Sets and Categories: What Foundational Approaches Tell Us About Mathematical Thought

MSc Thesis *(Afstudeerscriptie)*

written by

Robin Martinot
(born December 4th, 1995 in Nijmegen, Netherlands)

under the supervision of **Dr. Luca Incurvati**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense:
August 30th, 2019

Members of the Thesis Committee:
Prof. Dr. Ariana Betti
Dr. Luca Incurvati (supervisor)
Dr. Thomas Schindler
Dr. Ekatarina Shutova (chair)

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis explores the idea that set theory and category theory represent different ways of thinking. Adopting the perspectives of various foundational systems based in set theory and category theory, we investigate two common and informal conceptions about the distinction between set-theoretical and categorical thinking. One concerns the intuition that set theory and category theory respectively correspond to a bottom-up and a top-down approach to mathematics. The other captures the idea that category theory represents a higher level of abstraction than set theory. Our investigation brings us the two main results of this thesis. First, we argue that the bottom-up/top-down distinction is irrelevant to the distinction between set-theoretical and categorical thinking. Second, we claim that, while categorical foundations are generally characterized by a higher level of abstraction compared to set-theoretical foundations, this difference is more variable and more modest than generally thought.

In order to familiarize ourselves with the various foundational systems, we discuss set-theoretical foundations in Chapter 2, and categorical foundations in Chapter 3. Chapter 2 also treats the development of general category theory from set-theoretical foundations so as to better delineate the categorical way of thinking. The incorporation of a variety of systems in our approach is significant for the arguments leading to the main results in Chapter 4. Additionally, the arguments benefit from the new refinements we make to the bottom-up/top-down distinction and of a method of abstraction coming from computer science.

Contents

Acknowledgements

I would like to express some words of gratitude here. Luca, thank you for providing reliable and guiding supervision, and for always keeping a clear perspective on the development of the thesis. I was not very experienced in either philosophy or mathematics before starting the thesis, but you made it possible for me to write it in the exciting combination of these two fields — thank you for that!

I want to thank Liesbeth, Mark and Myrthe for believing in me. Although you may not have seen much of me in the period of writing my thesis, the knowledge of your support has always been a comfort.

Thanks to the numerous fellow MoL students (in particular hallway inhabitants Jelly, David and Nuno) for always providing a pleasant distraction from work. Thanks to Rachael, Dean and Zhuoye for letting me invade the Logic House with rants about life (and death) and (square) horses. Special thanks to Rachael for rubbing my arm when I needed it.

Finally, thank you Sam, for your steady supply of card games, your endless patience in listening to my thesis-related thoughts, your belief in me — but above all, thanks for being you.

Chapter 1

Introduction

This thesis occupies itself with the ways in which we can look at mathematics. In particular, we take on perspectives from foundations for mathematics, each of which sheds a fundamental light on what mathematics is about. We focus on the often-made comparison between foundations based in set theory and those based in category theory. Most of the philosophical debate has concentrated on what the nature of a genuine foundation for mathematics should be, and how set theory and category theory fit into this picture — however, this is not where our interests will lie. Instead, we aim to identify what it is that characterizes the distinction between a set-theoretical and a categorical way of thinking. Category theory has been advocated as a theory of mathematics that, compared to set theory, expresses a more natural and efficient approach towards mathematics (Lawvere, 1966). This intuition has not been properly formalized, however, and we need to understand the difference between set-theoretical and categorical thinking in order to verify it. Other than refining intuitions, it is generally relevant for the practice of mathematics to understand how foundational systems relate to each other. Investigating the ways of thinking represented by two different mathematical fields will eventually be relevant, as well, for the more general goal of understanding the nature of mathematical thought. This latter idea was introduced in (Mathias, 2001) and (Ernst, 2017), but has remained essentially unexplored.

Hence, in this thesis, we are interested in answering the following research question. *How (if so) can we distinguish categorical thinking from set-theoretical thinking?* As this is a rather broad research question, we restrict ourselves to the investigation of two common intuitions concerning its answer.

The first regards the informal distinction between a 'bottom-up' and a 'top-down' approach to mathematics. The idea, here, is that set theory and, respectively, category theory represent a bottom-up and a top-down manner of thinking, while regarding the same subject matter (this idea appears, for example, in (Awodey, 2004)). In this thesis, we formalize the bottom-up/top-down distinction, and we argue that it is in fact independent from the distinction between set-theoretical and categorical thinking.

Secondly, it is a broadly accepted claim that category theory is more abstract than set theory (remarks along these lines can be found in many articles, for instance in (Linnebo and Pettigrew, 2011) or (Landry, 2013)). This implies, contrary to the previous intuition, that the difference between set-theoretical and categorical thinking may find its roots in the fact that they enjoy distinct subject matters. We point out that the notion of 'abstraction' is consistently used in an informal way, and that (even if our intuition regarding this matter is true) it remains unclear exactly *how much* more abstract category theory is compared to set theory. We address both issues by borrowing (and slightly modifying) Floridi's *method of abstraction* described in (Floridi, 2013). This leads to our claim that, in general, categorical foundations embody a higher level of abstraction than set-theoretical foundations, although in several cases this difference is more modest than generally assumed.

In reaching the above-mentioned conclusions, we stress that our approach provides us with a faithful representation of mathematics from the perspective of set theory and category theory, by taking into account a variety of foundational systems based in either theory. Close regard for the differences between these foundations will be important for the development of our arguments. Our approach thus combines a rich perspective on foundations with new methods of distinguishing between set-theoretical and categorical thinking.

1.1 Structure of the thesis

In Chapter 2, we present an overview of the various set-theoretical foundations for category theory. This will remind the reader of the properties of existing set theories, and it will illustrate the means with which category theory can be incorporated into set theory. By including results from a range of set theories, we aim to get a good sense of what characterizes the set-theoretical approach to mathematics and its ability to found category theory. Knowledge of the potential of set theory to act as a foundation for category theory is a means to indirectly characterize properties of categorical thinking — this will be useful later, especially when categorical foundations themselves refuse to fully reveal their nature. Following this is Chapter 3, which treats a number of categorical foundations and their properties. This will, analogously, give us an idea of the mathematical characteristics of the categorical perspective on mathematics.

The relation between Chapter 2 and Chapter 3 will become apparent in Chapter 4. Here, knowledge from the previous two chapters is synthesized into the development of the two ideas outlined above, based on the informal distinctions between set-theoretical and categorical thinking. The first part of Chapter 4 motivates viewing the distinction between set theory and category theory separately from the bottom-up/top-down distinction. Subsequently, the second part of Chapter 4 will begin by introducing the formalization (and our modification) of the concept 'level of abstraction' that originates from computer science. We argue that a distinction between set theory and category theory based on abstraction levels tells part of the story. Nevertheless, we lie emphasis on our conclusion that the categorical abstraction of subject matter in the discussed foundations is more variable, and less strong, than generally thought. Our contributions to the debate on set-theoretical and categorical thinking can be summarized as follows.

1. One general recommendation of this thesis is that imprecise terms that are intended to link mathematical knowledge with philosophical reasoning benefit from being made explicit. This avoids misunderstandings when

such terms are used in different contexts. In this thesis, we contribute by making more explicit the terms *bottom-up*, *top-down* and *level of abstraction*, and we apply them to various set-theoretical and categorical foundations.

2. For the term *level of abstraction*, in particular, we call for a further improved framework that can capture true abstraction relations. To this end, we make a first modification to the method used in this thesis, but more rigorous changes are required.

3. Specific to the research question, we have made a start to characterize the distinction between categorical and set-theoretical thinking by singling out properties that are (not) important to this distinction. That is, we suggest that the bottom-up/top-down distinction is irrelevant to the research question, whereas differences in levels of abstraction do play a role — albeit less so than generally thought.

4. We relate the topic to the roles of foundations for mathematics. We suggest, through responding to (Landry, 2013), that the goal of foundations may be tied to their level of abstraction. The connection of set-theoretical thinking and categorical thinking to this relation should be further explored. Generally, however, our findings advocate a pluralistic view on the foundations for mathematics, where a situation is analyzed from the perspective of a foundation of an abstraction level suitable to the purpose on hand.

A remark. In this thesis we aim to discuss and compare a range of foundational systems, so that it is necessary to limit ourselves to the relevant mathematical properties of these systems that we use in our discussion. As such the level of mathematical detail will vary. We always aim to give a comprehensive view of the situation, and zoom in on technical notions when relevant. A more thorough and deep understanding of the subject would of course be given by fully exploring the mathematical properties of each system. We offer references to more in-depth mathematical treatments at various points in the thesis.

1.2 The basic notion of a category

Throughout this thesis we will assume that the reader has a basic familiarity with set theory and category theory. For an introduction to category theory, we refer to (Awodey, 2010), and for more advanced material covering topos theory to (Johnstone, 2002). Introductory and advanced material on set theory may be found in many textbooks, for example (Jech, 2013). For completeness we will state here the basic definition of a category as introduced by Eilenberg and Mac Lane in 1945 (we adapt the definition from (Awodey, 2010)). This definition is referred to as EM; besides being the main ingredient for category theory developed from set theory in Chapter 2, it will also come up in Chapter 4, as it has been proposed as a categorical foundation for mathematics. We will only briefly mention EM in Chapter 3, as the definition is already provided here.

Definition 1. A *category* consists of the following components:

1. Objects $A, B, C, \ldots$
2. Arrows $f, g, h, \ldots$
3. For every arrow f, there exist objects D and C called the *domain* and *codomain* of f, respectively (we write $f : D \to C$).
4. For every pair of arrows $f : A \to B$, $g : B \to C$, there exists the *composite* arrow $g \circ f : A \to C$.
5. For every object A, there exists the *identity* arrow $1_A : A \to A$.

These ingredients of a category are additionally required to satisfy the following laws:

1. *Associativity*: for arrows $f : A \to B$, $g : B \to C$ and $h : C \to D$, it holds that
$$h \circ (g \circ f) = (h \circ g) \circ f.$$
2. *Unit*: for every arrow $f : A \to B$, it holds that
$$f \circ 1_A = f = 1_B \circ f.$$

Thus, a category is taken to be anything satisfying this definition. We are now ready to embark on the next chapter, which will explore set-theoretical properties that affect the development of EM-category theory.

Chapter 2

Category theory from set theory

In this chapter, we give an overview of the set-theoretical systems that have been developed to provide a foundation for category theory, in order to better characterize EM-categorical thinking. Several main approaches to the founding of category theory can be identified, and their workings are quite well understood. As we will see, the key problem is to find a way to allow for the free construction of large categorical objects. This is not generally supported by standard set theory, as its axiomatization imposes restrictions on what qualifies as a set and, hence, a category (assuming the standard way of incorporating category theory into set theory — this will be illustrated later). Thus, we will see that in nearly every set-theoretical framework, some restriction to the development of category theory has to be imposed which we might like to avoid. Besides technical adequacy, some authors have also commented on the philosophical suitability of set-theoretical frameworks for category theory. For, not every set-theoretical system that allows for a large part of category theory is as intuitive and elegant as we would like. We begin the chapter by elaborating on criteria that can be imposed on set-theoretical foundations for mathematics, which will help us better familiarize ourselves with the way they work. In the next part of the chapter we will outline different set theories and their level of success in founding category theory. Rather than ZFC itself, several extensions of ZFC with inaccessibility assumptions or class axioms have proved to be successful in allowing for large categories. Alternatively, the incorporation of a particular reflection principle is a method worthwhile to pursue in several ways. We consider both variants of methods, where illustrations are often inspired by (Shulman, 2008), which provides an excellent mathematical discussion of the various approaches. We additionally discuss the set theory New Foundations (NF), which originates from type theory and regulates the formation of sets in a different way from ZFC. We will conclude the chapter by discussing the key patterns that arise among the different ways for set theory to found category theory.

2.1 Requirements for set-theoretical foundations

Let us consider for a moment what we would like a set-theoretical foundation for category theory to satisfy. This question has been asked and answered (differently) by several authors. The particular requirements imposed on set-theoretical systems tell us something about the way category theory is regarded. For instance, Feferman would like for category theory to show that it is not susceptible to Russell-like paradoxes. Consider his system of requirements (R) as described in (Feferman, 2011). For a foundational system S for category theory, Feferman proposes that S should do the following:

(R1) Allow us to construct the category of all structures of a given kind, e.g. the category **Grp** of all groups, **Top** of all topological spaces, and **Cat** of all categories.

(R2) Allow us to construct the category B^A of all functors from A to B, where A and B are any categories.

(R3) Allow us to establish the existence of the usual basic mathematical structures and carry out the usual set-theoretical operations.

(R4) S should be established to be consistent relative to a currently accepted system of set theory.

The first two requirements concern so-called *unrestricted* or *unlimited* categories, which are intended to capture truly *all* constructions of a given kind without somewhere restricting the universe that one works in. Thus, whereas the first two requirements are easily supported if we adopt Feferman's interest in unrestricted constructions, (R3) lacks a sense of formality that turns out to cost him its consistency with (R1). Namely, (Ernst, 2015) shows that Feferman's requirement that a foundation accommodates all categories of a given kind (R1) is inconsistent with (R3). Briefly, this is because (R1) lets us construct the category of all reflexive graphs — (R3) then implies that this category has several properties that allow for the construction of a particular contradictory arrow. The proof resorts to Cantor's diagonalization method in a category-theoretic context. Thus, (R) is not an obvious combination of requirements to be satisfied. In order to avoid contradictions within (R), (Enayat, Gorbow, and McKenzie, 2017) propose to weaken (R3) to the following.

(R3') Ordinary mathematics and category theory, along with its distinction between large and small, are naturally implementable.

While the notion of *natural implementability* is again rather informal, Enayat, Gorbow and McKenzie distinguish three levels of 'decreasing user-friendliness'. The friendliest option requires the underlying set theory to allow for category-theoretic notions with its usual membership relation, without requiring a restriction for category theory to a set or class. If restriction to a set or class is necessary, we find ourselves on the second level of user-friendliness. The least user-friendly variant involves the additional requirement that the implementation of category-theoretic notions comes with a (well-motivated) membership relation that is different from the membership symbol of the underlying language. Allowing for these restrictions in requirement (R3') sidesteps the contradictions that arise with Feferman's variant (R3). As we will see later, various set theories are able to satisfy such a weaker version of (R3). This asks for a bit more tolerance towards the set-theoretical implementation of usual category theory, however, in that its constructions should be representable in a slightly more indirect way than we are perhaps used to. Still, the results from (Enayat, Gorbow, and McKenzie, 2017) show that employing mathematically more nuanced requirements gives a good case for set-theoretic foundations for category theory.

On the other hand, we might consider category theory to be independent of its ability to accommodate unrestricted notions. Instead, we might take set theory as usually known by ZFC as the starting point for category theory, and formulate requirements from there that suit the development of category theory best. In that case, we might end up with a set of requirements (S) as follows (these are requirements that (Enayat, Gorbow, and McKenzie, 2017) implicitly extracted from motivations given in (Shulman, 2008)).

(S1) Ordinary category theory, along with its distinction between large and small, is naturally implementable.

(S2) The large/small distinction is relative, in the sense that for any x, there is a notion of smallness such that x is small.

(S3) ZFC is interpreted by $\in$, both when quantifiers are restricted to large sets, and when quantifiers are restricted to small sets.

Here, the large/small distinction is required to be implementable in such a way that any large object 'might as well be small'. That is, we want to be able to easily characterize the behaviour of large categories; the next section will elaborate on various ways of doing this. Furthermore, (S3) asks for the satisfaction of anything that ZFC proves, both for large and small sets. The preservation of ZFC in a set-theoretical foundation for category theory is useful if there is some benefit from adhering to this well-known set-theoretic context. Indeed, the advantages and disadvantages of using (extensions of) ZFC for category theory are well-known, and we mainly focus on such set theories in our comparison to category theory. However, it should be kept in mind that additional non-classical set theories may have interesting properties with respect to category theory, and should be explored in a similar fashion for a more complete picture.

In this chapter, we are interested in keeping track of the effects on category theory with respect to requirements like (R) and (S), when we consider different set theories. Hence, we do not explicitly take a stand on these requirements, although (R1) and (S3) are easy to motivate, as well as the requirement that there is a way to deal with large categories - the latter is necessary already for establishing familiar completeness results. Thus, we think it will be helpful to look at the ways in which set theory allows for different categorical constructions. When exploring unlimited, (R1)-like categories, we will assume a weaker variant of (R3) such as (R3'), in order to avoid previously mentioned contradictions. It will turn out that we can capture a lot of category theory already with a large/small distinction in our set-theoretical universe, while disregarding Feferman's unlimited categories.

The role of urelements

We take a brief excursion into explicit ways of accommodating (small) EM-categories within set theory. Commonly, a category is represented inside set theory as an ordered pair, containing a set of objects O, a set of arrows M, and a composition function $\circ$. While (Muller, 2001) and (Feferman and Kreisel, 1969) take such an ordered pair to be represented as $(O, M, \circ)$, (Feferman, 2011) expands this tuple to include explicit domain and codomain functions. Analogously to the way (Benacerraf, 1965) argues there is no 'one correct way' to define the successor operation for the natural numbers, the different ways of incorporating categories into set theory should all be taken as more or less correct. Hence, the particular method used to translate category theory into set theory is irrelevant for us — philosophically, there is the following relevant side to it, however.

It is characteristic of EM-category theory to say little about its objects and arrows except how they are related. When restricting category theory to a set-theoretical basis, then, one might want to add as little extra information as possible, in order to retain the apparent featurelessness of the objects and arrows in EM-categories. To this end, we suggest the inclusion of urelements in the reduction of category theory to set theory may be worthwhile. Where many philosophers claim that sets are 'constituted by their elements' (Linnebo, 2013; Boolos, 1971; and others), urelements are objects that do not have this property, although they can be elements of a set. As urelements possess no elements, but are not all equal to the empty set, they seem to have less features than usual sets, and may hence be more compatible with the neutral nature of categorical objects. Urelements could (in various ways) easily be incorporated into the above-defined implementation of categories.

Keeping this possibility in mind, we should now be ready to embark on the analysis of categorical constructions a little higher up inside various set theories, which will prove to elicit some interesting behaviours.

2.1 A notion of set

It has previously been noted that ZFC is not a very suitable foundation for category theory (see (Mac Lane, 1969), (Shulman, 2008), and others). We will briefly examine why this is so. With the previous section in mind, we see that the construction of a category inside set theory will rely on the collection of objects and arrows of the category being sets. Thus, when doing category theory in ZFC, we are able to construct any *small* category, as these simply consist of the relevant sets of objects and arrows. However, let us make a small counterpoint here. Suppose that we take a small category that consists of one object and its identity arrow, but its object is a large category (such as **Set** or **Cat**). Then the object must clearly be represented as a set in ZFC as well, even though it is a large category. In these situations we can simply take any set and identity function to *represent* the large category. However, we might wonder whether we find this implementation satisfying enough as a representation of our object and arrow. Note that where we argued in the previous section that set-theoretical representations of objects and arrows can say too much, in this case set theory seems to say too little. Hence, while ZFC is mathematically a perfectly fine foundation for small categories, the molding of categories into sets can result in a distorted picture.

In addition, ZFC can also handle working with a single large category. For example, we can build the category of all sets **Set** (as an informal class), with all sets as its objects, and all functions between sets as its arrows. This allows us to work inside **Set** using ZFC, but we will not be able to do anything with **Set** as a whole, as it is not a set itself. ZFC does not give us any tools to deal with **Set** in another way. Indeed, we also cannot build even larger categories or functors between large categories, which restricts the development of a large part of category theory.

However, there is a way in which we can define large categories (which can be informally thought of as non–set-like) while staying within ZFC (Shulman, 2008). For this we let *classes* be collections of sets defined by a property expressible in ZFC. Thus, we can express properties like 'X is a group' or 'X is a set', so that we can conceive of a class of all groups or a class of all sets (defined from the universe V). Then we now let a *large* category be a category such that its collections of objects and arrows are proper classes (i.e., classes that are not sets). Using previously defined 'class-properties' we can construct new properties, which allows us to perform basic operations on classes (and thus on large categories). For example, we can define the Cartesian product of two large categories, and we can prove that **Set** and other large categories are (co)complete. This will still not be enough, however, as ZFC does not possess any axioms for manipulating classes and lacks a way to quantify over them. Many theorems in category theory (such as the Adjoint Functor Theorem) involve quantification over large categories and are thus not even expressible in ZFC. The only way to deal with classes is via the properties expressed in the language of ZFC, which poses restrictions.

Thus, to define the collection of objects and arrows of a small category, a notion of set can be used, and ZFC provides this. However, its lack of machinery to deal with classes, which are merely implicit in the theory, provide serious drawbacks to the suitability of ZFC as a foundation for category theory.

2.1.1 Non–well-founded sets

We briefly consider a non–well-founded set theory developed in (Aczel, 1988). The motivation for this is based on (Incurvati, 2014), where it is argued that, contrary to general opinion, non–well-founded set theories embody a conception of set. Whereas ZFC promotes the *iterative conception of set*, Incurvati argues that the non–well-founded set theory ZFA embodies the *graph conception of set*. Indeed, we can think of sets as what is depicted by an arbitrary graph. This seems to be an interesting perspective for category theory, since a category is essentially a reflexive graph satisfying the axioms for composition (Ernst, 2015).

For the graph conception of set, we let nodes and (directed) edges represent sets and the converse membership relation, respectively. Then the edge $a \to b$ denotes that the set represented by b is an element of the set represented by a. Furthermore, a graph is *pointed* if there is a unique, distinguished top node (the point) that represents the set that the graph depicts. Finally, a graph is *accessible* if each node can be reached by some finite path starting from the point. Then we restrict our attention to directed accessible pointed graphs (*apgs*): this allows us to, from a given graph, uniquely recognize the set that it depicts. Generally, non–well-founded set theories are constructed from ZFC by eliminating its Foundation axiom, and adding in a particular Anti-Foundation axiom. There exist multiple such axioms that can be incorporated, and each specifies a slightly different requirement on which non–well-founded sets may exist.

However, the framework suffers from similar shortcomings to ZFC-like set theories considering its technical adequacy. Namely, every set in ZFA is ZF-bijective to a well-founded set in ZF, and every object in a model of ZF also exists in a model of ZFA (McLarty, 1993). Thus, ZFA does not allow for any new isomorphism types compared to ZFC. Secondly, even though the accommodation for non–well-founded sets seems promising perhaps for unrestricted categories, this is not the case. As ZFA, like ZFC, contains a restricted Comprehension axiom, the construction of an unlimited **Set** or **Cat** will still be blocked. Hence, non–well-founded set theory (ZFA, in particular) is no more successful as a foundation for category theory than ZFC.

(Incurvati, 2014) argues that the graph conception of set motivates most axioms of ZFA, and that it naturally gives rise to the existence of non–well-founded sets. The result that non–well-founded set theory motivates a different conception of set than ZFC, but does not differ in its ability to found category theory, will prove to be relevant for our characterization of categorical and set-theoretical thinking in Chapter 4.

2.2 A notion of largeness

There exist several ways of accommodating large categories in set theories based on ZFC, each of which has their own advantages and disadvantages for category theory. Nearly all of them impose a sense of size to the universe that we work in, and hand us a reference frame for the objects that we construct. The general approaches can be divided into (1) adding inaccessibility assumptions, (2) introducing a set-class distinction, and (3) making use of reflection principles. We finally consider a set-theoretical system that extends a theory built in New Foundations as well as ZFC.

2.2.1 Inaccessibility assumptions

First, we consider what happens when we add an axiom to ZFC that postulates the existence of a large cardinal. It has become convention to let this cardinal be *inaccessible* (that is, uncountable, a strong limit, and regular). If α is an inaccessible cardinal, then V_α is a model of the whole of ZFC and V_α is called a *Grothendieck universe*. Equivalently, a Grothendieck universe is a transitive set U that is closed under pairing, power sets and indexed unions. Indeed, the existence of such a U is equivalent to the existence of an inaccessible cardinal (Shulman, 2008; Feferman, 2011). If α is only a limit ordinal greater than ω, V_α need not be a model of the Replacement axiom of ZFC, and thus V_α is not necessarily a model of ZFC. However, it turns out that we do not have to go as big as inaccessible. We call a cardinal κ *worldly* if V_κ models ZFC (Incurvati and Löwe, 2016). Worldly cardinals are always strong limits, but they need not be regular. Hence, where our discussion of previous research extends ZFC with inaccessibles, we should keep in mind that the same result can be obtained with worldly cardinals. We will come back to this a little later on.

Now let **ZFC + I** be ZFC + "there exists an inaccessible cardinal". Note that ZFC alone cannot prove the existence of an inaccessible cardinal, so that this results in a theory strictly stronger than ZFC. Thus, pick an inaccessible cardinal κ; then V_κ will act as our frame of reference. That is, we define elements of V_κ to be *small sets*, and sets that are not necessarily elements of V_κ to be *large sets*. With this large/small distinction, we can construct many large categories, such as the functor category **[A,B]** for any two large categories **A** and **B**. This approach also allows us to distinguish between the size of large sets: **[A,B]** is larger than both **A** and **B**, which is not possible with class theories.

Still, it must be noted that this method does not allow us to be fully unrestricted in our constructions. In particular, it will not satisfy the requirements (R1), (R2) and (S2) that we mentioned in the beginning. By assuming the existence of an inaccessible cardinal κ, the universe V_κ will set a limit. Consider, for example, the category **Set**, which ideally should capture truly *every* set. What we can do first is construct the category **Set** of all small sets, which will live outside V_κ as a large category. However, clearly then **Set** does not include any large sets. In this case, we could introduce classes, that are even larger than large sets, or we could assume the existence of another large cardinal $\lambda > \kappa$. A redefinition of *large set* would then be 'element of V_λ', and we introduce the notion of *very large set* (quasi- or meta-category are also used as terms for very large categories) as a set that is not necessarily an element of V_λ. This would allow for the construction of a new '**Set**' that captures both small and large sets. Still, it is clear that we have not captured the very large sets by doing this. Thus, inaccessibles allow for safe construction of larger categories, but we will never be able to define the category of, say, *all* sets, in a larger universe than the one that we work in.

Another issue is that the distinction between large and small that we now have, might not always behave the way we want it to. For example, if we have an object that is large relative to V_κ, we may want to treat it as a small object, for which we would want to switch to a bigger universe than V_κ. Grothendieck proposed to add an axiom to ZFC that asserts the existence of arbitrarily large inaccessibles, so that the size of the universe may always be changed, and we can apply results about categories to larger categories. This gives us a denumerable number of inaccessible cardinals added to ZFC, resulting in the theory ZFC$^\omega$. In contrast, the single-universe system ZFCU proposed by Mac Lane (like the assumption of a single inaccessible) is usually sufficient for ordinary mathematics. However, the fact remains that care is required in dealing with results obtained in a particular universe. There exist categorical objects that depend on size considerations for a particular other category (e.g., for the category of *all* groups), meaning that the existence of these objects does in fact depend on the size of the universe. Hence, when changing universes one must be precise in making sure that arguments do not suffer from such a dependence.

Last, even though Grothendieck's ZFC$^\omega$ can satisfy (S2), it (and the other approaches based on inaccessible cardinals) presents a different worry. We are strictly strengthening our set theory, and this is not something that category theory desperately needs. In fact, it increases the risk that our theory is inconsistent, and we may make it incompatible with possible future axioms that could be useful for category theory (Enayat, Gorbow, and McKenzie, 2017). These objections are theoretical and perhaps pose a small risk in practice. However, it remains the case that these systems give us much more than is required. (Muller, 2001) reasons that category theory only talks about objects in the classes V, $\mathcal{P}(V)$, ..., $\mathcal{P}^n(V)$ (V being the class of all sets, and n a fixed finite number). Therefore, given that $\mathfrak{i}$ is the first inaccessible cardinal, we only need a denumerable sequence of increasing universes $V_{\mathfrak{i}_1+1}$, $V_{\mathfrak{i}_1+1}$, ..., $V_{\mathfrak{i}_1+\omega}$. However, ZFC+I, ZFCU, ZFC$^\omega$ and the theory ZMC (which will soon come up) give us far more than this, up to $V_{\mathfrak{i}_1^\omega}$ and further. It seems strange that we should assume so much more than what we really need. The existence of large constructions is not always directly necessary to prove statements about objects living in them. Hence, it may be that the assumption of inaccessibles can actually be eliminated from many arguments, making them excessive and not crucial to capturing what category theory needs.

2.2.2 Reflection principles

(Feferman and Kreisel, 1969) showed that, in fact, working with universes may be unnecessary for category theory. The size distinction between large and small can be made in a different way, one that eliminates the need for switching between universes. This requires the following notation. Let φ be any statement and M any structure, such as a potential model for a theory. Then we define φ^M as φ relativized to M, so that all quantifiers in φ are restricted to range only over elements of M. Furthermore, suppose that $M \subset N$, then $\varphi(x_1, ..., x_n)$ is *reflected* from N to M if

$$\forall x_1 \in M \cdots \forall x_n \in M(\varphi^N(x_1, \ldots, x_n) \Leftrightarrow \varphi^M(x_1, \ldots, x_n)).$$

Note that ZF proves the Levy Reflection Principle: for every formula φ, there is a set X such that φ is reflected from V to X. This principle is in fact equivalent to Infinity and Replacement (in the presence of the other axioms of ZF) (Incurvati, 2017). We here define a new reflection principle using a constant symbol $\mathbb{S}$. Indeed, $\mathbb{S}$ is added to our language and denotes the universe of small sets, as V_κ did before. Sets that are elements of $\mathbb{S}$ will now be called *small*, whereas sets not necessarily in $\mathbb{S}$ are called *large*. Then when adding the axiom "$\mathbb{S}$ is transitive and closed under subsets" and the instantiated reflection axiom

$$\forall x_1 \in \mathbb{S} \cdots \forall x_n \in \mathbb{S}(\varphi(x_1, \ldots, x_n) \Leftrightarrow \varphi^{\mathbb{S}}(x_1, \ldots, x_n))$$

to ZFC, we obtain the system ZFC/S. Here, in the translation of a formula φ to its relativized version $\varphi^{\mathbb{S}}$, each $(\exists x)$ and $(\forall x)$ is replaced by $(\exists x \in \mathbb{S})$ and $(\forall x \in \mathbb{S})$, respectively. Thus, instead of assuming a large infinity to make the distinction between large and small, we introduce a constant symbol and a reflection principle, which achieve the same result. This allows us to work with small objects, and use the reflection principle to apply our results to large objects. The use of a reflection principle then formalizes the intuition that the notion of smallness is a tool for legitimizing general constructions. Where we needed to be careful with our proofs for theorems in ZFC+I for them to be preserved in larger universes, this is unnecessary in ZFC/S. Unlike ZFC+I, furthermore, ZFC/S is a conservative extension of ZFC, so that we manage without strictly strengthening ZFC, although this comes with the following sacrifice.

The Replacement axiom (*if a class F is a function, then for any set X there exists a set $Y = F[X] = \{F(x) : x \in X\}$*) must be weakened for small sets by imposing the requirement that F is *small-definable*. This means that F has to be a subset of $\mathbb{S}$ of the form $\{s \in \mathbb{S} | \varphi(s)\}$ for a definable property $\varphi(s)$ that refers only to elements of $\mathbb{S}$. This turns out to restrict some of the category theory that we can do. For example, suppose that we want to prove that **Set** is a complete category. Here, **Set** equals **Set**$[\mathbb{S}]$, denoting that its objects are all sets in the universe $\mathbb{S}$. Then it should hold that **Set** has all limits for functors $F : \mathbf{A} \to \mathbf{Set}[\mathbb{S}]$ for small **A**. However, if F is not small, then it need not be small-definable, and $F(A)$ is not necessarily a set. Thus, the proof will not work, yet completeness of **Set** is something we would like to have. In fact, small-definable categories possess many of the properties that are desirable. This provides incentive to change the definition of properties such as completeness or of adjoints: for example, considering the functor category of *small* functors $[\![\mathbf{A}, \mathbf{Set}]\!]$ instead of the whole functor category $[\mathbf{A}, \mathbf{Set}]$ allows for successful proofs of adjoints. It must be noted that in practice, we will not suffer from many restrictions because most categorical objects that we are interested in are already (equivalent to) small-definable objects. Still, this remains tedious to establish, and so the small-definability restrictions are something we would like to avoid. Thus, the main issue with ZFC/S concerns the practical inconvenience of keeping track of the small-definability restrictions. Dealing with large cardinals and constructing ordinary category theory from this system is supported well, which perhaps makes the small-definability limitation a minor one.

It turns out that strengthening ZFC/S further to ZMC/S by adding the axiom "$\mathbb{S} = V_\kappa$ for some inaccessible cardinal κ" solves the small-definability problems (Shulman, 2008). This can easily be seen by the fact that V_κ will model ZFC and hence also Replacement. Every functor with small domain will thus automatically be small again. Repeatedly applying the reflection principle knowing that κ is inaccessible tells us that there must be infinitely many small inaccessibles and also arbitrarily large inaccessibles. ZMC/S is conservative over ZFC + "any finite set of formulas is reflected in some Grothendieck universe", so we have actually strengthened the reflection principle (Shulman, 2008). Hence, even though we do not have a category of all sets (which Feferman would still desire), this set theory allows us to always pick a category of sets large enough for any purpose that it might as well contain all of them. Thus ZMC/S is quite an aesthetic solution to the small-definability problems of ZFC/S, but we do lose conservativity over ZFC.

Alternatively, we can solve the problems from ZFC/S by using *indexed categories*. Indexed categories are formally defined in (Johnstone, 2002), and their precise relation to ZFC/S is explored in (Shulman, 2008). We will not discuss them in detail, however, as it concerns the incorporation of topological tools in ZFC/S, which is not inherently characteristic of the way set theory founds category theory. Additionally, their definition is quite intricate and not very relevant for our purposes here. However, their general effect can be described as follows. Essentially, indexed categories provide a more aesthetic way to deal with small-definable categories, which makes it natural to work with them. This is done by defining an *indexing* relative to our universe of sets; every small-definable category will give rise to such an indexed category. An *indexed category* is a family of categories consisting of a particular category for each object in the universe. Where in ZFC/S we had to check whether categories were small-definable at first, the indexing machinery will now do this for us. Depending on our prior set theory, different types of indexings are required. Once indexed, the categories that were small-definable will contain the well-behaved objects, such that our Replacement axiom can deal with them. It turns out that small-definability restrictions arose because, for example in ZFC/S, we implicitly used the wrong type of indexing such that our (weakened) Replacement axiom could not deal with this. Indexing corrects this 'mistake' so that categories will again naturally fit the properties of, for example, completeness, that we want.

2.2.3 Class theories

The first class-set theory that we discuss has been proposed as a way of overcoming some of the previously encountered difficulties, and is an extension of Ackermann set theory described by (Muller, 2001). Class-set theories can introduce a notion of largeness by the incorporation of classes as new types of objects. Unlike the class theories that we describe next, classes in Ackermann set theory are allowed to be elements of other classes. Like ZFC/S, Ackermann set theory is a conservative extension of ZFC (Shulman, 2008). (Muller, 2001) introduces an extension of Ackermann set theory that he intends to not be susceptible to at least conceptual limitations to founding category theory. First, the axioms of A (a slight adaptation of Ackermann set theory) are as follows:

1. *Axiom of Extensionality*: classes are identical if they have the same members.
2. *Completeness*: the class of all sets V is complete, i.e. all classes contained or included in sets are sets.
3. *Class Separation*: for any predicate $\varphi(\cdot, \mathcal{Y})$, where $\mathcal{Y}$ stands for any finite number of class-parameters, and for every class $\mathcal{Z}$, there exists a class $\mathcal{A}$ which contains exactly those members of $\mathcal{Z}$ for which $\varphi(\cdot, \mathcal{Y})$ holds.
4. *Set Existence*: for any *safe* predicate $\psi(Y, \cdot)$ where Y stands for any finite number of set-parameters, if the only classes for which it holds are sets, then these sets form a set.

A predicate ψ is called *safe* if it only contains set-parameters and the class V of all sets (which is *unsharply delineated*, as Ackermann calls it) does not occur in it. (Muller, 2001) regards this notion as better than, for example, requirements of predicativity or stratification (which we discuss later), as it still allows us to pick elements from the whole class V, only with minor limitations. Predicativity and stratification, on the other hand, actively prevent us from picking elements from V if they already occur (in a particular way) in the predicate ψ. Now (Muller, 2001) extends A by adding *Regularity* (R) and *Choice* (C) for sets, resulting in the theory **ARC**. Muller argues that ARC is philosophically more satisfying than the previous set theories that we discussed, as it does not suffer from the objections that a universe much larger than needed is defined, and still unrestricted categories are not supported. However, Ackermann set theory as a foundation of category theory turns out to have similar problems to ZFC/S and is in fact strictly weaker in its statements about classes (Shulman, 2008). Recall that for ZFC/S we had to weaken the Replacement axiom and we ended up with small-definability restrictions. As ARC merely adds a Regularity and Choice axiom, we see that this will not avoid the same restrictions.

Last, we mention two theories that also introduce classes as a new type of object, but more quickly turn out to have limitations when it comes to constructing categorical objects. Nonetheless, it is useful to see where things go wrong. First, *von Neumann-Bernays-Gödel* set theory (NBG) includes axioms to deal with classes, where only sets can be elements of sets or classes. NBG is conservative over ZFC. It allows familiar axioms such as Extensionality and Foundation to be applied to classes, and unlike in ZFC, we can quantify over classes. This facilitates dealing with large categories, which we can clearly now incorporate as classes. However, it can be proven that the Class Comprehension scheme (*if* φ *is a property, then there exists a class* $Y = \{x : \varphi(x)\}$) holds in NBG only for a limited collection of formulas φ, namely those formulas which quantify only over set variables (Shulman, 2008). This turns out to be problematic, as certain basic mathematical principles cannot be verified. For example, while ZFC can prove the induction principle for all statements $\varphi(n)$, NBG can only do this if φ does not quantify over classes. Consequently, there may be statements about large categories for which a natural proof would use induction, but which cannot be done in NBG. Additionally, we cannot form every (functor) category in NBG (Ernst, 2017), and hence requirements (R1) and (R2) cannot be met. We can only construct **Cat** or **Set** for *small* categories and sets. Similarly, what we can do is form the functor category **[A, B]** for *small* **A** and *large* **B**. Namely, take a functor $F : \mathbf{A} \to \mathbf{B}$, then Replacement will ensure that the functor is a set. Then **[A,B]** will consist of all such functors, and will hence be a class. However, suppose that **A** is also large. Then **A** is not a set, so that Replacement cannot be applied and the existence of **[A,B]** is not guaranteed.

An adaptation of NBG results in *Morse-Kelly set theory* (MK), which supports full Class Comprehension. Indeed, this means that the notion of mathematical induction is fully available again, and statements about large categories are more easily proved. Unlike NBG, MK is strictly stronger than ZFC. Furthermore, even though we can do more with MK for category theory than with NBG, arbitrary functor categories are still unattainable. To be able to define the collection of functions between classes seems a reasonable request, and hence this restriction is conceptually dissatisfying, certainly for Feferman's requirements (R). Additionally, not being able to have **[A,B]** for large **A** and **B** prevents results obtained from a perspective of higher category theory. We will not mention details here, but we might want to say something about the cohomology theories that satisfy the Eilenberg-Steenrod axioms, or about the 2-category of Grothendieck toposes (see (González, 2018) for an application). This will not be possible without arbitrary functor categories. Even though most results in ordinary mathematics will not depend on these constructions, we see that there do exist mathematical results depending on such categories, and hence we would like to accommodate them.

2.2.4 New Foundations

The set-theoretical foundations we have seen so far allow for the construction of many desired large categories, but we are still not equipped to accommodate truly unrestricted categories. With this in mind, Feferman proposes a set-theoretical system based in Quine's New Foundations (NF) that should allow for all unlimited categorical constructions (Feferman, 2011; Feferman, 2013). This set theory uses the concept of type to enforce a well-foundedness in the description of sets that ZFC achieves in the constitution of sets with its restricted Comprehension axiom. The language of NF is first-order, with a single sort variable (which Feferman refers to as 'classes') and two relation symbols $\in$ and $=$. Extending NF, Feferman constructs the system S* that, in order to impose a sense of largeness, adds the class of all sets as a constant symbol, as we will see. To build this system, we first define the system NFU, that allows for multiple urelements in NF by using a weakened form of Extensionality (Ext'). This results in the following axiomatization for NFU:

(Ext') $(\exists X(X \in A) \wedge \forall X(X \in A \leftrightarrow X \in B)) \rightarrow A = B$

(SCA) $\exists A \forall X(X \in A \leftrightarrow \varphi)$ where φ is stratified and the variable A does not occur in φ. (SCA) consists of all universal closures of this axiom.

(SCA) describes a Stratified Comprehension axiom: this is what avoids the usual paradoxes. To see this we need to consider the following definition.

Definition 2. Let φ be a sentence of first-order logic that contains no relation other than $=, \in$. Then φ is *stratified* if it is possible to assign a nonnegative integer (a *type*, or $tp(x)$) to each variable x in φ such that:

1. Each variable has the same type wherever it appears.
2. In each atomic sentence $x = y$ in φ, $tp(x) = tp(y)$.
3. In each atomic sentence $x \in y$ in φ, $tp(y) = tp(x) + 1$.

Thus, (SCA) allows for a universal class of all sets V (and we have $V \in V$). (Feferman, 2013) applies the following extensions to make the framework suitable as a foundation for category theory. First, he includes a pairing axiom (P) $((X_1, X_2) = (Y_1, Y_2) \rightarrow X_1 = Y_1 \wedge X_2 = Y_2)$ that results in the theory NFU(P). The formation of arbitrary categories of a given kind (such as the category of all groups, or all topological spaces) and that of arbitrary functor categories is already possible in NFU(P). For example, using the implementation of categories in NFU(P) as described in the beginning of this chapter, one can define the class of all categories (as the axiomatic definition of a category can be described in a stratified way) and so construct the category with this class as its objects. This results in the category of all categories.

The stratification condition, however, results in some type-shifting problems when trying to define arbitrary Cartesian products and functions on equivalence classes. To account for these problems in a better way, (Feferman, 2011) constructed the system S^*, which is an extension of NFU(P) and of ZFC. The language of S^* introduces set variables next to class variables, and a constant symbol V_0 that represents the class of all sets. Then S^* extends the axioms (Ext') and (SCA) with the following axioms.

Sets and classes	$\forall x \exists X(x = X)$ $X \in V_0 \leftrightarrow \exists x(x = X)$ $X \in x \rightarrow X \in V_0$
Empty set	$\exists! z \forall y(y \notin z)$
Operations on sets	$\{x, y\} \in V_0$ $\bigcup x \in V_0$ $\mathcal{P}(x) \in V_0$ $(x, y) = \{\{x\}, \{x, y\}\}$
Infinite set	$\exists a[\exists z(z \in a \wedge \forall y(y \notin z)) \wedge \forall x(x \in a \rightarrow x \cup \{x\} \in a)]$
Replacement	$\forall x, y_1, y_2[\psi(x, y_1) \wedge \psi(x, y_2) \rightarrow y_1 = y_2] \rightarrow$ $\forall a \exists b \forall y[y \in b \leftrightarrow \exists x(x \in a \wedge \psi(x, y))]$
Foundation	$\exists x \psi(x) \rightarrow \exists x[\psi(x) \wedge \forall y(y \in x \rightarrow \neg\psi(y))]$, where $\psi(x)$ is any L^* formula not containing y
Universal Choice	$\exists C \left[\begin{array}{l} \forall X, Y_1, Y_2((X, Y_1) \in C \wedge (X, Y_2) \in C \rightarrow Y_1 = Y_2) \\ \wedge \forall X(\exists Y(Y \in X) \rightarrow \exists Y(Y \in X \wedge (X, \{Y\}) \in C)) \end{array} \right]$

Feferman also notes that MK with *Universal Choice* is interpretable in S^*. This system can deal with functions on equivalence classes in a stratified way, by working with representatives of equivalence classes. However, the problems with Cartesian products remain. Since the notion of Cartesian product is important to the construction of many structures and for checking properties like completeness, we do not want to give up on it.

This problem still has not been solved, although we mentioned before that it has to do with the inconsistency of Feferman's requirement (R3) with (R1). Since in the system $S^{(*)}$ unlimited collections are objects of the theory, it is not surprising that we cannot carry over all familiar laws of mathematics. However, there is no clear selection criterion for which part of (R3) should be preserved. In recent developments, (Thomas, 2018) comments on (McLarty, 1992), who showed that **Set** and **Cat** are not Cartesian closed in NF-style set theories. Thomas points out that these categories do still have a property approximating Cartesian closure. However, because of numerous limitations introduced by their stratification restrictions, Thomas still concludes that NF-style set theories are not a good foundation for category theory.

A final variant on the use of New Foundations in a foundational theory for category theory is discussed by (Enayat, Gorbow, and McKenzie, 2017). For this we need the following notion.

Definition 3. A set A is *cantorian* if there is a bijection from A to $\{\{a\} | a \in A\}$. A is *strongly cantorian* if there is a bijection from A to $\{\{a\} | a \in A\}$ that maps each $a \in A$ to $\{a\}$.

The stratification property of (SCA) makes sure that you cannot show that every cantorian set is strongly cantorian. Restricting to strongly cantorian classes solves the type-shifting problems (Feferman, 2011). Still, the collection of strongly cantorian classes does not form a class, so that (R1) cannot be satisfied. However, let us now define NFUA as NFU + Infinity + Choice + "every cantorian set is strongly cantorian". Then NFUA satisfies (R1) and (R2), and (R3) can be satisfied in a restricted way. (Enayat, Gorbow, and McKenzie, 2017) show that ZMC/S may be interpreted inside NFUA. This results in a better accommodation of (R3), while the requirements (S) are also fully solvable within this interpretation.

In the end, it seems to come down to the question whether the accommodation of unrestricted categorical notions, besides being conceptually pleasing, has other benefits for category theory. (Enayat, Gorbow, and McKenzie, 2017) suggest this might be the case, by showing that the category **Rel** has (co)products, and **Set** coproducts, that are indexed by the set of all singletons. In ZFC-like systems, however, the locally small versions of **Rel** and **Set** do not yield this result for the set (or class) of all small singletons. Thus, NFU-based category theory seems to provide results that ZFC-based category theory cannot. (Enayat, Gorbow, and McKenzie, 2017) conclude from this that the fully unrestricted (R1)-categories may have categorical relevance other than philosophical, in that they can induce relevant mathematical results. Hence, it may still be useful to continue to explore systems like NFU, NFUA and S*.

2.3 Discussion

Summing up. We have described the set theories that are most used as a foundation for category theory, and we analyzed their strengths and limitations in doing so. Neither ZFC nor its non–well-founded variants can deal well with large categories, calling for the implementation of more creative methods to allow for categorical results relying on a distinction between large and small sets. In pursuance of such a distinction, we looked at ways to capture the universe of sets we work in, so that objects living outside this universe can be regarded as large. Taking the universe of sets to be defined by V_κ for κ inaccessible (or a worldly cardinal) is one way; instantiating a reflection principle with a 'universe'-constant $\mathbb{S}$ is another. Both methods allow for additional incorporation of bigger inaccessibles and constants, respectively. By letting $\mathbb{S} = V_\kappa$ for κ inaccessible, the combination of these methods results in a powerful candidate for a foundation for category theory. Class theories like NBG and MK, on the other hand, incorporate explicit axioms with which to manipulate large objects. Although this method is effective, it lacks the power of the aforementioned approaches. We should remain aware, though, that large cardinal assumptions and reflection principles, the latter without strengthening ZFC, bring along some technical inconveniences. Finally, the Stratified Comprehension Axiom of NFU allows for unlimited constructions that are independent from universe restrictions. The introduction of a distinction between large and small objects on top of this axiom, however, helps to partly solve type-shifting problems, although Cartesian closedness of **Set** can not yet be attained. Still, systems like S* and NFUA have been shown to make some (recent) strides in satisfying Feferman-like requirements for a foundation of category theory.

Take away. This chapter has brought to light several things that we can take with us. First, it should be clear that it remains debatable what a foundation for category theory should precisely be. Categorical results from a set-theoretical perspective are dependent on size constraints; hence, each set-theoretic system provides a somewhat different account of category theory. Setting requirements like (R) and (S) may give us something to hold on to when trying to evaluate set-theoretic foundations, but they also tend to oversimplify the situation. Such requirements namely have to be rather general, as justifications for strict conditions are absent. Consider (R3), which demands that the 'usual basic mathematical structures' exist and we can carry out the 'usual set-theoretical operations'. This allowed for the finding of a contradiction with (R1) in (Ernst, 2017), because the general nature of both requirements lead to a specific instance where they clashed. Instead of taking such requirements very literally when investigating foundations for category theory, we argue that they should simply be taken as a general set of guiding principles. The relevance of the quest of finding 'the one foundation' for category theory is thereby weakened, as different requirements are desirable in different contexts. Depending on what mathematical or philosophical glasses one puts on, it is relevant to include or exclude requirements such as the accommodation of unlimited categories. In this thesis, we deliberately take various foundations into account so that a colourful picture of set-theoretic category theory can be sketched.

Furthermore, the results from different foundations tell us that for the practical purposes of a mathematician, foundations such as ZFC/S, ZFC+I and class theories suffice. Issues like universe-juggling and having to carry out small-definability checks, however, make foundations philosophically less attractive. In picking a foundation, then, a trade-off must be made between simplicity or elegance of the system, and its strength. Importantly, this has shown us that we cannot simply take the term 'set-theoretic foundation for category theory' to denote one thing. Set theories vary substantially in their axioms and how these affect category theory. For example, more tolerant Comprehension axioms in class theories allow for the construction of more complex categories - here, Stratified Comprehension requires the least number of restrictions by dealing with the syntax of a class-describing formula. On the other hand, we saw that Foundation has little to do with the construction of category theory. The heart of the matter seems to really lie in sufficiently strong axioms that allow for a distinction between familiar small sets and more external large sets.

What is next? We should keep in mind the set-theoretical systems we discussed in this chapter and which properties affect the construction of category theory the most. They will come back in Chapter 4, where we base our argument concerning the distinction between set-theoretical and categorical thinking on these results. First, however, we will examine purely categorical implementations of a foundation for mathematics in Chapter 3. This will tell us more explicitly how mathematics based in these foundations is developed, and will allow us later to see where they differ from set-theoretical foundations. It will additionally be informative to see whether this comparison is consistent with the behaviour of category theory within set theory, as we analyzed it in the current chapter.

Chapter 3

Categorical foundations

Following on the previous chapter, we will here discuss purely categorical implementations of a foundation for mathematics. That is, we will look at systems with axioms that take categorical notions as primitive, such as arrows and a composition operation. In this chapter, we discuss four attempts at founding mathematics from a categorical system. Some systems rather explicitly incorporate set theory in a categorical way, for example by finding analogues for set-theoretical relations such as membership, or by resembling axioms from set-theoretical foundations. For other systems, the possibility of including set-theoretical notions is only of secondary relevance to the system itself. We will see that several attempts have been quite successful, and that there is still ongoing research regarding the development of some more intricate systems. Note that we also take EM as a categorical foundation: it is recommended as a suitable structuralist foundation for mathematics in (Landry, 2013). Although EM perhaps seems an unlikely candidate for a foundation for mathematics at first sight, the argument supporting the view that it is relies on the idea that foundations for mathematics come with different purposes. This will be touched on briefly in the next section, and we will refine this argument in Chapter 4. Before discussing the other systems, we will elaborate on some of the difficulties which exist in determining the exact requirements on a categorical foundation.

3.1 Requirements for categorical foundations

In order to get a better understanding of different categorical foundations, we would like to discuss some possible requirements for categorical systems. We will often remain neutral and limit ourselves to outlining the issues we treat, but it is helpful to be aware of the main proposed requirements.

First, an obvious requirement and analogue of which we treated for set-theoretical foundations of category theory, is that a categorical foundation should allow for the 'usual mathematical constructions'. That is, we should be able to develop with our foundation the mathematical fields that are relevant for practical purposes. Similarly to its set-theoretical analogue, however, this will remain a rather general condition that should not be taken too strictly.

Second, the role of set-theoretical notions in a categorical foundation is not obvious. Various categorical foundational systems and particular axioms are explicitly inspired or motivated by set-theoretical analogues. The question then arises whether these statements are inherently set-theoretical (making the justification of categorical foundations 'dependent' on set theory), or whether they stand independently as generally desired properties for mathematics. Related to this is the position one takes regarding set theory as a foundation: if we take set theory (in general, or a particular foundation like ZFC) to constitute the 'official' foundation for mathematics, we may care about shaping categorical foundations similarly to set-theoretical ones. In this case, it is also relevant to consider strengthening a categorical system in order that it can obtain a strength equal to ZFC — even though it is arguably not the case that every axiom of ZFC is necessary to get a powerful theory that supports a great deal of mathematics. Although we do not concern ourselves in this thesis with the justification of individual foundations, we do mention that we do not regard the possible conceptual dependency of categorical foundations on set theory reason enough to discard such set-theory–like systems. Note that all such theories are formulated entirely in categorical terms, and they should provide an informative perspective on any possible encompassing distinction between set-theoretical thinking and categorical thinking that we find in this thesis. Hence, they are useful for us to include in our discussion. Concerning the comparison of the strength of categorical systems to set theories, we regard this as a useful method that tells us more about how to relate theories, but we will not maintain anything like it as a strict requirement.

Furthermore, because the set-theoretical foundations for mathematics of the previous chapter seemed to pursue an accomodation of larger and larger — even unlimited — categories, we mention this requirement here, too. We note that explicit size is of less importance in categorical foundations than in set-theoretical ones, as sometimes there is simply no way to evaluate this from the categorical theory itself. When it is possible, however, the categorical approach does not magically provide the means to incorporate very large objects. In fact, the category of sets (as axiomatized by ETCS) is similar to ZFC regarding the way it deals with size. For the category of categories, it remains unclear whether it can be a truly unrestricted category (incorporating *all* categories). More often than not, size restrictions are not explicitly discussed for categorical systems, because these systems appear more flexible regarding the specific models of the theory, and this would narrow down the possible models too strongly. However, as the matter is generally not discussed for foundational systems, we recommend that it should be better investigated in the future which theories allow for unlimited categories and why. Naturally, we do not take the incorporation for unlimited categories as a requirement for our categorical foundations.

Finally, it should be noted that authors differ on what they think may be called a foundation for mathematics, regarding the structure of the axioms. The axioms of set-theoretical systems are generally assertoric: ZFC, for example, asserts the existence of an empty set, an infinite set, and from these sets the existence of infinitely many sets can be guaranteed. This way, a universe of sets can be constructed explicitly from the axioms. It is sometimes thought (for example, by (Hellman, 2003)) that such existence axioms are necessary for a system that aspires to be a foundation of mathematics, and so also for category theory. The categorical systems that we will discuss, however, differ in how much their axioms assert. The category of sets as axiomatized by ETCS will prove to ensure the existence of many objects, like standard set theory itself. The axioms for algebraic set theory and the category of categories, on the other hand, secure the existence of fewer constructions. This leads to results more often being derived as hypotheticals instead. Note that EM, in fact, does not provide any assertoric axioms at all. The question, then, is whether the fact that categorical foundations do not allow us to explicitly construct all objects that we are interested in is really problematic for their role as a foundation of mathematics. In (Landry, 2011) and (Landry, 2013), Landry has argued that, depending on our objectives, we can indeed do without any assertoric axioms. Namely, it is argued that foundations with and without assertoric axioms carry different roles as a foundation for mathematics. This justifies the different formulations of foundations which differ in their purposes. In Chapter 4, we will come back to Landry's argument concerning the purpose of foundations more elaborately, and we will discuss how it relates to our findings. For now, recall that we aim to take into account a variety of factors in this thesis that may turn out to be important in the distinction between the set-theoretical and categorical approach to mathematics. Hence, we include both assertory and non-assertory theories in our discussion of categorical foundations.

Now that we are aware of various possible requirements for categorical foundations (and their complications), we are ready to embark on the discussion of the first categorical foundational system.

3.2 The category of sets

We here consider a categorical axiomatization of set theory. Indeed, since set theory has proven to be a successful foundation of mathematics, it makes sense to construct a categorical analogue of it. In particular, we will look at an axiomatization of the category of sets that is intended to be an analogue of ZFC. This theory, called the Elementary Theory of the Category of Sets (ETCS), was developed by Lawvere in 1964. It arose from the idea that set theory should not take the element relation as primitive, but rather a relation that captures isomorphism-invariant structure. Thus, the only properties that the 'elements' of ETCS-sets have are those relating them to other elements. This motivates the fact that quantifiers in ETCS range over mappings only. Note that composition is taken as a primitive relation for pairs of mappings f, g such that the codomain of f is the domain of g. Before we list the axioms of ETCS, note that we need a notion of set-theoretical elementhood that category theory does not possess. We define the following.

Definition 4. (a) x is an element of A, denoted $x \in A$, if and only if it is an arrow $x : 1 \to A$.

(b) a is a subset of A if and only if a is a monomorphism with codomain A.

This makes sense, as every subobject of an object A in **Set** is represented by a set in $\mathcal{P}(A)$. Now consider the axioms of ETCS (adapted from (Lawvere, 2005)), which describe the properties of the category **Set**.

1. The usual axioms for an abstract category (as given in the introductory chapter of this thesis).

2. All finite roots exist. Equivalently, it can be assumed that there exist a terminal object 1 and initial object 0; for every pair of objects A, B the coproduct $A+B$ exists; and for any pair of maps $f, g : A \to B$ the equalizer $k : E \to A$ and coequalizer $q : B \to C$ exist.

3. For any pair of objects A, B, the exponential B^A exists.

4. *('Axiom of Infinity')* There exists a natural numbers object N.

5. *(Well-pointedness)* 1 is a generator. This means that if $f, g : A \to B$, then $f \neq g$ implies $\exists a(a \in A \wedge af \neq ag)$.

6. *(Axiom of Choice)* If the domain of f has elements, then there exists a g such that $fgf = f$.

7. Every object other than 0 has elements.

8. Every element of a coproduct is a member of one of the injections.

9. There is an object with more than one element.

The axioms above describe a particular kind of topos: a well-pointed topos with an NNO and satisfying Choice. Well-pointedness and Choice separately ensure that our topos is Boolean, and that ETCS is in fact equivalent to Bounded Zermelo set theory with Choice (BZC) (Shulman, 2008). BZC is exactly ZFC without Replacement and with bounded Separation. Thus, if the axioms of ETCS are intended to resemble ZFC, it is reasonable to consider strengthening ETCS with a categorical Replacement axiom. As Replacement implies full Separation, this would suffice. Obtaining the strength of ZFC may be attractive, as we lack the tools to prove induction for any formula and perform transfinite constructions on functors in BZC (Shulman, 2008). For example, we cannot iterate the power set functor which sends n to $\mathcal{P}^n(\omega)$ even ω times, as without Replacement we can have $V_{\omega \cdot 2}$ as our universe, but $|\mathcal{P}^\omega \omega| = \beth_\omega \notin V_{\omega \cdot 2}$.

Shulman makes a distinction between two versions of categorical Replacement that have been put forward. The first one comes from (McLarty, 2004) and is the following:

> Let X be an object and $\varphi(y, Z)$ a definable property such that for any 'element' $x : 1 \to X$ there exists an object S_x unique up to isomorphism with $\varphi(x, S_x)$. Then there exists a morphism $S \to X$ such that for any x there is a pullback square
>
> $$\begin{array}{ccc} S_x & \longrightarrow & S \\ \big\downarrow & & \big\downarrow \\ 1 & \xrightarrow[x]{} & X \end{array} \tag{3.1}$$

Equivalent versions have been described by other authors (Osius, 1974; Lawvere, 2005). Adding this axiom (R) indeed leads to equiconsistence of ETCS+R with ZFC. Intuitively, it is a categorical way of saying that S is the disjoint union of each S_x, and hence S is the image of the fibers $f^{-1}(x)$. This axiom is strongly dependent on Well-pointedness, as we are representing elements of X as arrows.

We may not be willing to let Replacement be so dependent on Well-pointedness, however, as none of the other axioms of ETCS require this. Therefore, another possibility is to add what has been called the *categorical axiom of iterative replacement* in (Taylor, 1999). This axiom uses a categorical definition of ordinals (as well-founded extensional coalgebras) and indexed families of functor iterates are considered. The axiom then states that there exists a family for each ordinal α and α-indexed functor T such that the former is isomorphic to the colimit of T. This essentially says that we can have transfinite constructions on functors (Shulman, 2008). The effects of adding this axiom to ETCS are largely unknown.

It is good to keep in mind that for the purposes of reconstructing mathematics, we do not have to assume a Replacement axiom or, in fact, many other axioms of ETCS. Any elementary topos with an NNO will do for reconstructing a good deal of mathematics already. Namely, in any elementary topos we can develop an internal logic, in order to model a set-like theory. Here, for any formula $\varphi(x)$, we can define a subobject $[\![\varphi]\!] \rightarrowtail A$; we can think of this subobject as a 'set' $\{x \in A : \varphi(x)\}$. Taking subobjects to be predicates, we can develop a notion of truth, falsity and the usual logical connectives. Hence, this internal logic allows us to reason with objects in our topos as if they are abstract sets.

However, this internal logic is in general intuitionistic (as we lack axioms such as Well-pointedness and Choice). In an intuitionistic setting, categorical Replacement (like the version from (McLarty, 2004)) no longer implies unbounded Separation, and we lose other results such as transfinite induction the way we know it (Shulman, 2008). Namely, this categorical Replacement was heavily dependent on Well-pointedness. Additionally, when constructing mathematics in an intuitionistic topos, translating results obtained from classical logic is not always obvious, so care must be taken in doing this. However, there do exist structuralist foundations for intuitionistic mathematics, such as CETCS as described in (Palmgren, 2012). CETCS even includes an intuitionistic version of Well-pointedness, and to strengthen Replacement there is an axiom of Collection. Thus, in the presence of stronger axioms we might obtain a workable intuitionistic foundation. Still, we will here restrict our attention to the better-known setting of classical logic.

3.3 The category of classes

A different approach is called *algebraic set theory* (AST). It uses the language of category theory to axiomatize a category $\mathcal{C}$. The axioms for $\mathcal{C}$ resemble those for an elementary topos, and motivate the interpretation of $\mathcal{C}$ as a category of classes. This idea starts from the fact that algebraic properties are strongly related to set-theoretical ones, which allows for a characterization of set theory based on arrows, instead of membership. The axioms imposed on $\mathcal{C}$ ensure that $\mathcal{C}$ is a Heyting pretopos; hence, $\mathcal{C}$ will possess an internal logic. Additional axioms in AST will assert the existence of a subcategory $\mathcal{S}_{\mathcal{C}}$ of $\mathcal{C}$, which will give rise to a notion of set. The internal logic of the categories, then, will allow for ordinary set-theoretic reasoning. We will furthermore incorporate a power class axiom, that also gives rise to the notion of *universes*. Let us begin by defining our *Heyting category* $\mathcal{C}$. The following conditions should hold, as described in (Awodey, 2008):

1. $\mathcal{C}$ has all finite limits, including a terminal class 1.

2. $\mathcal{C}$ has all finite coproducts, including an initial class 0.

3. $\mathcal{C}$ has kernel quotients. This means that for every arrow $f : C \to D$, the pullback of f against itself $k_1, k_2 \colon K \to C$ (called the kernel pair) has a coequalizer $q \colon C \to Q$.

$$\begin{array}{ccccc} K & \overset{k_1}{\underset{k_2}{\rightrightarrows}} & C & \xrightarrow{q} & Q \\ & & {\scriptstyle f}\downarrow & & \\ & & D & & \end{array} \tag{3.2}$$

4. $\mathcal{C}$ has dual images. To see what this entails, note that for any arrow $f : C \to D$, we have a pullback functor $f^* \colon Sub(D) \to Sub(C)$ (with $Sub(D)$ the category of subobjects of D). Then the axiom ensures that f^* has a right adjoint $f_* \colon Sub(C) \to Sub(D)$, so that it holds that, for any $U \leqslant C$ and $V \leqslant D$:
$$f^* V \leqslant U \text{ iff } V \leqslant f_* U$$

These requirements allow for our category to be described as 'regular', which leads to admittance of first-order logic. In particular, the axioms imply that for any arrow f as in (4), f^* will also have a left adjoint. These adjoints can represent existential and universal quantifiers and, like before, we can let operations on subobjects correspond to logical connectives. Additionally, it can be seen that for each object C, $Sub(C)$ is a Heyting algebra, so that we also have the Heyting implication '$\Rightarrow$'. In short, $\mathcal{C}$ models intuitionistic, first-order logic with equality.

At this point, we would like to obtain a notion of smallness, in order to denote which objects are sets. This is done by axiomatizing a collection $\mathcal{S}$ containing the *small maps* of $\mathcal{C}$. The intuition as in (Joyal, Izak Moerdijk, Ieke Moerdijk, et al., 1995) is that an arrow $f : A \to B$ is small if each fiber $f^{-1}(b) \subseteq A$ is small. Thus, the axioms capture some properties of maps with small fibers, and they are stated as follows.

1. $\mathcal{S} \hookrightarrow \mathcal{C}$ is a subcategory and has the same objects as $\mathcal{C}$.

2. The pullback of a small map along any map is small.

3. Every monomorphism $m : C \rightarrowtail D$ is small.

4. If $f \circ e$ is small and e is a regular epimorphism, then f is small.

5. If $f : A \to C$ and $g : B \to C$ are small, then so is the copair $[f, g] : A + B \to C$.

Note that the first requirement ensures that every identity map and every composite of two small maps are small. The second condition expresses that the smallness property is indeed one of the fibers of the map, while the other conditions ensure that the small maps are closed under basic operations on classes. Now, it makes sense that we should call an object A small if, for some $f : A \to X$ and $x \in X$, we have a small fiber such that $f^{-1}(x) = A$. Therefore, we have the following definition:

Definition 5. $X \in \mathcal{C}$ is small if $X \to 1$ is a small map.

We furthermore let a relation $R \rightarrowtail C \times D$ be small if its second projection $R \rightarrowtail C \times D \to D$ is a small map, and a subclass $A \rightarrowtail C$ is small if the class A is small. We can now refer to the small classes as *sets*. With the axioms that follow we formalize that each class has a powerclass representing its sub*sets*. A powerclass comes with a membership and subset relation, and with the notion of powerset we can construct 'cumulative hierarchies' of sets built on any class. First, consider the powerclass axioms.

(P1) Every class C has a *powerclass*: an object $\mathcal{P}C$ with a small relation $\in_C \rightarrowtail C \times \mathcal{P}C$ such that, for any class X and for any small relation $R \rightarrowtail C \times X$ there is a unique arrow $\rho : X \to \mathcal{P}C$ such that the following is a pullback diagram:

$$\begin{array}{ccc} R & \longrightarrow & \in_C \\ \downarrow & & \downarrow \\ C \times X & \xrightarrow[1_C \times \rho]{} & C \times \mathcal{P}C \end{array} \tag{3.3}$$

(P2) The internal subset relation $\subseteq_C \rightarrowtail \mathcal{P}C \times \mathcal{P}C$ is a small relation.

Hence, the powerclass operation classifies the small subobjects of a class, and $\mathcal{P}C$ is small if C is. The first condition tells us that any small relation cRx can be written as $c \in_C \rho(x)$ for the unique $\rho : X \to \mathcal{P}C$ such that $\rho(x) = Rx$. The second requirement establishes the smallness of the powerclass $\{x : x \subseteq_C y\}$ of any set y. Now the powerclass operation gives rise to the notion of a *free $\mathcal{P}$-algebra*, for which we use the following definition.

Definition 6. A *ZF-algebra* is a partially ordered class A that has all small joins and is equipped with a successor operation $s : A \to A$.

The free ZF-algebras have been shown to coincide with the free $\mathcal{P}$-algebras (Awodey, 2008). A $\mathcal{P}$-algebra, then, is a pair (A, α) with object A and map $\alpha : \mathcal{P}A \to A$ (where $\mathcal{P}$ is the powerclass operation extended to a functor). A *free* $\mathcal{P}$-algebra FC on a class C, furthermore, is a $\mathcal{P}$-algebra $f : \mathcal{P}(FC) \to FC$ with a map $\mu : C \to FC$ that has a universal mapping property. Lambek has proven that f and μ give rise to an isomorphism $\mathcal{P}FC + C \cong FC$. With this, elements of FC can be regarded as, either, a subset of FC, or an element of C. Thus, the intuition is that FC can be thought of as containing all 'powerset-levels' of elements.

A free $\mathcal{P}$-algebra for C is hence similar to the cumulative hierarchy $V(C)$ built on C. Then we let $V(C)$ for a given C be the *universe* on C, and we assume that for each class, such a universe exists. That means that $V(0)$ exists, and we call it the *(initial) universe* V, representing the class of all sets. Within this universe we can have a membership relation by defining $a \in b$ if and only if $s(a) \leqslant b$ (Awodey, 2008). Similarly in terms of ZF-algebras, V is the free ZF-algebra on 0, with the successor operation $s : a \to \{a\}$ and taking unions as joins. Then $(V, \in)$ models IZF. Here, it is possible to assume an axiom of Collection to strengthen the axiom of Replacement. Furthermore, we can assert that $\mathcal{C}$ is Boolean, for example by adding axioms of Well-pointedness and Choice. Then clearly $(V(0), \in)$ will model classical ZF. We could finally assert an axiom of Infinity. This can be done by asserting that there exists a small object I

together with a monomorphism $1 + I \rightarrowtail I$, with the intuition that I is Dedekind infinite (Simpson, 1999). This is equivalent to asserting that $\mathcal{C}$ has a small NNO.

Any model of NBG is also a model of AST if we assert that smallness corresponds to being a set (Shulman, 2008). The other way around does not work, however, as AST allows for the existence of larger classes than NBG. By asserting Well-pointedness, Choice and Infinity for $\mathcal{C}$, the subcategory $\mathcal{S}$ of sets becomes a WPTNC like ETCS was, and $\mathcal{S}$ will also satisfy the categorical Replacement we saw before. Working in AST with these additional axioms closely corresponds to working in NBG. We can furthermore assume that $\mathcal{C}$ is a topos, which gives us equivalence to BZC+I. Finally, if we assume a Replacement axiom for $\mathcal{C}$ as well, we obtain equivalence to ZFC+I (Shulman, 2008).

3.4 The category of categories

An intuitive idea for a categorical foundation of mathematics, furthermore, is to axiomatize the category of categories. We will refer to the theory that does this as the Category of Categories and Adjoint Functors (CCAF), a term introduced by (Lawvere, 1966). We call the intended model *CAT* of this theory a metacategory, which has categories as its objects, and functors as its arrows. This term is mainly intended to avoid confusion with the terms 'object' and 'arrow', which are only used for the categories inside *CAT*. CCAF is presented in a two-sorted first-order logic, with function symbols dom, cod, 1 (identity) and $\circ$ (composition) satisfying the usual requirements. Furthermore, categories $\varnothing, \mathbf{1}, \mathbf{2}, \mathbf{3}, \mathbf{E}$ are introduced, as well as functors $c, d, \alpha, \beta, \gamma, e_1, e_2$, and several sorted operators. Defining these particular categories and functors will allow for a notion of 'object' and 'arrow' of a particular category in *CAT*. We present the axioms, adapted from an improvement of the original theory by (McLarty, 1991), and then clarify the role of the constant categories and functors.

1. *CAT* has all finite limits and colimits, with initial category $\varnothing$ and terminal category $\mathbf{1}$, and it is Cartesian closed. $\varnothing$ is not isomorphic to $\mathbf{1}$.

2. The constants $\mathbf{2}$, d and c are such that $d : \mathbf{1} \to \mathbf{2}$ and $c : \mathbf{1} \to \mathbf{2}$.

3. The functor $\binom{d}{c} : \mathbf{1} + \mathbf{1} \to \mathbf{2}$ is not an epimorphism. Equivalently, if we define a category $\mathbf{E}$ and arrows e_1, e_2 as the pushout of $\binom{d}{c}$ along itself, the axiom amounts to requiring $e_1 \neq e_2$. The diagram below depicts $\mathbf{E}$.

$$d \underset{e_2}{\overset{e_1}{\rightrightarrows}} c \tag{3.4}$$

4. Here we introduce the constants $\mathbf{3}$ and α, β and γ. This is a pushout:

$$\begin{array}{ccc} \mathbf{1} & \xrightarrow{c} & \mathbf{2} \\ {\scriptstyle d}\downarrow & & \downarrow{\scriptstyle \alpha} \\ \mathbf{2} & \xrightarrow[\beta]{} & \mathbf{3} \end{array} \tag{3.5}$$

   Additionally $\gamma : \mathbf{2} \to \mathbf{3}$ with $\gamma \circ d = \alpha \circ d$ and $\gamma \circ c = \beta \circ c$.

5. For every first-order expression $R(f, g)$, with f, g variables of functor type, it holds: if $R(f, g)$ defines a functorial relation from arrows of **A** to those of **B**, then there is a functor **F**: **A** → **B** such that $R(f, g)$ iff $\mathbf{F}f = g$.

6. Some category has a non-identity isomorphic arrow.

7. There is a dual operator $\_^{op}$ that preserves identity functors, domains, codomains and composites of functors, and, for all categories **A** and functors **F**, $(\mathbf{A}^{op})^{op} = \mathbf{A}$ and $(\mathbf{F}^{op})^{op} = \mathbf{F}$.

8. $d^{op} = c : \mathbf{1} \to \mathbf{2}$.

Indeed, we start with the terminal category **1**, which contains only one object. Then for any category **A** in *CAT*, the objects of **A** are defined as functors $f : \mathbf{1} \to \mathbf{A}$. Instead, for arrows we need the category **2**. As axiom 2 specifies, **2** has two objects, d and c. Then arrows are intuitively functors $f : \mathbf{1} \to \mathbf{A^2}$, and such a functor corresponds to a functor $g : \mathbf{2} \to \mathbf{A}$. Indeed, then $g \circ d : \mathbf{1} \to \mathbf{A}$ equals $dom(g)$ and $g \circ c : \mathbf{1} \to \mathbf{A}$ equals $cod(g)$.

To account for composition of arrows within a category, we resort to axiom 4. Note that α, β and γ are arrows of **3**. Furthermore, $cod(\alpha) = cod(\beta)$ and the equalities in the axiom tell us that $dom(\gamma) = dom(\alpha)$ and $cod(\gamma) = cod(\beta)$. That is, $\gamma = \beta \circ \alpha$. This allows us to define composites $b \circ a$ for any arrows a, b in a category **A** such that $a \circ c = b \circ d$. Namely, we have the existence of the pushout of axiom 4. Then as we have another commutative square with $a, b : \mathbf{2} \to \mathbf{A}$ instead of $\alpha, \beta : \mathbf{2} \to \mathbf{3}$, there exists a unique functor $t : \mathbf{3} \to \mathbf{A}$, where $t \circ \alpha = a$ and $t \circ \beta = b$. Then we can define the composition $b \circ a$ as the arrow $t \circ \gamma : \mathbf{2} \to \mathbf{A}$, which will have the domain of a and the codomain of b.

It might help to consider the depictions of the categories **1**, **2** and **3** as follows, based on (Ernst, 2017).

$$1_1 \qquad\qquad d \xrightarrow{1_2} c \qquad\qquad \begin{array}{ccc} \alpha \circ d & \xrightarrow{\alpha} & \alpha \circ c \\ & \searrow^{\gamma} & \downarrow \beta \\ & & \beta \circ c \end{array} \tag{3.6}$$

The CCAF axioms plus an axiom of Infinity are equiconsistent with the axioms for a well-pointed topos with a natural numbers object and satisfying Separation (McLarty, 1991). We can consider strengthenings of CCAF to obtain a more complete foundation of mathematics. For example, we can assert the existence of a category **A** such that **A** is a model of ETCS. This means we could work with a topos of sets and obtain the more higher-order results we can get from there. In fact, we may assert the existence of any category or functor, should that benefit the theory. Alternatively, axioms concerning *CAT* itself could be added, in order to make CCAF a more encompassing theory. For example, it is unclear from this system whether *CAT* should be an object of itself.

Besides the categories $\varnothing$, **1**, **2**, **3** and **E**, the axioms of CCAF prove the existence of several other finite categories. For this, it needs to be assumed that **E** has no more than the four arrows that it is required to have (McLarty, 1991). With this assumption, a theorem can be proven that implies the existence of particular categories. Out of these categories, the ones that have all finite products are equivalent to lattices, and the ones that are Cartesian closed are equivalent to Heyting algebras. In the theory, sets are then taken to be discrete categories.

(McLarty, 1991) does not think that the CCAF axioms prove that every category has a maximal discrete subcategory (that is, a set of objects), hence that there likely will be large categories.

However, apart from these particular categories, no other category has to exist according to the CCAF axioms. This means that the usual categorical results must mainly be derived as hypotheticals. Hence, this gives us results in a more general form, instead of providing specific instances.

3.5 Higher categories

A perspective discussed in (Shulman, 2008) is that for large sets, and similarly, large categories, it often suffices to know them up to isomorphism or equivalence (respectively). For smaller constructions, a more accurate representation is sometimes desired. In order to provide this level of generalization for large categories, however, we need to go further than axiomatizing a category of categories. The notion of categorical equivalence requires that of natural transformations, and hence we need to work in a 2-category. Thus, we may be interested in axiomatizing the 2-category of large categories **CAT**, as this represents more accurately the way we might want to deal with large categories.

There is no universally agreed upon axiomatization of **CAT**. Attempts have been made, however: for example, (Weber, 2007) gives a 2-categorical generalization of an elementary topos, and investigates properties of the 2-category **CAT**, where **CAT** has an internal category of sets. (Shulman, 2009) informally lists several desired axioms for the 2-category. Here, it is suggested that **CAT** should be a 2-pretopos, which gives us an internal logic, and it could assert the existence of certain exponentials. Additionally, following (Weber, 2007), a category of sets could be defined as a 'classifying discrete opfibration', satisfying suitable axioms. Finally, an axiom of Well-pointedness would ensure that our 2-category is indeed one of categories. Now the object of sets allows for the usual categorical constructions with the help of 2-categorical limits and the internal logic. This gives us an idea of what the axiomatization of **CAT** would look like; we will not go into any more detail here.

Once we are willing to axiomatize 2-categories, we could similarly assemble 2-categories (of large categories) in a 3-category, and so on. It is unclear how this would be axiomatized. However, there exist two theories, relating to each other, that are inspired by higher-categorical ideas. We shortly wish to discuss both here. The first, called FOLDS, was developed in (Makkai, 1998). Makkai expressed the first understanding that Martin-Löf type theory can be used to formalize (higher) categorical mathematics. Makkai looks for the *universe* and the *language* for a structuralist foundation of abstract mathematics, as category theory does not provide these notions clearly. That is, he aims to find a metatheoretic description of a category of categories. The idea is that *First-Order Logic with Dependent Sorts* (FOLDS) suits to be the language, and the universe should contain something called weak n-categories. First, a set A is defined as a type A, and 'elements' of A possess a dependent type $E(A)$. Dependent types allow for many more constructions in this way. Furthermore, an isomorphism principle is asserted, saying that properties of objects in a category are invariant under isomorphism. This principle requires a change in the definition for a functor, and this is given by the notion of a *saturated anafunctor*, of which we omit the technical definition here. By replacing functors with saturated anafunctors, we get from the 2-category **CAT** the bicategory **SanaCat —** an additional modification gives us a so-called anabicategory.

The point is that (Makkai, 1998) suggests that for a 2-category **CAT** we should take a saturated anabicategory. If we keep on increasing the dimension (and hence, generality), then for an n-category we should take the general notion of a *weak n-category*, a concept that is still being researched. The idea is, then, that the universe of higher-dimensional categories (or weak n-categories) is an alternative to the set-theoretic cumulative hierarchy. Now with FOLDS, Makkai introduces the notion of a *one-way category*, which is a category where each object can only be the domain of finitely many arrows and there cannot be any infinite paths; a one-way category corresponds to a partially ordered set. This is formalized by the notion of an **L**-structure, which is a functor from a one-way category into **Set**. A one-way category **L** is called a similarity type, and for each **L** there will exist the concept of **L**-equivalence of **L**-structures. That is, FOLDS-expressible properties over **L** are invariant under **L**-equivalence. This expresses the idea that the identity relation is something that is derived from a context, instead of being provided a priori.

Now, Univalent Foundations (UF) followed from FOLDS as a further development in the search for a way to syntactically restrict what we can express about mathematical objects, and category-theoretic language is used in defining this theory. Indeed, Tsementzis (2017) argues that the "real culprit" of the accommodation for more than just the essential properties of mathematical objects lies in "the availability of a global (untyped) identity predicate". Essential to the theory, then, is the idea that we should treat isomorphic mathematical structures as being identical (as we already do this informally). Thus, it provides a richer notion of equality. This is formalized by the Univalence Axiom, which says that, given elements of the universe as types, equivalence between types is the same as equality between types (Kapulkin and Lumsdaine, 2012). Informally, it can be phrased as follows:

> (UA) Identity is isomorphic to isomorphism.

That is, for objects A, B (UA) says that $(A =_U B) \cong (A \cong B)$. This allows UF to succeed rather well in satisfying the by Tsementzis formulated property:

> (SFOM) Any theoretical context can be naturally formalized in [a foundational system] $\mathcal{S}$ in such a way that any grammatical property of an object in $\mathcal{S}$ is invariant under the relevant criterion of identity in that context.

Hence, in UF, given a set-, group- or category-theoretical identity criterion, we can restrict the language in such a way that isomorphic objects with respect to this criterion are made indistinguishable (even though we remain aware of the multiple representations of this entity). This property prevents us from being able to express properties that are not structural.

Tszementzis argues that this makes UF a better structural foundation than (ZFC and) ETCS. Namely, ETCS incorporates a global, untyped identity relation, so that it can still express differences between, say, the additive group of integers $\mathbb{Z}$ and the additive group of even integers $2\mathbb{Z}$, even though they are set-theoretically isomorphic. To see this, define the arrow $f : 1 \to \mathbb{Z}$ (i.e., $f \in \mathbb{Z}$),

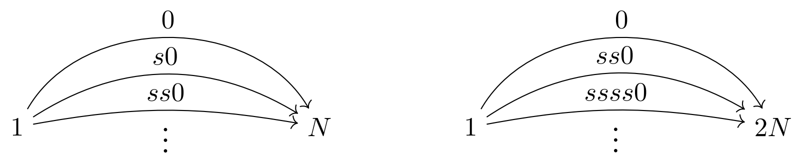

Figure 3.1: Two isomorphic yet different NNOs in ETCS

and let the property $\varphi(x)$ be that $cod(f) = x$. Then ETCS will always satisfy $\varphi(\mathbb{Z})$ (i.e., $cod(f) = \mathbb{Z}$), but not $\varphi(2\mathbb{Z})$ (i.e., $cod(f) = 2\mathbb{Z}$). An analogous situation for the natural numbers is illustrated by Figure 3.1. The problem is thus that ETCS (and ZFC) still allow for the existence of distinct isomorphic objects, and that they can detect this distinction. To solve this, we might for example formulate ETCS inside Makkai's system. However, this remains problematic for categories, as the identity criterion here is that of categorical equivalence. Thus, FOLDS will still express properties that are not invariant over equivalence.

We agree with Tsementzis that this argument makes ETCS a weaker candidate for a structuralist foundation. It remains unclear, however, to what extent categorical foundations are characterized by structuralist properties. The result given by Tsementzis currently gives us a property that ZFC and ETCS have *in common*; it is possible, then, that a 'structuralist' way of thinking is not inherent to categorical thinking compared to set-theoretical thinking. Of course, this a rather quick conclusion to reach here. However, it provides an interesting addition to the discussion concerning the relation between structuralism and category theory, and is worthy to be explored further. Still, even if it is the case that structuralism does not relate to an essential property of categorical thinking, it may still be the case that structuralist thinking is better accommodated for in a categorical setting, compared to a set-theoretical one. This raises the question when precisely something can be regarded as inherent to a certain theory (which is a very interesting matter, but unfortunately we lack the time to properly address it in this thesis).

Related to this (and returning to our purposes), is that in both FOLDS and UF, the main method to improve the suitability of categorical theories as structuralist foundations is by taking syntactical measures. However, we state that this is not a method inherent to category theory itself, as it concerns changing the way that identity is incorporated in the language of a theory. Again, such changes may bring very useful properties with them in a categorical setting, but they do not inform us about approaches inherent to categorical foundations themselves. This is something the reader should keep in mind — because the approach relates to the method of abstraction in Chapter 4, however, we will briefly cover the effect of the current approach there, focusing on UF.

3.6 Discussion

Summing up. We have seen different approaches towards the axiomatization of a categorical foundation for mathematics. Starting with an elementary topos is a common approach for defining some category of sets. The theory ETCS adds axioms for Well-pointedness, an NNO, Choice (and optionally Replacement), and thereby resembles standard set theory. These extra assumptions are merely optional for the category of classes of AST — adding them would turn the subcategory of small classes $\mathcal{S}_C$ into a WPTNC as well. However, AST alone provides a collection of axioms that introduces powerobjects and universes on top of the notion of smallness. A different approach is taken for CCAF, where we are merely told how categories are defined from functors in the meta-category and what basic properties they satisfy. Here, the theory allows for a rather free implementation of specific categories, which may be added as extra assumptions. Finally, FOLDS, and inspired by it, UF, pursue a foundation based in category theory that requires mathematical objects to be invariant under a suitable notion of isomorphism. To this end, the Univalence Axiom in UF explicitly alters the definition of identity.

Take away. From this chapter, we should take away that, similar to the set-theoretical foundations for category theory, a good deal of variation exists among categorical foundations. This shows that the categorical approach to mathematics possesses a certain flexibility with respect to its perspective on mathematics. In particular, categorical foundations thrive when pushed close to set-theoretical notions, but also appear in more neutral forms when sticking to basic categorical notions. We also see that, compared to set-theoretical foundations, categorical foundations vary more in the type of axioms they support. That is, whereas set-theoretical foundations for category theory all explicitly construct their objects, categorical foundations are less settled on the objects they assert. Furthermore, without the primitive notion of membership, categorical foundations are shown to inhabit various levels of analysis. Where ETCS and AST analyze the (set-representing) objects of a category, CCAF instead analyzes categories themselves. EM, on the other hand, defines the notion of a single category, but can clearly capture many of those with its definition. It should be kept in mind, then, that these factors individually cannot pin down the categorical perspective on mathematics. This will prove relevant in the following chapter.

A more coherent pattern which is seen in categorical foundations is that arrow-relations replace the role of (set-theoretical) membership. In ETCS, this is explicitly the case, as members of a set are directly taken to be arrows from the terminal object. In AST, however, a local element relation belonging to powerobjects is represented as a subobject, although specific members are not recognizable from this definition. Where CCAF does not incorporate explicit set-theoretical membership, it represents the 'elements' of specific categories (i.e., their objects and arrows) as particular functors. Thus, it seems that arrows are a deciding factor in a great deal of the relevant notions in categories, and they, too, are flexible in the role they can take on.

What is next? We hereby conclude the part of this thesis that concerns itself with the introduction of foundational systems based in set theory and category theory. That is, we are ready to apply our newly obtained knowledge of these systems to the exciting and slightly mysterious domain of mathematical thought.

Chapter 4

Mathematical thought

With the knowledge of the behaviour of category theory in set-theoretical foundations and the make-up of categorical foundations fresh in our minds, we are all set to address the main research question of this thesis in the current chapter. Recall that (Mathias, 2001) and (Ernst, 2017) have put forward the suggestion that category theory and set theory represent 'two distinct modes of thought'. Their idea is that, by considering which fields of mathematics are naturally captured by category theory (and set theory), and where categorical (and set-theoretical) methods fail, it might become clear that a different kind of mathematical thinking is at play in either theory. Although we consider this worthwhile to pursue, our starting point is the idea that ways of mathematical thinking are well-represented by foundations for mathematics — and that combining the perspectives from various foundations within category theory and set theory provides us with a realistic and complete view of what both theories are about.

Then, instead of deducing from applications to various mathematical fields what types of thinking are at play, we change the order. The approach from foundations lends itself quite naturally to outline, first, factors that describe a distinction in mathematical thinking, which can then be investigated with respect to foundational systems. This will avoid the uninformative situation where theories turn out to be suitable for mathematical fields that were developed inside them or with respect to them, and unsuitable for fields that have been developed in other theories. In that sense, by beginning with a conception about mathematical thought that does not rely on technical mathematical ideas, we hope to be able to investigate the distinction between set-theoretical and categorical thought slightly more fruitfully.

Indeed, there exist several conceptions about the distinction between category theory and set theory concerning mathematical thinking. In this chapter, we investigate two such conceptions: one focuses on the idea that category theory and set theory represent different ways of thinking about the same subject matter, while the other suggests that they are perhaps similar ways of thinking but have a varying subject matter. In particular, we focus first on the conception that the contrast between set theory and category theory corresponds to the distinction between the bottom-up and top-down approach, and second on the idea that category theory analyzes mathematics with a higher level of abstraction than set theory. If these factors turn out to be important for the distinction between set-theoretical and categorical approaches for mathematics, further research could investigate whether they indeed correspond to the mathematical fields that are strongly and weakly supported by category theory and set theory.

The chapter is structured as follows. We clear up the definition of a bottom-up and a top-down approach to mathematics first, after which we apply the result to set-theoretical as well as categorical foundations to argue that these terms do not carve up the conceptional space in terms of set theory and category theory. Second, we borrow a formalization of abstraction levels developed in (Floridi, 2013), and we argue it should incorporate an additional requirement for our purposes that helps capture true abstraction relations more directly. After applying the method to specific foundations for mathematics, we suggest that in several cases, categorical foundations approach mathematics from a higher level of abstraction than set-theoretical foundations. However, we also find that the differences in abstraction between set-theoretical and categorical foundations are less big than commonly thought, as supported by the lack of convincing satisfaction of our additional requirement. This leads us to the last section of this chapter, which will take into account particular roles of foundations.

4.1 Bottom-up and top-down

The distinction between a bottom-up and a top-down way of thinking appears regularly in mathematics. As the distinction is applicable to various (non-) mathematical processes, an informal description of the terms 'bottom-up' and 'top-down' is usually provided. Intuitively, both terms possess a dynamic component, where 'bottom-up' tells us we are working towards bigger things, while 'top-down' somehow starts with a bird's-eye view and zooms in. However, as our aim is to identify different types of thinking, we shall need a more precise definition of the terms. We will bring this about by responding to the use of the terms 'bottom-up' and 'top-down' in (Awodey, 2004) and (Landry, 2013). In both papers, a top-down approach is argued to fit a categorical way of working, while the bottom-up approach is (implicitly) associated with the well-known case for set theory. However, in this section we wish to argue that the conception of 'top-down' used for this distinction does not capture what our basic intuition tells us. Rather, we will see that Awodey's definition is more an ad hoc way of fitting it into a structuralist way of thinking. This, then, does not correspond to how 'top-down' is commonly used, so that 'top-down' is perhaps not a good indicator of the categorical way of thinking after all. With our cleared-up definition of 'bottom-up' and 'top-down', we will argue against their carving up of the conceptional space in terms of category theory and set theory.

First, consider the way in which Awodey and Landry use the distinction between bottom-up and top-down. We should note that the debate they focus on mainly addresses the suitability of category theory as a foundation for mathematics, a matter we do not take up ourselves. However, their association of different ways of working to category theory and set theory is relevant for us. In (Landry, 2013), Landry uses Awodey's conception of 'bottom-up' and 'top-down' to argue that they correspond to a constitutive and organizational role of foundations for mathematics, respectively. She argues, from there, that category theory, with its various systems, can be a foundation for mathematics in the sense that it organizes mathematics. Relating the distinction between top-down and bottom-up with this, Landry states that the category theory axioms can be taken as "top-down implicit definitions". She says:

> [...] the [category-theoretic] axioms, as schematic implicit definitions (as opposed to assertory truths), structure our mathematical concepts in terms of the relations that bear between them (as opposed to in terms of the "subject matter" of which they are constructed or constituted) so that the mathematical structuralist, as Awodey's distinction between top-down and bottom-up ways of working suggests, begins with the axioms. (Landry, 2013, p. 41)

Then let us take a closer look at Awodey's distinction. Awodey's notion of a 'bottom-up' approach seems intuitive and clear-cut, and we will adopt it as the basis of our own definition. As described in (Awodey, 2004), one here starts with a "specific range of specific "objects", presumed or constructed, but somehow fixed and given". A mathematical statement X is then interpreted inside this range of objects: it is taken to apply to the objects that X talks about and that are accessible to us from the given 'domain'. We add to this that, when a new entity is introduced, it is important that this happens along with an *increase in complexity*. That is, with the available objects and application of the rules at our disposal, via the bottom-up approach we construct more complex entities than we had before. Doing so will give us increasingly more information about the make-up and behaviour of a system as a whole.

For the top-down approach to mathematics, however, we wish to argue against Awodey's (and thereby also Landry's) conception. For Awodey, the top-down way of thinking does not presuppose a fixed universe of objects anymore. Similarly, top-down mathematical statements do not concern universal quantification over a given domain. Instead, a top-down statement X is "a *schematic* statement about a *structure* [...] which can have various *instances*" (Awodey, 2004). The 'schematic' property of top-down statements is meant to convey that they are if-then statements, where the range of the 'if' part does not need a restriction. This is because we have not specified a collection of things, and we can simply take anything that satisfies the if-requirement. Awodey explicitly relates his top-down approach to structuralism, and he argues that, by having its axioms 'go down' by specifying more and more structure of concepts (instead of actually constructing objects), category theory is very effective as a structuralist framework.

We wish to respond, however, that the conception of top-down as relating to structuralism does *not* do justice to the true meaning of 'top-down'. The informal nature of 'top-down' allows for multiple interpretations of the concept, including that of Awodey. However, this can lead to misunderstandings, and we encourage a consistent use of the term. In particular, we favor an intuitive definition that starts from the conception of 'bottom-up' and 'top-down' as opposites. Recall that for the bottom-up approach, we stated that a build-up of complexity was characteristic. This suggests that for the top-down approach, a breaking down or *decrease in complexity* is essential. That is, we start by 'externally' considering a system of mathematical structures, which we can break down into subsystems of lower complexity. This implies that (in well-founded cases) we in the end reach the most basic elements of our theory, that have lowest complexity. This requires that, as we are in general able to zoom in to the most basic elements we have, we should always have in mind, or at least decide at some point, what universe we are working in. This corresponds to the bottom-up way of working. In short, where the bottom-up-minded can build from available elements a structure of higher complexity, the top-down-minded can envisage a complex structure and 'fill in the gaps' with lower-complexity elements to ensure its existence. This, then, forms our distinction between bottom-up and top-down. Instead of ascribing the difference between bottom-up and top-down to the explicit or implicit treatment of objects, corresponding to a constructivist and structuralist approach, respectively, 'bottom-up' and 'top-down' should be taken as differing in how they pass between various complexity levels in mathematical structures.

Note that we, in turn, have not been fully explicit about the meaning of increasing and decreasing 'complexity'. This is because the exact interpretation of our sense of complexity is dependent on the particular objects one works with. Hence, when using this notion in our investigation of 'bottom-up' and 'top-down' ways of thinking in the various foundational systems, we will specify each time what we mean by it.

We hope now to have cleared up our conception of top-down and bottom-up. We will next argue that, by looking at the perspectives from categorical and set-theoretical foundations for mathematics, the distinction between (our sense of) top-down and bottom-up is not inherently related to the distinction between category theory and set theory. Let us clarify that we do not take the bottom-up/top-down distinction to be generally unhelpful: on the contrary, it allows for an enriched characterization of many concepts. However, we argue that the distinction is not helpful for discriminating a categorical from a set-theoretical way of thinking. Using results from the previous chapters, we will see that set theory and category theory allow for, and possess properties of, both ways of thinking. Hence, neither way seems to capture entirely the nature of set theory or category theory.

4.1.1 Relation to set theory

By considering set-theoretical foundations and how they form mathematical objects, we recognize two main properties that motivate the answer to whether they advocate a bottom-up or a top-down way of thinking. The first is well-foundedness, relating to the conception behind a set theory, and the second concerns the specific workings of particular axioms in a system. We argue that both allow for an upward and downward direction of complexity change. For the application of the term 'complexity' to sets, we would like to distinguish between two natural ways in which sets possess complexity. These do not define a universally acknowledged description of 'set-theoretic complexity', but they correspond to notions of membership-related structure that are well-known within set theory.

1. First, each set has a cardinality representing its size. It is simple and natural to think of a set X as being more complex than a set Y if $|X| > |Y|$.

2. Intuitively, the membership relation also expresses differences in *depth* (i.e., length of membership chains) in each set. For well-founded hereditarily finite sets, the longest membership chain can be determined by

taking the maximum depth in a graph representation of a set, such as the apg-representation described in (Incurvati, 2014) (recall the definition of an apg from Chapter 2). This notion happens to correspond exactly to the notion of *rank*. Whereas the rank is also an appropriate measure of complexity for well-founded non-hereditarily finite sets, it here no longer corresponds to the maximum depth of its apg-representation. For example, ω will only contain finite membership chains, but it has infinite rank. In the non–well-founded case, neither rank nor the depth of an apg-set is an appropriate measure of complexity.

Hence, we will take as our definite (and more encompassing) measure of complexity for an apg-representation of a set, the *number of its non-isomorphic sub-apgs*. This approach is advocated for and applied to general graphs by (Kim and Wilhelm, 2008); the situation is easily adapted to apgs.

The measure of non-isomorphic sub-apgs, combined with the cardinality of a set, is then what we adopt in order to evaluate the build-up and breakdown of complexity in set theories. This will reveal how different axiomatizations and conceptions of set theory approach the formation of objects in their theory.

Well-foundedness and non–well-foundedness. Starting with the property of (non–)well-foundedness, recall that we pointed out in Chapter 2 that the Foundation axiom (absent in non–well-founded set theories like ZFA) plays an important role in generating the iterative conception of set. In fact, it corresponds explicitly to this conception, as the requirement that there can be no infinite descending chain of sets is equivalent to saying that for every set X, there exists an ordinal number α such that $X \in V_\alpha$ (for V_α a level of the cumulative hierarchy V) (Moss, 2018). This allows us to argue that, depending on the incorporation of a Foundation axiom or not, set theory can accommodate both a bottom-up and a top-down conception towards its objects. Should we choose to incorporate Foundation, then the iterative conception of set reflects a bottom-up approach. For this, we have the empty set as a base element, and the Power Set axiom as the main complexity-increasing factor. Observe that the Power Set axiom increases both the rank and cardinality of sets (and the number of subgraphs of the apg-representation), although its effect on cardinality is notably stronger than on rank (if a set has n members, its power set has 2^n, whereas one application of the power operation adds just one extra 'level' in the membership chains). Then, when considering an individual set, it is natural to ask what level of the cumulative hierarchy it is part of, and to regard it as a product of a complexity-increasing operation. That is, a set is here something that is constructed from prior-existing, less complex sets.

This perspective is countered by the graph conception of set. The approach here is to begin with the whole conception of a set, and to repeatedly apply the reverse membership relation to unravel its members. Hence, the whole set 'exists first', and we are pursuing its less-complex members. Note that, while this unravelling of sets results in sets with membership chains that are no longer than before, we are not necessarily pursuing sets with fewer members than before, or equally many. In fact, if we start from any singleton $\{X\}$ containing a set X with more than one member, applying the reverse membership operation leaves us with a set with more members than before. Still, the measure of non-isomorphic sub-apgs provides a more consistent picture, as applying the reverse membership operation to an apg will always leave us with an apg that has no more non-isomorphic subgraphs than before.

In short, the graph conception of set then has a top-down nature, relying on the idea that "a set simply is [...] an object having a (hereditary) membership structure" (Incurvati, 2014). By allowing for downward infinity, non–well-founded set theory encourages the unravelling of sets, and by that a top-down approach to sets. If we only allow for upward infinity, the bottom-up view of sets is more intuitive. Thus, the perspective towards infinity we take affects the direction of complexity changes in the conception of set, and it even affects which objects may exist. Still, note that neither conception rules out occurrences of the other direction of complexity change in the theory.

Workings of the axioms. This relates to the second argument we make here, which concerns the particular workings of the axioms of a foundation. First, we claim that practically every set theory contains 'bottom-up axioms' as well as 'top-down axioms'. In ZFC, (and hence any extension of ZFC), the Pairing and Power Set axiom acting on available sets are clear examples of (non-strict) complexity-increasing operations (in both senses mentioned before). The Union axiom often (non-strictly) increases the cardinality of a set, suggesting it has bottom-up characteristics. For Union, the measure of non-isomorphic sub-apgs tells us, however, that the top node in the set-representing apg should discard all the first edges it is connected to, and create new edges to the nodes that were 'second in line'. This certainly corresponds to a (non-strict) decrease in the number of non-isomorphic subgraphs. Hence, the Union axiom correponds to a top-down approach to sets. The point here, then, is that, whereas the conception of set (iterative, graph, or something else) may advocate a general sense of top-down or bottom-up for a foundation, individual axioms often still allow for both approaches in the construction of sets.

Note that NFU is a bit of a special case here. If we take the system S^* described in Chapter 2, we see that it fits in with the previous paragraph as it extends ZFC. However, NFU by itself is axiomatized only by (weakened) Extensionality and Stratified Comprehension. While Extensionality only tells us when two sets are equal, we claim that Stratified Comprehension exhibits a top-down way of working. This axiom tells us that sets in NFU are simply 'predicates in extension', without a sense of which class exists prior to another. When applying the axiom, we create a set that contains every set satisfying some stratified property. We remark here that this presupposes a collection of *all* stratified sets that we can quantify over, and out of which we may select a subset. Clearly, this subset will be of lower complexity than the implicit universal collection of NFU-sets—hence, it advocates a top-down approach. We might argue that NFU allows for a simulation of a bottom-up construction of sets, by for instance applying Comprehension for a (stratified) property φ, a (stratified) property ψ, and subsequently forming a more complex set by taking the property $\varphi \wedge \psi$. However, we note that the more complex set is here constructed by a separate application of Comprehension, involving once more a complete quantification over all stratified sets, instead of building on the actual sets represented by φ and ψ. Thus, the method by which NFU forms sets is essentially always top-down, even though indirectly we can detect a build-up of complexity. Note that the Comprehension axiom also has a top-down working in the form of full Comprehension (in MK), restricted Comprehension (in NBG) and Separation (in ZFC and various extensions).

Furthermore, in NFU, Stratified Comprehension certainly allows for non–well-founded sets in the theory (such as a set of all sets), although note that the graph conception cannot be applied here. Similarly, one could envision an iterative conception of set by taking the urelements allowed by weakened Extensionality to form the basic elements of the theory. However, these elements do not at all form the basis of a complexity-increasing hierarchy induced by the axioms. The universe of New Foundations is untyped, so that there is no explicit idea of a bottom-up construction from these elements or a top-down unravelling of objects up to these urelements. Thus, NFU lacks a compelling conception of its sets that advocates either a top-down or bottom-up approach. The workings of the particular NFU-axioms themselves, however, represent a top-down way of working.

Finally, we recall from Chapter 2 that ZFC-classes can be defined implicitly by taking a ZFC-definable property. Through an informal use of (class) Comprehension, ZFC can here 'imagine' that classes exist somewhere outside its universe, and using ZFC they can be manipulated. As these classes are objects that can be thought of only as instances of (informal) Comprehension, we can only associate them with a top-down construction. This is in line with the fact that they cannot be reached by the power set operation on another set, and are hence not part of the (bottom-up) iterative conception of set. A similar argument works for NBG, where the informal use of Comprehension in ZFC is included explicitly in the theory of NBG to allow for the construction of proper classes. Thus here, we see a distinction between small and large (definable) sets that relates to the distinction between bottom-up and top-down.

4.1.2 Relation to category theory

The knowledge we obtained of categorical foundational systems will help form our claim that category theory, too, exhibits both bottom-up and top-down characteristics. As before, we will consider both the more conceptional approach to mathematics of a foundation, and its axiomatization. While conceptions of categorical foundations are more often top-down than bottom-up, most categorical axioms turn out to have a bottom-up nature (although explicit bottom-up construction of objects generally does not encompass all objects that categorical theories aim to talk about). Several foundations, additionally, take a neutral position regarding the bottom-up/top-down distinction.

In order to talk about complexity in categorical foundations, we should take into account the different types of objects that each foundation looks at: each will require a different sense of complexity. Below, we will describe a semi-formal measure of complexity for each foundation. It is not the purpose of this thesis to dive into the precise complexity of various categorical objects, so that we remain a little general here at times. Our purpose here, then, is to show that there exist measures of complexity for categorical foundations that we can relate to the notions of bottom-up and top-down approaches to mathematics. Thus, consider the following foundations and corresponding complexity measures.

CCAF The objects of study in CCAF are categories. The complexity of a category should intuitively capture the complexity of its structure in terms of objects and arrows. A rather simple way of formalizing this is to consider its size, i.e. the number of objects and arrows in a category. However, this tells us nothing about the complexity difference between two categories that have the same number of objects and arrows (even though we could imagine one has a simpler structure than the other). Therefore, we additionally take on the previously used approach from (Kim and Wilhelm, 2008). That is, we will take a category to have a higher complexity than one with the same number of nodes and edges, if it has more non-isomorphic subcategories than the other.

ETCS As a categorical set theory, the objects of study of ETCS are sets and their members, instead of entire categories. Thus, we would like to have a notion of complexity for categorical sets. As we retain the dependency of sets on their members in ETCS, we will adhere to the notion of membership for our complexity measure. We took a cardinality increase to be part of a complexity increase in set theory—the same requirement can be applied to ETCS. Indeed, we take sets X in ETCS for which there exist many arrows $x : 1 \to X$ to be more complex than sets for which there are fewer such arrows. However, as membership chains of length >1 are not formalized in ETCS, the apg-representation of an ETCS-set would only give us its cardinality without any further membership structure. Thus, as membership chains do not play a role here, we can simply take the cardinality of a set to represent its complexity.

AST In AST, membership does not let itself be captured by specification of individual members. However, the theory does impose a notion of smallness on some of its classes, which encompasses the objects and arrows with the most straightforward properties (such as being an identity arrow). Thus, we let small AST-classes be of lower complexity than proper AST-classes. Furthermore, the relation of the powerclass operation to its analogue in set theory, and the intention of a powerclass to represent the subsets of a class suggests that a powerclass $\mathcal{P}C$ should be taken to be of higher complexity than C (even for proper classes, as a powerclass $\mathcal{P}C$ is only small if C is). Similarly, a universe $V(C)$, intended to represent the cumulative hierarchy built on C, should be taken as more complex object relative to C. Finally, note that $\mathcal{C}$ has finite limits and finite coproducts, which are notions that should be taken as more complex than the objects from which they can be constructed.

It seems that we have to make do with these general and relative requirements for a sense of complexity in AST: we do not seem to be able to capture absolute complexity in the sense that we can always compare objects in terms of their complexity. The issue with this is, of course, that we are here already looking at the axioms and deriving from them notions of complexity. Our intention was, instead, to come up with an intuitive measure of AST-complexity first, and *then* apply this to the axioms. However, a natural candidate for this does not present itself in AST.

We can now apply these requirements to the relevant foundations. Starting with CCAF, its assertion that the 'metacategory' *CAT* has all finite limits and colimits, and is Cartesian closed, certainly increases the number of subcategories of *CAT*. Namely, the exclusion of limits and colimits already creates various non-isomorphic subcategories of *CAT*. The given properties of categories $\varnothing$, **1, 2, 3** and **E**, furthermore, provide explicit categories and functors to include or exclude in *CAT*, acting on its number of subcategories. Similarly, the CCAF axioms provide the existence of functors defined by a relation, and for every category it harbors they provide a dual (opposite) category. These axioms affect the complexity of the metacategory *CAT*, but note that the assumption of finite limits and colimits also acts on the complexity of individual categories: a limit category will be more complex than the categories it was constructed from. Furthermore, to directly affect the complexity of an individual category, the axioms should act on the construction of functors from **1** and **2** (i.e., the objects and arrows of a category, respectively). The only axiom that does this, however, asserts that some category has a non-identity isomorphic arrow (technically increasing the number of subcategories of this category).

Thus, CCAF does not have a lot of axioms that act on complexity, but the ones that do seem to increase it. This suggests that the CCAF axioms represent a more bottom-up than top-down approach to mathematics. Furthermore, we also maintain that categories made up of more arrows and objects (i.e., of a bigger size) are more complex. As the individual categories are defined from the metacategory, the conception of individual categories starts by conceiving the whole CCAF-universe. Only then can we focus on a particular category in the metacategory by considering a particular subset of functors that commute with a particular category. This 'zooming in' corresponds both to a size reduction and a decrease in non-isomorphic subcategories of *CAT*. In this sense, then, the conception of CCAF advocates a more top-down approach to the make-up of individual categories in the theory.

The axioms of ETCS say more about complexity changes. For example, they assert that the coproduct $A + B$ of two sets A, B always exists. $A + B$ will, through composition, have as many members as A plus B; taking the coproduct in ETCS hence increases the complexity of sets. It can also be shown that exponentials A^B (which are additionally ensured to exist for sets A, B) generally have more members than A (the situation is less clear for B). Hence, these axioms represent bottom-up ways of thinking. Additionally, objects like 0, 1 and an NNO are specifically assumed to exist in the theory and are suitable for applications of bottom-up axioms. The other axioms mostly provide general requirements for the make-up of the model of ETCS, and do not indicate complexity changes. Concerning the conception towards mathematics that categorical foundations employ, we see that the axioms for ETCS are similar to those of ZFC. We are told what a category is, followed by several existence claims and properties of categories. As we have no membership chains, however, the Axiom of Foundation is irrelevant and gives us no incentive to either adopt the iterative or the graph conception of set. A similar argument has been made in (Linnebo and Pettigrew, 2011), where it was concluded that the conception of the iterative hierarchy depends strongly on the membership relation between sets, which ETCS does not provide. Hence, we claim that the *conception* of ETCS is neutral in terms of the distinction between bottom-up and top-down. Then, by our observation that the axioms of ETCS express mainly bottom-up ways of constructing categorical sets, we conclude that the bottom-up approach is slightly better represented.

For AST, note that we already largely gave away the effect of the axioms on complexity. For its conception, we observe that the axioms first of all start with a category of classes, from which the notion of smallness is subsequently derived. Hence, this resembles a complexity-decreasing process. Furthermore, bottom-up axioms establish, for every class, the existence of a powerclass and a universe based on it. The axiom ensuring the existence of limits and finite coproducts also has a bottom-up nature. These are the main notions in the axiomatization of AST that we pointed out as possessing comparable complexity. It thus seems that AST combines bottom-up and top-down axioms acting on its objects, although bottom-up axioms are more common.

EM-category theory. We here separately consider the EM-axioms. EM tells us what a category is: there is no explicit existence assumption of a category, nor of an object or arrow. We could perhaps infer a measure of complexity, where objects that have more arrows going to and from them than other objects are relatively more complex; or we can take into account the number of objects that a particular object is related to. However, this gives us nothing to hold on to with respect to the EM-axioms, as it is entirely neutral about the specific structure of categories. That is, with its axiomatization, we claim that EM does not express a top-down or bottom-up approach. We mean here that EM does not tell us how to *approach* or *understand* the concept of category. Of course, the definition itself has many instances and it captures all structures that are categories. However, there is no sense of movement or direction of complexity in the axioms, something that the bottom-up and top-down approach both require. Thus, we conclude that EM-category theory is neutral with regard to this distinction.

Finally, we see that the type of set theory that successfully founds EM-category theory also does not characterize a sharp distinction between bottom-up and top-down. First, compared to ZFC, the assumption of an inaccessible is clearly one of increasing complexity. The reflection principle, however, provides a connection between objects of higher and lower complexity, while it remains of itself rather neutral of the direction of this movement. Thus, compared to ZFC, the more successful category theory founding capacities of ZMC/S rely mostly on a bottom-up axiom. On the other hand, NFU constructs its categories with the top-down Comprehension axiom—its success in founding (unlimited) category theory compared to ZFC can thus here be taken to rely on a top-down approach. Furthermore, the inclusion or exclusion of a Foundation axiom (that determines the bottom-up or top-down conception of set) does not affect at all how EM-category theory is supported. These claims are then in line with our suggestion that EM-category theory is neutral with respect to the distinction between bottom-up and top-down.

4.1.3 Taking stock

In this section, we fine-tuned the bottom-up/top-down distinction, after which we explored the representation of bottom-up and top-down approaches in different set theories and category theories. In both cases, the formation of mathematical objects can be done in a bottom-up as well as a top-down manner. Set theory allows in particular for a bottom-up and a top-down approach to the conception of sets by (not) incorporating the Foundation axiom. This shows that it is misleading to claim that set theory and category theory can be distinguished by the fact that the former represents a bottom-up, and the latter a top-down approach to mathematics. Differences *between*, but also *within* particular foundations are important to take into account when characterizing set theory and category theory as wholes. It seems that the distinction between bottom-up and top-down is a property of the formulation of axioms and the type of objects they allow for. However, it does not so much correspond to an intrinsic property of a foundational system itself. A complication in our application of the distinction between bottom-up and top-down to set theory and category theory is that there is no universal notion of 'complexity' for foundational theories, so that we had to adapt the definition for different foundations. A more precise and formal description of complexity of objects should shed more light on the topic.

Still, the current result give reason to believe that category theory and set theory do not represent different approaches to the same thing. Perhaps, then, they have a different subject matter, in the sense that their respective objects come with a different level of abstraction. This is what we will explore in the next section.

4.2 Abstraction of subject matter

Awodey's original definition of top-down resorts to a structuralist view, where entities exist independently of their instances, and theorems are 'schemata'. Although we argued against this characterization of top-down, it might come closer to what is really at play for category theory. Perhaps what really captures Awodey's perspective, then, is that category theory embodies a higher level of abstraction than set theory. The existence of a difference in the level of abstraction between category theory and set theory has (informally) been suggested various times (see (Linnebo and Pettigrew, 2011), (Landry, 2013), and others). Intuitively, an increase in the level of abstraction of some system is characterized by, for one, a process of ignoring details from lower levels that have become irrelevant, and secondly by identifying new concepts that capture the essential objects, relations or properties at the higher level. This intuition, however, is seldom made explicit, so that arguments concerning the levels of abstraction of set theory and category theory become rather informal. Similarly to the distinction between bottom-up and top-down, then, we maintain that a more formal notion of abstraction is desired and would provide a better understanding of the way set theory and category differ with respect to it.

In the interest of this, we will borrow a method developed in (Floridi, 2013) based in the field of philosophy of information. Here, levels of abstraction are formalized as expressing an increasing epistemic clarity, and levels are related to each other by means of so-called gradients of abstraction. In what follows, we will define the main concepts that we will use from this method, after which we will provide an addition to the method related to our purposes. The main part of this section will then show that, from our definition, categorical foundations generally reflect higher, but varying, levels of abstraction (that are still less high than generally thought) with respect to set-theoretical counterparts. This entails that the distinction between levels of abstraction cannot be completely associated with the distinction between set-theoretical and categorical thinking, either. Furthermore, we claim that set-theoretical systems that can harbor many large categorical constructions act on properties that are related to the method of abstraction, and we suggest that this may allow them to 'simulate' a higher level of abstraction. Finally, we argue that the purpose of a foundational system is connected to the level of abstraction it possesses. At this point it is time to introduce the main formal notions we will concern ourselves with.

4.2.1 A method of abstraction

The method described in (Floridi, 2013) involves a form of 'levelism' that is intended to be the main working method in the philosophy of information. It has broad applications to different fields of study. As Floridi's method encompasses various other methods and it supports an epistemological approach, it generally suits our purposes here. However, the role of the method in the philosophy of information is rather different from its role in the philosophy of mathematics: in particular, we will see that the notions of 'level of abstraction' in the two fields do not always coincide. Thus, we make our own additional requirement to the definition later on. For now, note that levels of abstraction should be considered as "levels of observation or interpretation of a system". This comes with the idea that such a level should not be considered independently, without a purpose or context. Surely, we can imagine that each foundational system for mathematics is created with a purpose in mind; the relation of levels of abstraction to this will come back in the last part of our argument. Consider, first, the following three-part definition adapted from (Floridi, 2013) that captures the notion of a level of abstraction.

Definition 7. (a) x is a variable of type X (written $x : X$) if x is a uniquely defined conceptual entity, and X is a set that comprises all values that x may take on.

(b) A typed variable $x : X$ together with a statement α that clarifies the feature of the relevant system that x represents (i.e. an 'interpreted' typed variable), is called an *observable*.

(c) L is a *level of abstraction (LoA)* if it is a finite and non-empty set of observables.

Thus, a LoA[1] is essentially a set of 'conceptual entities', each capturing an aspect of the system at hand. A distinction is made between *discrete* observables (observables whose type is a finite set) and *analogue* observables (otherwise). Whereas most of the types that we will define will be analogue, we need to be slightly careful with their characterization as sets. Namely, we will often be ranging over set-theoretical or categorical objects in the metatheory, so that the type for us then comprises an external *collection*. Furthermore, the *behaviour* of a system at a particular LoA is given by a predicate, that takes observables as values. Any instantiation of types for the observables that the predicate makes true is called a *system behaviour*. A LoA together with a behaviour is called a *moderated* LoA.

[1]We write 'a LoA' with the pronunciation of LoA as the non-existent word *loa* in mind, instead of 'an LoA', which presumes the separate pronunciation of the letters 'l', 'o' and 'a'.

Some examples. In order to provide a basic intuition for the concepts just formalized, we provide some examples used by (Floridi, 2013). Suppose that our object of study is a human: then we could define a variable *h* (representing *height*) of type $\mathbb{R}$ and interpreted by the unit of metres (making it an observable). The behaviour of the system is then given by the predicate $0 < h < 3$, as the length of humans is clearly bounded above and below. Alternatively, if we are evaluating wine, we could have observables for *colour*, *clarity*, *alcohol level*, *price*, and so on, each with its own type. In this case, different LoAs will consist of the observables that suit our purpose for the wine: for example, the price of the wine is important if we would like to purchase a wine, but the colour of a wine is more important for the purpose of tasting it. Of course, observables between such LoAs could overlap. There are many more examples, but it should be clear that there is a certain level of freedom in the implementation of LoAs with respect to the analyzed system. This is part of why the method is so widely applicable, but it also requires us to justify the way we choose to apply it to mathematics later on.

The last relevant notion that we adopt from (Floridi, 2013) allows for the connection of different LoAs by means of a relation. A *relation* between sets A and C is simply taken to be a subset of $A \times C$. If A has a predicate p for its observable, R relates it to the predicate $P_R(p)$ on C that holds just at those $c : C$ that are related by R to some $a : A$ satisfying p. With this, a system can be discussed at various LoAs, as follows.

Definition 8. A *gradient of abstractions (GoA)* consists of the following.

1. A finite set of moderated LoAs L_i $(i \leqslant n)$.
2. A family of relations $R_{i,j} \subseteq L_i \times L_j$ $(0 \leqslant i \neq j < n)$. The family of relations $R_{i,j}$ relates the observables from each pair L_i, L_j of distinct LoAs such that:
   (a) For $i \neq j$, $R_{i,j}$ is the reverse of $R_{j,i}$.
   (b) The behaviour p_j at L_j is at least as strong as the translated behaviour $P_{R_{i,j}}(p_i)$, i.e. p_j implies $P_{R_{i,j}}(p_i)$.
3. For each interpreted type $x : X, y : Y$ in L_i, L_j (respectively) such that $(x : X, y : Y) \in R_{i,j}$, a relation $Rxy \subset X \times Y$.

Thus, a GoA establishes an explicit connection between the observables at multiple LoAs, allowing one to unambiguously reveal the way aspects of a system correspond to each other. Although this method acquires elegance from its simplicity, we cannot completely justify this way of relating LoAs for our purposes. The next section will therefore elaborate on the way in which we will apply the method to foundational systems. Note that (Van Leeuwen, 2014) argues that, additionally, *annotations* (i.e. meta-data) should be added to a LoA in order to be able to express the "micro-structure" that exists between successive LoAs. This is meant to refine the gradient between such LoAs. The annotations should "describe how observables of an LoA are constrained or otherwise to be used, as a guide to a deeper insight or capability at this level". We will not explicitly use the method of annotations here, as we will provide our formalization with enough explanation that it will not give us any other benefits. This seems more relevant for looking at more concrete systems from different perspectives.

4.2.2 Applying the method

The method of abstraction we just described has mostly been illustrated with real-world examples, in order to clarify its workings. The subject matter of this thesis, however, is mathematics, and our aim is to analyze LoAs that are characteristic for set theory and category theory. Taking the angle from foundations, we support the view that each particular system based in set theory or category theory reveals some information about the way the respective theory looks at mathematics. In particular, we wish to avoid the treatment of each foundation simply as a representation of 'set theory' or 'category theory'. Instead, we maintain that knowledge about the differences between particular systems will allow us to form a more nuanced picture of their approaches to mathematics. Then, a natural idea is to formalize the various set-theoretical and categorical foundations as occupying a particular LoA. That is, we regard each foundation as individually selecting essential aspects of mathematics. The formalization of these aspects into observables will let us form a GoA between foundations that are related to each other by a difference in level of abstraction. Now let us consider the suitability of the discussed method of abstraction for this purpose.

We can envision several requirements that we need the method of abstraction to satisfy in order to provide reliable results. When we obtain a GoA between two foundations of mathematics, we would like for it to express the two intuitively characterizing properties for a change in abstraction. That is, the GoA should tell us that there is a loss of information when going from one foundation of mathematics to the other, and it should point out the higher-level concepts that are identified in the foundation of a higher level of abstraction. Furthermore, it should generally be the case that the defined observables are natural and their choice can be justified, i.e. they are a reasonable pick and capture characterizing aspects of the relevant system. The respective types of related observables should ideally capture properties that are only present at a particular LoA, so that we can tell what really characterizes it. If this holds, we would also like the (type) relations in the GoA to be strict, in the sense that they can only go one way. Namely, if the types allow for a two-way GoA, we surely cannot tell which observable is part of the higher LoA with respect to the other. To do this, then, we should also require some kind of 'naturality' in the relation. That is, we should not be able to define an arbitrary relation between observables; we really want to relate concepts to abstractions of themselves. This is hard to formalize, but we argue that in the relation between types it is already helpful to keep track of where information is lost—this should indicate the direction of the change in abstraction. To make sure that the type relation specifies an abstraction, furthermore, it is natural to require that the higher LoA identifies concepts that group together elements of the lower LoA. This we can better incorporate in the method by adding to Definition 8 one of the following requirements (we elaborate more on the justification of these requirements after defining them). Here we take type X to belong to an observable of a lower LoA than that of the observable with type Y.

> Take a relation $R_{i,j}$, and interpreted types $x : X \in L_i$, $y : Y \in L_j$ such that $(x : X, y : Y) \in R_{i,j}$.
>
> R1. We require the inverse of the relation $Rxy \subset X \times Y$ to be *(necessarily) injective*. That is, suppose that $y_1, y_2 \in Y$ and that

there exist $x_1, x_2 \in X$ such that Rx_1y_1 and Rx_2y_2. Then $y_1 \neq y_2$ implies $x_1 \neq x_2$.

R2. We require the inverse of the relation $Rxy \subset X \times Y$ to be injective. Additionally, for at least two elements $x_1, x_2 \in X$, there *(necessarily)* exists a $y \in Y$ such that $(x_1, y) \in Rxy$ and $(x_2, y) \in Rxy$.

More concretely, by not assuming either R1 or R2, we allow for a grouping of different elements $x_1, x_2 : X$ of type X under one element $y : Y$, as well as a grouping of $y_1, y_2 : Y$ under an element $x : X$. To us, this suggests that the observables are not well-defined, as the types do not seem to capture a consistent level of abstraction. In R1, we require such a grouping process, *if* it occurs, to only go one way. That is, R1 still allows for an inclusion GoA, where no grouping of elements takes place. This is because we do not wish to commit ourselves to R2, as there *can* be a one-to-one mapping that still expresses a change in abstraction. In particular, this could be a loss of information alone, leading to a creation of higher-level concepts without a collapsing or grouping process of elements — however, mere loss of information should generally of course not be taken to express a difference in the level of abstraction. Hence, we regard a strict R1-GoA to be weaker than an R2-GoA. Observe, furthermore, that we formulated R1 and R2 with 'necessarily' between brackets. This is because we would also like to know when R1- and R2-relations are accommodated, yet in a less convincing way. This will inform us about the naturality of the pairing of observables. Also, we do not wish to impose either the R1- or R2-relation on our GoAs, as our goal here includes getting a complete picture and characterizing possibly different relations between LoAs, as well. We think it is important, however, in the case where relations deviate from R1 and R2, to be clear about what they represent.

Enforcing the second requirement comes down to explicitly demanding an instance of a generalization in a GoA. Making this distinction will allow us to judge the significance of GoAs we establish a little better. Additionally, R2 allows us to explicitly recognize where concepts are grouped together (i.e., abstracted) to form higher-level concepts. This means that R2 identifies the creation of higher-level concepts that go directly together with a collapsing process of type instantiations. Plenty of concrete examples motivate this requirement: for instance, biologists can analyze the human body in terms of molecules, whole cells, organs, or as a whole. Each level can be analyzed on its own, without he or she needing, or in fact *having*, the complete information of the make-up of the relevant structure in terms of atoms. Relations between the levels are clearly R2-relations. A simple mathematical example is given by the equations (1) $1 \cdot 1 = 1$, (2) $2 \cdot 1 = 2$, (3) $3 \cdot 1 = 3$, and so on. We would like these to abstract to the equation $n \cdot 1 = n$, instead of a one-to-one relation of (1) to $n \cdot 1 = n$, (2) to $m \cdot 1 = m$, etcetera - clearly, $n \cdot 1 = n$ and $m \cdot 1 = m$ should be taken to be the same, and we want (1), (2), (3), and so on, to collapse onto the same abstraction. This allows us to more formally, and with more certainty, recognize abstraction relations.

Our added requirements lack in Floridi's approach, as the primary goal of his abstraction method is to capture associations between the way we (as humans) analyze a system with a particular goal in mind. For example, we may make a GoA from the *wine tasting* LoA to the *wine purchasing* LoA, by mapping similar observables such as *colour* to themselves, and relating the types with a (possibly) different interpretation to each other. Similarly, if a LoA is extended by adding extra observables, the inclusion relation R that sends each observable in the first LoA to itself in the second LoA, forms a GoA for Floridi. For us, it results in an R1-relation for the types of each observable. This makes sense to us, as adding an observable intuitively means that a different aspect of a system is analyzed, yet with the same level of detail. For Floridi, instead, the addition of the observable and its implied change in the purpose behind the analysis of the system directly stands for a change in abstraction level.

Several remarks on the reliability of the method of abstraction are still in order. The relations of a GoA expanded with our R1- and R2-requirement allow for a characterization of the generation of higher-level concepts on a higher level of abstraction. Namely, this shows directly and clearly how concepts are related to each other to form new ones. However, the relation process is heavily dependent on the interpretation of the system by the user of the method, as she may define observables with their interpreted types in nearly any way she likes. This large amount of freedom makes it tempting for the user to mold the observables and types exactly in such a way that the desired GoA is attained. This may be especially tricky in a theoretical subject such as mathematics, that does not provide us with a fixed subject matter with perceptible properties. Additionally, the loss of information aspect of an increase in abstraction remains implicit in the method. Hence, we will have to elaborate on this part separately for every GoA that we create. Even with our adaptations, then, the results from the method should be interpreted carefully.

In the following subsections, we explore several GoAs by first making the observables in the relevant foundations explicit, followed by the relations of the GoA. In particular, this will reveal an explicit GoA between ZFC and ETCS+R, and various possible ones between NBG and AST. Additionally, we discuss several ways to incorporate CCAF. After that, we will more generally discuss the relation of UF and EM to the method of abstraction. For the types of our observables, the specific implementation of (meta-)quantification can be chosen in various ways. We shortly outline the possibilities and their strengths and weaknesses; however, it should be stressed that the argument we make will work with any of these methods. We take as preparatory example the observable *members* that we will use the most. In short, this observable should, for any set, output a collection of its members. Thus, suppose we are working in ZFC, then the type ideally consists of all possible collections of sets. The motivation for this observable will be given in the next section. For now, consider the various possible implementations of the type of *members*. Various advantages and disadvantages of these methods specific to the LoAs will be discussed along with their introduction.

1. First, we can restrict to ZFC-definable sets. Here, we take a definable property φ to be a formula such that ZFC proves that there is a unique element x such that $\varphi(x)$ holds. Then, given a ZFC-definable set $\varphi(x)$, the observable *members* can tell us what its members are by outputting the collection of its ZFC-definable members. That is, $\varphi(x)$ is assigned $\{\psi(y) : \text{ZFC} \vdash \exists x, y(\varphi(x) \wedge \psi(y) \wedge y \in x)\}$ (where the brackets $\{\}$ should be understood as defining an 'external' collection). Thus, the type of the whole observable then becomes $\{X : X$ is a collection of ZFC-definable

members}.

2. Another way to go is to consider, for ZFC, all its models M and take their (class) union $\mathfrak{Z}$. Some might prefer this method over the first one as the quantification will be over actual sets, instead of formulas. It will turn out to be relevant to the argument to assume that the models in $\mathfrak{Z}$ are disjoint. We can easily see that we may do this (without loss of generality), by the simple fact that we can replace each model M by $M \times \{M\}$. That is, we can label all of our elements with their model (although other ways exist). Then, however, in the union of models there now exist multiple copies of the empty set, the natural numbers, and other sets, which are not of use to us.

3. Alternatively, we could assume that there is one canonical model of ZFC. This would get rid of multiple occurrences of sets as in the union of models, but it might be less clear if we have all the sets we want when comparing theories.

Besides these mathematical approaches, an option that requires a little more tolerance is one that considers a meta-universe of mathematical objects. The idea is that this universe is neutral in the sense that it does not belong to any foundational point of view. Hence, it can be shaped into a ZFC-accessible universe V (and thus capturing *all* objects its theory allows for), while other theories interpret the meta-universe in different ways. This approach ensures that we have all the objects we want, and from the ZFC-point of view we can talk about actual sets and their members. While this approach is similar to the general way of thinking in this thesis, it is admittedly not very well-defined. Perhaps if worked out a little more carefully, it is as viable an option as the above three, and we should keep in mind that it is a (possibly quite elegant) method. However, we will currently motivate our definitions of LoAs and GoAs with the more sharply delineated methods that we just discussed.

We add that the GoAs we create will require us to relate objects of a particular theory to objects of a different theory. In the definable case, it is not clear if we can relate definable objects in one theory canonically to definable objects in another theory — that is, it is unclear how the respective formulas should correspond. When considering the union of models of a particular theory, it is sometimes obscure how one should relate the models of particular type instances. For some theories there may exist a canonical model correspondence that leads to an obvious choice of models to relate. This type of relation in a GoA has our preference, as it can reveal abstraction relations directly between corresponding type instances. If this is not the case, however, an arbitrary choice of model and particular type instance will have to be made. This will work fine, too, but it is a little less elegant. As it is not always known whether a canonical correspondence between models of different theories exists, we will leave the exact choice between the two methods open when constructing our GoAs — this section may be seen as a template for making this choice with the proper knowledge.

At times the implementation of GoAs involves several technicalities. We outline the more general idea of the argument at the beginning of every GoA. After discussing our GoAs between set-theoretical and categorical foundations, we will elaborate briefly on the role of set-theoretical foundations for category theory with respect to the method of abstraction. Note that we do not claim to exhaustively specify the differences in LoAs between foundations, nor that our implementation of the method of abstraction is the only reasonable one. However, we regard our attempt as a start of the formalization of mathematical concepts such as abstraction, which can be used in the philosophy of mathematics to characterize the distinction between set theory and category theory.

Relating ZFC to ETCS+R

We are not the first to propose that there is a difference in the level of abstraction between set theory and ETCS(+R). McLarty says that

> The sets of ETCS are abstract structures in exactly [the following] sense. An element $x \in S$ in ETCS has no properties except that it is an element of S and is distinct from any other elements of S. (Lawvere, 2005, p. 3)

Thus, we see that McLarty's conception of abstraction more or less coincides with the process of losing irrelevant information. We have claimed, however, that identification of higher-level concepts is an additional (stronger) indication of abstraction. This is not obvious from McLarty's clarification. More generally, ETCS(+R) is commonly referred to as axiomatizing the 'category of abstract sets' ((Linnebo and Pettigrew, 2011), (Landry, 2013), and others). However, this characterization of ETCS(+R) has been used informally and has never been questioned. We argue that this use of abstractness conceals a more nuanced picture of the situation. In fact, we will show that with our definition of levels of abstraction, there does exist a difference in abstraction between ZFC and ETCS+R, yet it is markedly less big than generally assumed. We make this idea precise with the defined framework for abstraction, an approach which has not been taken before. First, we will motivate and establish a LoA for ZFC, followed by one for ETCS+R. Consequently, we can formulate a GoA between the two defined LoAs, which will allow for an R1- and an R2-relation, but not in a *necessary* way.

Defining the LoAs. In order to form an observable for ZFC, we should consider the perspective of ZFC on mathematics. An observable should tell us something about this point of view, and it should capture a property that can always be assessed. Just like determining the colour of wine or the height of humans gives us knowledge about the object of study, we are looking for properties that characterize mathematics from the point of view of ZFC. This quickly leads us to the notion of set, the only type of object ZFC believes exists. We note that sets have many properties, all of which could be formalized into separate observables. That is, we could have observables for cardinality, rank, the number of ordered tuples a set contains, its intersection with a particular other set, and so on. However, we observe that all of these properties can be inferred from the characterization of the members of a set (by formalizing the observables as described below). As the members of a ZFC-set are sets themselves, looking at membership will give use the entire internal structure of a set, from which we can deduce other properties. Then, it seems that we can make do with just one 'membership observable' $\mathbf{O}_{\text{ZFC}}$ that does exactly this. Hence, we define our LoA $\mathbf{L}_{\text{ZFC}} = \{\mathbf{O}_{\text{ZFC}}\}$.

We see that ETCS+R naturally gives rise to a similar LoA. The theory supports the two sorts *object* and *arrow*, with the former representing sets and the latter forming the members of sets as arrows from the terminal object. Unlike in ZFC, then, members should not be thought of as sets, and there exists no explicit notion of membership chains. Still, knowledge of the members of an ETCS+R-set will provide a thorough characterization of the set itself, as aside from their members, ETCS+R-sets are merely characterless 'points'. Hence, we can introduce the membership observable $\mathbf{O}_{\text{ETCS+R}}$ that, like $\mathbf{O}_{\text{ZFC}}$, for a given set outputs a collection of its members. Additional observables for properties such as cardinality are, like before, derivable from $\mathbf{O}_{\text{ETCS+R}}$. This allows us to define our second LoA as $\mathbf{L}_{\text{ETCS+R}} = \{\mathbf{O}_{\text{ETCS+R}}\}$.

The natural idea for the GoA between $\mathbf{L}_{\text{ZFC}}$ and $\mathbf{L}_{\text{ETCS+R}}$ is then to relate a ZFC-set with a particular cardinality to an ETCS+R-set with that same cardinality. Namely, this is the closest correspondence we get between the two types of sets, and it allows us to directly examine the change in abstraction between the two notions.

Formalizing the observables. Note that we can formalize the observables in the various ways described earlier (and perhaps even more). The approach from definable sets results in the following definition of the observables.

> $\mathbf{O}_{\text{ZFC}}$:= *members* : $\{X|X$ is a collection of ZFC-definable sets$\}$
> $\mathbf{O}_{\text{ETCS+R}}$:= *members* : $\{X|X$ is a collection of ETCS+R-definable arrows with domain 1 and codomain a particular object$\}$

However, this is not the most satisfactory approach here. This is because there exist only countably many formulas φ that can describe a definable set. This means for example that, if the observable outputs a collection of countably many definable sets as the members of a certain definable set, we do not know if the latter was indeed countably infinite, or if it was actually uncountable. We care about making this distinction, as we motivated our membership observable with the idea that knowledge of the members of a set gives us all other information we would like to know about it. This method then does not seem to suffice for this, if we wish to take uncountable sets into account. Of course, we could make the approach work by adding another observable for cardinality, but this seems to (partly) undo the motivation for the membership observable that it provides enough information by itself. Nevertheless, for finite sets this method works well enough.

Alternatively, then, we could use a model approach. Taking the (class) union $\mathfrak{Z}$ of models of ZFC and the union $\mathfrak{E}$ of models of ETCS+R, our observables look as follows.

> $\mathbf{O}_{\text{ZFC}}$:= *members* : $\{X|X$ is a collection of sets $x \in M$ for some $M \in \mathfrak{Z}\}$
> $\mathbf{O}_{\text{ETCS+R}}$:= *members* : $\{X|X$ is a collection of arrows with domain 1 and codomain a particular object $x \in N$ for some $N \in \mathfrak{E}\}$

Here, our sets and members are not dependent on size limitations, although recall that we assume disjointness of the models, so that multiple copies of various sets will exist in the model unions.

Finally, we stress that, although technical details of the quantification process may seem irrelevant, it is important to be aware of the possible formalizations. The purpose of this part of the argument is to show that we can compare observables in different theories and explicitly point to a process of abstraction. Whereas the idea for the relations between observables will often be quite intuitive, showing that it can be formally implemented considerably strengthens the argument, as it shows us that such an abstraction process can *actually* take place, and it is not just an unattainable idea confined to philosophical debates.

Now that we know how to make the type of our observables precise, we can relate them to constitute a GoA between $\mathbf{L}_{\text{ZFC}}$ and $\mathbf{L}_{\text{ETCS+R}}$.

Defining a GoA. A natural requirement for connecting the observables is to have the cardinality of collections of ZFC- and ETCS+R-sets match up. That is, we want to define $R_{\text{ZFC,ETCS+R}} = \{(\mathbf{O}_{\text{ZFC}}, \mathbf{O}_{\text{ETCS+R}})\}$ with the following property. Let $type_{\text{ZFC}}$ and $type_{\text{ETCS+R}}$ stand for the respective types we outlined above: then the relation $Rxy \subset type_{\text{ZFC}} \times type_{\text{ETCS+R}}$ should be such that for every collection of ZFC-members X, R picks a collection of ETCS+R-members Y such that $|X| = |Y|$. This requirement seems easy to justify, as we are essentially mapping ZFC-sets to their isomorphism class in ETCS+R.

There are multiple ways of formally implementing this, i.e. of deciding exactly which Y it is that R should pick. The idea is that we want to select the most natural way, and that this will tell us whether R is a weak or a strong abstraction relation. We observe that letting R pick an *arbitrary* ETCS+R-collection Y for every ZFC-collection of members X is not the most attractive way. This would require some application of a Choice principle, which cannot register the actual link between X and Y — this is, however, what a GoA is intended to do. Thus, we consider another way.

We here work from the approach that takes the elements of the types of both observables to come from the union of the relevant models. If x is a ZFC-set in model M, let X_M be the collection of its members. We have that "ETCS plus replacement is equivalent to ZFC, in the strong sense of an equivalence of models" (Shulman, 2008). That is, given our model M, there exists a model M' of ETCS+R that contains an analogue x' of x. M' is created by taking the sets of M as its objects, and letting the functions in M be additionally represented as the arrows (note that, for every ZFC-set y, this will automatically give rise to the corresponding number of 'elements' for y', i.e. arrows $f : 1 \to y'$). Then define $X'_{M'}$ to be the collection of arrows $1 \to x'$; let R relate X_M with $X'_{M'}$. This gives us a unique correspondence between collections of ZFC-sets and ETCS+R-sets.

Hence, we see that the most natural implementation of R is an R1-relation, and not an R2-relation. Of course, there do *exist* R2-relations between the observables of ZFC and ETCS+R. For example, if R maps ZFC-collection X to ETCS+R-collection Y in the way just defined, we may then modify the relation by letting all ZFC-collections with cardinality $|X|$ be related to Y. This might be worth considering, as (assuming $(X, Y) \in R$) it is unclear whether an ETCS+R-set $Z \neq Y$ such that $|X| = |Z|$, apart from the difference in name, would be an inherently different choice for X to map to. We saw in Chapter 3 that the NNOs N and $2N$, while having the same cardinality, can be distinguished by properties inherent to themselves.

Hence, although we can define an R2-relation, it is not so much a *neces-*

sary one. First, because the R1-relation defined above seems the most natural and avoids potential problems like the one Tsementzis points out; second, because we could change R in the opposite way as well, by letting it relate a ZFC-collection X to all ETCS+R-collections with cardinality $|X|$. Thus, where the implementation of a relation that preserves cardinalities can be done in several ways (including R2-relations), we conclude that the natural relation satisfies only R1, and is therefore not a strict abstraction relation in our sense. Additionally, it is not even the case that R1 is *necessary*, since there exist other possible (non-R1) relations. As relations in a GoA should go together with information loss, however, the latter possible relations are not completely justifiable (even though our method does not explicitly require an indication of information loss).

Our results show that ETCS+R treats members of sets largely on the same level of detail as ZFC, except that it forgets information about the internal structure of sets. Whereas this might intuitively correspond to an abstraction relation, our two-part definition of increased abstraction and our method show that the relation expresses a rather weak increase in abstraction. Still, note that the way ETCS+R considers mathematics compared to ZFC is affected. Even though ZFC-bijectivity does not translate to equality via the natural relation we defined above, ETCS+R supports a more local perspective on membership and hides some details in the description of sets that distract from their relevant properties. Concerning abstraction relations from ZFC, however, there does not appear to be a successful or unique identification of a higher-level concept of members in ETCS+R.

Taking stock. The employed method of abstraction has allowed us to form a GoA between ZFC and ETCS+R, but it has left us without confidence in the necessary nature of the R2- and even the R1-relations. Thus, the use of a framework of abstraction has helped verify, even weaken, the intuition that ETCS+R is more abstract than ZFC. We have explored several ways of relating the two membership properties of sets, which turned out to enjoy more freedom than desired. It should be kept in mind that, despite the absence of the creation of unique higher-level concepts, our GoA does come with a process of suppression of irrelevant details, which partly characterizes an increase in abstraction. Although this part is currently not well-represented in our framework of levels of abstraction, our use of the method has still contributed to a refined comparison between ZFC and ETCS+R with respect to abstraction. We conclude that a formal approach to the concept of abstraction is required to have an accurate sense of its characterization of these two foundational systems.

Relating NBG to AST

Whereas ZFC and ETCS+R supported a rather intuitive GoA that quickly showed its strengths and weaknesses, the situation becomes a little less precise here. We examine the relation between NBG and AST, both of which incorporate class-like objects. However, although membership in NBG plays a similar role to that in ZFC, we will see that AST assigns a different role to membership than ETCS+R, and we are forced to be more resourceful in the design of observables. Still, AST has been described as capturing abstract versions of set-theoretical notions.

Awodey mentions:

> AST thus separates two distinct aspects of set theory in a novel way: the limitative aspect is captured by an abstract notion of "smallness", while the elementary membership relation is determined algebraically. The second aspect depends on the first in a uniform way, so that by changing the underlying, abstract notion of smallness, different set theories can result by the same algebraic method. (Awodey, 2008, p. 4)

This already suggests to us that both the notion of smallness and the notion of power set (that induces a 'membership relation') play an important role in shaping the perspective of AST on mathematics. In this section, we will consider two GoAs between NBG and AST. The first is rather self-evident, while the second is a little more technical; both will turn out to have their own strengths and weaknesses. Let us consider the observables.

Defining the observables. As an extension of ZFC, NBG takes into account the same definable sets as ZFC. Additionally, however, it incorporates (proper) classes definable by a property that only ranges over sets. Thus, compared to before we are including new objects, so that NBG approaches mathematics from a broader perspective than ZFC. It seems then that what we are dealing with in terms of observables is a widening of the type of the ZFC-observable. It might also be argued that classes here essentially are predicates in extension instead of actual objects, for which we cannot really talk about members and for which the observable cannot explicitly output a collection of its members. However, we note that NBG does not necessitate that classes be definable by a predicate; there exist models of NBG where 'undefinable' classes exist. Hence, we do not wish to take up the distinction between sets and classes based on predicates here. Furthermore, the definability and model approach to the type of an observable both allow us to talk about and point to the members of classes, because we always start with a background of all of the available sets and classes. This means that we regard classes as having the same fundamental nature as sets here. Since this view relies in part on the methodology we use concerning observables, we think it is still philosophically interesting to analyze the idea that classes and sets may be distinguished otherwise further. The fact that our method allows us to regard classes similarly to sets tells us that the membership observable is suitable for the LoA we create for NBG. Hence, we will introduce a membership observable $\mathbf{O}_{\text{NBG}}$ and define $\mathbf{L}_{\text{NBG}} = \{\mathbf{O}_{\text{NBG}}\}$. The various implementations of this observable are thus essentially the same as its ZFC-counterpart. And here, too, knowing the members of a set or a class in NBG will provide us with other useful properties we are interested in. In particular, note that because of the available background of objects in the model approach, we will be able to tell from an arbitrary collection of members whether it is a class or set.

For this, recall that all models in the union of NBG-models $\mathfrak{N}$ may be assumed to be disjoint. Recall also that the observable will look as follows.

> $\mathbf{O}_{\text{NBG}}$:= *members* : $\{X|X$ is a collection of sets $x \in M$ for some $M \in \mathfrak{N}\}$

Like before, for any NBG-set or -class x, the observable will output the collection containing every member of x. If the observable outputs a collection X

of members, the set or class x with these members will belong to one of the disjoint models we started with, and we can deduce which one this is from the collection of members X. From this, then, we can deduce whether X represents a set or a class. All we have to do for this is check in the relevant model whether there exists a set s such that $x \in s$ (as NBG-sets are always members of some other set). This can be verified simply by the predicate $Set(x) := \exists y(x \in y)$. Still, we should note that this argument depends on the clever use of the disjointness property of the models and it is not generally true without this assumption. We maintain, though, that knowledge of the members of a set or class should come with the framework that it belongs to, as without this context we cannot really make sense of this knowledge at all. As in the examples of the colour of a wine, the length of a human, and so on, the system that is analyzed (the wine or human) is given before analyzing a property. Hence, we argue that knowledge of the model that a collection of members belongs to is reasonable to have. If we take the approach from NBG-definable sets and classes, it is a little more complicated to deduce the property of being a set or class. Still, we may regard that here as a weakness of the definability approach.

Hence, this allows us to leave out an observable in NBG that tell us whether an object is a set or a class. Of course, we still *can* include one, but we would like our LoAs to capture only the essential properties of a system, so that within a level of abstraction it is clear what is truly characteristic of the perspective the system embodies.

Turning now to AST, we observe that the situation is a little different. The properties that define classes in AST do not resort to explicit membership, and hence there is no absolute notion of size. Each powerobject does come with a (local) membership relation as a subobject, and membership within a universe (which is a ZF-algebra) can be defined in terms of the successor operation. However, both of these notions characterize membership more externally and locally without giving an explicit way to quantify over members of classes. Hence, classes in AST are not defined by their explicit members; instead, what characterizes a class in AST is first of all the size it takes on, in the form of something we cannot derive from the knowledge of members. This results in a first, simple observable as defined below.

$\mathbf{O}_{\text{AST1}} := \textit{size} : \{\textit{small}, \textit{large}\}$

Indeed, given an AST-class X (a definable class, or a class from a given model), $\mathbf{O}_{\text{AST1}}$ tells us whether X is small or large. There exist other properties, in addition, that further characterize classes in AST. Note that while the membership observable used before is very powerful in that it determines practically all properties of a set, analogous properties in AST are not connected with or determined by a membership-like concept. Instead, to better capture classes in AST, more observables are needed. As the powerobject (which represents the 'subsets' of a particular class) plays an important part in the theory, as well as the notion of universe, we define the following additional observables.

$\mathbf{O}_{\text{AST2}} := \textit{subsets} : \{X : X \text{ is of the form } \mathcal{P}C \text{ for an AST-class } C\}$

$\mathbf{O}_{\text{AST3}} := \textit{universes} : \{X : X \text{ is of the form } V(C) \text{ for an AST-class } C\}$

Given an AST-class C, $\mathbf{O}_{\text{AST2}}$ and $\mathbf{O}_{\text{AST3}}$ provide us with its powerobject $\mathcal{P}C$ and universe $V(C)$, respectively. Note here again that quantification over subsets and universes can be done according to the various discussed methods. As we are not relying on one of the observables for additional properties, the objection against the approach concerning definable classes becomes less strong here. We can now take $\mathbf{L}_{\text{AST}} = \{\mathbf{O}_{\text{AST1}}, \mathbf{O}_{\text{AST2}}, \mathbf{O}_{\text{AST3}}\}$.

Defining a GoA. Our GoA between NBG and AST will ideally consist of relations $R_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST1}})\}$, $P_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST2}})\}$ and $Q_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST3}})\}$. However, whereas $R_{\text{NBG,AST}}$ is easily implemented, possible relations for $P_{\text{NBG,AST}}$ and $Q_{\text{NBG,AST}}$ are harder to motivate. This suggests that the latter observables do not capture a direct abstraction of NBG-membership. Although they are essential for describing classes in AST, they do not group collections of members under one concept. We will elaborate on this later and the implications this has for our method of abstraction. First, consider the more obvious relation $R_{\text{NBG,AST}}$.

$R_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST1}})\}$ rather easily allows for connections between the respective types. Namely, we define the relation $Rxy \subset \mathit{type}_{\text{NBG}} \times \mathit{type}_{\text{AST1}}$ as follows. Take a collection of NBG-members X. Then if this collection is a proper NBG-class, we let R relate X to *large*. Otherwise R relates X to *small*. Clearly, this has to result in an R2-relation, as *small* and *large* cannot both be related to a particular NBG-set or proper class. Thus, this GoA appears to express a stronger difference in abstraction than the one between ZFC and ETCS+R. Of course, this idea is quickly undermined by the previously made observation that one can create a similar observable for *size* in NBG, with type $\{\mathit{set}, \mathit{class}\}$, which could have a reversed GoA mapping small AST-objects to *set* and large ones to *class*. However, we have noted before that the notion of size in NBG is derivable by the knowledge of its members. That is, NBG contains more information about its sets that makes it more natural (and more characteristic of $\mathbf{L}_{\text{NBG}}$) to create an encompassing observable for *members*, than it is to create one for the distinction between sets and classes, as this leaves out much information that NBG characteristically gives. Hence, we argue that the size observable is more suitable for AST. Still, the matching of *members* with *size* may not be completely natural — this will come to light even more with the next observables, and we explore this further there.

For the relations $P_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST2}})\}$ and $Q_{\text{NBG,AST}} = \{(\mathbf{O}_{\text{NBG}}, \mathbf{O}_{\text{AST3}})\}$, we will remain a little more informal here. Despite various possibilities for the relations, we have not been able to properly justify any of them. We will argue that this is because of the limits of the method of abstraction for our purposes, and that this gives incentive to add an additional requirement to make sure related observables match up. It is interesting to consider, however, possible relations between our current observables: consider the following options.

1. As each model of NBG gives rise to a model of AST when we take 'set' to mean 'small' (Shulman, 2008), the approach that defines the type as the union of all models of the respective theories will give us a natural way to relate models of NBG to those of AST. This gives rise to a relation as follows. If x is an NBG-set or -class in model M, let X_M be the collection of its members. Furthermore, let $x'_{M'}$ be the analogue of x in the AST-model M' induced by M. Then it is a natural choice to let $P_{\text{NBG,AST}}$ relate X to $\mathcal{P}x'_{M'}$.

Clearly, a similar relation works for $Q_{\text{NBG,AST}}$. Similarly to the situation concerning ZFC and ETCS+R, we can modify the relation (so that it satisfies R2) by making all NBG-sets map to the same powerobject or universe. However, if the AST-model induced by an NBG-model can tell the difference between all of its objects, this is not completely justifiable. In that case, this type of mapping gives us at most an R1-relation.

2. Alternatively, we can consider the definable sets and classes of our theories. Note that AST gives us the existence of the object 0, from which the theory may define a 'hierarchy' of objects $\mathcal{P}0$, $\mathcal{P}\mathcal{P}0$, and so on. These objects may be regarded as corresponding to the finite levels of the cumulative hierarchy (indeed, finite, as AST only guarantees finite coproducts). Then a straightforward $P_{\text{NBG,AST}}$-relation is defined as follows: if a collection X of NBG-sets is an element of V_α (with α finite) in the cumulative hierarchy in NBG, $P_{\text{NBG,AST}}$ maps X to the AST-object $\mathcal{P}^\alpha 0$. This gives us an R2-relation, as we are essentially mapping all sets of a particular hierarchy level to the same object. However, clearly this only works for the finite levels of the hierarchy, and we are left without an obvious correspondence for bigger NBG-sets. The case for $Q_{\text{NBG,AST}}$, concerning universes, is even less clear. There exist NBG-classes (like V, but also others) that act as universes — these classes could be justifiably related to AST-universes. Then, however, sets in NBG are here the questionable cases, as there is no definable universe that naturally corresponds to an NBG-set.

Matching up the observables. Note that in both possibilities for P and Q, we are relating the broad *members* observable of NBG to the more particular AST-observables of powerobjects and universes. That is, where every NBG-object can itself be a member, not all AST-objects are instances of powerobjects or universes. Furthermore, while the *members* observable, because of its detailed information, allows us to deduce properties of sets (like their powersets, or the hierarchy built on them) *from* their members, the loss of information about specific elements in AST forced us to separate these notions into new observables. This makes it rather unnatural to relate the *members* observable to only powerobjects or universes, as the abstraction relation should ideally capture abstractions of particular observables. That is, a GoA from NBG should capture an abstraction of the notion of 'member' (something that was clearly there in the relation between ZFC and ETCS+R). Hence, as our observables do not match up right now, the relations we defined are not very reliable. Note that the same applies to the GoA relating *members* to *size*. However, this GoA allows more naturally for a relation as, contrarily to the *subsets* and *universes* observables, every AST-object does take on an instance of *size*.

Still, the relations we managed to define here are informative to some extent. By mapping a collection of NBG-members (defined with detail) to a powerobject (of an arbitrary class or one representing the appropriate level of the cumulative hierarchy) or a universe, we are still given information about how our knowledge of an NBG-object differs with that of a particular AST-object. For example, whereas $\mathcal{P}$ in NBG explicitly picks out all possible subsets of a set, in AST it defines an object $\mathcal{P}C$ that helps define the 'universal small relation' on C, so that the internal logic can think of it as a set of subsets. Similarly, the free

$\mathcal{P}$-algebras give us properties that we associate with universes, such as that they capture all powersets of their elements.

Of course, this is better formalized by a match-up of observables. For ZFC and ETCS + R, the observables happened to match up, whereas the relation here requires an evaluation first. This suggests, then, that observables should be defined *after* deciding which two LoAs to connect. However, this is not an obvious requirement to make. For NBG then, we could make observables, and relate them to AST as follows.

1. We keep the notion of membership, so that we have our original observable *members*. Note that AST says that membership of a class is the universal small relation on it: each class C comes with a subobject $\in_C$. Then, for a collection of NBG-members X, we would like to relate it to the subobject $\in_C$ of a particular class C. As before, however, we do not have a natural way to pick C. And, as AST contains many of such subobjects, this will not naturally give rise to an R2-relation. We do see, however, that information is lost and we focus on the *essential* properties of membership. Instead of pointing to particular members, AST-membership is specified by a universal small relation (see Chapter 3); i.e. the relation satisfying just the important properties and nothing else.

2. As mentioned before, we could introduce a size observable *size* of type $\{set, class\}$. This leads to the one-to-one (R1-)relation mapping *set* to *small* and *class* to *large*. Still here, too, we note a loss of information: X is an NBG-set if there is no bijection $f : X \to V$ (V being the class of all sets); hence, this relies again on membership, in particular on cardinality. AST-smallness, on the other hand, is defined from the smallness of arrows towards 1. This relies on the uniqueness of this arrow, instead of absolute size, which is irrelevant.

3. Third, we could create the NBG-observable *powerset* (or similarly *universes*), essentially selecting out of the type of the *members* observable all collections that contain all subsets of a set. The relation to AST with this observable remains R1 at most. However, we do see that, while the members of NBG-powersets $\mathcal{P}(X)$ relate to the elements X in a unique way, this is also captured in AST. For a class C, $\mathcal{P}C$ should capture the idea that it contains all subsets of (i.e. monomorphisms to) C. This is done by letting $\in$ be a subobject to $C \times \mathcal{P}C$.

Thus, matching up of observables makes the relations in GoAs at least more justifiable, and it pinpoints where information is lost. In the case of NBG and AST, however, it still does not elicit strong abstraction relations.

Taking stock. The method of abstraction applied to NBG and AST has led us to an R2-relation concerning the size of objects. By moving from NBG to AST, we lose information about the precise cardinality of a set or class, and we introduce the higher-level concepts *small* and *large* that cannot even talk about members or cardinality. The notions of power set and universe, according to our method, do not elicit a true abstraction relation with the NBG-observable. Hence, the method has allowed us to pinpoint which aspects of AST contribute more and which contribute less to an increase in abstraction. Note, as well, that compared to set theory, the loss of information aspect of abstraction is present more strongly in AST than ETCS, as we lose information about the precise cardinality of sets (and classes) as well as their internal structure. This section has also given us incentive to reconsider our method of abstraction and encourage the matching up of observables.

Bringing CCAF into the picture

Other than ETCS and AST, which are theories that explicitly resort to set-theoretical notions, we would like to apply the method of abstraction to CCAF. McLarty states about this theory: "I regard these axioms as a background theory to be used with axioms on particular categories or functors" (McLarty, 1991), and he suggests one of these particular categories could model ETCS. As CCAF obtains full potential with such additional assumptions, we are interested in the observables of CCAF by itself, as well as those of ETCS incorporated in CCAF. Similarly, the inclusion of a model of AST will be of interest. We will see that the *members* observable of previous theories allows for a GoA with CCAF — although its further separation from set-theoretical notions leads to the 'matching-up problem' of observables we discussed before. The set-derived notions from AST face this issue as well. In contrast to the situation of the previous section, there is no intuitive way to remedy this problem. However, although CCAF does not reveal R2-relations from our GoAs, it does allow for a very flexible and neutral theory where the loss of information aspect of abstraction is clearly present. Let us consider how we can introduce observables for CCAF.

Defining the observables. Recall that CCAF provides the given categories $\varnothing$, **1**, **2**, **3** and **E**, which act as tools with which one can determine the objects and arrows of additional categories. The axiomatization of CCAF focuses on these categories, while specifying few specific properties of other categories, or their objects and arrows. Just like we took sets to be the primary object of study in ETCS, and classes in AST, we take categories (with their objects and arrows defined as particular functors) to be the primary type of object in CCAF. Then our observables should say something about categories. We could also imagine observables that act on particular functors (for example an observable for *sort* with type $\{object, arrow\}$), but we will here be consistent with the idea that the observables should be applicable to the primary subject matter of CCAF, i.e. entire categories. The following two observables then come to mind (illustrated by the approach of taking the union $\mathfrak{C}$ of all models of CCAF).

> $\mathbf{O}_{\text{CCAF1}}$:= *objects* : $\{X|X$ is a collection of functors with domain **1** and codomain a particular category $C \in M$ for some $M \in \mathfrak{C}\}$
>
> $\mathbf{O}_{\text{CCAF2}}$:= *arrows* : $\{X|X$ is a collection of functors with domain **2** and codomain a particular category $C \in M$ for some $M \in \mathfrak{C}\}$

That is, when we put in a particular category (represented in CCAF by a characterless point), these observables will give us the collections of its objects and arrows. Alternatively, we can combine these observables into one observable $\mathbf{O}_{\text{CCAF3}}$ called *structure* and let it output the collections of objects and arrows of a particular category simultaneously. Note that essentially, *structure* gives you the same information as *arrows* (as each (identity) arrow comes with a domain and codomain object already), only more explicitly. Hence, both $\mathbf{O}_{\text{CCAF2}}$ and $\mathbf{O}_{\text{CCAF3}}$ will give us a complete characterization of categories in CCAF, and we define without loss of generality $\mathbf{L}_{\text{CCAF}} = \{\mathbf{O}_{\text{CCAF3}}\}$.

Defining possible GoAs. We can create several GoAs relating $\mathbf{L}_{\text{CCAF}}$ to the set theory LoAs $\mathbf{L}_{\text{ZFC}}$ and $\mathbf{L}_{\text{NBG}}$, and to category theory LoAs $\mathbf{L}_{\text{ETCS+R}}$ and $\mathbf{L}_{\text{AST}}$. For each GoA to $\mathbf{L}_{\text{CCAF}}$, its observable $\mathbf{O}_{\text{CCAF3}}$ requires us to relate the type instances of the other LoA to a collection of objects and arrows specifying a whole category. With this in mind, let us consider the *members* observables of ZFC, NBG and ETCS+R first. Two possible GoAs come to mind: one relates a collection of set members to a category 'representing' that set — the other embeds the collection into a bigger category of sets representing the model it is part of. In the spirit of matching up our observables, however, our GoA should capture an abstraction in the representation of an individual set; hence, we omit the latter option. We treat all three *members* observables together. Hence, without loss of generality we will define $R_{\text{ZFC, CCAF}} = \{(\mathbf{O}_{\text{ZFC}}, \mathbf{O}_{\text{CCAF3}})\}$. We mention the following options for the implementation of the relation between types (note that other possibilities exist, but we feel these capture the most intuitive ones).

1. Let X be a collection of ZFC-members. Then R relates X to the discrete category whose collection of objects equals X and whose arrows are only identity arrows. Note that there could surely exist multiple discrete categories of the same cardinality in CCAF (whose objects are different), but (certainly in the union of models approach) there will always exist a category whose objects are 'named' precisely after the elements of X. The uniqueness of models guarantees the uniqueness of this category. Thus, R can map X to a natural unique category. Still, this will only leave us with an R1-relation.

   And once more, like in the case of ETCS+R, we can modify this relation to create a (non-neccessary) R2-relation, by relating all ZFC-sets X of a particular cardinality to the same discrete category of that cardinality. The similarity to ETCS+R is highlighted even more by considering what this CCAF-representation of sets means for membership. Suppose we represent a ZFC-set X by a discrete CCAF-category. Then $x \in X$ if and only if $x : \mathbf{1} \to X$, i.e., x is a functor in the metacategory *CAT*. This is of course precisely the ETCS+R-notion of membership, only on a metacategory level. This shows that we are not given a higher level of abstraction than ETCS+R by this representation of sets.

2. Of course, R could also send X to an ETCS+R-style representation of its set. Alternatively, X could be related to a category that forms the tree representing the membership structure of X; the latter representation is described in (McLarty, 2004). Both ways display the membership structure of X a little better. However, the former option 'borrows' a method that is arguably not inherent to CCAF — and the GoA becomes essentially the same as the one defined for ETCS+R. The second option is no different, as it will relate X to the unique tree-structured category that has objects named just like the members of X. Again, we are left with only a justifiable R1-relation.

Hence, the membership observables do not very naturally form GoAs that capture an abstraction relation. For option 1, we obtain a non-necessary R2-relation, where the CCAF-representation of an ETCS+R-set seems to be essentially a rewriting of notions that possess the same level of information. For ZFC- and NBG-sets, then, we lose knowledge of the internal structure of a set, but retain the cardinality. Still, we note that the categories representing membership are here represented as particular functors. Hence, while we lost information with other methods about the internal structure of sets and their cardinality, we keep knowledge about cardinality here. Note that in the metacategory *CAT*, sets themselves are represented as (characterless) points ('categories'), while their members are the functors from category **1**. Alternatively, through commutativity properties, the members correspond to the identity arrows on themselves, i.e. the corresponding functors from **2**. In that sense, compared to the ETCS+R-representation of sets, CCAF provides unnecessary information by supporting a double representation of members of sets with the inclusion of arrows.

For AST, we find ourselves in a little more trouble again. Individual representations as categories of the type instantiations *small*, *large* do not seem to make sense. Similarly, relating individual powerobjects or universes to categories does not seem very useful at all. While *small* and *large* do not belong to a particular set or even a particular model, powerobjects and universes have no internal structure or explicit cardinality, so that a representation of these as a category would result in a singleton category for each of them. This leaves us with few options to naturally obtain a relation to a category. The problem, of course, is that CCAF does not present itself with notions of size differences, power operations or universes.

We conclude that the correspondence of the AST-observables to the notion of a category is not a very fruitful one, while the relation of the *members* observable to CCAF reveals a level of information about sets similar to ETCS+R. Even if we incorporate a model of ETCS+R and AST as whole categories **Set** and **Class** (respectively) explicitly in CCAF, it is quickly seen that the 'functorization' of objects and arrows from ETCS+R and AST retains basically all properties they had before. In this case, as there exists exactly one functor for a member of an ETCS+R-set, the cardinality of such sets is retained. Similarly, we should be able to detect properties like smallness, powerobjects and universes with the representation of an AST-model in CCAF. The only point of interest is that ETCS+R *explicitly* calls particular arrows 'members', whereas in CCAF it is only the functors from **2** to our category **Set** that represent members, and it is not inherent to the theory itself. In that sense, CCAF is perhaps more flexible with respect to the notions that it can incorporate, but its axiomatization does not provide an abstraction of these theories.

Relation of CCAF to EM. The consideration of EM adds two things here. First, we would like to observe that EM, by simply providing the ingredients of a category with its axioms, is similar to CCAF, only does not consider several categories at once with relations between them. That is, each EM-category in a way exists by itself. This suggests that, as for ETCS+R and AST, its objects of study consist of the individual objects and arrows of a category (instead of the category as a whole). Still, the EM-observable that can give us most information provides, for a given object, the objects and arrows it is (not) related to, thereby characterizing the entire category. Hence, EM-observables can essentially provide the same information as the CCAF-observables; they are merely given from a different perspective. Then, the relations (and their implications) between the membership observables of ZFC, NBG and ETCS+R, as well as the other AST-observables, are similarly defined as in the CCAF case, and we will not repeat them here.

Second, if we want to relate CCAF and EM to each other with a GoA, the natural relation is one-to-one, mapping a CCAF-category to itself in EM. However, we may also relate these two theories to each other via a different route, to more explicitly capture their differences. For this we create an observable similar to one that we rejected before, because it acted on CCAF-functors instead of its categories. Here, however we implement it as an observable that acts not on categories, nor on functors, but on both of them. We define $sort_1 : \{category, functor\}$ for CCAF and $sort_2 : \{object, arrow\}$ for EM. Then an obvious (R2-)GoA from EM to CCAF relates both *object* and *arrow* (of a category, not a metacategory) to *functor*. Thus, we see here that CCAF seems to abstract the notion of category by representing the structure of categories merely with the notion of functor. Of course, among functors, we can still differentiate between objects and arrows, but it shows that we may think of categories as containing a single sort of object only. The loss of information for this GoA concerns the precise identity of 'objects' (of a *category*). Whereas in EM, objects have their own independent identity, in CCAF they are defined as compositions of arrows with the *domain* or *codomain* arrow. Indeed, a CCAF-object is defined by all of the arrows it is a codomain or domain of. The arrow that describes an object the most, then, is the identity arrow, for which the object is both the codomain and the domain. Still, this is a more subtle loss of information than, say, the cardinality of a set. Hence, even though our GoA is R2, the lack of clear information slightly weakens the strength of the abstraction relation.

Taking stock. The method of abstraction has here shown us that CCAF surely has the resources to capture as much detail about the objects of ETCS+R and AST as the respective theories themselves. That is, there is no necessary and natural relation that captures a true abstraction process for their observables. In the case of set theories like ZFC, the abstraction relations of members are similar to those of ETCS+R, and give us only justifiable R1-relations. We note, however, that CCAF is not intended to represent sets by itself (unlike ETCS+R and AST), and is not designed to provide abstractions of them. This gives us the same problem as before, where our observables are arguably not harmonized well enough to provide a suitable GoA between them. We finally found an R2-abstraction relation between the *sort* observables of EM and CCAF.

Thus, our method of abstraction suggests that CCAF shows variable behaviours when related to other theories. This tells us that foundations of mathematics cannot be taken to possess an absolute value with regard to their level of abstraction, but that this is a notion that should be considered relative to other foundations.

An application to UF

We shortly elaborate on the relation of the method of abstraction to UF. Recall that the Univalence Axiom says that 'identity is isomorphic to isomorphism'.

This may essentially be taken as implying that, if two objects have the same structure, they are identical. That is, mathematical objects *are* structures—this makes UF philosophically so attractive for structuralism. We saw before that we can take a suitable version of isomorphism in a particular context, and then allow for strict identity of objects to not be tracked. Awodey says the following about the implication of this for objects.

> Rather than viewing it as identifying equivalent objects, and thus collapsing distinct objects, it is more useful to regard it as expanding the notion of identity to that of equivalence. For mathematical purposes, this is the sharpest notion of identity available; the question whether two equivalent mathematical objects are "really" identical in some stronger, non-logical sense, is thus outside of mathematics. (Awodey, 2014, p. 10)

In order to define GoAs with respect to UF, of course, our method of abstraction encourages the 'collapsing' way of thinking. Take two objects $\{\varnothing\}$ and $\{\{\varnothing\}\}$ in ZFC, for example. Then ZFC tells us these are distinct: the collections X and Y containing $\varnothing$ and $\{\varnothing\}$, respectively, are two different instances of the type of the *members* observable. When relating these by means of a GoA to a *members*-like observable $\mathbf{O}_{\text{UF}}$ of UF, however, without working out the details we may suppose that both X and Y are (R2-)related to the same instance of the isomorphism-invariant type of $\mathbf{O}_{\text{UF}}$.

In that sense, UF has us lose information about the strict identity of objects, which is now a notion only present external to the theory. The new sense of identity can be seen as the introduction of a higher-level concept in UF. However, recall that the isomorphism-invariant notion of identity is caused by a restriction in the syntax of the theory. Hence, we argue that this R2-relation may not be viewed as inherently caused by category theory. We consider UF as a very interesting and promising theory, but perhaps not of use for making a distinction between set-theoretical and categorical ways of thinking. This claim, however, should be investigated further, something we lack the time for in this thesis.

4.2.3 EM-category theory from set theory

Various processes that induce a loss of information, as well as some identifications of higher-level concepts, have shown to characterize the categorical approach to mathematics compared to the set-theoretical approach. This result has been obtained by taking the perspective of foundations for mathematics, where it was regularly the case that categorical notions derivable from set theory were responsible for the increase in abstraction. The method was successful, then, in capturing abstraction relations between concepts that are already closely connected. On the one hand, this makes sense, as something can really only be abstract with respect to something else. On the other hand, this still leaves us without a good way to compare the level of abstraction of set theories to that of (parts of) AST, CCAF, higher categorical foundations and EM.

In order to provide additional support to the results obtained from categorical foundations, then, we once more take the general notion of a category as in Chapter 2. As we there investigated which properties of set theories allow for more category theory built from this basic definition, we can here investigate the level of abstraction of various set theories. Namely, if general EM-category theory (which provides the basis for most categorical foundations) reflects a higher level of abstraction compared to set theory (but our current method of abstraction is unable to show this), we should be able to detect a change in abstraction between set theories that support category theory in different ways. Indeed, if a set theory allows for a lot of category theory after the addition of particular axioms, these axioms might according to our method of abstraction add a new observable or change the type of an old observable. While we already argued that NBG adds objects to the type of the observable for ZFC, we also claim that the assumption of an inaccessible, and the NFU-based system S^* have a similar effect. Reflection principles, on the other hand, more strongly involve a change of the actual values of the type of *members*, instead of just adding to it. We suggest that the first process generally allows for increased complexity of objects, but not a higher level of abstraction. The second process, however, may represent a higher abstraction level compared to ZFC. We argue as follows.

1. Adding an inaccessible to ZFC explicitly increases the size of the universe, and hence of the type of the *members* observable. Thus, we can have the 'inclusion' R1-GoA between ZFC and ZFC+I, and similarly between ZFC and NBG or S^*. Note that we (as mentioned before) take classes to embody a similar type of object as sets concerning their members, as the different approaches we take towards the type of *members* explicitly tells us what members we have available — this allows for a similar treatment of sets and classes. These GoAs, however, do not come with either loss of information, or identification of higher-level concepts as indicated by an R2-mapping. What is new, instead, is the incorporation of new, larger objects that are built out of smaller sets. This indicates the introduction of new *complexity* levels; however, there is no sign of an abstraction process that generalizes the notion of membership. Still, this type 'enrichment' allows for the construction of larger categorical objects than before, and might be regarded as an 'illusion' of a higher level of abstraction. It can boost the capacities for supporting category theory by upping the limit on size tremendously. However, all the level of detail of sets and their members is retained.

2. The reflection principle in ZFC/S provides a new kind of independence from size, however, as 'anything we prove in ZFC/S about small objects is also true about large objects' (Shulman, 2008). That is, the small sets can be viewed as being inherently 'coupled' to the large sets. Hence, the type of *members* for ZFC/S is not increased by adding new instances; instead, we add some sort of meta-information to the types of small sets that tells us we can transfer results to large sets. In that sense, the type instantiations do embody a kind of higher-level object, because of the 'internal relating process' of sets of different sizes. This is more like an abstraction relation than just a size increase; we treat multiple different objects *similarly* with regards to their properties, and ZFC/S talks about a concept of set that is independent from strict cardinality. However, as ZFC/S does not literally define 'set' as this more encompassing concept, the reflection principle may be said here, too, to provide the illusion of increasing the

level of abstraction. In reality, however, ZFC/S-sets possess the same level of detail as ZFC-sets.

Recall that ZMC/S, which combines the inaccessibility assumption and use of a reflection principle, was arguably the most successful set-theoretical foundation for category theory. By using both ways of 'simulating' a higher level of abstraction, then, it seems that EM-category theory is accommodated better. The two processes of increase in type complexity, and an internal coupling of type instantiations, however, may be viewed as 'abstraction-like' processes that thereby assist further development of category theory. However, we have not seen any truly convincing instances of abstraction relations among the set-theoretical foundations for category theory. This supports the idea that set theories in general occupy a similar level of abstraction, whereas categorical foundations show more variation.

4.2.4 Taking stock

The results from the method of abstraction applied to categorical foundations and counterparts reveal several things. Almost none of the GoAs we constructed are necessary R2-relations, except the relation between the *size* observables of NBG and AST, and the one between the *sort* observables of EM and CCAF. Relating set-theoretic membership to ETCS+R and CCAF, furthermore, allowed for natural R1-relations and non-necessary R2-relations. Together with the knowledge of where information is lost, this gives us a newly found characterization of abstraction relations between set-theoretical and categorical foundations. Generally, then, we see that categorical foundations do find themselves on a higher level of abstraction than set-theoretical theories. However, the different axiomatizations of categorical foundations allow for quite some variety in the implementation and strength of GoAs. As the only explicit R2-abstraction relation between a set theory and a category theory here concerns that of small- and largeness, we conclude that the difference in abstraction between set-theoretical and categorical foundations is less big than generally thought. Most abstraction relations do not involve a necessary collapsing of type instances onto another type instance.

Furthermore, we have seen that the method of abstraction needs improvement in several ways. For one, it needs a better way of denoting the information that is lost by going from one LoA to another. Second, we saw that a better framework is called for in order to have GoAs only be defined between observables that are properly matched up. Possible ways to improve the method are further discussed at the end of this chapter. Right now, we will briefly discuss the purpose of set-theoretical and categorical foundations.

4.3 The role of foundations

We end this chapter by addressing the role of set-theoretical and categorical foundations, and how the method of abstraction relates to these. Instead of considering all perspectives on the goals of foundations from scratch, we focus mainly on the perspectives discussed in (Landry, 2013). Based on these views, we refine the idea that set-theoretical and categorical foundations may be distinguished by their purpose. We end by advocating the naturally arising position of foundational pluralism in mathematics.

Recall that Landry takes Awodey's distinction between bottom-up and top-down (that we redefined) and relates it to two different roles of foundations. Indeed, Landry argues that a (set-theoretical) bottom-up foundation for mathematics has a constitutive role, whereas a (categorical) top-down foundation for mathematics has an organizing role. According to Landry, a *constitutive* foundation is concerned with

> [...] constructing the structure of concepts by beginning with some fixed domain of facts as its constitutive subject matter (Landry, 2013, p. 41)

We note that this can be related to the (philosophical) goal of foundations described in (Shapiro, 2004) to create a suitable ontology for mathematics (although this goal does not necessarily pursue the 'one correct' ontology). That is, the constitutive role of a foundation gives us an account of exactly what our structures are made up of, which sheds light on the nature of mathematical objects. On the other hand, an *organizing* foundation

> [...] "takes place through" the axiomatic method, which itself aims to "structure" concepts in terms of their relations, and so organizes or founds "the facts" that fall under such concepts by beginning with the axioms. (Landry, 2013, p. 41)

The axiomatic method is a notion borrowed from Hilbert, who defines it in a way that reminds of Awodey's notion of top-down. That, is we start with the assumed existence of elements, and then relate these to each other by means of (implicit) axioms. We observe that this is perhaps more related to the (mathematical) goal of foundations treated in (Shapiro, 2004), which is to serve mathematics by providing insights into mathematical fields. In what follows, we will argue for the following claims. Since we argued against the description of 'top-down' as in (Awodey, 2004) and (Landry, 2013), we do not think that the described sense of an 'organizing' role of mathematics inherently relates to the term 'top-down'. Rather, we suggest that this 'organizing' role of mathematics may be more connected to levels of abstraction that are higher than needed for the constitutive role. We then note that this will still not allow us to fully associate categorical foundations with an organizing role and set-theoretical foundations with a constitutive role. Although we regard this association helpful for a *general* characterization of the current foundations, we argue that a literal interpretation resembles a rather hasty generalization that hides the variety in roles and levels of abstraction between individual foundations. In line with this, we will claim that the distinction between the organizational and constitutive role should not be thought of as *inherent* to the distinction between categorical and set-theoretical foundations.

First, note that we showed earlier in this chapter that neither the bottom-up, nor the top-down approach is exclusive to set-theoretical or categorical foundations. In line with this, we argue that they are both applicable to the constitutive role and to the organizational role of foundations. Namely, we argued that both 'bottom-up' and 'top-down' require knowledge of what their structures are made up of, but that they express a difference in the direction of changes in complexity. The constitutive role of foundations ensures that we know the precise nature of our objects in the sense of their construction; hence, with the top-down and bottom-up method, we can work our way up or down with particular axioms to characterize our objects. On the other hand, the organizational role of foundations does not require knowledge of the precise make-up of mathematical objects. Instead, the relations that the axioms imply between systems that instantiate these axioms are primary. This should quickly remind us of a difference in abstraction. An organizational foundation should highlight the essential properties and concepts belonging to mathematics, where these are understood without the exact notion of the construction of its instantiations. However, we maintain that both bottom-up and top-down ways of thought are *possible* for an organizational role, in the sense that axioms can express changes in complexity levels. Still, for organizing notions and capturing relations between them, full detail of their nature is not necessary. In that sense, a loss of information is taking place when going from constitutive foundations to organizing foundations, as the former by definition need full knowledge of their objects, while the latter do not. As the notions of an organizing foundation, then, are independent from their specific make-up, they 'collect' more constitutive notions under (imaginably) higher-level concepts. That is, we argue that the role of an organizing foundation should be recognized as approaching mathematics from a higher level of abstraction compared to the role of a constitutive foundation.

We take care, however, not to bluntly associate set theory or category theory with either side of this distinction. In the previous section, we saw that categorical foundations rather consistently demonstrate some kind of loss of information aspect when compared to set-theoretical counterparts. Categorical notions, then, automatically take on a more general nature. However, it does not always seem to be the case (based on our current method of abstraction) that there is an identification of higher-level concepts. This has led us to conclude that the categorical foundations we discussed find themselves on lower levels of abstraction than generally assumed, although still on higher levels than the various set theories. Here, then, we argue against Landry's association of set theory and category to a constitutive and organizational foundational role, respectively, as the distinction is simply not that clear. With respect to the current foundations we discussed, however, the *general* pattern seems to be that categorical foundations approach mathematics from a higher LoA than set-theoretical foundations—as a consequence, this entails that categorical foundations generally succeed better in organizing mathematics, and less in providing a constitutive account of mathematics, than set-theoretical foundations.

However, we stress that this general pattern found in this thesis should be taken to apply to the foundational systems we discussed here, but perhaps not to any set-theoretical and categorical foundation in general. That is, there should be more investigations into whether the better suitability of the organizational role to categorical foundations, and that of the constitutive role to set-theoretical foundations, is in fact intrinsic to the distinction between set theory and category theory. Take, for example, the EM-axioms, which Landry adopts as the ultimate structural (and organizational) foundation for mathematics. We might imagine that a similar axiomatization of a set theory is possible, where we start by defining the ingredients of a set (which are other sets), and add Extensionality (perhaps Foundation, too, or other suitable axioms) as laws. Without the explicit existence of elements, and the existence of power sets, unions, and so on merely as instantiations of the definition of 'set', we seem to come close to a set-theoretical foundation with an *organizing* role. Whether this is truly feasible is yet unclear, but it opens up the suggestion that the organizational and constitutive role of foundations may be more a property of the formulation of axioms, instead of the axiomatized theory itself. For now, however, we maintain the tentative conclusion that categorical foundations are more characterized by a higher level of abstraction than set-theoretical foundations, and are hence more susceptible to taking on an organizing role towards mathematics.

This perspective finally leads us to advocate a sense of foundational pluralism. Namely, taking one particular (categorical or set-theoretical) foundation for mathematics as the 'correct' one now seems to come down to saying that mathematics is associated with certain levels of abstraction and not with others, and that it has a preference over constituting its objects versus organizing them. This should not agree with our intuitions about mathematics: instead, what characterizes mathematics is the *combination* of various levels of abstraction and goals to analyze its objects, which enriches our understanding of them. This is where the purpose of a mathematician comes in, as well, as (s)he can then pick a foundation for mathematics, depending on the level of abstraction or the perspective (s)he is looking for. This chapter has shown us that variation is abundant among foundational systems, and that each system comes with its own strengths and weaknesses for different purposes. Then, provided one remains clear about their purposes and is aware of the relations between foundational systems, foundational pluralism should allow for a better adaptation of methods to the goal one is working towards.

4.4 Discussion

Summing up. In this chapter, we have explicitly cleared up the difference between a bottom-up and a top-down approach to mathematics, arguing against the use of 'top-down' employed by Awodey and Landry. Subsequently, we saw that both set-theoretical and categorical foundations allow for bottom-up and top-down ways of thinking. We concluded that this distinction is not inherent to the difference between set-theoretical and categorical thinking.

Next, we embarked upon an analysis of the levels of abstraction of various foundations, by using the method of abstraction (with an additional requirement) of (Floridi, 2013). Our results show that ETCS + R, CCAF and EM are more abstract than ZFC, but fail to present a true abstraction relation. Furthermore, AST seems partly more abstract than NBG, as shown by an R2-relation, although we lack proper ways of relating a different part of AST to NBG. Additionally, we found an R2-relation showing that CCAF is more abstract than EM. These relations were obtained from observables that naturally arise from the perspectives of the relevant theories on mathematics.

Finally, we suggested that the level of abstraction of various foundations corresponds to the role they intend to play. Where Landry has argued that the constitutive role and the organizational role of foundations for mathematics correspond to a bottom-up and a top-down approach to mathematics, respectively, we take a different stand. We propose that the constitutive role of foundations is more characterized by lower abstraction levels, whereas higher abstraction levels better capture the organizational role.

Take away. This chapter has taught us several things. First, we have seen that it is important to maintain clear definitions of concepts that can have multiple interpretations. The interpretation of a concept that one adopts can quite drastically change the philosophical ideas that arise from it. We have seen this in the first part of this chapter, where Awodey's structuralist interpretation of 'top-down' led to the idea (argued against by us) that category theory is not identifiable with bottom-up characteristics.

Second, we have seen that categorical foundations and set-theoretical foundations show varying behaviour with respect to the two conceptions we analyzed. Variation among set-theoretical and categorical foundations, that already became apparent in Chapter 2 and 3, has here shown to truly matter. It explicitly revealed that the categorical and set-theoretical way of thinking cannot be pinned down by the terms 'bottom-up' and 'top-down', and that categorical foundations cannot be regarded as expressing one level of abstraction. The *general* pattern of categorical foundations as expressing a higher level of abstraction than set-theoretical foundations, however, gains strength from our approach, as it persisted despite the varying properties of categorical foundations. Still, as this result is a generally accepted one, we simultaneously lay emphasis on the fact that the more nuanced picture shows variable abstraction levels and several rather weak abstraction relations with set-theoretical foundations.

Third, we have seen that the level of abstraction of theories is an interesting area that needs further research. We suggest that, to allow for a better comparison of set theories to categorical theories, the currently used method of abstraction be adapted and specialized to our purposes. The aim here is to formalize the method such that any GoA that can be defined with it provides a truly reliable and informative abstraction relation. Our added requirements R1 and R2 have already started to provide an indication of the strength of abstraction relations, so that there are more resources with which to describe and compare the significance of relations. Furthermore, we recommend that (among others) the following aspects are addressed. First, the method of abstraction should make sure that only observables which already display similarities are related by a GoA. To start, then, it could be required that the names of observables have the approximate same meaning. Second, the method should explicitly require its user to identify a loss of information along with a GoA-relation between types. This can simply be implemented as an annotation, or by using more formal methods. Finally, the method can be made more reliable in capturing the essence of a system, by providing guidance on how to introduce observables. Of course, the appropriate way of doing this differs per system that is analyzed. This entails that such guidance should be offered and justified separately for each system. For foundations of mathematics, it can for example be recommended to consider the individual axioms and, in a more determined way, select out of these essential properties of the structures they define.

What is next? We conclude this thesis by evaluating our results and relating them to future research in the final chapter.

Chapter 5

Conclusion

After an excursion into various set-theoretical and categorical foundations for mathematics and an analysis of two commonly held conceptions concerning the distinction between their approach to mathematics, it is time to come back to the original research question and evaluate our approach.

Answering the research question. This thesis concerned itself with characterizing the distinction between set-theoretical and categorical approaches to mathematics. We contribute to the attainment of this goal in two ways. First, we narrow down the scope of the research question by arguing that it should exclude the bottom-up/top-down distinction from its possible answers. The main arguments supporting this claim rely on the varying effects of set-theoretical axioms on the bottom-up or top-down conception of set, the varying (sometimes even neutral) ways that categorical axioms act on complexity, and the fact that the development of EM-category theory is not affected by changes in the bottom-up or top-down nature of the set theories taken as its foundation. We surveyed the necessary knowledge for these arguments in Chapter 2 and 3. Second, we put forward a candidate positive answer to the research question with the suggestion that categorical foundations possess a higher level of abstraction than set-theoretical foundations. This is the first time that this result has found support from a formalization of levels of abstraction. Furthermore, the method has revealed a general consistency in levels of abstraction among set-theoretical foundations, as opposed to a more fluctuating picture on the categorical side. This suggests that category theory has the resources to cover a range of abstraction levels; however, it is too early to tell whether this fundamentally characterizes the distinction between set-theoretical and categorical thinking. We additionally suggest that Landry's distinction between the constitutive and organizational role of foundations, instead of corresponding to the bottom-up/top-down distinction, should be associated with differences in levels of abstraction.

Relevance of this thesis. We maintain that elucidation of the motivations and justifications of foundations for mathematics is important for the interpretation of mathematical results from different fields. Being able to transfer a result between foundations helps shed light on the different perspectives towards it.

However, such a characterization of a result will only truly be understood if we can make the differences between the foundations themselves explicit. Our results, in particular, give reason to regard the differences in set-theoretical results and categorical results as obtained from a particular level of abstraction. Alternatively, it is relevant to be aware of the level of abstraction one is occupying while working towards a result, as this entails knowing the capacities and limits of the theory that is used. Of course, it remains to be seen if our interpretation is consistent with the successful and less successful applications of either theory.

Independently from the research question, this thesis has also contributed to the sharpening of terms with multiple interpretations towards definitions more usable in the philosophy of mathematics. We discovered that incautious use of terms such as 'bottom-up', 'top-down' and 'abstractness' can lead to overgeneralized and uninformative conclusions. Hence, we stress generally that concepts on the border of mathematics and philosophy benefit from being made explicit.

Open questions. Future research should continue the quest of distinguishing set-theoretical from categorical thinking. Different conceptions of ways of thinking can be analyzed and applied to various foundations — for example, the idea that categorical foundations represent a more structuralist way of thinking than set-theoretical foundations can be investigated. Furthermore, the idea from (Mathias, 2001) and (Ernst, 2017) of analyzing the strengths and weaknesses of set-theoretical and categorical foundations in various fields is a different approach to characterizing the distinction. Our suggestion that category-theoretic foundations embody a higher level of abstraction than set-theoretical foundations deserves additional support, which may be found here. Still, the possibility exists that the variation in abstraction indeed corresponds to the set theory/category theory contrast, but that it is not *characteristic* of it. That is, it could be more of a byproduct than an actual underlying cause of this distinction.

Furthermore, additional (set-theoretical and categorical, but also other foundational) theories should have a place in the hierarchy of levels of abstraction. Foundational systems based in homotopy type theory, for example, may fit into the picture, thereby further elucidating the way theories relate to each other via abstraction.

Last, we have seen that the method of abstraction can be improved in several ways. Future research should investigate whether there exists a method that reliably provides relations between mathematical levels of abstraction, and in what ways the current method can be developed further.

Finally. This thesis sets in motion the quest to understand what different foundations for mathematics can tell us about mathematical thought. We eagerly anticipate further developments and enlightenment in this intriguing area.

Bibliography

Awodey, Steve (2004). "An Answer to G. Hellman's Question" Does Category Theory Provide a Framework for Mathematical Structuralism?"" In:

— (2008). "A brief introduction to algebraic set theory". In: *Bulletin of Symbolic Logic* 14.3, pp. 281–298.

— (2010). *Category theory*. Oxford University Press.

— (2014). "Structuralism, invariance, and univalence". In: *Philosophia Mathematica* 22.1, pp. 1–11.

Benacerraf, Paul (1965). "What numbers could not be". In: *The Philosophical Review* 74.1, pp. 47–73.

Boolos, George (1971). "The iterative conception of set". In: *The Journal of philosophy* 68.8, pp. 215–231.

Enayat, Ali, Paul Gorbow, and Zachiri McKenzie (2017). "Feferman's Forays into the Foundations of Category Theory". In: *Feferman on Foundations*. Springer, pp. 315–346.

Ernst, Michael (2015). "The prospects of unlimited category theory: doing what remains to be done". In: *The Review of Symbolic Logic* 8.2, pp. 306–327.

— (2017). "Category theory and foundations". In: *Landry*, pp. 69–89.

Feferman, Solomon (2011). "Enriched stratified systems for the foundations of category theory". In: *Foundational Theories of Classical and Constructive Mathematics*. Springer, pp. 127–143.

— (2013). "Foundations of unlimited category theory: What remains to be done". In: *The Review of Symbolic Logic* 6.1, pp. 6–15.

Feferman, Solomon and G Kreisel (1969). "Set-theoretical foundations of category theory". In: *Reports of the midwest category seminar III*. Springer, pp. 201–247.

Floridi, Luciano (2013). *The philosophy of information*. OUP Oxford.

González, Julia Ramos (2018). "Grothendieck categories as a bilocalization of linear sites". In: *Applied Categorical Structures* 26.4, pp. 717–745.

Hellman, Geoffrey (2003). "Does category theory provide a framework for mathematical structuralism?" In: *Philosophia Mathematica* 11.2, pp. 129–157.

Incurvati, Luca (2014). "The graph conception of set". In: *Journal of Philosophical Logic* 43.1, pp. 181–208.

— (2017). "Maximality principles in set theory". In: *Philosophia Mathematica* 25.2, pp. 159–193.

Incurvati, Luca and Benedikt Löwe (2016). "Restrictiveness relative to notions of interpretation". In: *The Review of Symbolic Logic* 9.2, pp. 238–250.

Jech, Thomas (2013). *Set theory*. Springer Science & Business Media.

Johnstone, Peter T (2002). “Sketches of an elephant: A topos theory compendium-volume 1”. In: *Sketches of an Elephant: A Topos Theory Compendium-Volume 1, by Peter T Johnstone, pp. 562. Foreword by Peter T Johnstone. Oxford University Press, Nov 2002. ISBN-10: 0198534256. ISBN-13: 9780198534259*, p. 562.

Joyal, Andri, Izak Moerdijk, Ieke Moerdijk, et al. (1995). *Algebraic set theory*. Vol. 220. Cambridge University Press.

Kapulkin, Chris and Peter LeFanu Lumsdaine (2012). “The simplicial model of univalent foundations (after Voevodsky)”. In: *arXiv preprint arXiv:1211.2851*.

Kim, Jongkwang and Thomas Wilhelm (2008). “What is a complex graph?” In: *Physica A: Statistical Mechanics and its Applications* 387.11, pp. 2637–2652.

Landry, Elaine (2011). “How to be a structuralist all the way down”. In: *Synthese* 179.3, pp. 435–454.

— (2013). “The genetic versus the axiomatic method: Responding to feferman 1977”. In: *The Review of Symbolic Logic* 6.1, pp. 24–51.

Lawvere, F William (1966). “The category of categories as a foundation for mathematics”. In: *Proceedings of the conference on categorical algebra*. Springer, pp. 1–20.

— (2005). “An elementary theory of the category of sets (long version) with commentary”. In: *Reprints in Theory and Applications of Categories* 11, pp. 1–35.

Linnebo, Øystein (2013). “The potential hierarchy of sets”. In: *The Review of Symbolic Logic* 6.2, pp. 205–228.

Linnebo, Øystein and Richard Pettigrew (2011). “Category theory as an autonomous foundation”. In: *Philosophia Mathematica* 19.3, pp. 227–254.

Mac Lane, Saunders (1969). “One universe as a foundation for category theory”. In: *Reports of the Midwest Category Seminar III*. Springer, pp. 192–200.

Makkai, Michael (1998). “Towards a categorical foundation of mathematics”. In: *Logic Colloquium’95: Proceedings of the Annual European Summer Meeting of the Association of Symbolic Logic, held in Haifa, Israel, August 9-18, 1995*. Association for Symbolic Logic, pp. 153–190.

Mathias, Adrian RD (2001). “The strength of Mac Lane set theory”. In: *Annals of Pure and Applied Logic* 110.1-3, pp. 107–234.

McLarty, Colin (1991). “Axiomatizing a category of categories”. In: *The Journal of symbolic logic* 56.4, pp. 1243–1260.

— (1992). “Failure of Cartesian closedness in NF”. In: *The Journal of symbolic logic* 57.2, pp. 555–556.

— (1993). “Anti-foundation and self-reference”. In: *Journal of Philosophical Logic* 22.1, pp. 19–28.

— (2004). “Exploring categorical structuralism”. In: *Philosophia Mathematica* 12.1, pp. 37–53.

Moss, Lawrence S. (2018). “Non-wellfounded Set Theory”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2018. Metaphysics Research Lab, Stanford University.

Muller, Frederik A (2001). “Sets, classes, and categories”. In: *The British journal for the philosophy of science* 52.3, pp. 539–573.

Osius, Gerhard (1974). “Categorical set theory: a characterization of the category of sets”. In: *Journal of Pure and Applied Algebra* 4.1, pp. 79–119.

Palmgren, Erik (2012). “Constructivist and structuralist foundations: Bishop’s and Lawvere’s theories of sets”. In: *Annals of Pure and Applied Logic* 163.10, pp. 1384–1399.

Shapiro, Stewart (2004). “Foundations of mathematics: metaphysics, epistemology, structure”. In: *The Philosophical Quarterly* 54.214, pp. 16–37.

Shulman, Michael A (2008). “Set theory for category theory”. In: *arXiv preprint arXiv:0810.1279*.

— (2009). *Category of categories as a foundation of mathematics*. MathOverflow. (version: 2009-12-20). URL: https://mathoverflow.net/q/9412.

Simpson, Alex K (1999). “Elementary axioms for categories of classes”. In: *Proceedings. 14th Symposium on Logic in Computer Science (Cat. No. PR00158)*. IEEE, pp. 77–85.

Taylor, Paul (1999). *Practical foundations of mathematics*. Vol. 59. Cambridge University Press.

Thomas, Morgan (2018). “Approximating Cartesian Closed Categories in NF-Style Set Theories”. In: *Journal of Philosophical Logic* 47.1, pp. 143–160.

Tsementzis, Dimitris (2017). “Univalent foundations as structuralist foundations”. In: *Synthese* 194.9, pp. 3583–3617.

Van Leeuwen, Jan (2014). “On Floridi’s method of levels of abstraction”. In: *Minds and Machines* 24.1, pp. 5–17.

Weber, Mark (2007). “Yoneda structures from 2-toposes”. In: *Applied Categorical Structures* 15.3, pp. 259–323.