

What we talk about when we talk about numbers*

Richard Pettigrew
Richard.Pettigrew@bris.ac.uk

December 16, 2016

In this paper, I'd like to describe and motivate a new species of mathematical structuralism. In the philosophy of mathematics, structuralism is a genus of theses concerning the subject matter and ontology of mathematics, as well as the correct semantics for mathematical language. Each species that belongs to that genus is motivated by the observation that mathematicians are agnostic about the intrinsic or internal nature of the objects that they study. In this sense, structuralism is very much a philosophy of mathematics that is inspired by and guided by mathematical practice. Mathematicians are indifferent to the non-mathematical features of the objects they study. They care only about the so-called structural features of those objects. For instance, they care that 2 is less than 3 and that π is transcendental. They do not care whether 2, 3, or π is a set or a class of sets, a Dedekind cut in the rationals or an equivalence class of Cauchy sequences of rationals, a universal or a particular, an abstract object or a concrete one, a necessary existent or an entity that exists only contingently, and so on. But, while each species of structuralism agrees on this indifference, they differ significantly on the ontology and semantics of mathematics that best accommodates it.

Why offer a new species of structuralism when the genus is already so crowded? The subject matter of mathematics, together with the semantics of mathematical language, has an extensive job description. There are many boxes that any candidate ontology and semantics would ideally tick. As we will see, while each of the existing species of structuralism ticks many of these boxes, they all leave many untouched. I hope that my new version will tick all of the boxes.

My strategy is as follows: I will begin with a species of structuralism as

*I am very grateful to Bahram Assadian, Eddy Keming Chen, Trent Dougherty, Chris Menzel, Laurie Paul, Ian Phillips, Daniel Rubio, Ralph Wedgwood, Robbie Williams, and Jack Woods for their help with various parts of this paper. I dedicate the paper to the memory of my late PhD supervisor, John Penn Mayberry (1939-2016), a determined set-theoretic structuralist, a wonderful philosopher, and a great friend.

different from my final proposal as can be. Then I will raise an objection to that species that will lead us to formulate a new species that avoids the objection. But now I will note a objection to this new species, and I will formulate a further new species that avoids both objections. And so on. I will repeat this process until we arrive at our new version of structuralism, which avoids all objections.

Before we start, a disclaimer: this paper covers quite a lot of ground. Each new tickbox in the job description for the subject matter of mathematics deserves, and has received, much more detailed discussion than I am able to give it here. But this paper is programmatic — my purpose is to motivate moving to a new version of structuralism. So I hope readers will forgive me if I reject their favoured version of structuralism without the full discussion they would wish.

1 Structuralism and the axiomatic method

As John P. Burgess (2015) argues in detail, structuralism in the philosophy of mathematics is the inevitable response to the introduction of the axiomatic method as the fundamental methodology of mathematics towards the end of the nineteenth century. And the axiomatic method was, in turn, the inevitable conclusion of the quest for greater rigour in mathematics and the attempt to expel geometric, spatial, and other forms of intuition from mathematical proofs and definitions. According to the axiomatic method, each area of mathematics — real or complex analysis, probability or measure theory, group theory, number theory, graph theory, linear algebra, topology, and so on — is characterized by a set of axioms. These pick out the items of interest in that area — the real numbers, the complex field, the probability spaces, the groups, the natural numbers, the graphs, the vector spaces, the topological spaces, and so on. They do this by spelling out the properties shared by all of the items of interest.

If we take a set-theoretic approach, the items of interest are *systems*. In this context, a system consists of an underlying set or a family of underlying sets, perhaps equipped with some distinguished elements of those sets, distinguished functions involving those sets, and relations amongst the members of those sets. Thus, for instance, Cayley's group axioms characterise the subject matter of group theory (Cayley, 1854). They apply to systems $(G, e, *)$, where e is a distinguished element of the underlying set G , and $*$ is a binary function on G . Similarly, Dedekind's axioms for a complete ordered field characterise the subject matter of real analysis (Dedekind, 1872). They apply to systems $(R, 0, 1, +, \times, <)$, where $0, 1$ are distinguished elements of R , $+$ and \times are binary functions on R , and $<$ is a binary relation on R . And Peano's axioms for a vector space characterise the subject matter of linear algebra (Peano, 1888). They apply to systems

$(V, K, \mathbf{0}, +_V, 0, 1, +_K, \times_K, \cdot)$, where $\mathbf{0}$ is a distinguished element of V , $+_V$ is a binary function on V , $0, 1$ are distinguished elements of K , $+_K, \times_K$ are binary functions on V , and \cdot is a function defined on $K \times V$. And so on.

On the other hand, if we take a category-theoretic approach, the items of interest are *categories*, or *objects* in *categories*. There are (at least) four ways in which we might formulate the axiomatic method on this approach. On the first, which we might call the *category-based approach*, the axioms characterise a certain sort of category, saying that the items of interest in the area of mathematics in question are all and only the categories of this sort. This is close to the set-theoretic approach, except that the axioms are stated in terms of the behaviour of objects and morphisms in the categories in question, not in terms of underlying sets, distinguished elements, and so on. Thus, for instance, we can give axioms that say when a category can be viewed as a group.¹ On the second, which we might call the *object-based approach*, the axioms characterize the category whose objects are all and only the items of interest in the area of mathematics in question. Thus, the axioms specify how the morphisms between the items of interest behave. As Steve Awodey puts it:

“From Dedekind, through Noether, and to the work of Eilenberg and Mac Lane, the fact has clearly emerged that mathematical structure is determined by a system of objects *and their mappings*, rather than by any specific features of mathematical objects viewed in isolation.” (Awodey, 1996, 209)

For instance, we can give axioms that characterize the category Grp , whose objects are all the groups and whose morphisms are the group homomorphisms; and those axioms pay attention only to the behaviour of morphisms between groups — they say nothing of the internal nature of the groups themselves. On the third way, which we might call the *functor-based approach*, the axioms characterise a generic category in such a way that we can identify the items of interest in a particular area of mathematics with the functors from that generic category into the category of sets, Set . This is the strategy that Lawvere presented in his dissertation; the generic categories that he describes are known as Lawvere theories (Lawvere, 1963). On the final way, which we might call the *particular-object-based approach*, the axioms characterize a particular sort of object in a particular sort of category, saying that the items of interest are all and only those objects in those categories. For instance, we might give axioms that pick out the natural numbers object in a given topos (Lawvere, 1963).

¹Namely, when it contains one object, and all of the morphisms from that object to itself are left- and right-invertible. The elements of the group are then the morphisms, the group multiplication operation is the composition operation on morphisms, and the group identity element is the identity morphism on the single object in the category.

Note: on both the set-theoretic and the category-theoretic approaches, each area of interest comes equipped with a notion of *structure-preserving map*, or *homomorphism*. In the set-theoretic case, this is usually determined by the nature of the systems to which the axioms apply. Since a group is a system $(G, *_G, e_G)$, a group homomorphism is a function $\varphi : G \rightarrow H$ such that (i) $\varphi(e_G) = \varphi(e_H)$ and (ii) $\varphi(g *_G g') = \varphi(g) *_H \varphi(g')$ — that is, φ preserves the distinguished element e_G of G and the binary operation $*_G$ on G . However, there are cases in which the homomorphisms cannot be read off the nature of the systems. For instance, a topological space (X, T) is a set T of subsets of an underlying set X satisfying particular closure conditions — the subsets are called the open sets of the topological space. The structure-preserving mappings from one underlying space X to another X' are the continuous mappings, to wit, those for which the inverse image of an open set in X' is an open set in X . In the category-theoretic case, in contrast, the structure-preserving mappings are given for the category-based approach by the functors between the categories, for the object-based and particular-object-based approach by the morphisms in the category of all items of interest, and for the functor-based approach by the natural transformations between the functors. Thus, on the object-based approach, a group homomorphism is a morphism in the category Grp; a homomorphism between topological spaces is an morphism in the category Top. An invertible homomorphism is called an isomorphism. This notion will become important when we come to distinguish different species of structuralism.

This, then, is the axiomatic method, the methodology of mathematics that structuralism seeks to accommodate. But there are different ways in which we might accommodate this methodology, and these give rise to a number of different species of structuralism.

2 Extreme Structure Realism

Here is an extreme version of structuralism: Number theory is concerned with a single simply infinite system $(\mathbb{N}, \mathbf{0}, \mathfrak{s})$, which we call the natural number system. Real analysis is concerned with a single complete ordered field $(\mathbb{R}, \mathbf{0}, \mathbf{1}, +, \times, <)$, which we call the real number structure. Group theory is not concerned with just one system; it is concerned with many. But there is just one group for each isomorphism class: thus, there is just one Klein Vierergruppe V_4 , just one cyclic group Z_n of order n (for given n), just one permutation group S_n of n symbols (for given n). And so on for other areas of mathematics. In general, for every isomorphism class, there is just one privileged representative that belongs to the subject matter of mathematics. As Steve Awodey (2014, 1-2) puts it: “The following statement may be called the *Principle of Structuralism*: Isomorphic objects are identical”. This is not to deny that other systems exist. For instance, one might still

hold that the system of finite Zermelo ordinals and the distinct system of von Neumann ordinals are both simply infinite systems; or that both the system of Dedekind cuts in the rationals and the distinct system of equivalence classes of Cauchy sequences of rationals are complete ordered fields; or that the set of symmetries on a rectangle that isn't a square constitutes a Klein Vierergruppe, as does the distinct group comprising the elements $(\pm 1, \pm 1)$ under coordinate-wise multiplication. Extreme Structure Realism just says that those other systems do not belong to the subject matter of number theory, real analysis, or group theory, respectively.

On this view, what accounts for the indifference that mathematicians show towards the internal nature of their objects of study is that the objects posited here — *the* natural number structure and *the* natural numbers themselves, *the* real numbers and *the* real numbers themselves, *the* elements of the Klein Vierergruppe, and so on — have no internal nature. That is what distinguishes this view from old-fashioned platonism. As structuralists of this stripe sometimes put it, the objects of mathematical study are *incomplete* — they have only structural properties; they have no non-structural properties. Thus, 2 is less than 3 and π is transcendental, but it is neither true nor false that 2 is a set-theoretic member of 3, neither true nor false that 2 is a mereological part of 3, neither true nor false that π is a Roman emperor, and so on. As Dedekind puts it, when speaking of the elements of the unique complete ordered field that is the subject matter of real analysis: they have none of the properties that “one would surely attach only very unwillingly to the numbers” (Dedekind, 1888a). Or, as Stewart Shapiro puts it, when speaking of the elements of the simply infinite system that is studied by number theory:

“[O]ne can look into the identity between numbers denoted by different descriptions *in the language of arithmetic* [...] But it makes no sense to pursue the identity between a place in the natural-number structure [i.e. an element of the simply infinite system] and some other object, expecting there to be a fact of the matter” (Shapiro, 1997, 79).

3 Burgess' Incompleteness Objection

The first objection to Extreme Structure Realism that we will consider was raised by John P. Burgess (1999, 286) in his review of Shapiro's book-length exposition of the position. As we will see, this objection can be answered, but its answer will be illuminating for what follows, so I include it here.

Extreme Structure Realism takes the subject matter of number theory to be a privileged simply infinite system $(\mathbb{N}, \mathbf{0}, \mathfrak{s})$ whose elements are incomplete in some sense. There are various ways one might spell out this claim of incompleteness: the elements of \mathbb{N} have only their number-theoretic

properties; they have only their structural properties; they have only those properties ascribed to them by mathematicians; they have only the properties that they share with all corresponding elements of other simply infinite systems. But, as Burgess points out, however we put it, this incompleteness claim can't be true.

For instance, take the natural number 0 , which is the zero element of $(\mathbb{N}, 0, s)$. And suppose we make the incompleteness claim precise by saying that 0 has only those properties that it shares with all zero elements of simply infinite systems. That is, 0 has property Φ iff, for all simply infinite systems $(N, 0, s)$, 0 has property Φ . But now consider that very property, namely, the property of having only those properties that you share with all zero elements of simply infinite systems. That is, consider the property Ψ such that x has Ψ iff, for any property Φ , x has Φ iff, for all simply infinite systems $(N, 0, s)$, 0 has Φ . 0 has that property, but there are zero elements of other simply infinite systems that do not — in some simply infinite systems, for instance, you are the zero element, but you do not have property Ψ . Thus, if 0 is incomplete in this sense, then it is not incomplete in this sense. Made precise in this way, the incompleteness claim of Extreme Structural Realism leads to a contradiction. And similar objections can be raised for the other ways of making the incompleteness claim precise.

There are various ways we might try to weaken the incompleteness claim so that it does not lead to a contradiction (Shapiro, 2006; Linnebo & Pettigrew, 2014). We might, for instance, restrict attention to the *first-order* properties of systems and say that the first-order properties of a number are precisely the first-order properties it shares with the corresponding elements in all simply infinite systems. But 0 has the first-order property of being the number of unicorns in Bristol in 2016, while that first-order property is not shared by the zero elements of all simply infinite systems — again, you are the zero element in some simply infinite systems, but you are not the number of unicorns in Bristol in 2016. Or we might restrict our claim to the *intrinsic* properties of a system and say that the intrinsic properties of a number are precisely the intrinsic properties it shares with the corresponding elements in all simply infinite systems. But 0 has the intrinsic property of being abstract and the intrinsic property of not being a set, while there are zero elements that do not have those intrinsic properties — again, you are a zero element in some systems, but you are not abstract; and the set containing only you and the Eiffel Tower is the zero element in some systems, but that does not have the property of not being a set.

In both cases, what goes wrong is the same, and it is what Burgess originally identified: by trying to ensure that our numbers have the properties required of them by their role as the objects of number theory, and by trying to ensure that they have no more than these properties, we endow them with the properties — being the number of unicorns in Bristol in 2016, or the property of being abstract or not being a set — that accrue to them be-

cause they are the objects of number theory; and these properties are not shared by many of the corresponding elements in other simply infinite systems. This leads to the following suggestion (Linnebo & Pettigrew, 2014, 271-2). Some properties of objects are more fundamental than others. The property of being Scottish and the property of being a philosopher are more fundamental properties of me than the property of being a Scottish philosopher; the property of being tall is more fundamental than the property of being tall or extrovert. Similarly, the property of being the additive identity element of the natural numbers is a more fundamental property of $\mathbf{0}$ than being abstract or not being a set. So, we might say: a natural number has a property Φ fundamentally iff the corresponding position in every simply infinite system also has property Φ (whether fundamentally or not). This captures what the structuralist wants, but it also allows that a natural number can have properties — such as being abstract, or being the number of unicorns in Bristol in 2016 — that the corresponding element in some other simply infinite systems lack. It simply demands that the natural number does not have those other properties fundamentally. Rather, they are derived properties; properties that the natural number has in virtue of having the fundamental properties it does. This, I think, answers Burgess' objection.

4 Hellman's Permutation Objection

Next, we turn to Hellman's permutation objection against Extreme Structure Realism (Hellman, 2006, 545). According to Extreme Structure Realism, $(\mathbb{N}, \mathbf{0}, \mathfrak{s})$ is the unique privileged simply infinite system that provides the subject matter of number theory. When I quantify over all natural numbers, I quantify over the elements of \mathbb{N} ; when I talk of particular numbers, I talk of particular elements of \mathbb{N} , such as $\mathbf{0}$, $\mathfrak{s}(\mathbf{0})$, $\mathfrak{s}(\mathfrak{s}(\mathbf{0}))$, and so on. Now, suppose that $\pi : \mathbb{N} \rightarrow \mathbb{N}$ is a permutation of the elements of \mathbb{N} . And now consider the system $(\mathbb{N}, \mathbf{0}^\pi, \mathfrak{s}^\pi)$, where $\mathbf{0}^\pi = \pi(\mathbf{0})$ and $\mathfrak{s}^\pi(n) = \pi(\mathfrak{s}(\pi^{-1}(n)))$, for any n in \mathbb{N} . Then, since $(\mathbb{N}, \mathbf{0}, \mathfrak{s})$ is a simply infinite system, so is $(\mathbb{N}, \mathbf{0}^\pi, \mathfrak{s}^\pi)$. What's more, since the elements of the latter system are precisely the elements of the former system, the elements of the latter are incomplete in the way required by structuralism just in case the elements of the former are. Thus, there is nothing to tell between $(\mathbb{N}, \mathbf{0}, \mathfrak{s})$ and $(\mathbb{N}, \mathbf{0}^\pi, \mathfrak{s}^\pi)$ as the subject matter of number theory. Nothing in the practice of number theory could possibly tell between them. Thus, there are as many legitimate candidates for the subject matter of number theory as there are permutations of the natural numbers, that is, continuum-many.

It seems to me that this is devastating for the version of Extreme Structure Realism that I find in Dedekind's writings, and which I will call *Dedekind*

Structure Realism or *DSR*.² For Dedekind, every isomorphism class contains only one privileged member that belongs to the subject matter of mathematics. But he says no more about this privileged member than that it belongs to the isomorphism class and that it is incomplete, perhaps in the sense spelled out above in our response to Burgess' incompleteness objection. He says so little about this privileged member because he thinks there is no more to say. But precisely because he says so little about it, Hellman's objection shows that, if there is one simply infinite system that answers to what he does say, then there are continuum-many. And similarly for other isomorphism classes. Thus, there can be no non-arbitrary way to select the unique simply infinite system that is the subject matter of number theory.

A natural response to this is to concede that there are many incomplete simply infinite systems that have equal claim to be the subject matter of number theory, and to say that, while it might seem that number theory is talking only of one simply infinite system, it is in fact quantifying over all and only the incomplete simply infinite systems. Thus, take the following sentence of number theory, which states the Green-Tao theorem (Green & Tao, 2008):

For any n in \mathbb{N} , there are a, b in \mathbb{N} such that, for all k in \mathbb{N} , if $k \leq n$, then $ak + b$ is prime.

Then, according to this amended version of Dedekind's structuralism, it says:

For any incomplete simply infinite system $(N, 0, s)$, and for any n in N , there are a, b in N such that, for all k in N , if $k \leq_{(N,0,s)} n$, then $((a \times_{(N,0,s)} k) +_{(N,0,s)} b)$ is $\text{prime}_{(N,0,s)}$.

Call this species of structuralism *Amended Dedekind Structure Realism* or *DSR⁺*. This, I think, is the version of structuralism that Lawvere (1994) extracts from the use of category-theoretic language and techniques in mathematics. As Lawvere articulates this view, a set, which for the category theorist is simply a minimally structured system, is "a bag of dots which are devoid of properties apart from mutual distinctness" (Lawvere, 1994, 6). That is, the elements of a set have no internal nature — the only fundamental property of the set itself is its cardinality. What's more, Lawvere takes the same to be true of all mathematical systems, not just the sets — their elements have no internal nature.

"In the mathematical development of recent decades one sees clearly the rise of the conviction that the relevant properties of mathematical objects are those which can be stated in terms of their abstract structure rather than in terms of the elements

²For a different interpretation of Dedekind, see (Reck, 2003) and (Yap, 2009a).

which the objects were thought to be made of.” (Lawvere, 1966, 1)

On this, Lawvere and Dedekind agree. But Lawvere, unlike Dedekind, does not require there to be just one such system for each isomorphism class — he allows that there might be many, just as DSR^+ does. After all, category theory has a particular name for categories that contain just one representative of each isomorphism class — they are called *skeleton categories*. And Lawvere does not claim that all of mathematics takes place inside skeleton categories.

While Dedekind’s version of Extreme Structure Realism must be amended to accommodate Hellman’s objection, Stewart Shapiro’s version, which I call *Shapiro Structure Realism* or *SSR*, need not (Shapiro, 1997). It is natural to say that two isomorphic systems share something in common, namely, their structure. Shapiro reifies this shared structure by positing, for each isomorphism class, a universal that he takes to be what is shared by systems in that isomorphism class. That is, he says that, for each isomorphism class, there is a universal, and the particulars that participate in this universal are all and only the systems in that isomorphism class. Thus, there is a universal corresponding to the class of all complete ordered fields; it is this universal that they all share in common. Similarly, there is a universal corresponding to the class of Klein Vierergruppen, one corresponding to the class of all simply infinite systems, and so on (Shapiro, 1997, 74).

Now, for Shapiro, this universal is not itself just a system that belongs to the isomorphism class of its participants — just as the universal echidna is not itself an echidna. However, Shapiro does claim that there is a perspective from which we might view the universal — the perspective that he calls the *places-as-objects* perspective — such that, from that perspective, it is a system from the isomorphism class to which it corresponds. In this respect, it is not like the universal echidna, which is not an echidna from any perspective (Shapiro, 1997, 83-4).

For Shapiro, then, this system — namely, the one extracted from the universal corresponding to a given isomorphism class by considering it from the places-as-objects perspective — is the unique system belonging to that isomorphism class that belongs to the subject matter of mathematics. While any other system obtained from it using Hellman’s permutation trick will also belong to the isomorphism class, it will not belong to the subject matter of mathematics. And what’s more, this exclusion is not arbitrary in the way it would have to be if we were to make DSR true. By saying more than Dedekind about what the privileged member of an isomorphism class is — it corresponds to the universal that covers that class — and how it is obtained — it is extracted from that universal by considering it from a particular perspective — Shapiro avoids Hellman’s complaint.

However, in the endeavour of formulating mathematical structuralism,

saying more about the mathematical objects you posit comes with risks. As I noted at the beginning, the central motivation for mathematical structuralism is to describe a subject matter for mathematics and a semantics for mathematical language that accommodates and respects the indifference that mathematicians exhibit towards the internal nature of the entities they study. Extreme Structure Realists respond to this by trying to posit entities with no internal structure. They try to posit entities that have none of the properties towards which mathematicians are indifferent; that is, incomplete entities. As we saw in our response to Burgess' incompleteness objection above, they cannot do this on pain of contradiction. They can say that the objects of mathematics have none of the properties to which mathematicians are indifferent *fundamentally*; but they must accept that, by having only those properties *fundamentally*, and by being the objects of mathematics, those objects immediately accrue further properties *derivatively* — for instance, the property of being the number of unicorns in Bristol in 2016, or having only those properties *fundamentally* about which mathematicians are not indifferent. These are the properties that are forced on these objects by having such a sparse set of fundamental properties. However, Shapiro goes much further. He says that the objects of mathematics are a certain new sort of universal — a structural universal — that can be considered from different perspectives, sometimes as a system, sometimes not. In order to answer Burgess' incompleteness objection, Shapiro must say that these properties are not fundamental to these objects, but derived. But they go well beyond what is demanded of these objects by simply having the sparse set of fundamental properties that they have. These extra properties — the properties that come with being structural universals — are not forced on these objects by only having as their fundamental properties those towards which mathematicians are not indifferent. Dedekind's incomplete systems are the minimal objects that have only those fundamental properties.

Thus, Shapiro's proposal seems at odds with the Extreme Structural Realist's quest for ontology for mathematics composed of minimal objects. However, he might respond that, while his structural universals are not demanded by incompleteness considerations, they are demanded by uniqueness considerations, as Hellman's objection shows. If there is to be just one privileged member of each isomorphism class that will belong to the subject matter of mathematics, then it must have something that picks it out from all the isomorphic copies obtained by Hellman's permutation trick. And that will require it to have more than strictly follows from having the sparse set of fundamental properties about which mathematicians are not indifferent.

Thus, we have two options: SSR and DSR⁺. Both avoid Hellman's permutation objection. Shapiro retains a single privileged member of each isomorphism class, but does so at the expense of going well beyond what

mathematicians would endorse. The amended version of Dedekind's species of structuralism countenances many different members of an isomorphism class amongst the subject matter of mathematics, but none of them go beyond what mathematicians would endorse. How should we choose? As we will see, neither is ideal.

5 Weaver's representation theorem objection

Let us move now to the third objection to Extreme Structure Realism. This is due originally to George Weaver (1998), though Jessica Carter (2005, 2008) also presses a related point using more sophisticated examples from mathematical practice. The point is straightforward: contrary to the motivating claim of structuralism, mathematicians are sometimes not indifferent to the internal nature of the objects that they study. Weaver notes that this is particularly true in the case of representations theorem. The purpose of such theorems is to show that each mathematical object of one sort is isomorphic to some mathematical object of another sort. If there were but one representative of each isomorphism class within the subject matter of mathematics, these theorems would be trivially true. And, even if there were many such representatives, these theorems are only interesting if some of those representatives have an internal nature — the interest of a representation theorem is that one mathematical object with a certain sort of internal nature is isomorphic to a mathematical object with a different sort of internal nature. Thus, amongst the subject matter of mathematics, we must find systems whose elements have an internal nature.

The most striking example of a representation theorem is Cayley's Theorem in group theory. This says that, for every group G , there is a subgroup of the symmetric group on the underlying set of G , $\text{Sym}(G)$, that is isomorphic to G . The symmetric group on G , $\text{Sym}(G)$, is the group whose elements are the permutations of the elements of the underlying set of G , whose identity element is the identity mapping on the underlying set of G , and whose group multiplication operation is the composition operation on mappings. Thus, the elements of $\text{Sym}(G)$ have an internal nature. They are one-one mappings from a given set to itself. And this nature is crucial for the theorem. The theorem says that any group can be represented by a group whose elements have this specific sort of internal nature. But the feature that makes that particular sort of group distinctive is not something that can be stated in the language of group theory. It is not some fact about how the group multiplication operation behaves. Rather, it is a feature of the internal nature of the objects of the group.

There are two ways in which a structure realist — a proponent of SSR or DSR^+ — might respond to this objection. They might argue that we can interpret Cayley's Theorem without positing systems whose elements have

an internal nature; or they might accept that such systems are required to interpret Cayley's Theorem, and introduce them into the subject matter of mathematics along with the incomplete system(s) already included by SSR and DSR⁺.

Consider the first option. This is the idea behind a number of definitions in category theory. It is endorsed explicitly by certain category theorists, such as F. William Lawvere (1994) and Steve Awodey (1996, 2014), though it is worth noting that there is nothing in the mathematics of category theory that demands it. The point is that category theory is able to take mathematical statements that are ostensibly about entities with an internal nature and interpret them instead as statements that concern only the morphisms between different objects in a category, or the functors between different categories. The famous definition of a product of two objects in a category illustrates how this might be done in principle (Eilenberg & Mac Lane, 1945). Consider the claim, from group theory, that to every pair of groups, G and H , there corresponds a direct product $G \times H$, which is also a group. This is standardly taken to mean that there is a group $G \times H$ whose elements have a particular internal nature, and whose identity element and group multiplication operation are specified in terms of that internal nature: each element of $G \times H$ is an ordered pair (g, h) , where g is in G and h is in H ; the identity element is $e_{G \times H} = (e_G, e_H)$, and $(g, h) *_{G \times H} (g', h') = (g *_G g', h *_H h')$. In category theory, however, the claim is interpreted as concerning only the objects and the morphisms of the category, Grp , namely, the groups and the group homomorphisms. It says nothing about what the elements of those groups are. In category theory, a direct product of two groups G and H is an object $G \times H$ equipped with two morphisms, $\varphi_G : G \times H \rightarrow G$ and $\varphi_H : G \times H \rightarrow H$, known as projection morphisms, such that the following holds: if there is an object C and morphism $\psi : C \rightarrow G \times H$, then there are morphisms $\theta_G : C \rightarrow G$ and $\theta_H : C \rightarrow H$ such that (i) $\varphi_G \circ \psi = \theta_G$ and (ii) $\varphi_H \circ \psi = \theta_H$. By this sort of construction, category theory is able to capture the mathematically relevant content of the notion of a direct product group — which we would usually express set-theoretically in a way that makes explicit reference to the internal nature of the elements of a direct product group — without making reference to the internal nature of any of the objects in question. Another example: in Lawvere's axiomatization of the category of sets, he is able to state the power set axiom — which tells us that for every set there is another whose elements are precisely the subsets of the first — without saying anything about the internal nature of the elements of any set (Lawvere, 1965). And so on for other constructions that initially seem to refer to the internal nature of the elements of systems. In this way, a category-theoretic reading of structure realism avoids Weaver's objection. Let's call this position DSR*. Its ontology is that of DSR⁺, but it uses category-theoretic techniques, such as those just outlined, in order to interpret all mathemat-

ical statements as concerning just incomplete systems and the morphisms between them.

Now consider the second option. On this option, SSR and DSR⁺ must include in the subject matter of mathematics not only the privileged incomplete representative(s) of each isomorphism class, but also certain representatives whose elements have an internal nature — thus, not just the incomplete representative(s) of the isomorphism class of symmetric groups of order n (for any n), but also, for each set S of cardinality n (for any n), the group of permutations on S , for example — and the elements of these system have an internal nature. As Carter (2008, 199) puts it: “mathematics certainly deals with structures, but [...] structures may not be all there is to mathematics”.

Which other representatives must SSR and DSR⁺ include in the subject matter of mathematics? To find out, we can look to mathematical practice, for they must include enough other representatives to underwrite that practice — recall: structuralism is a philosophy inspired and motivated by mathematical practice. When do mathematicians feel able to construct a new sort of group or ring or probability space? It is exactly this question that Zermelo set out to answer when he laid down his original axioms for set theory (Zermelo, 1908). Nowadays, too often, we think of the iterative conception of sets as the fundamental conception of the set-theoretic universe — first, you have the null set; then you have its power set; then you have the power set of that; and so on. But this picture originates in Zermelo’s 1930 paper on the subject (Zermelo, 1930). In his 1908 paper, he sought to bring together the principles that had been laid down by Cantor (1883) and Dedekind (1888b). While Cantor was particularly interested in the theory of pure cardinality, Dedekind was interested in systematically enumerating the operations that mathematicians routinely use when they construct the systems that form the subject matter of their study. In §8 of *Was Sind Und Was Sollen Die Zahlen?*, for instance, Dedekind gives the definition of what we would now call the union of a family of sets and says that the union exists whenever the sets in the family exist. In §17, he introduces the notion of the intersection of a family of sets and says, again, that the intersection exists whenever the sets exist. In §25, he asserts that the image of a set under a mapping that is defined on each element in it is itself a set, an assumption that became Fraenkel’s and Skolem’s axiom of replacement (Fraenkel, 1922; Skolem, 1922). He does not mention a power set axiom nor an axiom of separation, but Zermelo makes good on these omissions as well as cleaning up Dedekind’s original formulations. Astonishingly, in this first attempt at such an enumeration, Zermelo (later supplemented by Skolem and Fraenkel) comprehensively listed the operations that underpin the constructions that are carried out within standard mathematics. Thus, SSR and DSR⁺ must at least add to the subject matter of mathematics those systems whose underlying sets can be constructed using the operations of

Zermelo-Fraenkel set theory — we will call these systems the *set-theoretic systems*.

Supplemented in this way, SSR becomes SSR^+ and DSR^+ becomes DSR^{++} . However, if we avoid Weaver's and Carter's objection by expanding the subject matter of mathematics in this way, SSR^+ and DSR^{++} face a new objection. They posit all the set-theoretic systems in a given isomorphism class as well as the incomplete systems that they posited originally. However, there is a version of structuralism that posits only the set-theoretic systems — this is *Set-Theoretic Structuralism* or *STS* (Bourbaki, 1970; Mayberry, 2000). According to this, the subject matter of mathematics consists of only the set-theoretic systems — the simply infinite system of von Neumann ordinals; the simply infinite system of the Zermelo ordinals; for any set-theoretic system that represents the rationals, the complete ordered field of the Dedekind cuts on that system and the complete ordered field of equivalence classes of Cauchy sequences on that system; and so on. Now, take a sentence that SSR would interpret as concerning the unique privileged incomplete simply infinite system $(\mathbb{N}, 0, s)$. STS, in contrast, would interpret that sentence as quantifying over all set-theoretic simply infinite systems — the von Neumann ordinals, the Zermelo ordinals, Hellman-style permutations of them, and countless others. Similarly, where SSR would interpret the claims of real analysis as concerning the unique, privileged complete ordered field whose elements have no internal nature, STS interprets them as quantifying over all set-theoretic systems that are complete ordered fields. And so on. In realist versions of structuralism, such as SSR or DSR^* , we account for any indifference that mathematicians exhibit towards the internal nature of their objects by saying that the objects they study have no internal nature. In STS, in contrast, we account for any such indifference by saying that, while the objects that those mathematicians study have internal natures, mathematicians often ignore that internal nature and talk only of the properties that any system will share with another that is isomorphic to it and whose only difference is in the internal nature of its elements.

6 Parsimony vs faithfulness

Thus, we have STS, which posits all the set-theoretic systems, *and nothing more*. And we have SSR^+ and DSR^{++} , which posit all the set-theoretic systems, *and also posit a raft of incomplete systems as well*. Considerations of parsimony would therefore seem to tell against SSR^+ and DSR^{++} , and in favour of STS. Unless there is something in favour of the former and against the latter that outweighs this consideration, it seems that we should reject SSR^+ and DSR^{++} . What might that be? Shapiro (1997, 2006) holds that one of the advantages of SSR/ SSR^+ over other species of structuralism (including, I presume, DSR^+ / DSR^*) is that it provides a subject matter for mathe-

matics that allows us to give a semantics for mathematical language that is faithful to its surface grammar. According to Shapiro, mathematicians use expressions such as ‘ \mathbb{N} ’, ‘ \mathbb{R} ’, ‘ V_4 ’, ‘the natural numbers’, ‘the real numbers’, ‘the Klein Vierergruppe’, ‘ π ’, ‘ e ’, ‘0’, ‘1’, etc. as singular terms. Thus, any faithful semantics for their language should identify, for each of these expressions, a single entity to which that expression refers. SSR/SSR⁺ does this: ‘ \mathbb{N} ’ refers to the unique, privileged incomplete simply infinite system, as does ‘the natural numbers’; ‘0’ refers to the zero element of that system (at least when that expression occurs in a paper on number theory); and so on. According to STS, in contrast, ‘ \mathbb{N} ’ and ‘the natural numbers’ are not singular terms, but free variables ranging over all simply infinite systems; and ‘0’ is a dependent variable that names the zero element of the simply infinite system in question. Thus, according to Shapiro, SSR/SSR⁺ provides a semantics that is faithful to the syntax of mathematical language; STS does not. And this, you might think, warrants the larger ontology posited by SSR⁺.

I’m not sure how we should weigh the faithfulness of a semantics against the parsimony of an ontology when we are adjudicating between rival philosophies of mathematics, but fortunately I don’t think this is necessary. As Pettigrew (2008) argues, the linguistic evidence does not favour treating ‘ \mathbb{N} ’, ‘the natural numbers’, etc. as singular terms rather than as a particular sort of free variable. Of course, Shapiro is right in thinking that an expression like ‘ \mathbb{N} ’ looks initially like a proper name, which is a singular term; and the definite article in ‘the natural numbers’ again suggests that the expression refers to a unique object, which would make it a singular term as well. But consider the following sentence from a chemistry textbook: ‘¹H is stable’. Or consider this sentence from Quine’s ‘On What There Is’: “McX never confuses the Parthenon with the Parthenon-idea” (Quine, 1980). Or this sentence from earlier in this very paper you are currently reading: “There are two ways in which the structure realist might respond to this objection”. From these three sentences consider the expressions ‘¹H’, ‘McX’, and ‘the structure realist’. Are they singular terms? As with ‘ \mathbb{N} ’ and ‘the natural numbers’, they initially seem to be. But a little thought makes clear that this is not how they function.³ There is no unique entity to which ‘¹H’ refers; Quine did not use ‘McX’ to pick out any particular person; and nor did I use ‘the structure realist’ to do so. Rather, in each case, they act as free variables that range over a class of entities — ‘¹H’ ranges over all atoms of the isotope of hydrogen with one proton and no neutrons; ‘McX’ ranges over all philosophers who hold the views that Quine ascribes to McX in his paper; ‘the structure realist’ ranges over all

³This is not to deny that one might give a semantics on which they are singular terms, perhaps using the ontology of Kit Fine’s theory of arbitrary objects (Fine, 1985). The claim here is only that this is not the natural reading.

structure realists. Thus, these expressions are what Pettigrew (2008) calls *dedicated free variables*. We use free variables often in mathematics: ‘Let r be a real number in the closed unit interval’; ‘Suppose m, n are natural numbers such that $\frac{m}{n} = \sqrt{2}$ ’; and so on. In each case, we introduce the free variable — ‘ r ’, ‘ m ’, ‘ n ’ — by stipulating the properties we are going to assume of it — being a real number in the closed unit interval; being natural numbers whose ratio is the square root of 2; and so on. Then we reason using this free variable, never assuming that what it ranges over has any properties other than those given by the stipulation. Finally, we reach some conclusion involving the free variable. And we can infer that this conclusion holds of any entity that satisfies the stipulation. In the case of ordinary free variables, such as ‘ r ’, ‘ m ’, ‘ n ’, the stipulation is required because each might be used with different stipulations in different mathematical contexts — I might use ‘ r ’ to range over all real numbers in the closed unit interval, but I might also use it to range over all transcendental real numbers or just over all real numbers or all non-negative real numbers, and so on. However, there are some free variables that are dedicated to being introduced by a particular stipulation. In these cases, it is not necessary to make the stipulation every time they are used. ‘ ${}^1\text{H}$ ’ is an example. It is always used to range over all atoms of the isotope of hydrogen with one proton and no neutrons. ‘ McX ’ is (now) another. Quine fixed its stipulation forever, at least amongst philosophers. ‘The structure realist’ is another. And, we might think, ‘ \mathbb{N} ’, ‘the natural numbers’, ‘ \mathbb{R} ’, ‘ 0 ’, ‘ π ’, etc. are others as well. Pettigrew does not argue that they are. Rather, he argues that there is no way to tell simply by looking at the grammar of the sentences that contain them, and the inferences between such sentences that mathematicians endorse, whether a given expression is a singular term or a dedicated free variable. Thus, it does not count in favour of SSR that it can provide a semantics for mathematical language on which these expressions are singular terms. STS provides a semantics for them on which they are dedicated free variables. And, as we have seen, no evidence from mathematical practice can tell between treating them one way and treating them another.

We began, in section 2 with an extreme version of structure realism, which then divided, in section 4, into SSR and DSR. In the light of Hellman’s objection, we replaced DSR with DSR^+ . In the light of Weaver’s objection, we replaced SSR with SSR^+ and we moved either to DSR^* — a categorial version of DSR^+ — or to DSR^{++} . In the light of the parsimony objection and the failure to rebut it by appealing to the faithfulness of the semantics provided, we now drop SSR^+ and DSR^{++} in favour of STS. We are thus left with DSR^* and STS.

7 The access problem

Which of these two positions should we favour, DSR* or STS? In their current guises, neither. However, as we will see, STS can be amended so that it provides a satisfactory ontology for mathematics and a satisfactory semantics for mathematical language. DSR* can be amended in the same way. But we will argue that, so amended, it isn't as satisfactory as an account of mathematics.

The problem with both STS and DSR* is the perennial access problem (Benacerraf, 1973; Field, 1989). STS includes sets in its ontology. It includes the everyday sets, such as the set of rabbits in Gloucestershire or the set of foxes in Somerset. These are sets of urelements, i.e., sets whose members are not themselves sets. Sets of urelements and in general sets with urelements in their transitive closure are known as impure sets. But STS must include more than just these. It must also include the so-called pure sets, i.e., sets whose transitive closure includes no urelements. If it does not, then the truth of various existential claims and the falsity of certain universal statements becomes contingent. In a universe with no urelements, there can be no impure sets. Thus, if there are no pure sets either, there are no sets at all, and no existential claims about sets can be true and no universal claims about sets can be false. Thus, STS posits pure sets as well as impure sets. But if the truths of mathematics are to be known, and if they are concerned with sets, as STS claims, we must know facts about those pure and impure sets. But, as we will see, we don't.

Similarly for DSR*. It does not include pure or impure sets in its ontology (though it does include Lawvere's sets of "lauter Einsen"). It does, however, posit a vast array of incomplete systems — systems whose elements have no internal nature. But if the truths of mathematics are to be known, and if they are concerned with these incomplete systems, as DSR* claims, we must know facts about them. But, again, as we will see, we don't.

In both cases — in the case of sets and in the case of the incomplete entities that compose the systems posited by DSR* — the epistemic problem is the same. As Justin Clarke-Doane (2016, 21) puts it, such objects are causally, counterfactually, and constitutively independent of us.⁴ And this, the argument goes, prevents us from having any knowledge concerning them.

⁴In fact, I won't assume here that impure sets are independent of us in any of these ways. After all, it's plausible that the UK Government has a causal influence on my life, and we may wish to identify the UK Government with the set of people it contains (Maddy, 1992). Also, if the set of members of my immediate family had contained only three elements, I would not have had a brother. So it doesn't seem to be counterfactually independent of me either. However, as we saw above, STS also posits pure sets, and they are independent of us in these ways.

The argument is given originally by Paul Benacerraf (1973), who assumed a causal condition on knowledge:

Causal S knows $p \Rightarrow S$ believes p and S 's belief in p was (partially) caused by p .

Since sets and incomplete systems are causally independent of us, no fact about them can cause my belief in that fact; and so, on the causal theory of knowledge, I cannot know that fact. But, the causal theory of knowledge is false — amongst other failures, it entails that we cannot know the future.

In the light of this failure, Hartry Field (1989, 26) reformulated the objection using the following alternative necessary condition, which he hoped was sufficiently weak that it would be acceptable to anyone, regardless of their favoured analysis of knowledge:

Explicable Reliability S knows $p \Rightarrow$ it is in principle possible to explain the reliability of S 's beliefs in propositions similar to p .

Now, if p is a proposition about the colour of a particular apple in front of me, it is clear that this condition is satisfied. My beliefs about the colour of fruit are reliable, and their reliability is explained by the way that my visual system interacts with light reflected from the surface of objects. However, if p is a mathematical proposition — one concerning impure sets or incomplete systems — it is not so clear that Field's condition is satisfied. How are we to explain the reliability of my beliefs about such objects? Since these objects are causally independent of me, there can be no causal explanation of the sort I have in the case of the apple. Is there another sort of explanation available?

Here's an attempt, which draws on an observation by Øystein Linnebo (2006, 559-562), though Linnebo ultimately rejects this approach. Each of my mathematical beliefs is derived deductively from a handful of basic axioms. If you take a set-theoretic approach and favour STS over DSR*, these are the axioms of Zermelo-Fraenkel set theory (ZFC). If you take a category-theoretic approach and favour DSR* over STS, these will be the axioms of some category-theoretic analogue to ZFC, such as the axioms of Lawvere's Category of Categories as a Foundation for Mathematics (CCAF), suitably fixed up to deal with the well-known problems (Lawvere, 1966; Isbell, 1967). Whichever you choose, these axioms tell you which systems there are and what basic properties they have. Thus, I explain the reliability of the vast set of mathematical beliefs I hold like this: I believe some basic axioms (those of ZFC or CCAF), and these axioms are true, so my beliefs in them are reliable; and I derive all other mathematical beliefs from these axioms using the method of deductive inference, which is conditionally reliable — given true inputs, it gives true outputs.

Is this explanation satisfactory? You might think that Field's demand for explanation simply re-emerges as a demand to explain the reliability of your beliefs in the basic axioms. But those axioms are few: on standard presentations, there are no more than ten axioms of ZFC or of CCAF. If my beliefs in these axioms are all true, it seems that their reliability — namely, the correlation between the truth of those axioms and my beliefs in those axioms — does not really call for explanation. What's more, the pressure to explain their reliability is reduced even further if, in fact, I derive my belief in the basic axioms from my belief in one single global axiom that entails them all, such as a reflection principle or other limitation of size principle in set theory. If I do this, then I explain the reliability of my mathematical beliefs by noting that I have a single true belief in the reflection principle in question, and I derive all my other mathematical beliefs from that. And this is a satisfactory explanation for the phenomenon for which Field's condition on knowledge, Explicable Reliability, demanded an explanation, namely, the reliability of my beliefs, or the correlation between what I believe about mathematical objects and what is true about those objects. To see this, consider an analogous situation. I notice that two quantities always match: the size of the US economy (presented in some unit) and the number of bacteria in the petri dish in my laboratory. This is a correlation that calls for explanation. And here is a satisfactory explanation: they started out matching at some earlier time, and they have grown at exactly the same rate ever since. In the explanation, I note that they matched at one point, and then note that the principle that determines the size of the US economy on the basis of its size at earlier times is the same principle that determines the number of bacteria in the dish at a given time on the basis of the number at earlier times — it is a particular growth function. Similarly, in my explanation of the correlation between what I believe about mathematics and what is true, I note two things: (i) they match at one point — I believe the single basic axiom and it is true; (ii) the principle that determines the truth of a proposition about mathematical objects is the same as the principle that determines whether I believe a given proposition about those objects — a mathematical proposition is true only if it follows from the single basic axiom, and I believe a mathematical proposition only if I can derive it from that single basic axiom. Now, though he formulates a similar response to Field's objection, Linnebo (2006, 562) argues that it fails, because the putative explanation doesn't explain the *connection* between the mathematical truths and the mathematical beliefs. But that is not what Field's objection demands. It demands only an explanation of the correlation of between those truths and those beliefs. And that we can provide.

So Field's version of the epistemological objection to mathematical realism fails. But there is an alternative, simpler version, and it succeeds. Benacerraf was right to think that a belief in p counts as knowledge only

if there is some connection between p and the belief in p . He was wrong only in thinking that the connection must be causal. Instead, the connection must be counterfactual. A belief in p counts as knowledge only if it is counterfactually related to p in the correct way. The two standard ways to spell out this counterfactual dependence are these:

Sensitivity S knows p only if, S formed her belief in p using method M , and if p were false, S would not believe p using method M (Nozick, 1981).

Safety S knows p only if, in nearby situations, S does not falsely believe p (Sosa, 1999).

Now, it seems to me that, while our mathematical beliefs may well be safe, they are not sensitive. That is, in nearby situations where we hold the same mathematical beliefs that we in fact hold, they are true; but in the nearest situations in which our actual mathematical beliefs are false, we would still have those beliefs.

The standard objection to this claim is that mathematical beliefs, if true, are necessarily true (Clarke-Doane, 2016, 26). If sets exist, they necessarily exist; and they have the same set-theoretic properties at every world. And similarly for the incomplete systems posited by DSR*. If that is right, then there are no nearby cases in which our actually true mathematical beliefs are false, and thus those beliefs are automatically sensitive. However, I see no reason to think that sets, if they exist, exist necessarily. And so I think that there is a genuine question whether our beliefs about the existence and nature of sets are sensitive — what's more, I will argue that the answer to that genuine question is that they are not, and thus do not constitute knowledge.

Why might we think that sets, if they exist at one world and have certain properties there, also exist at every other world and have those same properties there? (And a similar question for the incomplete systems of DSR*; I will focus here on STS and sets, but everything I say will hold, *mutatis mutandis*, for the ontology of DSR*.) Here is one argument: Suppose sets exist at the actual world with certain properties. Now consider another possible world. What could possibly be different about that other world that would make it the case that sets don't exist there, or exist but have different properties? What could be responsible for that difference between worlds? When we consider a concrete object, such as the Eiffel Tower, which exists at one world but not at another, there is usually some difference between the two worlds that is responsible for this existential difference. If the Eiffel Tower does not exist at some non-actual world, there will be some difference between that world and the actual world that explains the difference: perhaps in that other world, Gustave Eiffel, having graduated from the École Centrale des Arts et Manufactures, went to work

for his uncle in Dijon, rather than remaining in Paris. And even when we consider putative contingent abstract objects — such as the sort of fictional entities posited by creationists, or singular propositions concerning contingent concreta — there is again always some difference between a world in which the abstract object exists and a world in which it doesn't that is responsible for the difference: in the creationist case, it would be the fact that the author did not write the fiction in the latter world; in the case of singular propositions, it is just that the subject of the proposition does not exist at the latter world. But what could be responsible for the difference between a world at which sets exist and a world at which they don't? Or between a world at which they exist with one set of properties and a world at which they exist with another set of properties? There is nothing that could play that role, or so the argument runs. But I think this is mistaken. Here is one difference that play that role: at the first world, the sets exist; at the second, they don't! The point is that there is no reason why an existential difference cannot be a brute fact. When an object is at the centre of a causal nexus involving other objects and a variety of causal influences, then it makes sense to ask what is responsible for it failing to exist in another possible world. But sets are not part of such a causal nexus. So there is no reason to think that their existence or non-existence at a possible world cannot be simply a brute fact.

Here's another argument for the necessity of sets. (Again, similar points can be made about incomplete systems.) Mathematical truths are necessarily true; mathematical propositions are about sets; therefore, sets exist necessarily.⁵ Or this variation: Mathematical truths are necessarily true; so the truth-makers of mathematical truths are necessary existents; sets are constituents of the truth-makers of mathematical truths; the constituents of necessary existents are themselves necessary; therefore, sets exist necessarily.⁶ In each case, the problem lies with the first premise: mathematical truths are necessarily true. The problem is not that the premise is false — I will end up agreeing that it is true. The problem is that the mathematical realist — the proponent of STS or DSR* — cannot appeal to this premise at this point in the argument. To see this, consider why we think the first premise is true. Our intuition is driven by the following thought: a mathematical truth makes no demands on the world; there is no way that it requires the world to be; so there is no way the world could be that would fail to make it true. So a mathematical proposition, if true, is necessarily true. The problem is that, while these claims may seem plausible in advance of a precise account of the semantics of these mathematical truths, once we propose a semantics that makes these truths depends on some heavy-duty ontology, such as impure sets (or incomplete systems), we cannot know

⁵Thanks to Ralph Wedgwood for raising this possibility.

⁶Thanks to Daniel Rubio and Trent Dougherty for raising this possibility.

that they are necessary unless we can know facts about the existence and nature of these mathematical entities. But that is precisely what is currently in doubt. The problem is that, as soon as we make mathematical truths depend on this heavy-duty ontology, it becomes much less plausible that they are vacuous and make no demands on the world.

We have no reason, then, to think that sets, if they exist at all, exist necessarily. So suppose they don't. Suppose they exist at the actual world, but not at some other worlds. It then becomes a non-trivial question whether our beliefs about sets are sensitive. What's more, the answer to that question is that they aren't. Suppose I believe some propositions about sets; and suppose that they are true. I form these beliefs using some method: perhaps I deduce them from a single basic axiom, such as a reflection principle or a limitation of size principle; perhaps I look to mathematical practice, and I see what basic principles are required by that, and deduce my mathematical beliefs from the principles I thereby formulate, as Dedekind did; perhaps I form them on the basis of mathematical intuition. Now take the nearest worlds where these beliefs are false — the sets don't exist at these other worlds, perhaps, or they don't have the properties they actually have. In that situation, would I still believe them on the basis of the method I used? The answer, it seems, is that I would. None of the methods I use, nor indeed any method that is actually available to me as a concrete physical entity, take as their input anything that will be different at these nearby worlds. Thus, they will give the same output.

It seems, then, that, if STS or DSR* is true, and if the entities they posit are not necessary contingents, then we cannot know the truths of mathematics because our mathematical beliefs are not sensitive in the way required for knowledge. Without reason to rule out this possibility, it seems that we should then abandon STS and DSR*. But what should we put in their place?

8 A modal version of set-theoretic structuralism

STS says that there are sets, and it says that they have certain properties; and it says that mathematics studies the systems built out of these sets, so that, for instance, a proposition in number theory concerns all set-theoretic systems that satisfy the Dedekind-Peano axioms for a simply infinite system. But problems arise because we cannot know that there are such entities nor that they have the properties that we ascribe to them. A natural alternative is to say: we don't know whether there are sets and we don't know that they have the properties we take them to have, but we know that they might exist and that they might have those properties and we know what is true of them at worlds at which they do exist and do have those properties, and mathematics studies that, so that a proposition in

number theory concerns what is true in all set-theoretic systems that satisfy the Dedekind-Peano axioms *at those worlds at which they exist and have the properties we take them to have*.

Thus, we have a modal version of STS that takes a sentence $\Phi(\mathbb{N}, 0, s)$ of number theory, for instance, and interprets it as the conjunction of the following two propositions:

- (i) Necessarily, for all set-theoretic systems $(N, 0, s)$, if $(N, 0, s)$ is a simply infinite system, then $\Phi(N, 0, s)$.
- (ii) Possibly, there is a set-theoretic system $(N, 0, s)$ that is a simply infinite system.

The theory, which we might call *modal set-theoretic structuralism* or *MSTS*, is therefore reminiscent of Geoffrey Hellman's modal structuralism, but with sets taking the place of Hellman's second-order entities (Hellman, 1993), or logical structuralism, the version of structuralism that Audrey Yap (2009a,b) extracts from Dedekind. Is this a satisfactory account of mathematics? It is close, I think, but not quite there. The problem arises when we ask how mathematics is applied to the physical world.

One of the main attractions of STS is that it gives a straightforward account of the application of mathematics to the physical world. There are three ways in which mathematics is applied in the physical sciences. On the first, the scientist identifies a physical system and notes that, as it occurs in the physical world, this system simply is a certain sort of mathematical system, such as a group. The theorems that apply to that sort of system therefore apply straightforwardly to the physical system in question and we can deduce facts about the physical system. This is the standard way in which group theory is applied in chemistry, for instance: the symmetries of the normal mode of a molecule form a group; having identified that group, the chemist then applies the results of group theory to discover further facts about that molecule. On the second way of applying mathematics to the physical world, the scientist identifies a physical system in *another* possible world — perhaps a perfectly flat plane — that closely but not perfectly resembles some physical system in the actual world — perhaps a very but not completely flat plane. They note that the former simply is a particular sort of mathematical system — perhaps a Euclidean space — and they use a theorem about that sort of mathematical system to deduce a fact about the physical system in the other world. They then infer that something very similar will be true of the physical system in the actual world, since this system closely but imperfectly resembles the system in the other world. In these sorts of applications, we reason about an idealised version of the actual world for the sake of simplicity and because we know the mathematics that applies to those versions; and we infer that the actual world will approximate that idealised version to some extent.

Thirdly, the scientist might identify a physical system, but instead of arguing that it is itself a mathematical system of a particular sort or that there is an idealised version in another world that is a mathematical system of a particular sort, they say that it is related to such a system, perhaps by being embedded into that system in some way. They then use theorems about the mathematical system in question to derive facts about the physical system in question. Thus, for instance, a system of rods of different lengths might be mapped into a complete ordered field in a particular way; and we might appeal to theorems in real analysis to derive facts about that system. In each of these cases, STS gives a straightforward account of how the application works. In the first case, the only ontology that is required consists of the impure sets in the actual world that make up the actual physical system in question, and STS posits those; in the second case, the required ontology consists of the impure sets in the actual world that make up the actual physical system as well as the impure sets in the other possible world that make up the idealised system, and STS posits all of those; and in the third case, the required ontology consists of the impure sets that make up the physical system in question, along with the pure sets that make up the mathematical system or systems to which it is related, and again STS posits all of that. Problems arise, however, when we move to MSTS, the modal version of STS. After all, in all cases, the application relies on the physical phenomena, or some idealised phenomena that closely resembles it, inhabiting the same world as the sets. But the physical system inhabits the actual world; and MSTS is premised precisely on the observation that we might not be able to know whether the actual world contains sets. Thus, in the light of this observation, the straightforward account of the applicability of mathematics that STS offers is lost when we move to MSTS. How, then, might we respond?

9 Instrumental nominalism about set-theoretic structuralism

At this point, we appeal to *instrumental nominalism* (Melia, 1995, 2000; Rosen, 2001; Pettigrew, 2012). Instrumental nominalism is usually deployed as a response to the indispensability arguments in the philosophy of mathematics. According to these arguments, since we use mathematics in our best scientific theories, and since we are justified in believing our best scientific theories, we are justified in believing in the existence of certain mathematical objects. According to the instrumental nominalist, our best scientific theories do not in fact appeal to mathematical objects, even though it seems that they do. Rather, our best scientific theories use mathematical language to make claims about the physical world, but without committing to the entities to which that mathematical language seems to commit

them. Now, other forms of nominalism agree on these point. But those other versions hold that the mathematical language in our scientific theories is really just a shorthand that we use to express something that we might express at greater length without ever mentioning a mathematical object (Field, 1980). Instrumental nominalism, on the other hand, concedes that this might not always be possible. But it holds that it is possible to say what the non-mathematical content of a scientific theory is without effecting such a translation. The idea is that, when I use mathematical language in the statement of a scientific theory, I really say only that the physical world is *as it would be if* the mathematical objects I speak of were to exist and the statement I made that referred to or quantified over those mathematical objects were then true. Suppose, for instance, that my best scientific theory says:

The average star has 2.4 planets.

Some nominalists might try to translate this as follows:

There are 12 planets and 5 stars \vee There are 24 planets and 10 stars \vee There are 36 planets and 15 stars \vee

But this nominalising translation involves an infinite disjunction, and it is unclear how to specify all of the disjuncts without appealing to mathematical language, e.g.,

$\vee_{n=1}^{\infty}$ There are $12n$ planets and $5n$ stars.

The instrumental nominalist, in contrast, concedes that there is no good nominalising translation that removes all reference to mathematical objects, and instead translates it as follows:

There is a possible world w such that: (i) the physical part of w is qualitatively identical with the physical part of the actual world @, (ii) numbers exist at w , and (iii) at w , the average star has 2.4 planets.

As I said above, instrumental nominalism is usually introduced in response to the indispensability argument. But the semantics that it offers for mathematical statements has other appealing features. Consider, for instance, an instrumental nominalist version of set-theoretic structuralism, which I will call *INSTS*. It says this: Suppose there is a mathematical sentence that would be translated, according to STS, as p ; then INSTS translates that sentence as follows:⁷

⁷Pettigrew (2012) explores how to formulate instrumental nominalism without reference to possible worlds.

There is a possible world w such that: (i) the physical part of w is qualitatively identical with the physical part of the actual world @, (ii) sets exist at w , and (iii) at w , p holds.

Like MSTs, INSTs does not rely on the actual existence of mathematical entities. Thus, it avoids the access problem. But, unlike MSTs, it recovers the virtues of the STs account of the application of mathematics in science. Above, we saw that it is not straightforward to give an account of the application of mathematics on MSTs because the physical phenomena to which we apply mathematics exist at the actual world while the mathematical objects involved in the application may well exist only at other worlds. Now this is true also on INSTs. However, on INSTs, we can take the account of the applicability of mathematics from STs and apply it to the world w posited by INSTs at which the physical world is exactly as it is in the actual world, but where there are also sets. Thus, INSTs retains the advantages of MSTs, but adds to them the straightforward account of the applicability of mathematics enjoyed by STs. It is INSTs that I wish to propose as my favoured version of structuralism.

But what about DSR*? In the previous two sections, I have focussed on exploring and developing modal versions of STs, i.e., MSTs and INSTs. But surely it is also possible to create a modal version of DSR* (MDSR*) and an instrumental nominalist version of DSR* (INDSR*). And, if we can do that, why favour INSTs over INDSR*? The reason is this: a crucial component of the translation that instrumental nominalism suggests asserts that there is a possible world at which the mathematical objects in question exist. Thus, while we need not know that sets exist at the actual world in order to know the truths of mathematics that INSTs posits, we do need to know that they exist at some possible world — that is, that they are metaphysically possible. And similarly for INDSR*: we need not know that the incomplete systems that it posits exist at the actual world; but we do need to know that they exist at some possible world — that is, that they are metaphysically possible. How do we come to know this? I will argue that there is a satisfactory answer in the case of INSTs, but not in the case of INDSR*. Or, more precisely, I will argue that any satisfactory answer in the case of INDSR* must go via an answer for the case of INSTs — and, in that case, we would do as well to stick with INSTs. I will adapt the arguments given by Linnebo & Pettigrew (2011), which are intended to show that no category-theoretic foundation for mathematics, such as Lawvere's Elementary Theory of the Category of Sets or Category of Categories as a Foundation for Mathematics, can be epistemically autonomous — it must rely for its justification on a set-theoretic foundation.

Linnebo and Pettigrew argue that, while such a category-theoretic foundation might be formulated independently of set-theoretic notions, our justification for believing the axioms must go via our justification for believing

the axioms of set theory. Now, as we saw above, I don't think that we can know the axioms that constitute a set-theoretic foundation for mathematics, nor any category-theoretic foundation. However, I do think that we can know that there is a possible world at which the set-theoretic axioms are true; we can know that those axioms are metaphysically possible. And our access to this modal fact is exactly as Linnebo and Pettigrew describe it: it goes via the iterative conception of set (Linnebo & Pettigrew, 2011, §5.1).

The iterative conception of set consists of two claims: a claim about the structure of the set-theoretic universe, and a claim about the sets that occupy it. According to the structure claim: the universe of sets is divided into a well-ordered collection of levels; a set exists at one level only if its members all exist at lower levels, and if there is no lower level such that all of its members exist at levels below that. According to the claim about the sets that occupy it:

the iterative conception of set amounts to the following claim of set-theoretic plenitude: relative to the constraints on the hierarchy just stated, whenever a set *could* occupy a level of the hierarchy, it *does*. (Linnebo & Pettigrew, 2011, 245)

It is then straightforward to derive the usual axioms of set theory from this conception. Linnebo and Pettigrew claim that this allows us to know the axioms of set theory; it gives us knowledge of the actual universe of sets. As we saw in section 7 above, I disagree. But I do think that this conception allows us to know that the axioms are metaphysically possible; they describe a possible universe of sets. Most importantly, it convinces us that there are no contradictions lurking in the axioms. There is a conceivable structure to the universe in which they are true, and this gives us our route to modal knowledge. But of course it is just this modal knowledge that we require to know the mathematical truths, given INSTS. Thus, the iterative conception of set delivers us the modal knowledge required by INSTS.

The problem with INDSR* is that there is no analogous route to the modal knowledge that is required by that account of mathematical truth. There is nothing that plays the role of the iterative conception of set for the incomplete systems posited by DSR*. Now, of course, we might note that the axioms of any existing category-theoretic foundation have a model in the set-theoretic universe; and we might then leverage our modal knowledge of the possibility that those set-theoretic axioms are true to obtain knowledge of the possibility that the category-theoretic axioms are true; and from there we might infer that it is not only possible that there is a set-theoretic model of those axioms, but also possible that there are incomplete systems of the sort posited by DSR* that also satisfy those axioms. But clearly such an argument already relies upon, and goes beyond, our argument for the possibility of the set-theoretic universe that we obtained from the iterative conception. And, in that case, it would be safer to stick

with INSTS and the modal knowledge we obtain directly from the iterative conception of set.

10 Conclusion

So, it is INSTS that I wish to propose as my favoured version of structuralism. It makes no claim of incompleteness, and thus avoids Burgess' objection (section 3). It accommodates Hellman's objection because it counts within the subject matter of number theory, for instance, any set-theoretic simply infinite system, and so includes one of Hellman's permuted systems whenever it includes the system from which it was obtained (section 4). The systems it posits as the subject matter of mathematics have exactly the internal natures that mathematicians sometimes care about and investigate in their representation theorems, and so INSTS avoids Weaver's objection, though it can also account for the indifference that mathematicians exhibit at other times towards the internal nature of their objects by saying that, at those times, they quantify over all objects in a given isomorphism class, thereby ignoring their internal natures (section 5). It renders mathematical knowledge possible — it is just modal knowledge — and thereby avoids the access problem (section 7). And finally it preserves a straightforward account of the applicability of mathematics in science (sections 8 and 9).

References

- Awodey, S. (1996). Structure in Mathematics and Logic: A Categorical Perspective. *Philosophia Mathematica*, 3(4), 208–237.
- Awodey, S. (2014). Structuralism, Invariance, and Univalence. *Philosophia Mathematica*, 22(1), 1–11.
- Benacerraf, P. (1973). Mathematical Truth. *Journal of Philosophy*, 70, 661–679.
- Bourbaki, N. (1970). *Éléments de mathématique: Théorie des ensembles*, vol. 1. Springer.
- Burgess, J. P. (1999). Book Review of Stewart Shapiro (1997) *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, Oxford, 1997. xii + 279 pages. *Notre Dame Journal of Formal Logic*, 40(2), 283–291.
- Burgess, J. P. (2015). *Rigor and Structure*. Oxford: Oxford University Press.
- Cantor, G. (1883). *Grundlagen einer allgemeinen Mannigfaltigkeitslehre*. Leipzig: Teubner.
- Carter, J. (2005). Individuation of objects – a problem for structuralism? *Synthese*, 143(3), 291–307.

- Carter, J. (2008). Structuralism as a philosophy of mathematical practice. *Synthese*, 163(2), 119–131.
- Cayley, A. (1854). On the Theory of Groups, as depending on the symbolic equation $\theta^n = 1$. *Philosophical Magazine Series 4*, 7(42), 40–47.
- Clarke-Doane, J. (2016). What is the Benacerraf problem? In F. Pataut (Ed.) *Truth, Objects, Infinity: New Perspectives of the Philosophy of Paul Benacerraf*, vol. 28 of *Logic, Epistemology, and the Unity of Science*. Switzerland: Springer.
- Dedekind, R. (1872). *Stetigkeit und irrationale Zahlen*. Braunschweig: Vieweg.
- Dedekind, R. (1888a). Brief an Weber. In R. Fricke (Ed.) *Gesammelte Mathematische Werke*, vol. 3, (pp. 488–490). Braunschweig: Vieweg.
- Dedekind, R. (1888b). *Was Sind Und Was Sollen Die Zahlen?*. Braunschweig: Vieweg.
- Eilenberg, S., & Mac Lane, S. (1945). General theory of Natural Equivalences. *Transactions of the American Mathematical Society*, 58, 231–294.
- Field, H. (1980). *Science without Numbers*. Oxford: Oxford University Press.
- Field, H. (1989). *Realism, Mathematics, and Modality*. Oxford: Basil Blackwell.
- Fine, K. (1985). *Reasoning with Arbitrary Objects*. Oxford: Basil Blackwell.
- Fraenkel, A. (1922). Zu den Grundlagen der Cantor-Zermeloschen Mengenlehre. *Mathematische Annalen*, 86, 230–237.
- Green, B., & Tao, T. (2008). The primes contain arbitrarily long arithmetic progressions. *Annals of Mathematics*, 167(2), 481–547.
- Hellman, G. (1993). *Mathematics without Numbers: Towards a Modal-Structural Interpretation*. Oxford: Oxford University Press.
- Hellman, G. (2006). Structuralism. In S. Shapiro (Ed.) *Oxford Handbook of Philosophy of Mathematics and Logic*, (pp. 536–62). Oxford: Oxford University Press.
- Isbell, J. R. (1967). Review of (Lawvere, 1966). *Mathematical Reviews*, 34.
- Lawvere, F. W. (1963). Functorial Semantics of Algebraic Theories. *Proceedings of the National Academy of Science of the U.S.A.*, 50(5), 869–872.
- Lawvere, F. W. (1965). An elementary theory of the category of sets. *Proceedings of the National Academy of Science of the U.S.A.*, 52, 1506–1511.
- Lawvere, F. W. (1966). The Category of Categories as a Foundation for Mathematics. In *Proceedings of a Conference on Categorical Algebra (La Jolla, California, 1965)*, (pp. 1–20). New York: Springer.
- Lawvere, F. W. (1994). Cohesive Toposes and Cantor's 'lauter Einsen'. *Philosophia Mathematica*, 3(2), 5–15.

- Linnebo, Ø. (2006). Epistemological Challenges to Mathematical Platonism. *Philosophical Studies*, 129, 545–574.
- Linnebo, Ø., & Pettigrew, R. (2011). Category theory as an autonomous foundation. *Philosophia Mathematica*, 19(3), 227–254.
- Linnebo, Ø., & Pettigrew, R. (2014). Two Types of Abstraction for Structuralism. *Philosophical Quarterly*, 64(255), 267–283.
- Maddy, P. (1992). *Realism in Mathematics*. Oxford: Oxford University Press.
- Mayberry, J. (2000). *The Foundations of Mathematics in the Theory of Sets*. Cambridge: Cambridge University Press.
- Melia, J. (1995). On What There's Not. *Analysis*, 55(4), 223–229.
- Melia, J. (2000). Weaseling away the indispensability argument. *Mind*, 109(435), 455–479.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.
- Peano, G. (1888). *Calcolo Geometrico secondo l'Ausdehnungslehre di H. Grassmann preceduto dalle Operazioni della Logica Deduttiva*. Turin: Bocca Brothers.
- Pettigrew, R. (2008). Platonism and Aristotelianism in Mathematics. *Philosophia Mathematica*, 16(3), 310–332.
- Pettigrew, R. (2012). Indispensability Arguments and Instrumental Nominalism. *Review of Symbolic Logic*, 5(4), 687–709.
- Quine, W. V. O. (1980). On What There Is. In *From a Logical Point of View*, (pp. 1–19). Cambridge, Mass.: Harvard University Press, 2nd ed.
- Reck, E. (2003). Dedekind's structuralism: An interpretation and partial defense. *Synthese*, 137, 369–419.
- Rosen, G. (2001). Nominalism, Naturalism, and Epistemic Relativism. *Philosophical Perspectives*, 15, 69–91.
- Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford: Oxford University Press.
- Shapiro, S. (2006). Structure and Identity. In F. MacBride (Ed.) *Identity and Modality*. Oxford University Press.
- Skolem, T. (1922). Einige Bemerkungen zur axiomatischen Begründung der Mengenlehre. In *Matematikerkongressen i Helsingfors den 4-7 Juli 1922, Den femte skandinaviska matematikerkongressen*.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Noûs*, 33, 141–153.
- Weaver, G. (1998). Structuralism and Representation Theorems. *Philosophia Mathematica*, 6(3), 257–271.

- Yap, A. (2009a). Logical structuralism and Benacerraf's problem. *Synthese*, 171(1), 157–173.
- Yap, A. (2009b). Predicativity and Structuralism in Dedekind's Construction of the Reals. *Erkenntnis*, 71(2), 157–173.
- Zermelo, E. (1908). Untersuchungen über die Grundlagen der Mengenlehre I. *Mathematische Annalen*, 65(2), 261–281.
- Zermelo, E. (1930). Über Grenzzahlen und Mengenbereiche. *Fundamenta Mathematicae*, 16, 29–47.