f^*v = f^*w$, so f^*e must be an isomorphism. By (iii) we now know that e is an isomorphism, forcing $v = w$. This proves (i). □

Chapter 2

Spatial Cover

In Section 1.3 we generalized the notion of sheaves to include sheaves over small categories compatible with a certain Grothendieck topology. The purpose of this chapter is to construct a topological space $\mathbf{X}$ such that a Grothendieck topos $\mathcal{E}$ can be embedded into the topos of sheaves over this space $\mathbf{X}$. In order to be able to construct the space $\mathbf{X}$, we require the topos $\mathcal{E}$ to be equipped with an object G whose subobjects generate $\mathcal{E}$ and a set of points K such that this set is jointly conservative. In the following we will refer to such points as *small points*, and say that a topos has enough small points when such a set exists.

This restriction is not too harsh. As a sanity check, note that the topos of sheaves on a certain topological space X comes with the set of points arising from the actual points of X, and this set is jointly conservative by Corollary 1.2.3. Many (non-spatial) toposes have enough small points, by Corollary 7.7.7 of Johnstone (1997) it follows that any topos with enough points satisfies this requirement. It was proven by Deligne in an appendix to Exposé 7 of Artin et al. (1972) that each coherent topos has enough points. The latter are a central object of study in algebraic geometry.

The material in this chapter is based on Butz and Moerdijk (1999). This paper also aims to prove a result on cohomology; we do not. We use all of their results not related to cohomology. We deviate from this text in that we stay on the side of étale spaces until the very last moment, where we pass over to sheaves. Some inspiration is taken directly from the dissertation Butz (1996), although the formulation there differs quite a lot from ours. Some intuition could be taken from Moerdijk (1996).

The argument as given in this chapter can be decomposed in four major parts. The first part is in constructing a space $\mathbf{X}$, using the small points and partial enumerations of the inverse-image of G under these points. We can then construct an étale space to each object of $\mathcal{E}$ in a functorial way, and this will give us a functor $\mathbf{e}^* : \mathcal{E} \to \mathbf{Etale}(\mathbf{X})$. The first part is concluded by proving this functor to be left-exact and colimit preserving, from whence general theory shows this to be the inverse-image part of a geometric morphism. We will do all of this in Section 2.2. In order to familiarize the reader with these partial enumerations and their geometric structure, we treat those in Section 2.1. An important concept we introduce here is the enumeration space of the inverse-image of the chosen generating object under a point.

We then proceed to show that étale spaces over this enumeration space are related to étale spaces over $\mathbf{X}$ via the geometric morphism $\pi : \mathbf{Etale}(\mathbf{X}) \to \mathbf{Set}$ as discussed in great detail in Subsection 1.2.3. This will give us a commuting diagram, which satisfies the Beck–Chevalley condition. The thus obtained natural isomorphism will be crucial to obtain our end result.

The third part is in providing the left-adjoints to formulate the Beck–Chevalley condition of the aforementioned diagram. A left-adjoint $\pi_!$ to π_* has been provided in Theorem 1.2.5, we provide a left adjoint $\mathbf{e}_! : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$ in Theorem 2.3.1. This entails that $\mathbf{e}$ is an essential geometric morphism, so in particular, it is known to preserve both limits and colimits.

Finally, we need to show that $\mathbf{e}$ actually is an embedding, which comes down to proving that $\mathbf{e}^*$ is full and faithful. To this end we prove the Beck–Chevalley condition, and use this to prove the counit of the adjunction between $\mathbf{e}^*$ and $\mathbf{e}_!$ to be an isomorphism. Having done all this, we have constructed an essential geometric embedding of the topos $\mathcal{E}$ into $\mathbf{Etale}(\mathbf{X})$. We can of course compose this embedding with the equivalence of categories between $\mathbf{Etale}(\mathbf{X})$ and $\mathbf{Sh}(\mathbf{X})$, finishing the argument. This yields this chapter's final result in the form of Theorem 2.3.3.

SECTION 2.1 Enumeration Space

We first take a step back from the situation described above and consider the space of partial enumerations on a generic set. To this end, let I be some infinite set and let S be a set which is smaller than I, such that S injects into I. We first introduce the notion of a partial enumeration. In the following definition we speak of partial functions, in the understanding that a partial map $\alpha : D \subseteq I \to S$ is simply a set-theoretic function $\alpha : D \to S$. So a partial function from I to S is simply a function $I \to S$ defined only on a (not necessarily strict) subset of I. We denote the domain of such a partial function α, the set of points in I for which the map is defined, by $\operatorname{dom}\alpha$.

Definition 2.1.1 (Partial Enumeration). A partial enumeration is a partial map $\alpha : D \subseteq I \to S$ with infinite fibers. We call a partial map $\alpha : D \subseteq I \to S$ finite when $\operatorname{dom}\alpha := D$ is finite.

Given two partial functions $\alpha_i : D_i \subseteq I \to S$ one can take their union $\alpha_1 \cup \alpha_2$. This is again a partial function precisely if α_1 and α_2 agree on the intersection of D_1 and D_2. Moreover, it is enumeration (or finite) if both α_1 and α_2 are. We write $\alpha_1 - \alpha_2$ for the restriction of α_1 to $D_1 - D_2$. When we write $\alpha_1 \subseteq \alpha_2$, we mean to say that the graph of α_1 is a subset of that of α_2.

We wish to endow the space $\mathbf{En}(I, S)$ of all partial enumerations with a topology, where neighborhoods are to reflect agreement with finite partial functions. Let $u : D \to S$ be a finite partial function, we define the basic open V_u as below.[1]

$$V_u := \{\alpha \in \mathbf{En}(I, S) \mid u \subset \alpha\}$$

The topology on $\mathbf{En}(I, S)$ is defined to be the topology as generated by these basic open sets. Note that the empty partial function gives rise to $V_\emptyset = \mathbf{En}(I, S)$. It is also easy to see that the intersection of two of these basic opens is simply the union of the finite partial functions defining them, and empty when the previous does not make sense. A bit more formally, $V_u \cap V_w$ consists of those partial enumerations α such that both $u, w \subseteq \alpha$. Now if u and w coincide on the intersection, $u \cup w$ defines a new partial function and $\alpha \in V_{u\cup w}$ precisely if $\alpha \in V_u \cap V_w$. Yet is u and w do not coincide on their intersection, then no partial enumeration α can extend both u and w, so the intersection of V_u and V_w is simply the empty set. Consequently we see that sets V_u do yield a basis for a topology on $\mathbf{En}(I, S)$. The partial nature of the elements of $\mathbf{En}(I, S)$ makes it easy to prove this topology to be both connected and locally connected, which is but an easy corollary of the following lemma.

Lemma 2.1.1. *Let $u : D \to S$ be a finite partial function, now V_u is connected.*

Proof. Suppose that V_u is the union of open U_1 and U_2. For $i = 1, 2$ we can pick some element $\alpha_i \in U_i$ and pick u_i such that $\alpha_i \in V_{u_i} \subseteq U_i$. Remark that $u \subseteq u_i \subseteq \alpha_i$ holds.

We now define $\beta := (\alpha_1 - u_2) \cup u_2$ and $\gamma := (\alpha_1 - u_2) \cup u$, and see that these are well-defined. Obviously, these are elements of $\mathbf{En}(I, S)$. It is clear that $u \subset \gamma \subseteq \beta, \alpha_1$ holds, so both β and α_1 are contained in every neighborhood of γ. We know that $\gamma \in V_u$, so $\gamma \in U_i$ for $i = 1, 2$. In case that $i = 1$, then β is contained $U_1 \cap U_2$, and if $i = 2$ then α_1 is. This proves that the intersection of U_1 and U_2 is non-empty, hence V_u is connected. □

Corollary 2.1.1. *The space $\mathbf{En}(I, S)$ is connected and locally connected.*

Proof. As $V_\emptyset$ is a basic open, it is connected by the above lemma. Moreover, any neighbourhood around a point contains a basic open, which is connected by the above lemma. □

Consider the (unique) map of spaces $\pi : \mathbf{En}(I, S) \to \mathbf{1}$, and recall from Subsection 1.2.3 that this gives rise to a geometric morphism $\pi : \mathbf{Etale}(\mathbf{En}(I, S)) \to \mathbf{Set}$. Moreover, by Theorem 1.2.5 and the above corollary we see that this geometric morphism is essential. That is to say, the inverse image functor $\pi^* : \mathbf{Set} \to \mathbf{Etale}(\mathbf{En}(I, S))$ has a left adjoint $\pi_! : \mathbf{Etale}(\mathbf{En}(I, S)) \to \mathbf{Set}$ which maps an étale space over $\mathbf{En}(I, S)$ simply to the set of its connected components. Using again the above corollary and Proposition 1.2.1 we see that π is in fact an embedding. This will be used later on in Section 2.3.

[1]We write $u \subset \alpha$ instead of $u \subseteq y$ in the definition here, but as α is a partial enumeration and u is finite it is clear that $u = \alpha$ can never hold.

In the remainder of this chapter we keep the Grothendieck topos $\mathcal{E}$ fixed, together with an object G whose subobjects generate $\mathcal{E}$ and a jointly conservative set of points K. We will call the points contained in this set K *small*. We will now construct the space $\mathbf{X}$, a point of which will be represented by a small point p and a partial enumeration of $p^*(G)$. In order to consider enumerations, we need an index set I such that $p^*(G)$ injects into I for every small point p. From now on we fix a choice of I, note that choosing $I := \coprod_{p \in K} G_p$ suffices. Because I is fixed from now on, in the following we will write $\mathbf{En}(S)$ to mean $\mathbf{En}(I, S)$.

We now have the tools necessary to construct the space $\mathbf{X}$. Given $\langle p, \alpha \rangle$ and $\langle q, \beta \rangle$ with p and q small points and α and β partial enumerations of $p^*(G)$ an $p^*(G)$ in I respectively, we call these pairs equivalent when there exists a natural isomorphism $\theta : p^* \to q^*$ with $\beta = \theta_G \alpha$. Define $\mathbf{X}$ to be the space of equivalence classes of such pairs. Note that for a point $\langle p, \alpha \rangle \in X$ one has that $\alpha \in \mathbf{En}(p^*(G))$.

Next we need a collection of basis opens on $\mathbf{X}$. Let P be a subsheaf of G^n for some $n \in \mathbb{N}$ and let $\mathbf{i} \in I$, we define the set

$$U_{\mathbf{i},P} := \{\langle p, \alpha \rangle \mid \alpha(\mathbf{i}) \in p^*(P)\}.$$

Let us elaborate a bit on the syntax here. When writing $\mathbf{i} \in I$ we mean that $\mathbf{i} = i_1, \ldots, i_n$ is such that $i_\iota \in I$ for all $\iota = 1, \ldots, n$. Because α is a partial enumeration of $p^*(G)$, we can evaluate α at i_ι and obtain a value $\alpha(i_\iota)$ exactly when i_ι is in the domain of α. We can do this for all $\iota = 1, \ldots, n$ and obtain a sequence of values in G of length n. This is exactly the same as an element of G^n. For brevity we simply write $\alpha(\mathbf{i}) \in P$ to mean that $\alpha(i_\iota)$ is defined for all $\iota = 1, \ldots, n$ and $\big(\alpha(i_1), \ldots, \alpha(i_n)\big) \in p^*(P)$. Let us now inspect some elementary properties of the structure we just created. In the following we will very often use the fact that small points $p : \mathbf{Set} \to \mathcal{E}$ have an inverse image part $p^* : \mathcal{E} \to \mathbf{Set}$ which is left exact and colimit preserving, by their very definition. We will no longer explicitly mention this at every point, and tacitly assume this to be so.

Remark 2.2.1 (Removal of Repetitions). It is more convenient to work with sequences $\mathbf{i} \in I$ without repetitions. Of course, for us to be able to do this safely, we have to demonstrate that this does not affect the sets $U_{\mathbf{i},P}$ defined above. To this end, let $i_1, \ldots, i_{n+1} \in I$ be a sequence and suppose that a repetition occurs, say $i_\iota = i_\kappa$ for some $\iota < \kappa$. We can now consider the map $\delta : G^n \to G^{n+1}$ defined by

$$\delta = \big\langle \pi_1, \ldots, \pi_\iota, \ldots, \pi_{\kappa-1}, \pi_\iota, \pi_\kappa, \ldots, \pi_n \big\rangle,$$

which sends each component to itself, and copies the ι-th component to the κ-th place. Now we can form the pullback of $P \subseteq G^{n+1}$ along δ to obtain P', a subsheaf of G^n. As p^* preserves pullbacks, it is clear that $\alpha(i_1, \ldots, i_{n+1}) \in p^*(P)$ precisely if $\alpha(i_1, \ldots, i_{\kappa-1}, i_{\kappa+1}, \ldots, i_n) \in p^*(P')$. We thus always assume sequences $\mathbf{i} \in I$ to be without repetitions, as the open sets they yield are equal.

Remark 2.2.2 (Powers of Enumeration). Given a point $\langle p, \alpha \rangle$ of X, we know α to be a partial enumeration of $p^*(G)$. One readily deduces that

$$\alpha^n : \langle i_1, \ldots, i_n \rangle \mapsto \langle \alpha(i_1), \ldots, \alpha(i_n) \rangle$$

is a partial enumeration of $p^*(G)^n = p^*(G^n)$.

Lemma 2.2.1. *The sets $U_{\mathbf{i},P}$ for $\mathbf{i} \in I$ and $P \subseteq G^n$ with $n \in \mathbb{N}$ form a basis for a topology on $\mathbf{X}$.*

Proof. Given $\mathbf{i}$, $\mathbf{j}$, P and Q it is easy to compute that

$$U_{\mathbf{i},P} \cap U_{\mathbf{i},Q} = U_{\mathbf{i};\mathbf{j},P \times Q},$$

where $\mathbf{i};\mathbf{j}$ denotes the concatenation of $\mathbf{i}$ and $\mathbf{j}$. Indeed, $\langle p, \alpha \rangle \in \mathbf{X}$ is such that $\alpha(\mathbf{i}) \in p^*(P)$ and $\alpha(\mathbf{j}) \in p^*(Q)$ precisely if $\alpha(\mathbf{i};\mathbf{j}) \in p^*(P \times Q) = p^*(P) \times p^*(Q)$. So we only need to show that the totality of basic opens covers the entire space $\mathbf{X}$. Take $\langle p, \alpha \rangle$ an arbitrary (representative of a) point of $\mathbf{X}$. Take the empty sequence $\emptyset$ and see that $\alpha(\emptyset) \in p^*(G^0) = \mathbf{1}$, proving $\langle p, \alpha \rangle \in U_{\emptyset,\mathbf{1}}$. □

The goal of this section is to construct a geometric morphism $\mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$. In order to attain this goal, it suffices to construct a functor $\mathbf{e}^* : \mathcal{E} \to \mathbf{Etale}(\mathbf{X})$ which preserves colimits and is left exact. We thus assign an étale space over $\mathbf{X}$ to each object (sheaf) in $\mathcal{E}$ in a functorial fashion. Let F be an object in $\mathcal{E}$, we now need to construct a topological space, which we will denote by $\mathbf{e}^*(F)$ for convenience, with a local homeomorphism down to $\mathbf{X}$. We define $\mathbf{e}^*(F)$ as a topological space to be the equivalence classes of triples $\langle p, \alpha, x\rangle$, where $\langle p, \alpha\rangle$ is (the representative of a) point in X and $x \in p^*(F)$. This equivalence relation extends the equivalence relation on representatives of points on X in a natural manner to these triples, so $\langle p, \alpha, x\rangle$ is equivalent to $\langle q, \beta, y\rangle$ if $\langle p, \alpha\rangle$ is equivalent to $\langle q, \beta\rangle$ via θ and $\theta_F(x) = y$.

In order to define a topology on $\mathbf{e}^*(F)$ we define the sets

$$U_{\mathbf{i},P,f} := \left\{\langle p, \alpha, x\rangle \mid \langle p, \alpha\rangle \in U_{\mathbf{i},P} \text{ and } p^*(f)\big(\alpha(\mathbf{i})\big) = x\right\},$$

with $\mathbf{i} \in I$ and $f : P \to F$ an arrow in $\mathcal{E}$. Recall that as $P \subseteq G^n$ for some n, and $\mathbf{i} = i_1, \dots, i_n$, the statement $x = p^*(f)(\alpha(\mathbf{i}))$ actually means that x equals the $p^*(f)$ image of the tuple $\alpha(i_1), \dots, \alpha(i_n)$, which is an element of $p^*(P)$ by the first constraint.

Lemma 2.2.2. *The sets $U_{\mathbf{i},P,f}$ with $\mathbf{i}$ a sequence in I, P a subobject of a finite power of G and $f : P \to F$ an arrow in $\mathcal{E}$ form the basis for a topology on $\mathbf{e}^*(F)$.*

Proof. Given two basic sets $U_{\mathbf{i},P,f}$ and $U_{\mathbf{j},Q,g}$ we can form the pullback below.

$$\begin{array}{ccc} C \times_F D & \xrightarrow{\pi_Q} & Q \\ {\scriptstyle \pi_P}\downarrow & & \downarrow{\scriptstyle g} \\ P & \xrightarrow{f} & F \end{array}$$

Realize that when we set $g\pi_Q = h = f\pi_P$, then $U_{\mathbf{i};\mathbf{j},P\times_F Q,h} = U_{\mathbf{i},P,f} \cap U_{\mathbf{i},Q,g}$.

We now need to prove that the basic open sets cover $\mathbf{e}^*(F)$. First, by Remark 1.3.1 we know that

$$\coprod_{f:P\to F, P\subseteq G^n} \operatorname{dom} f \xrightarrow{e} F$$

is an epimorphism. Let $\langle p, \alpha, x\rangle \in \mathbf{e}^*(F)$ be arbitrary. Now as an epimorphism can be expressed as a pushout, and p^* preserves colimits, we know that

$$\coprod_{f:P\to F, P\subseteq G^n} \operatorname{dom} p^*(f) \xrightarrow{p^*(e)} p^*(F)$$

is a surjection. Consequently, there must be some $P \subseteq G^n$, some $f : P \to F$ and some $c \in p^*(P)$ such that $p^*(f)(c) = x$. By Remark 2.2.2 we know that α^n is a partial enumeration as well. Consequently, there is some $\mathbf{i}$ such that $\alpha(\mathbf{i}) = c$. One can now readily deduce that $\langle p, \alpha, x\rangle \in U_{\mathbf{i},P,f}$. □

Due to this lemma, we can endow $\mathbf{e}^*(F)$ with the topology as generated by the basic open sets $U_{\mathbf{i},P,f}$. There is a projection map

$$\pi : \mathbf{e}^*(F) \to X, \qquad \langle p, \alpha, x\rangle \mapsto \langle p, \alpha\rangle\,.$$

To make $\mathbf{e}^*(F)$ an étale space, we show that this map is continuous and a local homeomorphism. Continuity follows from the fact that the fiber above $\langle p, \alpha\rangle$ is simply all of $p^*(F)$, so we derive

$$\pi^{-1}(U_{\mathbf{i},P}) = \bigcup_{f:P\to F} U_{\mathbf{i},P,f}\,.$$

Lemma 2.2.3. *The map $\pi : \mathbf{e}^*(F) \to X$ is a local homeomorphism.*

Proof. For each basic open set there is a section

$$\sigma : U_{\mathbf{i},P} \to U_{\mathbf{i},P,f}, \qquad \langle p, \alpha \rangle \mapsto \langle p, \alpha, p^*(f)(\alpha(\mathbf{i})) \rangle .$$

It is easy to see that $\sigma\pi = \mathrm{id}_X$. Moreover, for $\langle p, \alpha, x \rangle \in U_{\mathbf{i},P}$ we know $p^*(f)(\alpha(\mathbf{i})) = x$, so $\sigma(\langle p, \alpha \rangle) = \langle p, \alpha, x \rangle$ proving $\pi\sigma = \mathrm{id}_{\mathbf{e}^*(F)}$. We now know σ to be surjective, from whence continuity easily follows. Indeed, if $U_{\mathbf{j},Q}g \subseteq U_{\mathbf{i},P,f}$ then we know its pre-image to simply be all of $U_{\mathbf{j},Q}$. ☐

Remark 2.2.3 (Stalks and Points). Recall from Subsection 1.2.4 that a point $\langle p, \alpha \rangle \in \mathbf{X}$ gives rise to a geometric morphism $\langle p, \alpha \rangle : \mathbf{Set} \to \mathbf{Etale}(\mathbf{X})$. By Corollary 1.2.3 we know that the totality of the inverse image functors $\langle p, \alpha \rangle^* : \mathbf{Etale}(\mathbf{X}) \to \mathbf{Set}$ is jointly conservative. This motivates us to know what the composite $\langle p, \alpha \rangle^* \mathbf{e}^* : \mathcal{E} \to \mathbf{Set}$ looks like. To see this, fix a (representative of a) point $\langle p, \alpha \rangle$ of $\mathbf{X}$ and a sheaf F of $\mathcal{E}$. From the description in Subsection 1.2.4 we know that $\langle p, \alpha \rangle^* \mathbf{e}^*(F)$ simply equals the pre-image of $\langle p, \alpha \rangle$ under the map π. It is easy to see that $\langle q, \beta, x \rangle$ lies in the pre-image of $\langle p, \alpha \rangle$ under π precisely if $\langle p, \alpha \rangle = \langle q, \beta \rangle$. But having fixed a representative of the left-hand side, these triples are fully determined by the elements $x \in p^*(F)$. Consequently we see that $\langle p, \alpha \rangle^* \mathbf{e}^*(F)$ is isomorphic to $p^*(F)$. This isomorphism is natural in F, but this statement is meaningless until we make $\mathbf{e}^*$ into a functor.

Earlier we stated that the assignment $\mathbf{e}^*$, which maps objects (sheaves) in $\mathcal{E}$ to étale spaces over X, is functorial. Let us now give substance to this claim by constructing a map $\mathbf{e}^*(h)$ to a map of sheaves $h : E \to F$. We define this map as below

$$\mathbf{e}^*(h) : \qquad \langle p, \alpha, x \rangle \in \mathbf{e}^*(E) \qquad \mapsto \qquad \langle p, \alpha, h_p(x) \rangle \in \mathbf{e}^*(F).$$

It is clear that this map respects the projections. To show continuity it suffices to cover the image of $\mathbf{e}^*(h)$ with basic open sets, and prove that their inverse image is again open. We thus pick some $\langle p, \alpha, x \rangle$ in $\mathbf{e}^*(E)$, and consider a basic open $V' := U_{\mathbf{i},P,f}$ around its image $\langle p, \alpha, h_p(x) \rangle$. We now will construct an open $W \subseteq \mathbf{e}^*(E)$ and refine V' to an open V still containing $\mathbf{e}^*(h)(x)$ such that $\mathbf{e}^*(h)^{-1}(V) = W$.

First form the pullback of f along h, and consider the picture below.

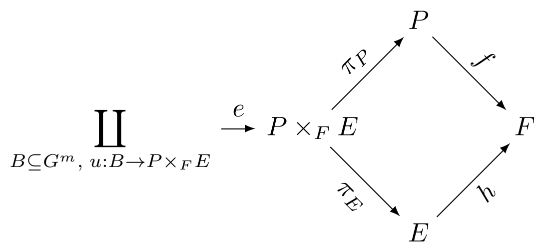

The object on the far-left is the coproduct of all arrows into the pullback $P \times_F E$. As before, it follows that there must be some subobject $B \subseteq G^m$, some arrow $u : B \to P \times_F E$ and some element $b \in B_p$ such that $u_p(b) = \langle \alpha(\mathbf{i}), x \rangle \in (P \times_F E)_p$. Now as B is a subobject of G^m we know that there exist $\mathbf{j} \in J$ such that $\alpha(\mathbf{j}) = b$. Define

$$\mathrm{graph}\left(B \xrightarrow{u} P \times_F E \xrightarrow{\pi_P} P \right) =: D \subseteq B \times P \subseteq G^m \times G^n = G^{n+m}.$$

We set $W := U_{\mathbf{j};\mathbf{i},D,\pi_E u \pi_B}$ and $U_{\mathbf{j};\mathbf{i},G^m \times C, f\pi_C}$ and claim that $\mathbf{e}^*(h)^{-1}(V) = W$. Remark that $\langle q, \beta, y \rangle \in V$ precisely if $\langle q, \beta, y \rangle \in V'$ and $\mathbf{j} \in \mathrm{dom}\,\beta$. Let $\langle q, \beta, y \rangle$ be a point of $\mathbf{e}^*(E)$. By spelling out the definitions we see that this is an element of W exactly when (i) $\mathbf{i}, \mathbf{j} \in \mathrm{dom}\,\beta$, (ii) $\beta(\mathbf{ji}) \in D$ and (iii) $y = (\pi_E u \pi_B)_p(\beta(\mathbf{ji}))$. We expand the definitions in (ii) to see that this holds precisely if $u\beta(\mathbf{j}) = \langle \beta i, z \rangle$ for some $z \in E_q$, subject to the condition that $h_q(z) = f_q(\beta(i))$. Using this knowledge in rewriting (iii) we see that $y = \pi_E u(\beta(\mathbf{j})) = z$. From here we immediately derive that $\langle q, \beta, y \rangle \in W$ if and only if $\langle q, \beta, h_q(y) \rangle \in V$, proving continuity. This makes $\mathbf{e}^*$ a bonafide functor $\mathcal{E} \to \mathbf{Etale}(X)$.

Now that we know $\mathbf{e}^* : \mathcal{E} \to \mathbf{Etale}(X)$ to be a functor, we can revisit the discussion in Remark 2.2.3. Consider a map $h : E \to F$ as above, and consider the diagram below, all the while fixing a representative $\langle p, \alpha \rangle$ of a point of $\mathbf{X}$.

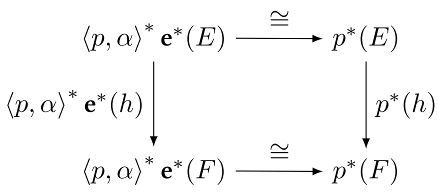

Now let $\langle q,\beta,x\rangle$ be an element of $\langle p,\alpha\rangle^* \mathbf{e}^*(E)$, which is to say, $\langle q,\beta\rangle = \langle p,\alpha\rangle$. Via the isomorphism on top this triple is sent to x, which is then mapped to $p^*(h)(x)$. The other composite maps this to $\langle q,\beta,p^*(h)(x)\rangle$, and then to the very same element $p^*(h)(x)$. This shows that the diagram commutes, proving that the stated isomorphism is in fact natural. We can further rephrase the above into saying that the diagram below commutes.

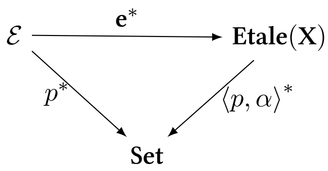

We use this fact to prove Lemma 2.2.5. Before we tackle this lemma, let us first contemplate the meaning of limits. Given a diagram $D : \mathcal{I} \to \mathcal{E}$, one can form its limit. This limit consists of a limiting cone $\Delta E \to D$, one often writes $\lim D$ for E.[2] Note that when we have two diagrams $A, B : \mathcal{I} \to \mathcal{E}$ and a natural transformation $f : A \to B$, then this gives rise to a map between their limits as depicted in the diagram below on the left. As a consequence, we can form the functor $\Delta \lim : \mathcal{E}^{\mathcal{I}} \to \mathcal{E}^{\mathcal{I}}$, which gives rise to the natural transformation $\lambda : \Delta \lim \to \mathrm{id}_{\mathcal{E}^{\mathcal{I}}}$ whose naturality square is depicted below on the right. Commutativity follows immediately from the diagram on the left.

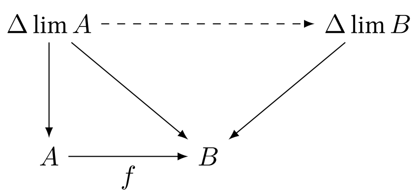

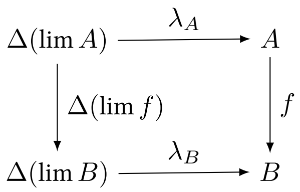

We wish to show that $\mathbf{e}^* : \mathcal{E} \to$ **Etale**(**X**) preserves finite limits. This will follow from the following, more general lemma. The key is to understand that when $D : \mathcal{I} \to \mathcal{E}$ is a diagram and $F : \mathcal{E} \to \mathcal{D}$ is a functor, then one can form a canonical arrow $F(\lim D) \to \lim FD$.

Lemma 2.2.4. *Suppose that $F : \mathcal{E} \to \mathcal{D}$, $G : \mathcal{D} \to \mathcal{C}$ and $H : \mathcal{D} \to \mathcal{D}$ are functors such that there is a natural isomorphism $\upsilon : GF \to H$. Assume that G and H preserve (finite, small) limits, and that $\mathcal{E}$, $\mathcal{D}$ and $\mathcal{C}$ all have these limits. Let $D : \mathcal{I} \to \mathcal{E}$ be a (finite, small) diagram and consider the limiting cones $\lambda_D : \Delta \lim D \to D$ and $\lambda_{FD} : \Delta \lim FD \to FD$. This gives rise to the unique map as indicated in the diagram below. Its G-image is an isomorphism.*

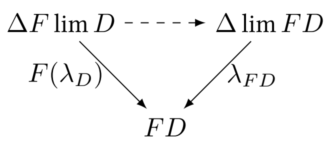

[2] This makes sense when we have a canonical choice for the limit of any diagram. This is the case in **Set**, so it is also the case in any Grothendieck topos, as limits are taken point-wise. In any case, vertices of limiting cones are canonically isomorphic.

Proof. The best way to prove this is to inspect the commutative diagram below. The dashed diagonal arrow are the unique arrows which make their respective triangles commute. The two squares are naturality of λ. The dashed rounded arrows are the unique arrows which make their triangles commute. Now as the center vertical arrows are all isomorphisms, we know that the outer diagram commutes. Consequently, as the bottom and diagonal rightmost arrows are isomorphisms by assumption, the desired follows immediately.

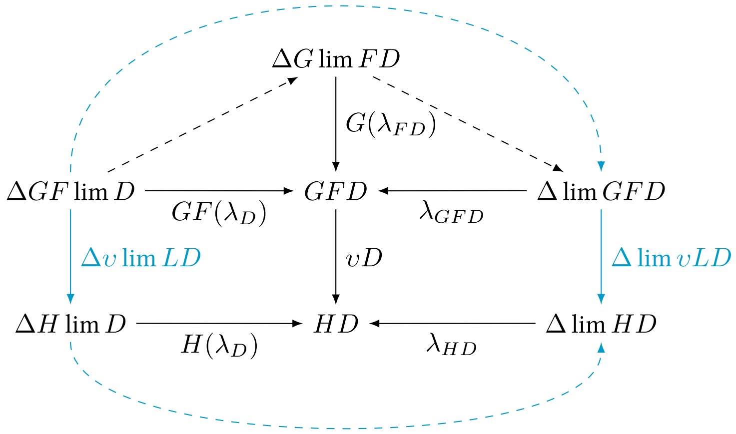

Lemma 2.2.5. *The functor* $\mathbf{e}^* : \mathcal{E} \to \mathbf{Etale}(\mathbf{X})$ *preserves colimits and finite limits.*

Proof. We know that $\mathbf{Etale}(\mathbf{X})$ has enough points, indeed, the set of all points of $\mathbf{X}$ yields a jointly conservative family as proven in Corollary 1.2.3. To show that $\mathbf{e}^*$ preserves finite limits, it thus suffices to check whether the canonical map $\mathbf{e}^*(\lim D) \to \lim \mathbf{e}^* D$ as explicated above is sent to an isomorphism by $\langle p, \alpha \rangle^*$ for every point $\langle p, \alpha \rangle$ of $\mathbf{X}$. This is an immediate consequence of Lemma 2.2.4, with $F = \mathbf{e}^*$, $G = \langle p, \alpha \rangle^*$ and $H = p^*$. We know G to be the geometric morphism as arising from a point, so it preserves finite limits by Theorem 1.2.2 and colimits because it has a right adjoint as proven in Theorem 1.2.4. The functor p^* preserves finite limits and colimits, because it is a point by assumption. The required diagram was shown to commute, or rather, we have shown the existence of a natural isomorphism $\langle p, \alpha \rangle^* \mathbf{e}^* \cong p^*$. An analogous lemma can be formulated for colimits. □

The above lemma proves that $\mathbf{e}^*$ is left exact, and preserves colimits. From this we can derive that $\mathbf{e}^*$ gives rise to a geometric morphism by means of general theory, as we explicate below.

Theorem 2.2.1. *The functor* $\mathbf{e}^*$ *is the inverse-image part of a geometric morphism* $\mathbf{e} : \mathbf{Etale}(X) \to \mathcal{E}$.

Proof. This follows from Lemma 2.2.5 and the Special Adjoint Theorem. A proof of of this is given by Mac Lane (1997), the proof is dual to that of Theorem 5.8.2. □

SECTION 2.3 A Beck–Chevalley Condition

Given a small point $p : \mathbf{Set} \to \mathcal{E}$ we can form a continuous map $\mathbf{i}_p : \mathbf{En}(p^*(G)) \to X,\ \alpha \mapsto \langle p, \alpha \rangle$. Indeed, one can easily compute

$$(\mathbf{i}_p)^{-1}(U_{\mathbf{i},P}) = \{\alpha \in \mathbf{En}(p^*(G)) \mid \alpha(\mathbf{i}) \in p^*(G)\} = \bigcup_{\mathbf{c} \in p^*(G)} U_{\mathbf{i} \mapsto \mathbf{c}},$$

where we write $\mathbf{i} \mapsto \mathbf{c}$ for the map which maps i_ι to c_ι, proving continuity. To our further discussion of this map it is crucial to note that $U_{\mathbf{i} \mapsto \mathbf{c}}$ is a connected component. It is connected by Lemma 2.1.1. To see that it is a component, suppose that it is contained in some connected open. Now as each open is the union of basic open sets, and as each basic open set is connected, we know this to be a basic open set, say U_u. Clearly, u must be undefined on anything but $\mathbf{i}$ in order to include $U_{\mathbf{i}\mapsto\mathbf{c}}$. Furthermore, it must be the case that $u(\mathbf{i}) = \mathbf{c}$, forcing $U_u = U_{\mathbf{i}\mapsto\mathbf{c}}$ as desired.

This map $\mathbf{i}_p$ gives rise to a geometric morphism $\mathbf{i}_p : \mathbf{Etale}(\mathbf{En}(p^*(G))) \to \mathbf{Sh}(X)$. We also know of the unique map $\pi : \mathbf{En}(p^*(G)) \to \mathbf{1}$, which gives rise to the map $\pi : \mathbf{Etale}(\mathbf{En}(p^*(G))) \to \mathbf{Set} \cong \mathbf{Etale}(\mathbf{1})$, as discussed in Subsection 1.2.3. We will now show, in Lemma 2.3.1 below, that the diagram

$$\begin{array}{ccc} \mathcal{E} & \xrightarrow{\mathbf{e}^*} & \mathbf{Etale}(\mathbf{X}) \\ \downarrow{\scriptstyle p^*} & & \downarrow{\scriptstyle \mathbf{i}_p^*} \\ \mathbf{Set} & \xrightarrow{\pi^*} & \mathbf{Etale}(\mathbf{En}(p^*(G))) \end{array} \tag{2.1}$$

commutes up to a natural isomorphism $\zeta : \mathbf{i}_p^*\mathbf{e}^* \to \pi^*p^*$. In Theorem 2.3.1 we prove that $\mathbf{e}$ is an essential geometric morphism, and π is essential for reasons explicated in Corollary 2.3.2. This gives us functors $\mathbf{e}_!$ and $\pi_!$, respectively left-adjoint to $\mathbf{e}^*$ and π^*, from whence we can form the composite below.

$$\pi_!\mathbf{i}_p^* \xrightarrow{\pi_!\mathbf{i}_p^*\eta^{\mathbf{e}}} \pi_!\mathbf{i}_p^*\mathbf{e}^*\mathbf{e}_! \xrightarrow{\pi_!\zeta\mathbf{e}_!} \pi_!\pi^*p^*\mathbf{e}_! \xrightarrow{\epsilon^\pi p^*\mathbf{e}_!} p^*\mathbf{e}_! \tag{2.2}$$

In the above composite we make use of the unit $\eta^{\mathbf{e}} : \mathrm{id}_{\mathbf{Etale}(\mathbf{X})} \to \mathbf{e}^*\mathbf{e}_!$ of the adjunction between $\mathbf{e}_!$ and $\mathbf{e}^*$, and of the counit $\epsilon^\pi : \pi_!\pi^* \to \mathrm{id}_{\mathbf{Set}}$ of the adjunction between $\pi_!$ and π^*. We prove that the above square (2.1) satisfies the Beck–Chevally condition in Lemma 2.3.5, which is to say, the above composite (2.2) is in fact an isomorphism.

Lemma 2.3.1. *There exists a natural isomorphism $\zeta : \mathbf{i}_p^*\mathbf{e}^* \to \pi^*p^*$, the natural isomorphism up to which (2.1) commutes.*

Proof. Let E be a sheaf of $\mathcal{E}$ and inspect the composites $\mathbf{i}_p^*\mathbf{e}^*(E)$ and $\pi^*p^*(E)$. The latter is simply the product $\mathbf{En}(p^*(G)) \times p^*(E)$, whereas the former is the pullback of $\mathbf{e}^*(E)$ along $\mathbf{i}_p : \mathbf{En}(p^*(G)) \to \mathbf{1}$. Look at the map below.

$$\langle x, \alpha\rangle \in p^*(E) \times \mathbf{En}(p^*(G)) \mapsto \langle \alpha, \langle p, \alpha, x\rangle\rangle \in \mathbf{En}(p^*(G)) \times_X \mathbf{e}^*(E)$$

Take a point $\beta \in \mathbf{En}(p^*(E))$, we now set out to prove that the map above is an isomorphism when restricted to the pre-image of β under the projection $p^*(E) \times \mathbf{En}(p^*(G)) \to \mathbf{En}(p^*(G))$, which is simply $p^*(E)$. At the other side, the pre-image of β are pairs $\langle \alpha, \langle q, \alpha, x\rangle\rangle$ where $\beta = \alpha$ and $\langle q, \alpha\rangle = \langle p, \alpha\rangle$, so these are fully determined by an element of $p^*(E)$. The map is thus bijective, proving the desired. □

With the lemma above we can prove an elementary property of the enumerations associated to a points within the same open of $\mathbf{X}$. The lemma states that sections of the étale space $\mathbf{e}^*(E) \to \mathbf{X}$ take the same value on nearby points in a neighbourhood $U_{\mathbf{i},P}$ when the associated enumerations take the same value on $\mathbf{i}$. This will help us to prove that the left-adjoint we construct to $\mathbf{e}^*$ is well-defined.

Corollary 2.3.1. *Let $\sigma : U_{\mathbf{i},P} \to \mathbf{e}^*(E)$ be a section. For any pair of points $\langle p, \alpha\rangle$ and $\langle p, \beta\rangle$ in $U_{\mathbf{i},P}$ it follows that if $\alpha(\mathbf{i}) = \beta(\mathbf{i})$ then $\sigma(\langle p, \alpha\rangle) = \sigma(\langle p, \beta\rangle)$.*

Proof. If $\alpha(\mathbf{i}) = \beta(\mathbf{i})$, then α and β lie in the same open $V := U_{\mathbf{i}\mapsto c}$ of $\mathbf{En}(p^*(G))$. By the above lemma we now know σ to be constant on $\mathbf{i}_p(V)$, from whence the desired is immediate. □

We now work to construct the left-adjoint $\mathbf{e}_! : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$ to $\mathbf{e}^*$. First realize that each object in $\mathbf{Etale}(\mathbf{X})$ is the colimit of basic opens $U_{\mathbf{i},P}$. Indeed, each topological space is simply the union of its opens, and each étale space above $\mathbf{X}$ is such that its (sufficiently small) opens are homeomorphic to the basic opens $U_{\mathbf{i},P}$. If we want $\mathbf{e}_!$ to be left-adjoint it has to preserve colimits, so we might as well only define it on basic opens $U_{\mathbf{i},P}$, and extend it via the colimit description as above. On a basic open $U_{\mathbf{i},P}$ we simply define $\mathbf{e}_!(U_{\mathbf{i},P} \to \mathbf{X}) := P$. We thus need to prove the existence of an isomorphism

$$\mathcal{E}(P,\, F) = \mathcal{E}\left(\mathbf{e}_! \begin{pmatrix} U_{\mathbf{i},P} \\ \cup\!\!\mid\!\downarrow \\ \mathbf{X} \end{pmatrix},\, F\right) \cong \mathbf{Etale}(\mathbf{X}) \begin{pmatrix} U_{\mathbf{i},P} & & \mathbf{e}^*(F) \\ \cup\!\!\mid\!\downarrow & , & \pi\downarrow \\ \mathbf{X} & & \mathbf{X} \end{pmatrix}. \tag{2.3}$$

To be able to do this, we first gather some information about the relation between open covers of basic opens in $\mathbf{X}$ and epimorphic families in $\mathcal{E}$. Given a finite sequence $i_1, \ldots, i_m \in \{1, \ldots, n\}$, one can form the map

$$\chi : S^n \to S^m,\ \langle \pi_{i_1}, \ldots, \pi_{i_m} \rangle .$$

An instance of the above arises when one assumes $\mathbf{i}$ to be a subsequence of $\mathbf{j}$. That is to say, $i_k = j_{a_k}$ for $1 \leq k \leq m$ and $\mathbf{a} \in \{1, \ldots, n\}$, where $\mathbf{j}$ is a sequence of length m and $\mathbf{i}$ of length n. This gives rise to a projection map $\chi : G^m \to G^n$.

Lemma 2.3.2. *Let $U_{\mathbf{i},P}$ and $U_{\mathbf{j},Q}$ be basic open sets of X and suppose that the latter is non-empty. The following are equivalent:*

(i) $U_{\mathbf{j},Q} \subseteq U_{\mathbf{i},P}$*;*

(ii) $\mathbf{i}$ *is a subsequence of* $\mathbf{j}$ *and the projection map* $\chi : G^m \to G^n$ *maps* B *into* C*.*

Proof. Suppose (i) holds. To prove that $\mathbf{i}$ is a subsequence of $\mathbf{j}$ we proceed by contradiction. First, pick a point $\langle p, \alpha \rangle \in U_{\mathbf{j},Q}$ and suppose that i_k does not occur amongst $\mathbf{j}$. Now define $\alpha' := \alpha - \{i_k\}$ and realize that $\langle \alpha', p \rangle \in U_{\mathbf{j},Q}$, but then $\langle \alpha', p \rangle \in U_{\mathbf{i},P}$ can not hold. As a consequence, we know $\mathbf{i}$ to be a subsequence of $\mathbf{j}$.

Let χ be the associated projection. We claim that $p^*(\chi)(p^*(Q)) \subseteq p^*(P)$ holds for every point small point p. Now as the set of small points p is assumed to be jointly conservative, this is sufficient to conclude that $\chi(Q) \subseteq P$. So let p be a small point and let $\langle g_1, \ldots, g_m \rangle \in p^*(Q) \subseteq p^*(G^m)$.[3] Define α' by sending j_k to g_k for $1 \leq k \leq m$, and extend it to a partial enumeration α of $p^*(Q)$. Note that this would in general not be possible were $\mathbf{j}$ to contain repetitions, but as we can safely assume it not to, everything works out. It now follows that $\langle p, \alpha \rangle \in U_{\mathbf{j},Q}$, because α is defined on $\mathbf{j}$ and $\alpha(\mathbf{j}) = \langle g_1, \ldots, g_m \rangle \in p^*(Q)$. Consequently, $\langle p, \alpha \rangle \in U_{\mathbf{i},P}$ follows by (i), from whence we derive $\alpha(\mathbf{i}) \in p^*(P)$. As this equals $\chi(\langle g_1, \ldots, g_m \rangle)$ we have proven (ii). The converse is a matter of spelling out definitions. ☐

Lemma 2.3.3. *Let $U := U_{\mathbf{i},P}$ be a basic open set and let $\mathcal{U} := \{U_\iota := U_{\mathbf{j}_\iota, P_\iota}\}_{\iota \in I}$ be a family of non-empty basic open subsets of U. To each U_ι we have the projection map $\chi_\iota : P_\iota \to P$ as guaranteed by Lemma 2.3.2. Now (ii) and (iii) are equivalent, and both are implied by (i).*

(i) U *is covered by* $\mathcal{U}$*;*

(ii) $\{\chi_\iota : P_\iota \to P\}_{\iota \in I}$ *is an epimorphic family;*

(iii) for all small points p *it holds that* $\coprod_{\iota \in I} p^*(\chi_\iota(P_\iota)) = p^*(P)$*.*

Proof. We first show that (ii) and (iii) are equivalent. Recall that (ii) holds exactly when

$$\coprod_{\iota \in I} P_\iota \xrightarrow{\coprod_{\iota \in I} \chi_\iota} P$$

is an epimorphic map. Note that the p^*-image of the above, for any point p, gives rise to a surjection. Now as a map is surjective precisely when its codomain is its image, and p^* preserves colimits, we obtain the equation as in (iii). From this the implication from (ii) to (iii) is immediate. The other direction follows from the assumption that the set of small points is jointly conservative.

[3] Every element of $p^*(Q)$ is of this form, as elements of B are tuples of elements of G, and p^* preserves finite limits by virtue of being left-exact.

Suppose that (i) holds, we wish to show that (ii) follows. So we let $c \in p^*(P)$ be arbitrary and aim to prove that $c \in p^*\left(\chi_\iota(P_\iota)\right)$ for some $\iota \in I$. Pick $\alpha \in \mathbf{En}(p^*(G))$ such that $\alpha(\mathbf{i}) = c$. We now know that $\langle p, \alpha\rangle \in U_{\mathbf{i},P} = U$. This gives us some $\iota \in I$ such that $\langle p, \alpha\rangle \in U_{\mathbf{j}_\iota, P_\iota}$. By assumption we know $\mathbf{i}$ to be a subsequence of $\mathbf{j}_\iota$. Consequently we obtain $b = \alpha(\mathbf{j}_\iota) \in p^*(P_\iota)$, from which we may derive $p^*(\chi_\iota)(b) = \alpha(\mathbf{i}) = c$, proving (ii) to hold. □

With the above lemma's we are now sufficiently equipped to construct the isomorphism (2.3).

Lemma 2.3.4. *For every basic open $U_{\mathbf{i},P}$ in $\mathbf{X}$ and every sheaf F of $\mathcal{E}$ we have the following isomorphism, natural in both.*

$$\mathcal{E}(P, F) = \mathcal{E}\left(\mathbf{e}_!\begin{pmatrix} U_{\mathbf{i},P} \\ \cup\!\!\mid\!\downarrow \\ \mathbf{X}\end{pmatrix}, F\right) \cong \mathbf{Etale}(\mathbf{X})\begin{pmatrix} U_{\mathbf{i},P} & & \mathbf{e}^*(F) \\ \cup\!\!\mid\!\downarrow & , & \pi\downarrow \\ \mathbf{X} & & \mathbf{X}\end{pmatrix}.$$

Proof. First consider a map $f : P = \mathbf{e}_!(U_{\mathbf{i},P}) \to F$, and recall that this defines a canonical section

$$\sigma_f : U_{\mathbf{i},P} \to U_{\mathbf{i},P,f} \subseteq \mathbf{e}^*(F), \qquad\qquad \langle p, \alpha\rangle \mapsto \langle p, \alpha, p^*(f)(\alpha(\mathbf{i}))\rangle\,.$$

To construct a mapping in the other direction, consider a section $\sigma : U_{\mathbf{i},P} \to \mathbf{e}^*(F)$. By Remark 1.2.1 we know that there exists an open cover $\{U_\iota := U_{\mathbf{j}_\iota, P_\iota} : \iota \in I\}$ of $U_{\mathbf{i},P}$ for some index set I such that $\sigma \restriction U_\iota$ is a canonical section, i.e. the inverse to $\pi : \mathbf{e}^*(F) \to \mathbf{X}$ on a small basic open. In a bit more detail, this gives us a family $\{f_\iota : Q_\iota \to E\}_{\iota\in I}$ such that σ restricted to U_ι equals

$$\sigma_{f_\iota} : U_\iota \to \mathbf{e}^*(F), \qquad\qquad \langle p, \alpha\rangle \mapsto \langle p, \alpha, p^*(f_\iota)(\alpha(\mathbf{j}_\iota))\rangle\,.$$

Due to Lemma 2.3.3 we now know that

$$\coprod_{\iota\in I} Q_\iota \xrightarrow{e := \coprod_{\iota\in I}\chi_\iota} C$$

is an epimorphism. As limits are computed point-wise in the category of sheaves on a site, we know that by evaluating on objects $C \in \mathcal{C}$ the resulting maps are surjections. This will allow us to send $x \in P(C)$ to $f_\iota(b)$ for the ι and $b \in Q_\iota(X)$ such that $\chi_{\iota C}(b) = x$, defining a map $P(C) \to F(C)$. Naturality of this map follows directly from naturality of the χ_ι and f_ι. We are left with the task of proving that the map is well-defined. We need to show that any two 'lifts' of elements in $P(C)$ get sent to the same element of F. This boils down to showing that for the pullback displayed on the left, the diagram on the right commutes.

$$\begin{array}{ccc} Q_\iota \times_C Q_\kappa & \xrightarrow{\pi_\iota} & Q_\iota \\ {\scriptstyle\pi_\kappa}\downarrow & & \downarrow{\scriptstyle\chi_\iota} \\ Q_\kappa & \xrightarrow[\chi_\kappa]{} & P \end{array} \qquad\qquad \begin{array}{ccc} Q_\iota \times_P Q_\kappa & \xrightarrow{\pi_\iota} & Q_\iota \\ {\scriptstyle\pi_\kappa}\downarrow & & \downarrow{\scriptstyle f_\iota} \\ Q_\kappa & \xrightarrow[f_\kappa]{} & F \end{array}$$

It suffices to check commutativity on stalks of small points p. Let $b \in p^*(Q_\iota \times_P Q_\kappa)$ be arbitrary and choose $\alpha_\lambda \in \mathbf{En}(p^*(G))$ such that $\alpha_\lambda(\mathbf{j}_\lambda) = \pi_\lambda(b)$ for $\lambda = \kappa, \iota$. As a consequence we know $\langle p, \alpha_\lambda\rangle \in U_\lambda = U_{\mathbf{j}_\lambda, Q_\lambda}$, from which we can derive

$$\big(f_\lambda\pi_\lambda\big)(b) = f_\lambda\big(\alpha_\lambda(\mathbf{j}_\lambda)\big) = \sigma\big(\langle p, \alpha_\lambda\rangle\big).$$

The desired follows from Corollary 2.3.1, which is applicable because $\chi_\iota\pi_\iota(b) = \chi_\kappa\pi_\kappa(b)$ entails $\alpha_\iota\pi_\iota(b) = \alpha_\kappa\pi_\kappa(b)$.

We still need to prove that these mappings are mutually inverse. Given a map $f : P \to F$ this is mapped to the canonical section $\sigma_f : U_{\mathbf{i},P} \to U_{\mathbf{i},P,f} \subseteq \mathbf{X}$, which clearly is a *canonical section*. We can thus cover $U_{\mathbf{i},P}$ with the cover consisting only of itself, associated to the map $f : P \to F$. The resulting map $P \to F$ is defined on components C as sending $x \in P(C)$ to $f_C(x)$, so this is exactly the same thing as the map f. On the other hand, a map $\sigma : U_{\mathbf{i},P} \to \mathbf{e}^*(F)$ yields a map $f : P \to F$ as described above. We know that σ_f is the canonical section associated to the basic open $U_{\mathbf{i},P}$. By Remark 1.2.1 we know that on a basic open each section is equal to a (unique) canonical section. This entails that $\sigma_f = \sigma$, proving the desired. □

Theorem 2.3.1. *The functor* $\mathbf{e}_! : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$ *is left adjoint to* $\mathbf{e}^* : \mathcal{E} \to \mathbf{Etale}(\mathbf{X})$*, that is to say, the following is an isomorphism natural in both* $p : E \to \mathbf{X}$ *and* F*.*

$$\mathcal{E}(P, F) = \mathcal{E}\left(\mathbf{e}_!\left(\begin{array}{c} E \\ p\big\downarrow \\ \mathbf{X}\end{array}\right), F\right) \cong \mathbf{Etale}(\mathbf{X})\left(\begin{array}{cc} E & \mathbf{e}^*(F) \\ p\big\downarrow, & \pi\big\downarrow \\ \mathbf{X} & \mathbf{X}\end{array}\right).$$

Proof. This is but an immediate corollary of the above lemma, for each $p : E \to \mathbf{X}$ is a colimit of basic opens. Thus a map f from p to $\mathbf{e}^*(F)$ is fully determined by the collection of maps $U_{\mathbf{i},P} \to (p : E \to \mathbf{X}) \to \mathbf{e}^*(F)$. To these we apply the previous lemma, and get the desired result. □

By the above theorem we know of a left-adjoint to the functor $\mathbf{e}^*$. The inverse-image functor π^* has a left adjoint as well, due to Corollary 2.3.2.

Corollary 2.3.2. *The map* $\epsilon^\pi : \pi_!\pi^* \to \mathrm{id}_{\mathbf{Set}}$ *is an isomorphism.*

Proof. This is a consequence of the fact that $\mathbf{En}(p^*(G))$ is locally connected, as proven in Corollary 2.1.1, and the more general Corollary 1.2.2. □

Let us now revisit the composite (2.2), as displayed below. To prove that the square (2.1) satisfies the Beck–Chevalley condition, we need to show that this composite is an isomorphism. Now, we know the center arrow in the composite below to be an isomorphism due to Lemma 2.3.1. By Corollary 2.3.2 it is clear that the right-most arrow also is an isomorphism, too. It thus is both necessary and sufficient to prove that the left-most arrow $\pi_!\mathbf{i}_p^*\eta^{\mathbf{e}}$ below is an isomorphism.

$$\pi_!\mathbf{i}_p^* \xrightarrow{\pi_!\mathbf{i}_p^*\eta^{\mathbf{e}}} \pi_!\mathbf{i}_p^*\mathbf{e}^*\mathbf{e}_! \xrightarrow{\pi_!\zeta\mathbf{e}_!} \pi_!\pi^*p^*\mathbf{e}_! \xrightarrow{\epsilon^\pi p^*\mathbf{e}_!} p^*\mathbf{e}_!$$

Lemma 2.3.5 (Beck–Chevalley Condition). *The composite* $\pi_!\mathbf{i}_p^* \to p^*\mathbf{e}_!$ *displayed above is an isomorphism.*

Proof. As mentioned above, we need to show that $\pi_!\mathbf{i}_p^*\eta^{\mathbf{e}}$ is an isomorphism. We need only to prove this on basic opens $U_{\mathbf{i},P}$. Let us first take a closer look at the map $\eta^{\mathbf{e}} : \mathrm{id}_{\mathbf{Etale}(\mathbf{X})} \to \mathbf{e}^*\mathbf{e}_!$, which on basic opens is defined by

$$\eta^{\mathbf{e}}_{U_{\mathbf{i},P}} : U_{\mathbf{i},P} \to \mathbf{e}^*\mathbf{e}_!(U_{\mathbf{i},P}) = \mathbf{e}^*(P), \qquad \langle p, \alpha\rangle \mapsto \langle p, \alpha, \alpha(\mathbf{i})\rangle .$$

The $\mathbf{i}_p^*$-image of $U_{\mathbf{i},P}$ is simply the pre-image of this space along the map $\mathbf{i}_p : \mathbf{En}(p^*(G)) \to \mathbf{X}$, which decomposes into the connected components $U_{\mathbf{i}\mapsto c}$ for all $c \in p^*(P)$. Note that on the other hand, by commutativity up to natural isomorphism ζ of (2.1), the connected components of $\mathbf{i}_p^*\mathbf{e}^*(P)$ are precisely those of $\pi^*p^*(P)$. Now as $\{x\}$ for $x \in p^*(P)$ is a connected component of $p^*(X)$ as a topological space, and as $\mathbf{En}(p^*(G))$ is connected by Corollary 2.1.1, we know that these connected components are precisely $\{x\} \times \mathbf{En}(p^*(G))$ for $x \in p^*(P)$.

To show that $\mathbf{i}_p^*(\eta_{U_{\mathbf{i},P}})$ is surjective on connected components, pick $x \in p^*(X)$ and consider the connected component $\zeta^{-1}(\{x\} \times \mathbf{En}(p^*(G)))$. Simply choose some α such that $\alpha(\mathbf{i}) = x$, and note that α is mapped to this connected component. Moreover, any other such map β lies in the open $U_{\mathbf{i}\mapsto x}$, which is a connected component. Hence $\pi_!\mathbf{i}_p^*(\eta_{U_{\mathbf{i},P}})$ is a bijection, proving the desired. □

We now know that $\mathbf{e} : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$ is an essential geometric morphism, and that the diagram (2.1) satisfies the Beck–Chevalley condition. Using this we can show that $\mathbf{e}$ is an embedding of toposes. That is to say, its inverse image functor is both full and faithful. As in Corollary 1.2.2, this amounts to proving that the counit $\epsilon^{\mathbf{e}} : \mathbf{e}_!\mathbf{e}^* \to \mathrm{id}_{\mathcal{E}}$ is an isomorphism. The desired is then immediate from Mac Lane (1997, Theorem 4.3.1). We prove this in the following theorem, by using that the topos $\mathcal{E}$ has a set of points which is jointly conservative and constructing an explicit inverse on the inverse-image of these points.

Theorem 2.3.2. *The geometric morphism* $\mathbf{e} : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$ *is an embedding.*

Proof. We need to show that $\mathbf{e}^*$ is full and faithful, which is equivalent to proving that the counit $\epsilon^{\mathbf{e}} : \mathbf{e}_!\mathbf{e}^* \to \mathrm{id}_{\mathcal{E}}$ of the adjunction between $\mathbf{e}_!$ and $\mathbf{e}^*$ is an isomorphism. As before, it suffices to check that the inverse-image of this map under each small point $p : \mathbf{Set} \to \mathcal{E}$ is an isomorphism. We thus wish to find an inverse to the natural transformation $p^*\epsilon^{\mathbf{e}}$. The inverse is constructed in the diagram below as the composite of the indicated arrows. This inverse is the composite of isomorphisms as given by Lemma 2.3.5 and Lemma 2.3.1 and Corollary 2.3.2. Consequently, it suffices to show that this isomorphism is a one-sides inverse to the map $p^*\epsilon^{\mathbf{e}}$.

We are done when we can show the diagram below to be commutative, so let us prove this. The triangles (1) and (2) commute clearly, whereas the triangle (3) is one of the triangle identities for the adjunction between $\mathbf{e}_!$ and $\mathbf{e}_*$. The square (4) is naturality of ζ and the square (5) is naturality of ϵ^π. Commutativity of (6) is obvious, proving the desired.

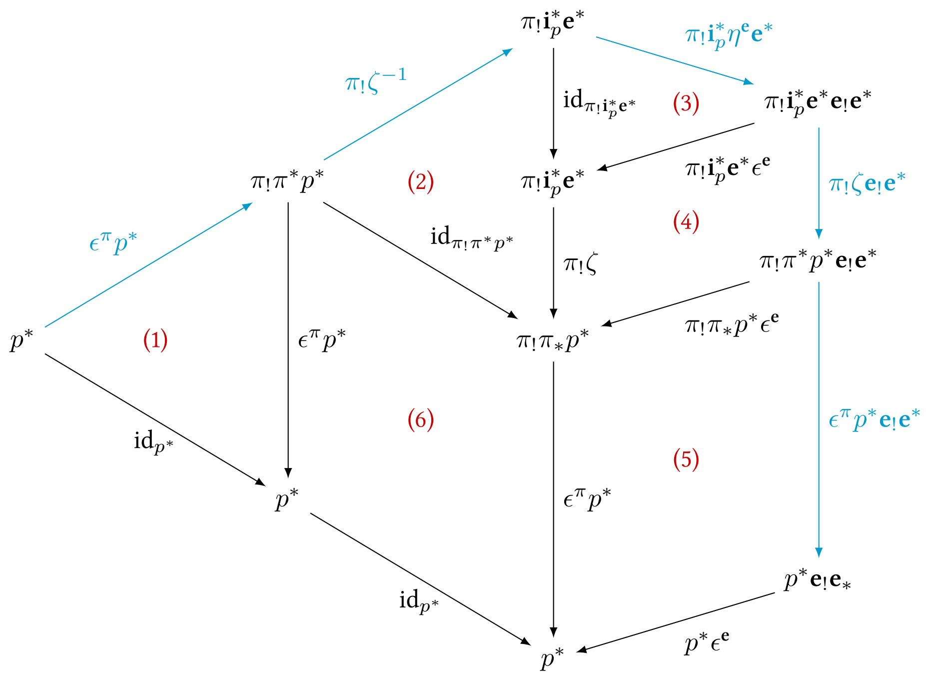

∎

Recall that $\mathbf{s} : \mathbf{Sh}(\mathbf{X}) \to \mathcal{E}$ is the geometric morphism as obtained through the equivalence of categories between $\mathbf{Etale}(\mathbf{X})$ and $\mathbf{Sh}(\mathbf{X})$ discussed in Subsection 1.2.1 and the geometric morphism $\mathbf{e} : \mathbf{Etale}(\mathbf{X}) \to \mathcal{E}$. In symbols, the inverse-image map $\mathbf{s}^*$ is simply the composite

$$\mathcal{E} \xrightarrow{\quad \mathbf{e} \quad} \mathbf{Etale}(\mathbf{X}) \xrightarrow{\quad \Gamma \quad} \mathbf{Sh}(\mathbf{X}) \ .$$

As a consequence, $\mathbf{s}$ is an essential geometric morphism because $\mathbf{e}$ is, as proven in Theorem 2.3.1. In particular, the inverse image functor $\mathbf{s}^*$ preserves all limits and colimits. Furthermore, it is full and faithful, which by Lemma 1.4.1 in particular entails that it is conservative, i.e. reflects isomorphisms. We thus created a spatial topos in which to embed our topos $\mathcal{E}$. The following theorem sums up the achievements of this chapter.

Theorem 2.3.3 (Spatial Cover). *Let* $\mathcal{E}$ *be a Grothendieck topos such that there exists a set of points of* $\mathcal{E}$*, of which the totality of inverse-image functors is jointly conservative. There exists a topological space* $\mathbf{X}$ *such that there is an essential embedding* $\mathbf{s} : \mathbf{Sh}(X) \to \mathcal{E}$*.*

Chapter 3

λ-Calculus

The λ-calculus was introduced by Church (1932), and intended as a foundation for logic. His original system was found to be inconsistent, as proven amongst others by Curry (1942). The subsystem concerned only with λ-terms and their conversions however turned out to be quite interesting, and this is what we today call λ-calculus. We will cover a certain typed variant of this system, our presentation based heavily on that of Lambek and Scott (1988).

We introduce λ-calculus via a slight detour, focussing first on *cartesian closed categories*. This in contrast with the quite common approach via computation. We wish to emphasize the relation between λ-terms and arrows in a cartesian closed category constructed in a 'canonical' fashion both early and often. Cartesian closed categories of particular interest to us are the category of sheaves on a topological space, or more generally, the category of sheaves on a site. Having covered these two types of categories extensively in Chapter 1, with as most prominent example of the former type the category $\mathbf{Sh}(\mathbf{X})$ of Chapter 2, we trust the reader to be quite familiar with them.

In order to neatly relate the λ-calculus to cartesian closed categories we follow a similar line of reason as Lambek and Scott (1988). To this end, let us revisit our foundations and look at the very definition of a category. The prevalent definition of categories, as given by Mac Lane (1997), makes use of three types of structures. First, it starts with the notion of a *graph*. Then follows what Lambek and Scott (1988) call a *deductive system*; a graph together with the ensured existence of an *identity arrow* id_X to each vertex X, and the operation of *composition*.[1] This allows one to *construct* arrows using the operations and arrows from the graph. One can formalize by means of the following two inference rules. The left-hand rule is reminiscent of the Cut-rule from sequent calculus. The right-hand rule has no assumptions (so it is an *axiom*) and bears obvious resemblance to the identity-axiom of zeroth order logic.

$$\frac{X \xrightarrow{f} Y \qquad Y \xrightarrow{g} Z}{X \xrightarrow{gf} Z}\ (\textsc{cut}) \qquad\qquad \frac{}{X \xrightarrow{\mathrm{id}_X} X}\ (\textsc{id})$$

Finally, a *category* is a deductive system subject to the laws as displayed below.

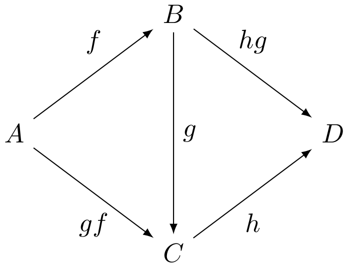

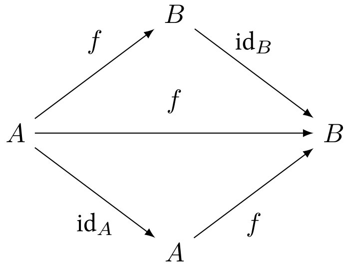

[1] Other authors use other names, Mac Lane (1997) speaks of a *metagraph* and *metacategory* respectively. We will use the term 'deductive system', as it more cleary conveys the relation to inference systems we wish to emphasize.

We introduce a cartesian closed category in Section 3.1 below by means of a certain deductive system. We then reformulate this deductive system into a more familiar form, resembling closely the inference system of zeroth order logic with positive connectives.

Motivated by giving names to proofs built using the inference system of zeroth order logic, we introduce the λ-calculus. The usual equalities are introduced, and motivated by equalities present in generic cartesian closed categories. In other texts these equations are motivated by a sense of computation, as for instance is the case in Barendregt and Barendsen (1988) and Sørensen and Urzyczyn (2006).[2] The former is a good reference, and provides more intuition for the computationally minded. The latter covers in great detail the so-called Curry–Howard correspondence,[3] which provides a correspondence between the typed λ-calculus and mathematical logic. In Section 3.2 we cover the simply typed λ-calculus, outfitted with the connectives for conjunction, truth and implication. Our goal here is the introduction of the definition of λ-calculus, as used by Awodey (2000) and originally given by Lambek and Scott (1988). We use the λ-calculus covered in this section as a motivation for this more general definition.

The λ-calculus provides an equational theory, equating terms of a given type. One might wonder whether there exists some model for this theory. In Section 3.3 we first introduce the type of a model we wish to consider, which assigns to a term an arrow in a given cartesian closed category. We are particularly interested in models that are *sound*, *complete* and *functionally complete*. The former means that the model is such that two terms are equated in the model when they are equal in the λ-calculus, and a model is *complete* when the converse holds. A model is said to be *functionally complete* when an arrow in the model from the terminal object to the interpretation of a type arises as the interpretation of a unique closed term in the λ-calculus, up to its equality. Not yet having provided rigid definitions of all terms in this paragraph, these statements remain quite vague until they come up again in our main Theorem 3.3.1. This theorem was originally proven by Awodey (2000) and we use this paper as an inspiration for many of the formulations in this chapter.

SECTION 3.1 Cartesian Closed Categories

Given a diagram $D : \mathcal{D} \to \mathcal{C}$ the limit $\lim D$ of this diagram is defined only up to isomorphism. As in the category of sets one can give a 'canonical' description for pullbacks and terminal objects, there is a bonafide functor $\lim : \mathcal{C}^{\mathcal{D}} \to \mathcal{C}$ for any finite category $\mathcal{D}$, defined by taking these predetermined choices. Now as limits and colimits in the category of presheaves over a category are taken point-wise, such a choice extends to this setting as well. In this section we define precisely what it means for a category to have chosen products, and chosen *exponents* as well. The exponent $A \Rightarrow B$ of B to A is the categorical generalization of the *relative complement* in partially ordered sets with greatest lower bounds, and as such yields a right adjoint $A \Rightarrow (-)$ to the product $(-) \wedge A$. In the following definition, as promised, we make good use of the language of deductive systems. We later on will show that a category which adheres to the definition below in fact has finite products and a (chosen) right adjoint to the product, which is the more common content of the term 'cartesian closed'.

Definition 3.1.1 (Cartesian Closed Category). A category $\mathcal{C}$ is said to be a cartesian closed category (CCC) when it is endowed with a chosen object $\mathbf{T}$ and operations $(-) \wedge (-)$ and $(-) \Rightarrow (-)$ which send a pair of objects to a new object, closed under the rules of Figure 3.1 such that the following diagrams commute, along with the diagrams for a category given above.

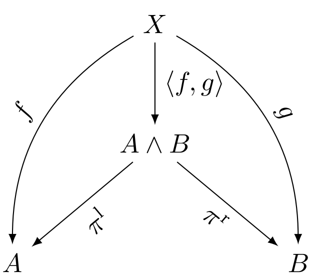

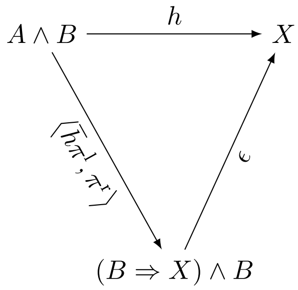

[2] We heavily rely on the latter to provide accurate historical references, such as given in this very introduction.

[3] Others call this the Curry–Howard–de Bruijn correspondence, emphasizing that de Bruijn not only recognizes the same ideas around the time Howard's original paper was written, but also applied them in his theorem prover Automath described in de Bruijn (1970).

Axioms

$$\frac{}{X \xrightarrow{\mathrm{id}_X} X} \qquad \frac{}{A \wedge B \xrightarrow{\pi^{\mathrm{l}}} A} \qquad \frac{}{X \xrightarrow{!} \mathbf{T}} \qquad \frac{}{A \wedge B \xrightarrow{\pi^{\mathrm{r}}} B}$$

$$\frac{}{(Y \Rightarrow X) \wedge Y \xrightarrow{\epsilon} X}$$

Inference Rules

$$\frac{X \xrightarrow{f} A \qquad X \xrightarrow{g} B}{X \xrightarrow{\langle f,g \rangle} A \wedge B} \qquad \frac{X \wedge Y \xrightarrow{f} Z}{X \xrightarrow{\overline{f}} Y \Rightarrow Z}$$

Cut Rule

$$\frac{X \xrightarrow{f} Y \qquad Y \xrightarrow{g} Z}{X \xrightarrow{gf} Z} \ (\textsc{cut})$$

Figure 3.1: Economic Rules of Inference Defining a CCC

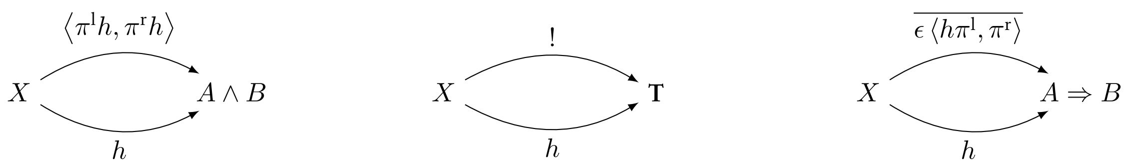

In the above definition we are given a mapping on objects and on arrows, the former denoted by $(-) \wedge (-)$ and the latter by $\langle (-), (-) \rangle$. These can be combined to form a mapping

$$(-) \wedge (-) : \mathcal{C} \times \mathcal{C} \to \mathcal{C}, \quad X \xrightarrow{f} A, Y \xrightarrow{g} B \mapsto X \wedge Y \xrightarrow{\langle f\pi^{\mathrm{l}}, g\pi^{\mathrm{r}} \rangle} A \wedge B.$$

Through some rudimentary computations one can show this to be *functorial*, making it an actual functor. The diagrams on the left in Definition 3.1.1 now express naturality of the collections of maps π^{l} and π^{r}. Likewise, we can form a functor

$$(-) \Rightarrow (-) : \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}, \quad X \xrightarrow{f} A, Y \xrightarrow{g} B \mapsto X \Rightarrow Y \xrightarrow{\overline{g\epsilon \langle \mathrm{id}_{X \Rightarrow Y}, f \rangle}} A \Rightarrow B.$$

With this knowledge, the bottom right hand diagram in the definition expresses naturality of $\epsilon : (-) \Rightarrow Y \wedge Y \to \mathrm{id}_{\mathcal{C}}$. We also have a natural transformation $\overline{(-)} : \mathcal{C}\left((-) \wedge Y,\, Z\right) \to \mathcal{C}\left((-),\, Y \Rightarrow Z\right)$, naturality of which is easy to see when using the top-right diagram and the distributive property $\langle f, g \rangle\, h = \langle fh, gh \rangle$ which follows from the left-hand diagrams. It readily follows that $\overline{(-)}$ in fact is an isomorphism, where the inverse mapping sends $h : X \to A \Rightarrow B$ to the map $\epsilon \left\langle h\pi^{\mathrm{l}}, \pi^{\mathrm{r}} \right\rangle$. The mappings indeed are mutually inverse, as follows from the right hand diagrams.

Due to the existence of the natural isomorphism $\overline{(-)}$ we know $(-) \wedge Y : \mathcal{C} \to \mathcal{C}$ to be left adjoint to $Y \Rightarrow (-)$. The natural transformations π^{l} and π^{r}, together with the top left diagram prove that $X \wedge Y$ is the product of X and Y. Finally, one readily sees $! : \mathrm{id}_{\mathcal{C}} \to \Delta\mathbf{T}$ to be natural due to the bottom center diagram, and from these two facts it follows that $\mathbf{T}$ is terminal. We have thus derived that any CCC is a CCC in the more popular sense, for it has all finite products and a right-adjoint to the functor $(-) \wedge C$ for all C.

Using the data ensured by the definition, we can form the deductive system of Figure 3.2. Note that in this system we write η to mean the unit of the adjunction given above, which is thus given on components by $\overline{\mathrm{id}_{(-)\wedge Y}}$. We automatically get the triangle equalities relating η and ϵ. The commutative diagrams we inferred above can be

Axioms

$$\frac{}{X \xrightarrow{\mathrm{id}_X} X}\ (\textsc{id}) \qquad\qquad \frac{}{X \xrightarrow{!} \mathbf{T}}\ (\mathbf{T}\mathrm{I})$$

Inference Rules

$$\frac{X \xrightarrow{f} A \quad X \xrightarrow{g} B}{X \xrightarrow{\langle f,g\rangle} A \wedge B}\ (\wedge \mathrm{I}) \qquad \frac{X \xrightarrow{h} A \wedge B}{X \xrightarrow{\pi^{\mathrm{l}} h} A}\ (\wedge \mathrm{E}_{\mathrm{L}}) \qquad \frac{X \xrightarrow{h} A \wedge B}{X \xrightarrow{\pi^{\mathrm{r}} h} B}\ (\wedge \mathrm{E}_{\mathrm{R}})$$

$$\frac{X \wedge Y \xrightarrow{f} Z}{X \xrightarrow{(Y \Rightarrow f)\eta} Y \Rightarrow Z}\ (\Rightarrow \mathrm{I}) \qquad \frac{X \xrightarrow{f} A \Rightarrow B \quad Y \xrightarrow{g} A}{X \wedge Y \xrightarrow{\epsilon(f \wedge g)} B}\ (\Rightarrow \mathrm{E})$$

Cut Rule

$$\frac{X \xrightarrow{f} Y \quad Y \xrightarrow{g} Z}{X \xrightarrow{gf} Z}\ (\textsc{cut})$$

Figure 3.2: Rules of Inference in a CCC

interpreted in terms of detour elimination. Naturality of the projection map π^{l} translates to saying that the two inferences in (3.1) are equal.

$$\frac{\dfrac{X \xrightarrow{f} A \quad X \xrightarrow{g} B}{X \xrightarrow{\langle f,g\rangle} A \wedge B}\ (\wedge \mathrm{I})}{X \xrightarrow{\pi^{\mathrm{l}}\langle f,g\rangle} A}\ (\wedge \mathrm{E}_{\mathrm{L}}) \qquad X \xrightarrow{f} A \qquad \frac{\dfrac{X \xrightarrow{g} A \quad X \xrightarrow{f} B}{X \xrightarrow{\langle g,f\rangle} A \wedge B}\ (\wedge \mathrm{I})}{X \xrightarrow{\pi^{\mathrm{r}}\langle f,g\rangle} A}\ (\wedge \mathrm{E}_{\mathrm{R}}) \tag{3.1}$$

The inferences of (3.2) are equal too, as follows form the commutative diagram displayed beneath them. The square is naturality of η, and the triangle is one of the triangle equalities. Now think for a moment about the category **Set**. This category is a CCC according to our definition, interpreting the operations in the usual sense. One easily sees that ϵ simply is function evaluation. Now take some map $f : X \to A \Rightarrow B$ in **Set**, and consider the composite defined by the derivation below, though some menial computations one can rewrite the expression obtained from this deduction to

$$\Big(x \in X \mapsto \big(a \in A \mapsto f(x)(a)\big)\Big).$$

As **Set** is a CCC we know that the proof below applies to this equation too, so this function and f must be equal. We view functions extensionally in **Set**, so it makes perfect sense that they are in fact equal. The derivation on the left can be seen as a proof with a detour, first eliminating an implication, only to re-introduce it by means of an assumption. On an intuitive level, this ought not to change the meaning of the proof, 'morally' justifying that these derivations are in fact equal. So in short, ($\Rightarrow$E) follows by ($\Rightarrow$I) does not alter the original proof.

$$\frac{\dfrac{X \xrightarrow{f} A \Rightarrow B \quad A \xrightarrow{\mathrm{id}_A} A}{X \wedge A \xrightarrow{\epsilon(f \wedge \mathrm{id}_A)} B}\ (\Rightarrow \mathrm{E})}{X \xrightarrow{(A \Rightarrow \epsilon(f \wedge \mathrm{id}_A))\eta} A \Rightarrow B}\ (\Rightarrow \mathrm{I}) \qquad X \xrightarrow{f} A \Rightarrow B \tag{3.2}$$

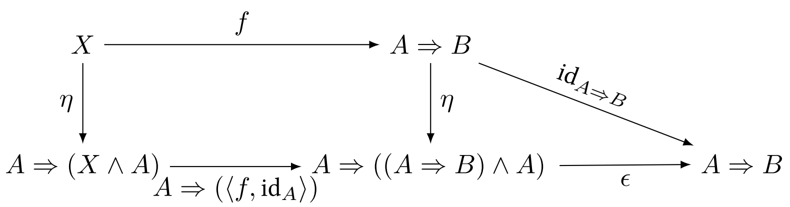

Finally, we note that the following two proofs of (3.3) are equal. Intuitively (insofar the following might feel intuitive), the derivation g is 'plugged into' the derivation f to change the domain from $X \wedge A$ to $X \wedge Y$. This equation is easier to understand when one takes g to simply be the identity map. Then it simply states that performing (⇒E) after (⇒I) does nothing to the original proof.

$$\dfrac{\dfrac{X \wedge A \xrightarrow{f} B}{X \xrightarrow{(A\Rightarrow f)\eta} A \Rightarrow B}\,(\Rightarrow\mathrm{I}) \qquad Y \xrightarrow{g} A}{X \wedge Y \xrightarrow{\epsilon((A\Rightarrow f)\eta)\wedge g} B}\,(\Rightarrow\mathrm{E}) \qquad\qquad X \wedge Y \xrightarrow{f\langle \mathrm{id}_X, g\rangle} B \tag{3.3}$$

The above proofs are equal due to the commutative diagram displayed below, which hinges mostly on one of the triangle equalities.

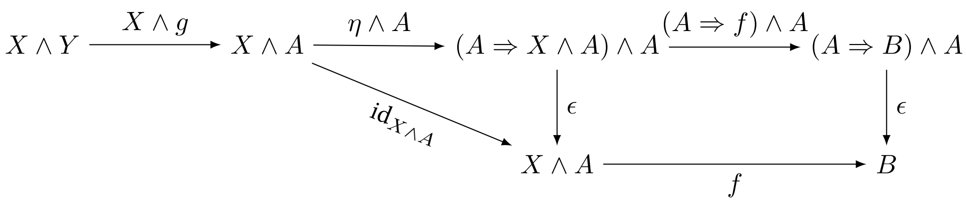

It is clear that a category closed under the system of Figure 3.2, endowed with the functors $(-)\wedge(-)$ and $(-) \Rightarrow (-)$ and subject to the mentioned naturality requirements and triangle identities, also satisfies the definition of a ccc as given earlier. Indeed, the projections can be obtained from (∧El) and (∧Er) applied to the identity map and the evaluation map is simply (⇒E) applied to the two sensible identity maps.

Let us now consider some basic examples of cartesian closed categories. The most clearcut example of a ccc is **Set**, as mentioned above. The maps η and ϵ can easily be described in an explicit fashion

$$\epsilon : B^A \times A \to B,\ \langle f : A \to B, a \in A\rangle \mapsto f(a), \quad \eta : B \to (B \times A)^A,\ b : B \mapsto \big(a \mapsto \langle b, a\rangle\big),$$

where ϵ is simply function evaluation and η the 'constant function-function'. The adjunction between $(-) \wedge C$ and $C \Rightarrow (-)$ gives rise to the natural isomorphism below.

$$f : X \times C \to Y \mapsto \Big(x \in X \mapsto (c \in C \mapsto f(\langle x, c\rangle))\Big)$$

In reminiscence of the method used to treat functions with multiple arguments in the Funktionenkalkül of Schönfinkel (1924) this could be called 'Schönfinkelisation', instead Christopher Strachey coined the term 'Currying' in 1974 and it became the prevalent name for this phenomenon.

To consider a situation completely different from **Set**, let $\mathcal{C}$ be a partially ordered set $\langle P, \leq\rangle$ considered as a category. We now know that $x \wedge y$ for $x, y \in P$ is the meet or greatest lower bound of x and y. In a generic poset one can define the complement of x relative to y as the greatest r such that $x \wedge r \leq y$ holds. It can be easily proven that s lies below the complement of x relative to y precisely when the meet of s and x lies below y. From this fact it is clear that the relative complement of x to y is given by $x \Rightarrow y$, and thus exists for each pair of objects. Note that this category satisfies all 'positive' requirements imposed on a *Heyting algebra*. That is to say, if we ignore any reference to falsity and disjunction in the definition of Heyting algebra we are left with a poset which is a ccc.

The following proposition provides us with an important ccc, which we will use later on in this chapter.

Proposition 3.1.1. *The category* **PSh**$(\mathcal{C})$ *is cartesian closed for any small category* $\mathcal{C}$.

Proof. We prove that **PSh**$(\mathcal{C})$ adheres to Definition 3.1.1. To this end, let P and Q be presheaves on $\mathcal{C}$. We define the presheaf $P \wedge Q$ point-wise, in the sense that $(P \wedge Q)(C) = P(C) \wedge Q(C)$, of course making use of the cartesian closed structure on **Set**. The projection maps π^{l} and π^{r} are defined in the sensible manner. We now set $\mathbf{T} := \mathbf{yT} = \mathcal{C}((-), \mathbf{T})$ and the maps $! : P \to \mathbf{T}$ are defined point-wise in the only way possible.

The presheaf $P \Rightarrow Q$ is defined as $\mathbf{PSh}(\mathcal{C})\,(\mathbf{y}(-) \wedge P,\, Q)$, which sends the object C of $\mathcal{C}$ to the set of natural transformations between $\mathcal{C}\,((-),\, C) \wedge P$ and Q. We now define evaluation on components as

$$\epsilon : \mathbf{PSh}(\mathcal{C})\,(\mathbf{y}(-) \wedge P,\, Q) \wedge P \to Q, \quad \big(\langle \alpha : \mathcal{C}\,((-),\, C) \wedge P \to Q, x \in P(C)\rangle \mapsto \alpha_C(\mathrm{id}_C, x)\big)_{C \in \mathcal{C}},$$

which is clearly natural. Finally, given an arrow $\phi : P \wedge R \to Q$ we define the arrow $\overline{\phi} : P \to R \Rightarrow Q$ on components $Y \in \mathcal{C}$ as below

$$p \in P(Y) \mapsto \Big(\langle f : X \to Y, r \in R(X)\rangle \in \mathcal{C}\,(X,\, Y) \wedge R(X) \mapsto \phi_X\big(P(f)(p), r\big) \in Q(X)\Big)_{X \in \mathcal{C}}.$$

Again, this mapping can easily be seen to be natural.

To finish the proof, we need to show commutativity of the five diagrams in the definition. Three of these, regarding products, ought to be clear from the point-wise definition. We now check that for a map $h : P \wedge Q \to R$ it is in fact the case that $\epsilon \left\langle \overline{h}\pi^{\mathrm{l}}, \pi^{\mathrm{r}} \right\rangle = h$. It suffices to check this on components, to this end, fix an object $C \in \mathcal{C}$. On components this is but an equation between two functions in **Set**, so simply take $p \in P(C)$ and $q \in Q(C)$ to compute

$$\begin{aligned}
\left(\epsilon \left\langle \overline{h}\pi^{\mathrm{l}}, \pi^{\mathrm{r}} \right\rangle\right)(p, q) &= \epsilon\big(\overline{h}(p), q\big) \\
&= \epsilon\Big(\Big(x \mapsto \big(\langle f : X \to Y, r \in R(X)\rangle \mapsto h_Y(P(f)(x), r)\big)_{X \in \mathcal{C}}\Big)(p), q\Big) \\
&= \epsilon\Big(\big(\langle f : X \to Y, z \in P(X)\rangle \mapsto h_Y(P(f)(p), r)\big)_{X \in \mathcal{C}}, q\Big) \\
&= h_C\Big(P\big(\mathrm{id}_Y\big)(p), q\big) = h_C(p, q)
\end{aligned}$$

proving that the desired indeed holds. The other diagram can be shown to commute in a similar fashion. □

We wish to form a category of CCC's, where the objects are CCC's and arrows are defined in some sensible manner. It makes sense that the arrows here are functors, preserving the structure in some meaningful way. Due to our strict definition of a CCC we have an explicit choice of products, so it stands to reason that the arrows should respect this assignment. Accordingly, all related structure ought to be preserved as well.

Recall from the previous chapter that a functor $F : \mathcal{C} \to \mathcal{E}$ *preserves* the limit of a certain diagram $D : \mathcal{D} \to \mathcal{C}$ precisely when the canonical arrow $F(\lim D) \to \lim FD$ is an isomorphism. In the case of binary products, this comes down to the demand that the canonical morphism

$$\big\langle F(\pi^{\mathrm{l}}), F(\pi^{\mathrm{r}})\big\rangle : F(A \wedge B) \to F(A) \wedge F(B)$$

is an isomorphism. We say that F *preserves binary products* precisely when this canonical map is an isomorphism. Given maps $f : X \to A$ and $g : X \to B$ there is also a clear relation between the maps $F(\langle f, g\rangle)$ and $\langle F(f), F(g)\rangle$. Indeed, we can compute that

$$F(f) = F(\pi^{\mathrm{l}})F(\langle f, g\rangle) = \pi^{\mathrm{l}} \left\langle F(\pi^{\mathrm{l}}), F(\pi^{\mathrm{r}})\right\rangle F(\langle f, g\rangle),$$

which proves that $F(\langle f, g\rangle)$ equals, up to a canonical structural change, the arrow $\langle F(f), F(g)\rangle$.

In any case, there is a canonical map $! : F(\mathbf{T}) \to \mathbf{T}$. Like above, we say that F preserves the nullary product, or rather, F *preserves the terminal object* whenever this map is an isomorphism.

It requires a bit more effort to see how one could canonically preserve exponents. We wish to relate $F(X \Rightarrow Y)$ and $F(X) \Rightarrow F(X)$, and seek an intrinsic bond. There is of course the natural transformation $\epsilon : (X \Rightarrow Y) \wedge X \to Y$, and assuming that F preserves products as described above, we can form the composite below.

$$F\,(X \Rightarrow Y) \wedge F(X) \xrightarrow{\left\langle F(\pi^{\mathrm{l}}), F(\pi^{\mathrm{r}})\right\rangle^{-1}} F\,((X \Rightarrow Y) \wedge X) \xrightarrow{\quad F(\epsilon) \quad} F(Y)$$

One can take the transpose of this composite, and obtain an arrow $F\,(X \Rightarrow Y) \to F(X) \Rightarrow F(Y)$. We say that F *preserves exponents* when this composite is an isomorphism. Given this arrow, there is a relation between $F(\epsilon)$

and ϵ, after suitable structural changes. Consider the following diagram, where $F(\epsilon)$ is equal to the composite map from the top right to the bottom right, and the indicated arrow is ϵ after structural changes. The diagram commutes, because the downwardly sloped triangle is one of the triangle identities, and all squares are naturality squares. A similar diagram can be constructed for $\overline{(-)}$ of $F(f)$ for any $f : X \wedge Y \to Z$.

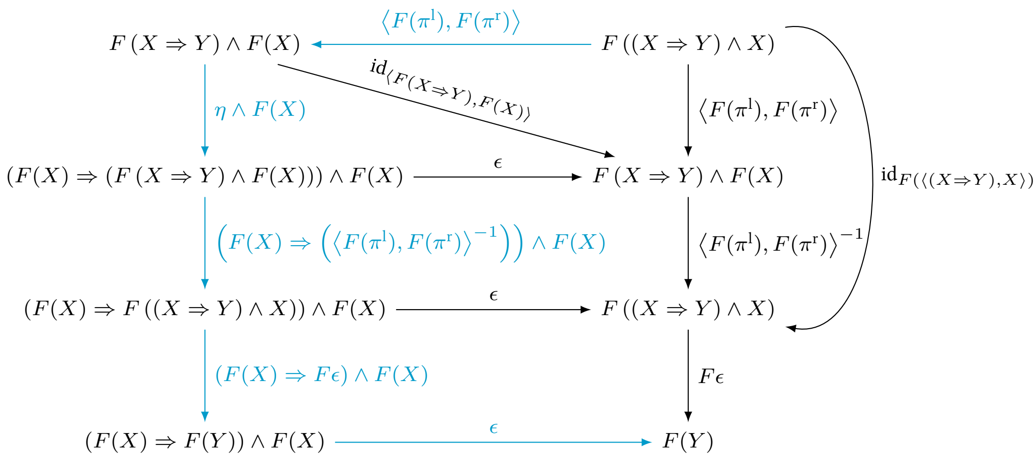

Let us now combine these notions of preservations into the definition of a structure-preserving map between cartesian closed categories.

Definition 3.1.2 (Cartesian Closed Functor). A *cartesian closed functor* $F : \mathcal{C} \to \mathcal{D}$ between CCC's $\mathcal{C}$ and $\mathcal{D}$ is a functor from $\mathcal{C}$ to $\mathcal{D}$ which preserves the cartesian closed structure. More precisely, the canonical maps below are isomorphisms.

$$F(\mathbf{T}) \cong \mathbf{T}, \quad F(A \wedge B) \cong F(A) \wedge F(B), \quad F(X \Rightarrow Y) \cong F(X) \Rightarrow F(Y),$$

When the canonical isomorphisms are identity morphisms, we call F a *strict cartesian closed functor*.

We are now ready to define the *category of cartesian closed categories* **Cart** as mentioned above. It can be proven through tedious computations that the composition of functors satisfying the constraints of Definition 3.1.2 satisfy them as well. Now as the identity functor is a cartesian closed functor too, **Cart** is a bonafide category. As a first example of a cartesian closed functor, we consider the Yoneda-embedding.

Proposition 3.1.2. *The Yoneda-embedding* $\mathbf{y} : \mathcal{C} \to \mathbf{PSh}(\mathcal{C})$ *is a cartesian closed functor for any small category* $\mathcal{C}$.

Proof. From the cartesian closed structure on $\mathbf{PSh}(\mathcal{C})$ as mentioned in Proposition 3.1.1 we see that $\mathbf{y}$ clearly preserves (binary and nullary) products. Consider the following two composites.

$$\mathbf{y}(X \Rightarrow Y) \xrightarrow{\eta} \mathbf{y}X \Rightarrow (\mathbf{y}(X \Rightarrow Y) \wedge \mathbf{y}X) \xrightarrow{\cong} \mathbf{y}X \Rightarrow \mathbf{y}((X \Rightarrow Y) \wedge X) \xrightarrow{\mathbf{y}X \Rightarrow \epsilon} \mathbf{y}X \Rightarrow \mathbf{y}Y$$

$$\mathcal{C}((-), X \Rightarrow Y) \xrightarrow{\cong} \mathcal{C}((-) \wedge X, Y) \xrightarrow{\cong} \mathbf{PSh}(\mathcal{C})(\mathbf{y}((-) \times X), \mathbf{y}Y) \xrightarrow{\cong} \mathbf{PSh}(\mathcal{C})(\mathbf{y}(-) \times \mathbf{y}X, \mathbf{y}Y)$$

The top composite is the map which we need to show to be an isomorphism in order for $\mathbf{y}$ to preserve exponents. The bottom composite consists only of isomorphisms, and it is easy to see that these composites have the same domain and codomain. It is now but an easy exercise to compute that both are simply the mapping given below, finishing the proof.

$$\alpha : C \to X \Rightarrow Y \mapsto \left(\left\langle D \xrightarrow{f} C, D \xrightarrow{g} X\right\rangle \mapsto D \xrightarrow{\langle \alpha f, g\rangle} (X \Rightarrow Y) \wedge X \xrightarrow{\epsilon} Y\right)_{D \in \mathcal{C}} \qquad \square$$

The spatial cover of a topos $\mathcal{E}$ with a set of enough points, introduced in Chapter 2, gives a geometric morphism $\mathbf{s} : \mathbf{Sh}(\mathbf{X}) \to \mathcal{E}$. The inverse-image part of this geometric morphism is an example of a cartesian closed functor, as we will prove in Theorem 3.1.1. In order to reach this result, we first prove the following more general lemma. Note that in the remainder of this section we are not as precise as we perhaps should be. Instead of proving explicitly that the specifically designated morphism is an isomorphism, we prove the existence of an isomorphism, which comes to us in such a natural manner that it really should be this particular morphism.

Proposition 3.1.3. *Let $f : \mathcal{C} \to \mathcal{D}$ be an essential geometric morphism, with full and faithful inverse-image part. If the composite*

$$f_! \left(C \wedge f^* D\right) \xrightarrow{\left\langle f_!(\pi^{\mathrm{l}}),\, f_!(\pi^{\mathrm{r}})\right\rangle} f_! C \wedge f_! f^* D \xrightarrow[\mathrm{id}_{f_! C} \wedge \epsilon^f]{} f_! C \wedge D \tag{3.4}$$

is an isomorphism for all objects C of $\mathcal{C}$ and D of $\mathcal{D}$, then $f^* : \mathcal{E} \to \mathcal{C}$ is a cartesian closed functor.

Proof. We know by assumption that f^* preserves finite limits, hence binary products and the terminal object in particular are preserved. Consider the following chain of natural isomorphisms.

$$\begin{aligned} \mathcal{C}\left(F,\, f^*\left(A \Rightarrow B\right)\right) &\cong \mathcal{D}\left(f_! F,\, A \Rightarrow B\right) \cong \mathcal{D}\left(f_! F \wedge A,\, B\right) \\ &\cong \mathcal{D}\left(f_! F \wedge f_! f_* A,\, B\right) \cong \mathcal{D}\left(f_!\left(F \wedge A\right),\, B\right) \\ &\cong \mathcal{C}\left(F \wedge A,\, f^* B\right) \cong \mathcal{C}\left(F,\, f^* A \Rightarrow f^* B\right) \end{aligned}$$

The isomorphisms on the first and last row exist because $f_!$ is left adjoint to f_* and both $\mathcal{C}$ and $\mathcal{D}$ are both cartesian closed. The first isomorphism on the second line exists because f_* is full and faithful, so $\epsilon : f_! f_* \to \mathrm{id}_{\mathcal{E}}$ must be an isomorphism. The second isomorphism is due to the assumed isomorphism (3.4). When we take the image of $\mathrm{id}_{f^* A \Rightarrow f^* B}$ we get an inverse to the canonical map $f^*\left(A \Rightarrow B\right) \to f^* A \Rightarrow f^* B$, which can be verified through tedious calculations. □

Corollary 3.1.1. *Let X be a locally connected and connected topological space. There is an essential geomtric morphism $\pi : \mathbf{Sh}(X) \to \mathbf{Set}$ and the composite (3.4) is an isomorphism. Moreover, the inverse-image part π^* is cartesian closed.*

Proof. This is but an immediate consequence of the above Proposition 3.1.3 and Lemma 1.2.12 □

Theorem 3.1.1. *Let $\mathcal{E}$ be a Grothendieck topos with a set of enough points. There exists a topological space $\mathbf{X}$ and an essential geometric morphism $\mathbf{s} : \mathbf{Sh}(\mathbf{X}) \to \mathcal{E}$ such that the inverse-image functor $\mathbf{s}^* : \mathcal{E} \to \mathbf{Sh}(\mathbf{X})$ is a cartesian closed functor.*

Proof. By Theorem 2.3.3 we get the space $\mathbf{X}$ and the essential geometric morphism $\mathbf{s} : \mathbf{Sh}(\mathbf{X}) \to \mathcal{E}$. Recall the space $\mathbf{En}(p^*(G))$ constructed in Section 2.1, which is locally connected and connected, thus by Corollary 3.1.1 endowed with an essential geometric morphsim $\pi : \mathbf{Sh}(X) \to \mathbf{Set}$ whose inverse-image part preserves exponents.

The diagram (2.1) commutes up to natural isomorphism by Lemma 2.3.1 and satisfies the Beck-Chevalley condition by Lemma 2.3.5. This respectively gives rise to isomorphisms $\mathbf{i}_p^* \mathbf{e}^* \cong \pi^* p^*$ and $\pi_! \mathbf{i}_p^* \cong p^* \mathbf{s}_!$. Also remember that $\mathbf{i}_p^*$ is left-exact, and so is p^* for any point $p := \langle p, \alpha \rangle$ of $\mathbf{X}$.

Note that we are done when we can prove that the composite of (3.4) is an isomorphism. As the topos $\mathbf{Sh}(\mathbf{X})$ has enough points by Corollary 1.2.3, it suffices to check that its p^*-image is an isomorphism. We now sloppily indicate why this is true, taking little care to argue why the following composite is in fact the p^*-image of (3.4). One ought to be able to prove this through computations involving primarily naturality and the triangle equalities.

$$\begin{aligned} p^* \mathbf{s}_!(E \times \mathbf{s}^* F) &\cong \pi_! \mathbf{i}_p^*(E \times \mathbf{s}^* F) \cong \pi_!(\mathbf{i}_p^* E \times \mathbf{i}_p^* \mathbf{s}^* F)) \cong \pi_!(\mathbf{i}_p^* E \times \pi^* p^* F)) \cong \pi_! \mathbf{i}_p^* E \times p^* F \\ &\cong p^* \mathbf{s}_! E \times p^* F \cong p^*(\mathbf{s}_! E \times F) \square \end{aligned}$$

SECTION 3.2 The λ-Calculus

We now will define propositional positive intuitionistic logic by means of a natural deduction system. To this end, we need to introduce some basic notions from proof theory.

The logic we work with here is only concerned with positive statements, that is, statements involving conjunction, implication and truth. We formally specify the set $\mathcal{L}_{\text{PIP}}$ of such positive formulae by means of the BNF below.[4]

$$\mathcal{L}_{\text{PIP}} ::= \mathbf{T} \mid \mathcal{L}_{\text{PIP}} \wedge \mathcal{L}_{\text{PIP}} \mid \mathcal{L}_{\text{PIP}} \Rightarrow \mathcal{L}_{\text{PIP}} \tag{3.5}$$

An element of the above defined set $\mathcal{L}_{\text{PIP}}$ is called a *formula*. While constructing a proof, one keeps in mind a (finite) number of assumptions, which form the 'context' of reasoning. We need to formally define what we mean by a *context*, that is, the 'environment' in which these assumptions are kept. It depends on the proof system what sort of structure the context ought to be, in the situation here we take a context to be a *finite set* of formulae. Another way to think of these contexts is as finite lists, up to order and multiplicity. As such it makes some sense to write Γ, Δ to mean the union of the contexts Γ and Δ, and to write ϕ instead of $\{\phi\}$. Of course, the notion of a judgement as defined below can be applied *mutatis mutandis* to any type of inference system, taking in mind the appropriate notion of formula and context.

Definition 3.2.1. A *judgement* is a statement of the form $\Gamma \vdash \phi$, to be read as 'Γ proves ϕ', where Γ is a *context* and ϕ is a *formula*.

Whenever the context within which we work is empty we write $\vdash \phi$ to mean $\emptyset \vdash \phi$. We now precisely define what we mean by a 'proof' using the notion of a labeled tree, following the explanation of Troelstra and Schwichtenberg (1996) or Sørensen and Urzyczyn (2006). It might help to take the dendritic terminology quite serious in reading proofs; trees start with their root on the ground and branch upwards. In the following definition we keep open the inference system with respect to which the proofs are to be considered. We do this because the notion can be applied to a wide variety of inference systems without any modification, and this formulation emphasizes that fact.

Definition 3.2.2 (Proof). A *proof* (with respect to a system of inference rules $\mathcal{S}$) is a finite *tree* T whose nodes are labeled with judgements (as appropriate for $\mathcal{S}$) subject to the laws:

(i) the labels of the immediate successors of a node v in T are the premisses of an inference rule of $\mathcal{S}$ whose conclusion is the label of v;

(ii) the labels of the leaves of T are axioms of $\mathcal{S}$.

We call the label of the root of a proof its *conclusion*. A judgement is called *provable* (in $\mathcal{S}$) if there is a proof (with respect to $\mathcal{S}$) whose conclusion is this very judgement.

Axioms

$$\frac{}{\Gamma, \phi \vdash \phi}\ (\text{ID}) \qquad\qquad \frac{}{\Gamma \vdash \mathbf{T}}\ (\text{TI})$$

Logical Rules

$$\frac{\Gamma \vdash \phi \qquad \Gamma \vdash \psi}{\Gamma \vdash \phi \wedge \psi}\ (\wedge\text{I}) \qquad \frac{\Gamma \vdash \phi \wedge \psi}{\Gamma \vdash \phi}\ (\wedge\text{E}_\text{L}) \qquad \frac{\Gamma \vdash \phi \wedge \psi}{\Gamma \vdash \psi}\ (\wedge\text{E}_\text{R})$$

$$\frac{\Gamma, \phi \vdash \psi}{\Gamma \vdash \phi \Rightarrow \psi}\ (\Rightarrow\text{I}) \qquad \frac{\Gamma \vdash \phi \Rightarrow \psi \qquad \Delta \vdash \phi}{\Gamma, \Delta \vdash \psi}\ (\Rightarrow\text{E})$$

Figure 3.3: Inference Rules of PIP

[4] See Knuth (1964) for information about the formal meaning of a BNF.

Now consider the inference system PIP in natural deduction style of Figure 3.3. It is easy to see that this bears great resemblance to the system of Figure 3.2. It is in fact a (variant of a) standard result that a judgement is provable in the system PIP exactly when it is 'true' in every partially ordered set that is also a CCC. There are two striking differences. First and foremost, all derivations were labelled in Figure 3.2. Secondly, there was a rule for the 'composition' of derivations, a consequence of the fact that a CCC is in particular a category.

We will now construct a labeling for proofs in PIP, analogous to the labels we had on arrows in a generic CCC as constructed using the deductive system of Figure 3.2. The language we choose our labels from is that of the λ-calculus, which we now introduce as a purely syntactic structure. The definition below inductively defines λ-terms, elements of the language of relevance here. We are however not interested in all λ-terms, for many of these do not correspond to proofs at all. We thus provide an inference system in Figure 3.4 which specifies exactly those terms we care for. This inference system is, of course, simply the system PIP as defined in Figure 3.3, where the derivations are labeled with λ-terms. We work with Church-style λ-terms, that is, λ-terms which encode within their syntax the formula they ought to be related to. This in contrast to Curry-style λ-terms, which do not have this property.

Definition 3.2.3 (Positive λ-Terms). Given an infinite set **V** we define the set of *λ-terms* as follows:

$$\begin{array}{rcll} \mathcal{L}_\Lambda & ::= & x^\phi & \text{variable of type } \phi, \text{ where } x \in \mathbf{V} \text{ and } \phi \in \mathcal{L}_{\text{PIP}} \\ & | & \lambda x^\phi . \mathcal{L}_\Lambda & \text{abstraction} \\ & | & \mathcal{L}_\Lambda \ \mathcal{L}_\Lambda & \text{application} \\ & | & \langle \mathcal{L}_\Lambda, \mathcal{L}_\Lambda \rangle & \text{pairing} \\ & | & \pi^{\mathrm{l}}_{\phi\wedge\psi}(\mathcal{L}_\Lambda) & \text{left projection} \\ & | & \pi^{\mathrm{r}}_{\phi\wedge\psi}(\mathcal{L}_\Lambda) & \text{right projection} \\ & | & * & \text{nil} \end{array}$$

An element of $\mathcal{L}_\Lambda$ will often be denoted by upper-case latin characters like M, N, P.

Before we can meaningfully discuss these terms we first need to define which variables are 'free' and which are 'bound' within a term.

Definition 3.2.4 (Free Variables). We define the *free variables* of a λ-term inductively as follows

$$\begin{array}{rcl} \mathrm{FV}\left(x^\phi\right) & = & \{x : \phi\}, \\ \mathrm{FV}\left(\lambda x^\phi . M\right) & = & \mathrm{FV}(M) - \{x : \phi\}, \\ \mathrm{FV}(M\ N) & = & \mathrm{FV}(M) \cup \mathrm{FV}(N), \\ \mathrm{FV}(\langle M, N\rangle) & = & \mathrm{FV}(M) \cup \mathrm{FV}(N), \\ \mathrm{FV}(\pi^{\mathrm{l}}_{\phi\wedge\psi}(M)) & = & \mathrm{FV}(M), \\ \mathrm{FV}(\pi^{\mathrm{r}}_{\phi\wedge\psi}(M)) & = & \mathrm{FV}(M), \\ \mathrm{FV}(*) & = & \emptyset. \end{array}$$

A variable x in a term M is said to be *bound* when $x \notin \mathrm{FV}(M)$. We say that abstraction over x *binds* x and that abstraction is a *binder*. A λ-term is said to be *closed* when all its variables are bound.

The name of bound variables is irrelevant, hence renaming them can be done without dangerous consequences.[5] Details involving this so-called 'α-equivalence' are tricky and not relevant to our discourse, the interested reader is referred to the introductory texts Barendregt and Barendsen (1988) and Sørensen and Urzyczyn (2006) for more details.[6] As a last point of attention concerning names we assume the name of a bound variable to be 'fresh' or 'previously unused' at introduction to prevent 'name clashes'.[7]

We must update our notion of *context* to this new situation.[8] Whereas it sufficed to simply store all assumptions in a finite set when working with the system PIP, we now need to annotate these assumptions with labels. This should

[5]This is similar to how the name of a local coordinate is irrelevant in geometry, that $R[x]$ and $R[y]$ are isomorphic rings whenever both x and y are transcendental or more down to earth, that in calculus the expressions $\int_0^1 x \, \mathrm{d}x$ and $\int_0^1 y \, \mathrm{d}y$ mean exactly the same thing.

[6]The latter texts covers some of these subtleties, and provides pointers to further literature at the end of the first chapter.

[7]As Barendregt and Barendsen (1988) put it, 'for reasons of hygiene'. Note that this would have been problematic were **V** finite, or more generally, when there can be more terms than variables.

[8]What we call a context, is what Mitchell and Moggi, 1991 call an 'environment'.

be intuitively clear, for a proof which uses some assumption ϕ twice is quite different from a proof which uses two distinct assumptions ϕ. In order to be able to distinguish these, one needs to know which assumption is used, hence the labeling. In the situation of PLC, a context is a finite set of labeled assumptions of the form $x : \phi$, where x is a variable (element of **V**), and ϕ a formula (element of $\mathcal{L}_{\text{PIP}}$). In order to annotate the derivations of Figure 3.3 we need to fix a notation to say 'within the context Γ, the λ-term M denotes a proof of ϕ'. More precisely, we will annotate judgements with λ-terms, and will write $\Gamma \vdash M : \phi$ to intuitively mean what we just described.

From the above it ought to be clear that the context in which a proof occurs must always contain the assumptions made in the proof, which is reflected by the fact that $\mathrm{FV}(M)$ is always a subset of Γ whenever $\Gamma \vdash M : \phi$ is derivable within PLC. The free variables correspond to assumptions, so it is clear that they should be contained within the context. We now have sufficient machinery to consider the system PLC of Figure 3.4.

$$\frac{}{\Gamma, x : \phi \vdash x : \phi}\ (\text{ID}) \qquad\qquad \frac{}{\Gamma \vdash * : \mathbf{T}}\ (\mathbf{T}\text{I})$$

Logical Rules

$$\frac{\Gamma \vdash M : \phi \qquad \Gamma \vdash N : \psi}{\Gamma \vdash \langle M, N\rangle : \phi \wedge \psi}\ (\wedge\text{I}) \qquad \frac{\Gamma \vdash M : \phi \wedge \psi}{\Gamma \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}(M) : \phi}\ (\wedge\text{El}) \quad \frac{\Gamma \vdash M : \phi \wedge \psi}{\Gamma \vdash \pi^{\mathrm{r}}_{\phi\wedge\psi}(M) : \psi}\ (\wedge\text{Er})$$

$$\frac{\Gamma, x : \phi \vdash M : \psi}{\Gamma \vdash \lambda x^{\phi}.M : \phi \Rightarrow \psi}\ (\Rightarrow\text{I}) \qquad \frac{\Gamma \vdash M : \phi \Rightarrow \psi \qquad \Delta \vdash N : \phi}{\Gamma, \Delta \vdash M\ N : \psi}\ (\Rightarrow\text{E})$$

Figure 3.4: Inference Rules of PLC

We now say that a judgement $\Gamma \vdash M : \phi$ is *valid* when it can be derived using PLC. Given this inference system, we revisit the equations we derived between inferences or deductions back in the previous section. All rules of Figure 3.2, except the composition rule (CUT) which a CCC has because it is a category, have their counterpart in the rules of PLC in Figure 3.4. It thus makes sense to write down the equations (3.1), (3.2) and (3.3) we had in the setting of a CCC, and see what they amount to here. First we consider the two equalities on conjunctions (or products), as depicted below, bearing great resemblance to the equalities in (3.1).

$$\frac{\dfrac{\Gamma \vdash M : \phi \qquad \Gamma \vdash N : \psi}{\Gamma \vdash \langle M, N\rangle : \phi \wedge \psi}\ (\wedge\text{I})}{\Gamma \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}(\langle M, N\rangle) : \phi}\ (\wedge\text{El}) \qquad \Gamma \vdash M : \phi \qquad \frac{\dfrac{\Gamma \vdash N : \psi \qquad \Gamma \vdash M : \phi}{\Gamma \vdash \langle N, M\rangle : \psi \wedge \phi}\ (\wedge\text{I})}{\Gamma \vdash \pi^{\mathrm{r}}_{\psi\wedge\phi}(\langle N, M\rangle) : \phi}\ (\wedge\text{Er})$$

Note that one can prove that to a certain annotated judgement, that is, a statement of the form $\Gamma \vdash M : \phi$, there is at most one derivation in the inference system PLC. This is clear because the syntax of the λ-terms neatly reflect the structure of the inference rules. The correspondence would have been a little bit less tidy when we would have worked with Curry-style λ-terms instead of Church-style λ-terms. Indeed, because we write $\pi^{\mathrm{l}}_{\phi\wedge\psi}(N)$ instead of simply $\pi^{\mathrm{l}}(N)$, there is a unique formula χ such that $\Gamma \vdash N : \chi$ is the antecedent of the rule ($\wedge$El) forming the term $\pi^{\mathrm{l}}_{\phi\wedge\psi}(N)$. Of course, this formula χ must be equal to $\phi \wedge \psi$. This would not at all have been clear when the formulae ϕ and ψ were not encoded within the term, which would have been the case for Curry-style terms.

Due to the above, we might as well not write the proof trees above, and simply write down the equation

$$\pi^{\mathrm{l}}_{\phi\wedge\psi}(\langle M, N\rangle) = M = \pi^{\mathrm{r}}_{\psi\wedge\phi}(\langle N, M\rangle), \tag{3.6}$$

in the understanding that $\Gamma \vdash N : \psi$ and $\Gamma \vdash M : \phi$. We now inspect the equation between a deduction with a detour on eliminating and then (re-)introducing an implication, akin to (3.2).

$$\frac{\dfrac{\Gamma \vdash M : \phi \Rightarrow \psi \qquad x : \phi \vdash x : \phi}{\Gamma, x : \phi \vdash M\ x : \psi}\ (\Rightarrow\text{E})}{\Gamma \vdash \lambda x^{\phi}.M\ x : \phi \Rightarrow \psi}\ (\Rightarrow\text{I}) \qquad\qquad \Gamma \vdash M : \phi \Rightarrow \psi$$

When one thinks of a λ-term as a description for computation, it makes sense that the equality above should hold. The equation between λ-terms simply becomes

$$M = \lambda x^{\phi}.M\ x, \tag{3.7}$$

and so this intuitively reads as the statement that the function M and $x \mapsto M(x)$ are equal. In the context of cartesian closed categories we made this argument precise for the large CCC **Set**, which stated exactly that a function $f : X \to Y$ is the same thing as a function $x \in X \mapsto f(x)$. When one simply inspects the proof on the left hand side above, it is clear that this detour is utterly irrelevant to the further structure of the proof. One can thus identify these λ-terms in good conscience.

Finally, we consider the proof with a detour made by introducing and then eliminating an implication, similar to the equation of (3.3). Recall that in (3.3) one could take the 'proof' g to be the identity, to make the equation more obvious. We first inspect this basic form in the setting here, where the implication is eliminated by means of a very trivial proof.

$$\dfrac{\dfrac{\Gamma, x : \phi \vdash M : \psi}{\Gamma \vdash \lambda x^{\phi}.M : \phi \Rightarrow \psi}\,(\Rightarrow\text{I}) \qquad \dfrac{}{y : \phi \vdash y : \phi}\,(\text{ID})}{\Gamma, y : \phi \vdash \left(\lambda x^{\phi}.M\right)\ y : \psi} \qquad\qquad \Gamma, y : \phi \vdash M\left[x \mapsto y\right] : \psi$$

Again writing this equality purely on the λ-terms, we obtain $\left(\lambda x^{\phi}.M\right)\ y = M\left[x \mapsto y\right]$.[9] In this equation, the expression $M\left[x \mapsto y\right]$ denotes the λ-term which is exactly the same as M, except that all occurrences of x are replaced by y. The detour here is obvious, and it is clear that the structure of the proof on the left is in no important way different than the structure on the right. Most importantly, it is clear that the right-hand side can be derived if and only if the left-hand side can be derived. More generally we get the equality below.

$$\dfrac{\dfrac{\Gamma, x : \phi \vdash M : \psi}{\Gamma \vdash \lambda x^{\phi}.M : \phi \Rightarrow \psi}\,(\Rightarrow\text{I}) \qquad \Delta \vdash N : \phi}{\Gamma, \Delta \vdash \left(\lambda x^{\phi}.M\right)\ N : \psi}\,(\Rightarrow\text{E}) \qquad\qquad \Gamma, \Delta \vdash M\left[x \mapsto N\right] : \psi$$

Here we note that $M\left[x \mapsto N\right]$ simply denoted the λ-term M, where each x is replaced by N. It is not as clear as in the above case that $\Gamma \vdash M : \psi$ is valid if and only if $\Gamma, \Delta \vdash M\left[x \mapsto N\right] : \psi$ is valid, both under the assumption that $\Delta \vdash N : \phi$ is valid. This is a basic result of λ-calculus, which we will not prove here. In literature aimed at an audience of mathematicians or logicians, this is often called *subject reduction*. Computer scientists, such as Pierce (2002), instead speak of *preservation*, placing more emphasis of the computational aspect of λ-terms. Indeed, one could place a direction on these equalities in the direction of detour elimination and obtain a *rewrite system*, modeling computation. Subject reduction now simply means that performing a detour-elimination step 'preserves' the theorem the λ-term proves. Assuming that subject reduction does in fact hold, the thus sensible equation now becomes

$$\left(\lambda x^{\phi}.M\right)\ N = M\left[x \mapsto N\right]. \tag{3.8}$$

On an intuitive level, thinking of a λ-term as a function from its context to the formula it is a proof of, this equation means that the expression $f(x)$ is the same thing as the result of computing the evaluation of f at x. As such, computation and 'description of what to compute' are identified.

With the above we have ample motivation to consider the following abstract definition. Let us first go over the items in this definition; unexplained they might appear as coming out of the blue. The first two rules (i) and (ii) are a generalization of the definition of $\mathcal{L}_{\text{PIP}}$, specifying that the class of types (or formulae) ought to be closed under the type-forming operations relevant here. In (iii) we ensure the existence of at least infinitely many variables, a necessity to be able to safely choose fresh variables. In (iv) we specify a way to construct terms out of existing terms, generalizing the above, where terms were exactly those elements of $\mathcal{L}_{\Lambda}$ which could be inferred from Figure 3.4. Finally, we introduced the equations (3.6), (3.7) and (3.8) above. There is also an equation to make sure that $*$ is the only proof of $\mathbf{T}$, and an 'extensionality' equation on products. These basic equations are included in Figure 3.5. The other rules in Figure 3.5 specify that the equations behave well with respect to the structure of λ-terms, analogous

[9] Do note that it is understood implicitly that y occurs nowhere in the judgement $\Gamma \vdash M : \phi$. One should simply read y here as 'some variable, not yet used in the construction of M'.

to the equations on cartesian closed categories discussed in the previous section. The definition below is taken immediately from Lambek and Scott (1988), and it is also used by Awodey (2000).[10]

Definition 3.2.5 (Simply Typed λ-Calculus). A *simply typed λ-calculus* Λ is a triple $\langle F, T, E\rangle$ of classes, where T is indexed by F. An element of F is calles a *type*, an element of T_ϕ a *term* of type ϕ and an element of E is called an *equation*. The following laws govern these types, terms and equations.

(i) the symbol **T** is a type;

(ii) if ψ, ψ are types, so are $\phi \wedge \psi$ and $\phi \Rightarrow \psi$;

(iii) for each type ϕ there is a term (a variable) $x_i{}^\phi$ of type ϕ per $i \in \mathbb{N}$;

(iv) the class of terms is generated inductively from the *basic terms*, the aforementioned variables and the inference rules of Figure 3.4;[11]

(v) equations are of the form $\Gamma \vdash M \equiv N$, where Γ is a finite set of variables containing the free variables of both M and N;

(vi) the relation $\Gamma \vdash (-) \equiv (-)$ satisfies the rules of Figure 3.5;

We write $M : \phi$ to mean that M is an element of T_ϕ, and say that M is of type ϕ. We call two λ-terms M and N equal precisely when they are of the same type and $\Gamma \vdash M \equiv N$ can be derived, with $\Gamma = \mathrm{FV}(M), \mathrm{FV}(N)$.

Note that a variable encodes within its very syntax the type it has, or more mnemonically, one can syntactically determine the assumption from the variable associated to it. This allows us to speak of the type of a term, even when the term is not closed. Suppose for a moment that we did not carry this type information along in our variables. Then given some term M, say for instance $\lambda x^\phi . f\ x$, there is no way to (in general) uniquely assign a type to this term. Indeed, the type of M depends on that of f, the type of which we in general do not know. Yet in the λ-calculi we consider, we would instead write f^ϕ, making it abundantly clear which type the λ-term has.

Remark 3.2.1 (Free Variables). In a λ-calculus there may be terms whose existence is not necessitated by the rules stipulated above. Think for instance of a λ-calculus which has an additional type $\mathbb{N}$, and terms $0 : \mathbf{T} \to \mathbb{N}$, $\mathbf{s} : \mathbb{N} \to \mathbb{N}$. It is important to realize that the notion of free variables, the mapping $\mathrm{FV}((-))$ in particular, needs to be extended to the class of terms associated to the λ-calculus. In extending this mapping, it is understood that the only way to bind a variable is in using abstraction. In a way, one can think of these new terms as constants, and use them to build new terms. The λ-calculus just mentioned for instance contains the term $\lambda x^{\mathbb{N}}.\mathbf{s}\ (\mathbf{s}\ x)$, which has no free variables at all.

Given some set **A** of basic types, we can compute the closure $\mathcal{L}_\mathbf{A}$ of **A** under the type formation rules from the above definition. Now form the λ-calculus which has only those terms and equations needed to be a λ-calculus according to the above definition. See that this gives us an infinite supply of variables $x : \phi$ for each $\phi \in \mathcal{L}_\mathbf{A}$, and λ-terms which can be built from these using the inference rules of PLC. The special case where **A** is simply the empty set is what we treated above. In some sense, this λ-calculus is the most basic kind of λ-calculus one can form. We wish to make this explicit, and can do this via the following definition.

Definition 3.2.6 (Morphism of λ-Calculi). A morphism (or *translation*) $\Phi : \Lambda \to \Lambda'$ is a mapping which sends the terms of Λ to terms of Λ' and types of Λ to types of Λ' such that

(i) the enumeration of variables is respected, that is, $x_i \in \phi$ is mapped to $x_i \in \Phi(\phi)$;

(ii) the type associated to a term is respected by the mapping, in symbols, $M : \phi$ is mapped to $\Phi(M) : \Phi(\phi)$;

(iii) the amount of binders is not increased, in symbols, $\mathrm{FV}(\Phi(M)) \subseteq \mathrm{FV}(M)$;

[10] In this definition, (iv) might seem somewhat off. In Lambek and Scott (1988) there is no explicit distinction between basic terms and those terms constructed inductively thereof. Such a distinction however is implicitly assumed in Awodey (2000), so we make this differentiation here as well. It is no real problem, as most of the interesting λ-calculi are constructed in an inductive fashion like this anyway. The major benefit is that we can easily say that basic terms do not have any free variables, which helps us a bit when defining the interpretation of a generic λ-calculus in Section 3.3.

[11] These rules are chosen such that the set of variables taken into consideration ('context') is manageble. There is a rule to enlarge this amount, so this poses no additional difficulty.

(iv) type formation operations are preserved 'on the nose';

(v) term formation operations are preserved up to the the governing equalities;

(vi) equalities are preserved.

Definition 3.2.7 (Category of λ-Calculi). There is a category **Calc** of λ-calculi, where the objects are λ-calculi and the morphisms are as defined above. Two morphisms are considered equal when they map equal terms to equal terms. Morphisms are composed in the obvious manner.

We can now state that the basic λ-calculus, whose types, terms and equalities are only those forced by the definition of a λ-calculus, is an initial object in the category **Calc**. There is an obvious translation, acting as an identity on terms and types, from any λ-calculus to this basic λ-calculus. Any translation of this basic λ-calculus into a λ-calculus must respect the variables and all terms built out of them, and hence the totality of terms in the basic λ-calculus. Similarly, the type formation operators are preserved, and as all types in the basic λ-calculus are made solely out of these formation rules, all types are preserved. This proves there to be but one unique translation from the basic λ-calculus into any λ-calculus, so it indeed is an initial object.

It is possible to make a category out of each λ-calculus. This category is in particular cartesian closed, and will provide us with a model, as explained in the next section. Let us first formally define this 'syntactic category'.

Equivalence Relation

$$\text{FV}(M) \vdash M \equiv M \qquad \frac{\Gamma \vdash M \equiv N}{\Gamma \vdash N \equiv M} \qquad \frac{\Gamma \vdash M \equiv N \qquad \Gamma \vdash N \equiv P}{\Gamma \vdash M \equiv P}$$

Enlargement of Context

$$\frac{\Gamma \vdash M \equiv N \qquad \Gamma \subseteq \Delta}{\Delta \vdash M \equiv N}$$

Compatibility

$$\frac{\Gamma \vdash M \equiv N \qquad \Gamma \vdash R \equiv R}{\Gamma \vdash M\ R \equiv N\ R} \qquad \frac{\Gamma \vdash L \equiv L \qquad \Gamma \vdash M \equiv N}{\Gamma \vdash L\ M \equiv L\ N}$$

$$\frac{\Gamma \vdash M \equiv N \qquad \Gamma \vdash R \equiv R}{\Gamma \vdash \langle M, R \rangle \equiv \langle N, R \rangle} \qquad \frac{\Gamma \vdash L \equiv L \qquad \Gamma \vdash M \equiv N}{\Gamma \vdash \langle L, M \rangle \equiv \langle L, N \rangle}$$

$$\frac{\Gamma \vdash M \equiv N}{\Gamma \vdash \pi^{\text{l}}_{\phi\wedge\psi}(M) \equiv \pi^{\text{l}}_{\phi\wedge\psi}(N)} \qquad \frac{\Gamma \vdash M \equiv N}{\Gamma \vdash \pi^{\text{r}}_{\phi\wedge\psi}(M) \equiv \pi^{\text{r}}_{\phi\wedge\psi}(N)}$$

$$\frac{\Gamma \vdash M \equiv N}{\Gamma - \{x\} \vdash \lambda x^{\phi}.M \equiv \lambda x^{\phi}.N}$$

Basic Equalities

$$\begin{array}{rcrcl}
\text{FV}(M) & \vdash & \lambda x^{\phi}.M\ x & \equiv & M \\
(\text{FV}(M) \cup \text{FV}(N)) - \{x\} & \vdash & (\lambda x^{\phi}.M)\ N & \equiv & M[x \mapsto N] \\
\text{FV}(M) \cup \text{FV}(N) & \vdash & \pi^{\text{l}}_{\phi\wedge\psi}(\langle M, N \rangle) & \equiv & M \\
\text{FV}(M) \cup \text{FV}(N) & \vdash & \pi^{\text{r}}_{\phi\wedge\psi}(\langle M, N \rangle) & \equiv & N \\
\text{FV}(M) & \vdash & \left\langle \pi^{\text{l}}_{\phi\wedge\psi}(M), \pi^{\text{r}}_{\phi\wedge\psi}(M) \right\rangle & \equiv & M \\
\text{FV}(M) & \vdash & * & \equiv & M
\end{array}$$

Figure 3.5: Inference Rules for Equations

Definition 3.2.8 (Syntactic Category). Let Λ be a λ-calculus. We now define the *syntactic category* $\mathbf{T}(\Lambda)$ of Λ as the category whose objects are the types of Λ, and where an arrow $\phi \to \psi$ is a valid judgement $x : \phi \vdash M : \psi$. We say that two such arrows $x : \phi \vdash M : \psi$ and $y : \phi \vdash N : \psi$ are equal whenever $y : \phi \vdash M[x \mapsto y] \equiv N$ is provable in Λ. The identity arrow is given by the judgement $x : \phi \vdash x : \phi$ and the composition of two arrows $x : \phi \vdash M : \psi$ and $y : \psi \vdash N : \chi$ is given by the judgement $x : \phi \vdash N[y \mapsto M] : \psi$. It is easy to see that this in fact defines a category.

It is usually said that the arrows of this syntactic category $\mathbf{T}(\Lambda)$ are pairs $\langle x : \phi, M\rangle$ where M is a λ-term with x^ϕ as only free variable, with a similar notion of equality. Of course, this comes down to exactly the same thing as we defined above. The notation chosen here conveys more clearly the idea that these arrows are terms of a certain type within a context, and it is visually more appealing.

Recall from Remark 3.2.1 that one has to be careful with the free variables of a term in a generic λ-calculus. Think again of the λ-calculus described in that remark, and see for instance that $x : \phi \vdash \mathbf{s} : \mathbb{N} \Rightarrow \mathbb{N}$ is an arrow $\phi \to \mathbb{N} \Rightarrow \mathbb{N}$ in its syntactic category for any x and ϕ, as there are no variables in the term **s**. Another example of an arrow is $x : \mathbb{N} \vdash \mathbf{s}\,x : \mathbb{N}$, whose domain and codomain is $\mathbb{N}$.

We now wish to show that the syntactic category is in fact cartesian closed.

Proposition 3.2.1. *Let Λ be a λ-calculus. The syntactic category $\mathbf{T}(\Lambda)$ is cartesian closed.*

Proof. In order to prove the desired, it suffices to show that the structure of Figure 3.1 is present and that the necessary equations hold. Indeed, we already know of the type-forming operations demanded in Definition 3.1.1, so let us now simply provide the required arrows and operations thereon. Consider the arrows provided below, which provide the axioms of Figure 3.1.

$$\begin{aligned}
\mathrm{id}_X &:= x : \phi \vdash x : \phi \\
\pi^{\mathrm{l}} &:= p : \phi \wedge \psi \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}(p) : \phi \\
\pi^{\mathrm{r}} &:= p : \phi \wedge \psi \vdash \pi^{\mathrm{r}}_{\phi\wedge\psi}(p) : \psi \\
! &:= x : \phi \vdash * : \mathbf{T} \\
\epsilon &:= p : (\psi \Rightarrow \phi) \wedge \psi \vdash \pi^{\mathrm{l}}_{(\psi\Rightarrow\phi)\wedge\psi}(p)\ \pi^{\mathrm{r}}_{(\psi\Rightarrow\phi)\wedge\psi}(p) : \phi
\end{aligned}$$

There are two inference rules we need to provide, for in the above definition we already demonstrated that $\mathbf{T}(\Lambda)$ is a category. Let us first construct the rule for pairing, so assume we are given arrows $x : \gamma \vdash M : \phi$ and $x : \gamma \vdash N : \psi$. We can now construct the arrow below, as required.

$$p : \gamma \vdash \langle M, N\rangle : \phi \wedge \psi$$

Now suppose we are given an arrow $p : \phi \wedge \psi \vdash M : \chi$. We construct the following arrow.

$$x : \phi \vdash \lambda y^\phi . M[p \mapsto \langle x, y\rangle] : \psi \Rightarrow \chi$$

To complete the proof, we need to show that the five diagrams of Definition 3.1.1 commute. The first diagram is concerned with projections, at it is not very hard to see that

$$\left(x : \gamma \vdash M : \phi\right) = \left(x : \gamma \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}(\langle M, N\rangle) : \phi\right) = \pi^{\mathrm{l}}\,\langle x : \gamma \vdash M : \phi, x : \gamma \vdash N : \psi\rangle$$

holds. A similar equation can be made for the right projection. The second diagram is a bit more intricate, so let us construct the three arrows involved in this diagram step by step. First consider the top and right arrows as below.[12]

$$\begin{aligned}
h &= p : \phi \wedge \psi \vdash M : \chi \\
\epsilon &= z : (\psi \Rightarrow \phi) \wedge \psi \vdash \pi^{\mathrm{l}}(z)\,\pi^{\mathrm{r}}(z) : \phi
\end{aligned}$$

[12] Note that we suppress the type information in $\pi^{\mathrm{l}}(q)$, simply to prevent the terms from becoming too large to display properly. The type annotation should be understood to be there anyway.

From here we can compute $\overline{h}\pi^{\mathrm{l}}$ as the arrow given by

$$
\begin{aligned}
&= \left(a \,:\, \phi \vdash \lambda b^{\psi}.M\left[p \mapsto \langle a, b\rangle\right] \,:\, \psi \Rightarrow \chi\right) \circ \left(q \,:\, \phi \wedge \psi \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}\left(q\right) \,:\, \phi\right)\\
&= q \,:\, \phi \wedge \psi \vdash \lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right] \,:\, \psi \Rightarrow \chi
\end{aligned}
$$

With these ingredients at hand we can now compute $\epsilon\left\langle \overline{h}\pi^{\mathrm{l}}, \pi^{\mathrm{r}}\right\rangle$, given as below

$$
q \,:\, \phi \wedge \psi \vdash \pi^{\mathrm{l}}\left(\left\langle \lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right], \pi^{\mathrm{r}}(q)\right\rangle\right)\, \pi^{\mathrm{r}}\left(\left\langle \lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right], \pi^{\mathrm{r}}(q)\right\rangle\right) \,:\, \chi
$$

One can now derive, using the equalities on λ-terms, that the term involved actually equals h. Indeed, the projection equalities ensure that first equality below holds. The second equality follows essentially from (3.8), and the third is due to 'extensionality' of the product. The final equation is due to our convention on variable names; as neither p nor q are bound in M, they can be exchanged freely without changing the meaning of M. From this description it should be clear that, although we only prove the equality with partial rigour, it can be proven formally using the inference rules of Figure 3.5. The other equalities can be proven in a similar fashion.

$$
\begin{aligned}
& \pi^{\mathrm{l}}\left(\left\langle \lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right], \pi^{\mathrm{r}}(q)\right\rangle\right)\, \pi^{\mathrm{r}}\left(\left\langle \lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right], \pi^{\mathrm{r}}(q)\right\rangle\right)\\
= \;& \left(\lambda b^{\psi}.M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), b\right\rangle\right]\right)\,\left(\pi^{\mathrm{r}}(q)\right)\\
= \;& M\left[p \mapsto \left\langle \pi^{\mathrm{l}}(q), \pi^{\mathrm{r}}(q)\right\rangle\right]\\
= \;& M\left[p \mapsto q\right]\\
= \;& M
\end{aligned}
$$

□

It is possible to make **T** into a functor **T** : **Calc** → **Cart**. Given a transformation of λ-calculi, simply transform the terms associated to objects in the term model of the domain into terms associated to objects in the term model of the codomain, via the transformation present.

SECTION 3.3 Models of λ-Calculus

Consider the basic λ-calculus, which consists only of those terms, types and equalities it is required to have. In the previous section we informally saw that proofs in zeroth-order logic correspond directly with λ-calculus terms. Furthermore, we saw that two proofs equal up to detours got mapped to 'equal' λ-terms. This suggests that one could interpret λ-terms, informally, as proofs.

This basic λ-calculus is constructed in such a manner that all arrows in its syntactic category are reflected in some manner in every CCC. To be a bit more precise, to each CCC $\mathcal{C}$ there is a cartesian closed functor from the term model of the basic λ-calculus to $\mathcal{C}$. This functor maps objects of the basic λ-calculus, which are constructed using the type forming operations as in (3.5) alone, to an object of $\mathcal{C}$ using the functors $(-) \wedge (-)$ and $(-) \Rightarrow (-)$ and the distinguished object **T** to reflect these type-forming operations. We chose the names of the type-forming operations to reflect the functors, so the expression $(\phi \Rightarrow \psi) \wedge \chi$ both denotes a formula (element of $\mathcal{L}_{\text{PIP}}$) and an object of the term model of the basic λ-calculus. Terms are mapped, along their inductive definition as according to Figure 3.4, to maps $\mathcal{C}$ via the corresponding deductions of Figure 3.2.

It is important to remark that the above described functor indeed is a cartesian closed functor. First see that in the above description one maps $p \,:\, \phi \wedge \psi \vdash \pi^{\mathrm{l}}(p) \,:\, \phi : \phi \wedge \psi \to \phi$ to $\pi^{\mathrm{l}} : \phi \wedge \psi \to \phi$, interpreting the type-forming operations as the appropriate functors. As a consequence, products are preserved. In a similar manner one can see that exponents are preserved.

An interpretation, as defined below, will be an assignment sending types to objects and terms to arrows, satisfying some constraints. These constraints are meant to ensure that the interpretation $[\![(-)]\!]$ gives rise to a cartesian closed functor on the term model, when mapping arrows $x \,:\, \phi \vdash M \,:\, \psi : \phi \to \psi$ to $[\![x \,:\, \phi \vdash M \,:\, \psi]\!]$.

Definition 3.3.1 (Model of the λ-Calculus). An *interpretation* $\mathcal{M}$ of the λ-calculus Λ is a pair $\langle \mathcal{C}, [\![(-)]\!]\rangle$, where $\mathcal{C}$ is a CCC and $[\![(-)]\!]$ is a mapping from types to objects and judgements to arrows. More precisely, it maps a type ϕ of Λ to some object $[\![\phi]\!]$ of $\mathcal{C}$, and is extended in some fixed way to finite sets of types, such that $[\![\Gamma]\!]$ is defined as the product $\prod_{\gamma\in\Gamma}[\![\gamma]\!]$. Moreover, it maps judgements $\Gamma \vdash M \,:\, \phi$ to arrows $[\![\Gamma \vdash M \,:\, \phi]\!] : [\![\Gamma]\!] \to [\![\phi]\!]$.

This assignment is subject to the laws of Figure 3.6 regarding the interpretation of judgements, and the equalities regarding the interpretation of types below. The arrows π and Δ are the obvious and canonical projection and diagonal maps respectively. The isomorphisms described in the definition are such that when restricting to the term model as described above, these are simply the canonical arrows arising from this functor.

$$[\![\phi \Rightarrow \psi]\!] \cong [\![\phi]\!] \Rightarrow [\![\psi]\!], \quad [\![\phi \wedge \psi]\!] \cong [\![\phi]\!] \wedge [\![\psi]\!], \quad [\![\mathbf{T}]\!] \cong \mathbf{T}$$

We write $\mathcal{M} \models_\Gamma M = N$ when $[\![\Gamma \vdash M : \phi]\!] = [\![\Gamma \vdash N : \phi]\!]$ for terms M and N with both $\Gamma \vdash N : \phi$ and $\Gamma \vdash M : \phi$ valid. The interpretation $\mathcal{M}$ is said to be a *model* of the λ-calculus Λ if for all terms M and N one has that $\mathcal{M} \models_\Gamma M = N$ if the the equality $\Gamma \vdash M \equiv N$ holds in Λ. To mean that $\langle \mathcal{C}, [\![(-)]\!]\rangle$' is an interpretation of Λ, we will often say that $[\![(-)]\!]$ is an interpretation of Λ in $\mathcal{C}$.

$$\frac{}{[\![\Gamma, x : \phi]\!] \xrightarrow{\pi} [\![x : \phi]\!] = [\![\Gamma, x : \phi \vdash x : \phi]\!]} \qquad \frac{}{[\![\Gamma]\!] \xrightarrow{!} \mathbf{T} \cong [\![\mathbf{T}]\!] = [\![\Gamma \vdash * : \mathbf{T}]\!]}$$

$$\frac{\Gamma \vdash M : \phi \qquad \Gamma \vdash N : \psi}{[\![\Gamma]\!] \xrightarrow{\langle [\![\Gamma \vdash M : \phi]\!], [\![\Gamma \vdash N : \psi]\!]\rangle} [\![\phi]\!] \wedge [\![\psi]\!] \cong [\![\phi \wedge \psi]\!] = [\![\Gamma \vdash \langle M, N\rangle : \phi \wedge \psi]\!]}$$

$$\frac{\Gamma \vdash M : \phi \wedge \psi}{[\![\Gamma]\!] \xrightarrow{[\![\Gamma \vdash M : \phi\wedge\psi]\!]} [\![\phi \wedge \psi]\!] \cong [\![\phi]\!] \wedge [\![\psi]\!] \xrightarrow{\pi^{\mathrm{l}}} [\![\phi]\!] = [\![\Gamma \vdash \pi^{\mathrm{l}}_{\phi\wedge\psi}(M) : \phi]\!]}$$

$$\frac{\Gamma \vdash M : \phi \wedge \psi}{[\![\Gamma]\!] \xrightarrow{[\![\Gamma \vdash M : \phi\wedge\psi]\!]} [\![\phi \wedge \psi]\!] \cong [\![\phi]\!] \wedge [\![\psi]\!] \xrightarrow{\pi^{\mathrm{r}}} [\![\phi]\!] = [\![\Gamma \vdash \pi^{\mathrm{r}}_{\phi\wedge\psi}(M) : \phi]\!]}$$

$$\frac{\Gamma, x : \phi \vdash M : \psi}{[\![\Gamma]\!] \xrightarrow{\eta} [\![\phi]\!] \Rightarrow ([\![\Gamma]\!] \wedge [\![\phi]\!]) \xrightarrow{[\![\phi]\!] \Rightarrow [\![\Gamma, x : \phi \vdash M : \psi]\!]} [\![\phi]\!] \Rightarrow [\![\psi]\!] \cong [\![\phi \Rightarrow \psi]\!] = [\![\Gamma \vdash \lambda x^\phi.M : \phi \Rightarrow \psi]\!]}$$

$$\frac{\Gamma \vdash M : \phi \Rightarrow \psi \qquad \Delta \vdash N : \phi}{[\![\Gamma, \Delta]\!] \xrightarrow{\Delta} [\![\Gamma]\!] \wedge [\![\Delta]\!] \xrightarrow{\langle [\![\Gamma \vdash M : \phi \Rightarrow \psi]\!], [\![\Delta \vdash N : \phi]\!]\rangle} ([\![\phi]\!] \Rightarrow [\![\psi]\!]) \wedge \phi \xrightarrow{\epsilon} [\![\psi]\!] = [\![\Gamma, \Delta \vdash M\ N : \psi]\!]}$$

Figure 3.6: Equalities for an Interpretation of the λ-Calculus.

The above definition which specifies when $\mathcal{M} = \langle \mathcal{C}, [\![(-)]\!]\rangle$ can be called a *model* is quite similar to the usual notion of a model. One says that an interpretation (or *structure*) is a *model* of a certain theory T when the interpretation of some expression in the relevant language is *valid* in the model whenever it is *provable* in the theory. We know that a λ-calculus provides a theory of equality on its terms, with inference rules as given in Figure 3.5. The language we work with here is that of equations of λ-terms, so expressions are of the form $\Gamma \vdash M \equiv N$, where $\Gamma \vdash M : \phi$ and $\Gamma \vdash N : \phi$ are both provable in the term-inference system of Figure 3.4. This expression was said to be valid exactly when $[\![\Gamma \vdash M : \phi]\!]$ and $[\![\Gamma \vdash N : \phi]\!]$ are equal in $\mathcal{C}$. In short; a model of the λ-calculus interprets formal equations as actual equations in a given CCC, such that provable equalities are sent to equalities that hold in the CCC.

From the above definition, it is clear that the interpretation must respect the inductive structure of Figure 3.4. Moreover, the cartesian closed structure of $\mathcal{C}$ as ensured by Figure 3.2 is used in a very straightforward way. For idealistic reasons, the equality $[\![\Gamma \vdash * : \mathbf{T}]\!] =!$ is included in Figure 3.6, but it of course holds anyway by virtue of $\mathcal{C}$ being CCC.

Let Λ be some λ-calculus, and consider the CCC $\mathbf{T}(\Lambda)$. There exists an interpretation of Λ in its term model $\mathbf{T}(\Lambda)$, which simply sends terms along their inductive construction to arrows in $\mathbf{T}(\Lambda)$. Indeed, each term is either a basic term (a constant), in which case it can be mapped to itself (paired with an appropriately typed variable for the context), or inductively constructed using the inference rules of Figure 3.4. Each such inductive construction can be mapped along its inductive structure, using the CCC counterpart of the λ-calculus construct at hand. For instance, one has

$$[\![x : \phi, y : \psi \vdash x^\phi : \phi]\!] = p : \phi \wedge \psi \vdash \pi^{\mathrm{l}}(p) : \phi,$$

as is forced to hold by the equations in Figure 3.6. To give another example, consider

$$[\![x : (\phi \Rightarrow \psi) \wedge \phi \vdash \pi^{\mathrm{l}}(x)\, \pi^{\mathrm{r}}(x) : \phi]\!] = \epsilon,$$

which can be verified by spelling out the construction of the term on the left-hand side. First see that the interpretation of the projections equals the projections, in particular $[\![x : (\phi \Rightarrow \psi) \wedge \phi \vdash \pi^{\mathrm{l}}(x) : \phi \Rightarrow \psi]\!] = \pi^{\mathrm{l}}$. This forces

$$\begin{aligned}[\![x : (\phi \Rightarrow \psi) \wedge \phi \vdash \pi^{\mathrm{l}}(x)\, \pi^{\mathrm{r}}(x) : \psi]\!] &= \epsilon \left\langle [\![x : (\phi \Rightarrow \psi) \wedge \phi \vdash \pi^{\mathrm{l}}(x) : \phi \Rightarrow \psi]\!], [\![x : (\phi \Rightarrow \psi) \wedge \phi \vdash \pi^{\mathrm{r}}(x) : \phi]\!] \right\rangle \\ &= \epsilon \left\langle \pi^{\mathrm{l}}, \pi^{\mathrm{r}} \right\rangle = \epsilon.\end{aligned}$$

It is an immediate consequence of the definition of the syntactic category that this interpretation is in fact a model. The following proposition summarizes the above.

Proposition 3.3.1 (Term Model). *Let Λ be some λ-calculus. There is an obvious interpretation of Λ in $\mathbf{T}(\Lambda)$, and this interpretation is in fact a model. We call this interpretation the term model of Λ.*

We now have the necessary machinery to precise define what we mean by 'complete' and 'functionally complete'. It is clear from the construction of the term model that it is both complete and functionally complete.

Definition 3.3.2. A model $[\![(-)]\!]$ of Λ in the category $\mathcal{C}$ is said to be *complete* when given valid judgements $\Gamma \vdash M : \phi$ and $\Gamma \vdash M : \phi$ one has that

$$[\![\Gamma \vdash M : \phi]\!] = [\![\Gamma \vdash N : \phi]\!] \qquad \text{if and only if} \qquad \Gamma \vdash M \equiv N.$$

The model is called *functionally complete* when each arrow $\mathbf{T} \to [\![\psi]\!]$ arises as the interpretation of a judgement $[\![\emptyset \vdash M : \phi]\!]$, unique up to equality.

Interpretations combine well with cartesian closed functors. To be a bit more precise, let $[\![(-)]\!]$ be an interpretation of Λ in $\mathcal{C}$ and let $F : \mathcal{C} \to \mathcal{D}$ be a cartesian closed functor. By 'composing' the interpretation in $\mathcal{C}$ with F we obtain $F[\![(-)]\!]$, which is an interpretation of Λ in $\mathcal{D}$. Indeed, a cartesian closed functor respects the type-forming operations, and as it respects the structure maps as well, the equalities of Figure 3.6 hold for the interpretation $F[\![(-)]\!]$. Moreover, it is obvious that if $[\![(-)]\!]$ is a model of Λ in $\mathcal{C}$, then $F[\![(-)]\!]$ is a model of the same λ-calculus in the category $\mathcal{D}$. The following lemma is easily seen to be true.

Lemma 3.3.1. *Let $F : \mathcal{C} \to \mathcal{D}$ be a cartesian closed functor, and let $[\![(-)]\!]$ be a model of Λ in $\mathcal{C}$. Now the following statements about $F[\![(-)]\!]$ hold:*

(i) if F is faithful, then the new model is complete whenever the old one was;

(ii) if F is full, then the new model is functionally complete whenever the old one was;

We now have built sufficient machinery to tackle our main objective of constructing a spatial model for any λ-calculus. Before we proceed, let us first prove the following minor technical lemma, which provides the category of presheaves over a small category with a set of enough points. With this final ingredient we prove our main theorem with almost shameful ease.

Lemma 3.3.2. *Let $\mathcal{C}$ be a small category. For each object $C \in \mathcal{C}$ there is a point C of $\mathbf{PSh}(\mathcal{C})$, and the set of all such points is jointly epimorphic.*

Proof. Given an object $C \in \mathcal{C}$ we wish to define the point of $\mathbf{PSh}(\mathcal{C})$, which we denote by a slight abuse of syntax as $C : \mathbf{Set} \to \mathbf{PSh}(\mathcal{C})$. To this end, we define its inverse-image and direct-image parts on objects as below, and make the sensible definition on arrows.

$$C^* : \mathbf{PSh}(\mathcal{C}) \to \mathbf{Set},\ P \mapsto P(C), \qquad C_* : \mathbf{Set} \to \mathbf{PSh}(X),\ S \mapsto \mathbf{Set}\left(\mathcal{C}\left(C, (-)\right), S\right)$$

It is fairly easy to compute that C^* is indeed left exact, that is to say, preserves finite limits. This is a direct consequence of the fact that limits are computed point-wise in the category $\mathbf{PSh}(\mathcal{C})$. We now need to show that C^* is left adjoint to C_*, which amounts to proving the existence of a natural isomorphism

$$\mathbf{Set}\left(C^*(P),\, S\right) \cong \mathbf{PSh}(\mathcal{C})\left(P,\, \mathbf{Set}\left(\mathcal{C}\left(C,\, (-)\right),\, S\right)\right).$$

In order to do this we define the following mappings, and check that they are mutually inverse.

$$f : P(C) \to S \mapsto \Big(p \in P(X) \mapsto \big(g : C \to X \mapsto f(P(g)(p))\big)\Big)_{X \in \mathcal{C}}$$
$$\alpha : P \to \mathbf{Set}\left(\mathcal{C}\left(C,\, (-)\right),\, S\right) \mapsto \Big(p \in P(C) \mapsto \alpha_C(x)(\mathrm{id}_c)\Big)$$

In the one direction, it is easy to compute that given $f : P(C) \to S$ one obtains the map

$$p \in P(C) \mapsto \Big(x \mapsto \big(g \mapsto f(P(g)(x))\big)\Big)(p)(\mathrm{id}_C) = f\big(P(\mathrm{id}_C)(p)\big) = f(p),$$

which clearly equals f. The other direction requires a slightly longer computation. It is easy to see that the equation below holds, which follows mostly from the the naturality of $\alpha : P \to \mathbf{Set}\left(\mathcal{C}\left(C,\, (-)\right),\, S\right)$.

$$p \in P(X) \mapsto \Big(g : C \to X \mapsto \big((x \mapsto \alpha_C(x)(\mathrm{id}_C))(P(g)(p)) = \alpha_C(P(g)(p))(\mathrm{id}_C) = \alpha_X(p)(g)\big)\Big).$$

We have now proven C to be a point, so we are done when we can show that the totality of functors C^* for objects $C \in \mathcal{C}$ is jointly epimorphic. This is nothing but obvious, for if $g : P \to Q$ is an arrow in $\mathbf{PSh}(\mathcal{C})$, then $C^*(g) = g_C : P(C) \to Q(C)$. As a natural transformation is an isomorphism exactly if it is an isomorphism on all components, the desired follows immediately. ☐

Theorem 3.3.1. *Let Λ be some λ-calculus such that the classes of types and terms are sets. There exists a topological space X and a model $[\![(-)]\!]$ on $\mathbf{Sh}(X)$ which is both complete and functionally complete.*

Proof. Consider the syntactic category $\mathbf{T}\left(\Lambda\right)$ and the Yoneda embedding $\mathbf{y} : \mathbf{T}\left(\Lambda\right) \to \mathbf{PSh}(\Lambda)$. This is a full and faithful functor, which is cartesian closed by Proposition 3.1.2.

By Example 1.3.2 we know that $\mathbf{PSh}(\mathbf{T}\left(\Lambda\right))$ is a Grothendieck topos. Now apply Theorem 3.1.1 to this Grothendieck topos, and know that we have a set of enough points by Lemma 3.3.2. Due to the theorem we know of a topological space $\mathbf{X}$ and an essential geometric morphism $\mathbf{s} : \mathbf{Sh}(\mathbf{X}) \to \mathbf{PSh}(\mathbf{T}\left(\Lambda\right))$, whose inverse-image part is a full and faithful cartesian closed functor. We can thus form the composite $\mathbf{s}^*\mathbf{y} : \mathbf{T}\left(\Lambda\right) \to \mathbf{Sh}(\mathbf{X})$, which also is full, faithful and cartesian closed.

Now consider the term model $[\![(-)]\!]$ in $\mathbf{T}\left(\Lambda\right)$ of Proposition 3.3.1 as described in this section. We know this model to be both complete and functionally complete. It now follows from the above Lemma 3.3.1 that the model $\mathbf{s}^*\mathbf{y}[\![(-)]\!]$ is complete and functionally complete, as desired. ☐

Bibliography

Artin, Michael et al., editors (1972). *Théorie des Topos et Cohomologie Etale des Schémas.* Volume 270. Lecture Notes in Mathematics. Springer. DOI: 10.1007/BFb0061319 (cited on page 31).

Awodey, Steven (2000). "Topological Representation of the Lambda-Calculus". In: *Mathematical Structures in Computer Science* 10.1, pages 81–96 (cited on pages 4, 44, 55).

Barendregt, Henk P. and Erik Barendsen (1988). "Introduction to Lambda Calculus". In: *Aspenæs Workshop on Implementation of Functional Languages, Göteborg.* Programming Methodology Group, University of Göteborg and Chalmers University of Technology. URL: citeseer.ist.psu.edu/barendregt94introduction.html (cited on pages 44, 52).

Butz, Carsten Horst Erich (Nov. 1996). "Logical and Cohomological Aspects of the Space of Points of a Topos". PhD thesis. Utrecht University (cited on page 31).

Butz, Carsten Horst Erich and Ieke Moerdijk (Sept. 1999). "Topological Representation of Sheaf Cohomology of Sites". In: *Compositio Mathematica* 118.2, pages 217–233. DOI: 10.1023/A:1001169215774 (cited on pages 4, 31).

Cartan, Henri (1950-1951). "Faisceaux sur un Espace Topologique". In: *Espaces Fibrés et Homotopie.* Volume 3. URL: http://www.numdam.org/item?id=SHC_1950-1951__3__A14_0 (cited on page 12).

Church, Alonzo (1932). "A Set of Postulates for the Foundation of Logic". In: *The Annals of Mathematics* 33.2, pages 346–366 (cited on page 43).

Curry, Haskell B. (1942). "The Inconsistency of Certain Formal Logics". In: *The Journal of Symbolic Logic* 7.3, pages 115–117 (cited on page 43).

De Bruijn, Nicolaas Govert (1970). "Symposium of Automatic Demonstration". In: volume 125. Lecture Notes in Mathematics. Springer. Chapter The mathematical language AUTOMATH, its usage, and some of its extensions, pages 29–61. DOI: 10.1007/BFb0060623 (cited on page 44).

De Vrijer, Roel (1987). "Exactly estimating functionals and strong normalization". In: *Indagationes Mathematicae (Proceedings)* 90.4, pages 479 –493. ISSN: 1385-7258. DOI: 10.1016/1385-7258(87)90012-6. URL: http://www.sciencedirect.com/science/article/B7GXJ-4DT5R97-8X/2/a0f786461124e93f79fac81f56c73bf2 (cited on page 3).

Diamond, Fred and Jerry Shurman (2005). *A First Course in Modular Forms.* Edited by S Axler et al. Volume 228. Graduate Texts in Mathematics. Springer (cited on pages 7, 14).

Forster, Otto (1981). *Lectures on Riemann Surfaces.* Edited by P.R. Halmos et al. Springer (cited on pages 5, 14, 18).

Friedman, Harvey (1975). "Logic Colloquium". In: volume 453. Lecture Notes in Mathematics. Springer Berlin / Heidelberg. Chapter Equality between functionals. ISBN: 978-3-540-07155-6. DOI: 10.1007/BFb0064870 (cited on page 3).

Grothendieck, Alexandre (1955). *A General Theory of Fibre Spaces with Structure Sheaf.* Technical report. University of Kansas (cited on page 12).

Hartshorne, Robin (1977). *Algebaric Geometery.* Edited by S Axler et al. Volume 52. Graduate Texts in Mathematics. Springer (cited on pages 6, 7).

Johnstone, Peter T. (1982). *Stone Spaces.* Edited by D. J. H. Garling et al. Volume 3. Cambridge studies in advanced mathematics. Cambridge University Press (cited on page 16).

— (1997). *Topos Theory.* Edited by P. M. Cohn and G. E. H. Reuter. Volume 10. London Mathematical Society Monographs. Acedemic Press Inc. (cited on pages 5, 31).

— (2002a). *Sketches of an Elephant: A Topos Theory Compendium 1.* Edited by Dov M. Gabbay et al. Volume 43. Oxford Logic Guides. Oxford University Press (cited on page 5).

Johnstone, Peter T. (2002b). *Sketches of an Elephant: A Topos Theory Compendium 2*. Edited by Dov M. Gabbay et al. Volume 44. Oxford Logic Guides. Oxford University Press (cited on pages 5, 24).

Knuth, Donald E. (1964). "Backus normal form vs. Backus Naur form". In: *Communications of the ACM* 7.12, pages 735 –736 (cited on page 51).

Lambek, Joachim and Philip J. Scott (1988). *Introduction to Higher-Order Categorical Logic*. Edited by D. J. H. Garling et al. Volume 7. Cambridge studies in advanced mathematics. Cambridge University Press. ISBN: 0-521-35653-9 (cited on pages 3, 43, 44, 55).

Lang, Serge (1999). *Complex Analysis*. Edited by S Axler et al. Volume 103. Graduate Texts in Mathematics. Springer. ISBN: 0-387-98592-1 (cited on page 6).

Mac Lane, Saunders (1997). *Categories for the Working Mathematician*. Edited by S. Axler et al. Second. Graduate Texts in Mathematics. Springer. ISBN: 0-387-98403-8 (cited on pages 18, 24, 37, 41, 43).

Mac Lane, Saunders and Ieke Moerdijk (1991). *Sheaves in Geometry and Logic*. Edited by S Axler et al. Springer. ISBN: 0-387-97710-4 (cited on pages 5, 9, 17, 28, 29).

Mitchell, John C. and Eugenio Moggi (1991). "Kripke-style models for Typed Lambda Calculus". In: *Annals of Pure and Applied Logic* 51.1, pages 99–124 (cited on pages 4, 52).

Moerdijk, Ieke (1996). "Classifying spaces for toposes with enough points". In: *Milan Journal of Mathematics* 66.1, pages 377–389. DOI: 10.1007/BF02925366 (cited on page 31).

Pierce, Benjamin C. (2002). *Types and Programming Languages*. The MIT Press. ISBN: 0-262-16209-1 (cited on page 54).

Plotkin, G.D. (1973). *Lambda-definability and logical relations*. Technical report. School of Artificial Intelligence, University of Edinburgh (cited on page 2).

Schönfinkel, Moses (1924). "Über die Bausteine der mathematischen Logik". In: *Mathematische Annalen* 92.3-4, pages 305–316. DOI: 10.1007/BF01448013 (cited on pages 1, 47).

Scott, Dana S. (1980). "To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism". In: edited by Jonathan P. Seldin and J. Roger Hindley. Acedemic Press Inc. Chapter Relating Theories of the Lambda-Calculus, pages 403–450 (cited on page 3).

Selinger, Peter (2007). "Lecture Notes on the Lambda Calculus". Department of Mathematics and Statistics, Dalhousie University, Halifax, Canada (cited on page 2).

Serre, Jean-Pierre (1955). "Faisceaux Algebriques Coherents". In: *The Annals of Mathematics* 61.2, pages 197–278. URL: http://www.jstor.org/stable/1969915 (cited on page 12).

Silverman, Joseph H. (1994). *The Arithmetic of Elliptic Curves*. Edited by S Axler and K.A. Ribet. 2nd edition. Volume 106. Graduate Texts in Mathematics. Springer (cited on pages 7, 8).

Sørensen, Morten Heine and Pawel Urzyczyn (2006). *Lectures on the Curry-Howard Isomorphism*. Edited by S. Abramsky et al. Studies in Logic and the Foundations of Mathematics. Elsevier. ISBN: 0-444-52077-5 (cited on pages 1, 3, 44, 51, 52).

Terese (2003). *Term Rewriting Systems*. Edited by Marc Bezem et al. Volume 55. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press (cited on page 3).

Troelstra, Anne Sjerp and Helmut Schwichtenberg (1996). *Basic Proof Theory*. Edited by S Abramsky et al. New York, NY, USA: Cambridge University Press. ISBN: 0-521-57223-1 (cited on page 51).