Paper contrib. to the sections of: GAP. 6 Philosophie - Grundlagen und Anwendungen, Berlin, 11.-14.9.2006, pp. 106-119

Wolfgang Spohn*,†

How are Mathematical Objects Constituted? A Structuralist Answer

In my view, structuralism as presented by Shapiro (1991, 1997), Resnik (1991), and elsewhere offers the most plausible philosophy of mathematics: Mathematics is about structures, indeed it is *the science of pure structures*. Structures have no mysterious ontological status, and hence mathematics is not ontologically mysterious, either. Again, it is no mystery how we can acquire knowledge about structures and thus mathematical knowledge. We find structures everywhere. Hence, if mathematics is about structures, we can apply mathematics everywhere. In this way, structuralism promises to offer straightforward answers to the most pressing problems in the philosophy of mathematics.

However, there are not only structures, there are also mathematical objects, numbers, pairs, triangles, sets, etc. Concerning their nature, structuralism tends to metaphorics, the most preferred metaphor being that mathematical objects are *places* in mathematical structures. Maybe it is not really necessary to say more, since it is only the structures that really matter. Still, I think one should be explicit and precise about mathematical objects, and this is what this paper is intended to achieve. I tend to think that the amendment it adds to structuralism is both trivial and obligatory. Maybe, though, it is contested and hence of substantial interest.

In order to say what mathematical objects are one needs to have a conception of general ontology. I shall present such a conception very sketchily in section 1, just as much as to able to explain my preferred version of Leibniz' principle in section 2, which will become important when I am going to present and defend in section 3 how I think that mathematical objects are constituted.

* Department of Philosophy, University of Konstanz, 78457 Konstanz.

† I am indebted to Wolfgang Benkewitz for discussions about the general essentialistic framework, to Wolfgang Freitag for discussions about Leibniz' principle, and in particular to Christopher von Bülow for discussions and comments concerning mathematical structuralism.

1 Sketchy Remarks on General Ontology

General ontology, which embraces mathematical ontology, must start somewhere, and I find it most natural to start with objects and properties. Maybe one can start elsewhere, with processes, or tropes, or possible worlds, and maybe one can prove that the various starts lead to equivalent results in some specifiable sense. Anyway, here I shall talk about a class D of all possible objects and class $\mathbf{P}$ of all possible n-place properties ($n \geq 0$) (taken in a wide sense so as to comprise relations). And the task is to axiomatically characterize $\mathbf{P}$ and D in a substantial way (in principle formalizable in second-order logic). Thus we gather all we can say about objects and properties in general. It is clear that we thereby characterize *possible* objects; which of them are actual is usually a matter of contingency and not of philosophy. Hence I am certainly a realist about possibilities.

I am not at all clear about all the details of axiomatization. Let me sketch only some basic assumptions. First, the class $\mathbf{P}$ of possible properties must have an algebraic structure. Thus, I assume that each class $\mathbf{P}_n$ of all n-place properties is a *Boolean* algebra and thus comprises conjunctive, disjunctive, and negative properties. Moreover, I assume that each $P \in \mathbf{P}_n$ has a *range of applicability* $A(P) \subseteq D^n$; properties are usually applicable to certain kinds of objects and not to others (stones or numbers are neither awake nor asleep, for instance). Hence, one needs to say how the Boolean structure interacts with the ranges of applicability. Moreover, there are various assumptions relating properties having a different number of places. The most important is the application operation α that is defined for a property $P \in \mathbf{P}_n$ and an object $d \in D$ only when d is in $A(P)$ and that fills the n-th place of P thus turning P into a n-1-place property $\alpha(P, d)$. If all places are filled we get a state of affairs; i.e., $\mathbf{P}_0$ forms the important class of all possible states of affairs. (Hence, the application of a property to an object creates a new item in our ontology and does not *assert* that the property actually applies to the object.) Let us say that a class of properties is *algebraically closed* if it is closed under this vaguely sketched set of algebraic operations without the application operation α and *applicationally closed* if it is moreover closed under α.

$\mathbf{P}$ contains two special properties, the *identity* relation $\equiv$ ($=$ is reserved for metalinguistic identity) and the property E of *actual existence* (or subsistence in Meinong's sense). (The property of possible existence is tautological; every possible object possibly exists.) An important notion is that of a *relational* property; I think it is arguable that it is just the application operation α that creates relational properties and that all the other algebraic operations do not lead outside the realm of relational properties. All the other properties not at least once generated by the application operation α are *non-relational.* $\equiv$ and E are non-relational. But they are not qualitative. Hence, I postulate a class $\mathbf{Q}$ of qualitative properties or *quali-*

ties that is algebraically closed and does neither contain ≡ and E nor any 0-place properties (there are no purely qualitative states of affairs). The algebraic closure of $\mathbf{Q} \cup \{\equiv, \mathrm{E}\}$ is the class of non-relational properties, and the applicational closure of $\mathbf{Q} \cup \{\equiv, \mathrm{E}\}$ is the class $\mathbf{P}$ of all possible properties. Often one would like to avoid including identity and existence among the properties. Hence, I define the class $\mathbf{P}^*$ of *proper properties* as the applicational closure of $\mathbf{Q}$.

A basic notion of general ontology is the *actuality operator* @ that applies to states of affairs and states them to be *facts*; of course, it has to conform to the algebraic behavior of states of affairs.

The final notion to be introduced is perhaps the most important one, namely that of *ontological necessity* basically represented by an operator $\mathbf{N}$ that applies to a one-place property $P \in \mathbf{P}_1$ and an object $d \in D$ and says that P is an *essential property* of d. The conjunction of all the essential properties of d forms the *essence* of d. Corresponding to the distinctions above we can also speak of the qualitative, non-relational, and proper essence of an object. The idea is that an object can't exist or keep its identity without its essence; any object not having the essence of d must be different from d. Thus, objecthood is fundamentally a modal notion. The further conjecture not to be defended here is that this operator $\mathbf{N}$, i.e. the ascription of essential properties, is the foundation for all our talk of ontological necessity. Hence, my essentialism is certainly close in spirit to Fine (1994), though not in details. In particular, an operator $\Box$ for states of affairs may be introduced that asserts their ontological necessity.

2 Leibniz' Principle

I just made the familiar, but somewhat sloppy assertion that an object can't exist or keep its identity without its essence. This raises the issue of an adequate, i.e., a substantial and acceptable formulation of Leibniz' principle; only by thinking about this issue can we gain a proper understanding of objecthood or of the so-called constitution of objects.

Leibniz' principle has two parts; it asserts the indiscernibility of the identical and the identity of the indiscernible. The first part was always the trivial one. It does not hold for intensional contexts; but this only means that intensional contexts do not express properties of the objects denoted by the expressions occurring in these contexts—as Quine has emphasized many times (though his aim was to argue that there is at best necessity *de dicto*). The issue is about the second part. Indiscernibility implicitly quantifies over one-place properties; thus the crucial question is, what is the domain of quantification?

First, it is clear that identity must be excluded. Inclusion of self-identity runs empty, and, as is well known, inclusion of identity $\alpha(\equiv, d)$ with some object d renders Leibniz' principle true, but trivial. The same holds for existence. Possible existence is empty again. And actual existence E should be excluded as well, because it cannot be the only feature distinguishing two objects. Try to imagine the contrary: two objects sharing all qualitative or relational proper properties and still one existing and the other not. This seems impossible to me; the reason, though, will lie in the argument for the positive Leibniz' principle I am defending below.

This leaves us with the proper properties. However, we cannot say that $d = d'$ iff they share all proper one-place properties. As a rule, many proper properties will apply to d only contingently, and hence there is no saying whether or not these contingent properties apply to d as such; only d's essential properties apply to d in any case. This does not seem to be a deep objection, though. The idea was rather to say that $d = d'$ iff they actually share all proper properties, i.e., iff $@(\alpha(P, d))$ exactly if $@(\alpha(P, d'))$ for all proper one-place properties P. However, this won't do, either. The reason is that in this version Leibniz' principle would apply only to actual objects (since $@(\alpha(P, d)$ entails $@(\alpha(\mathrm{E}, d))$, whereas it is intended to apply to all possible objects d.

Thus, one might want to extend the idea that $d = d'$ iff they exist in the same possible worlds and share all their proper properties in those possible worlds. This is plausible, no doubt. However, we cannot yet express this extended idea, because in the ontological order envisaged in section 1 possible worlds emerge quite late as complex and in some sense complete configurations of possible objects and possible properties, i.e., states of affairs. (These would then be Wittgensteinian possible worlds. Of course, there are also Lewisian possible worlds, i.e., (analogously) spatiotemporally maximal possible objects. General ontology should acknowledge both kinds of worlds and clarify their relations.)

Should we wait then with fixing Leibniz' principle till we have reached these constructive stages? No, we should know beforehand about the relation between identity and indiscernibility; otherwise, we won't be able to say what possible worlds might be. For instance, part of saying this will be to explain so-called transworld identity, i.e., to which extent different worlds contain the same individuals. This extent may be null, as Lewis (1986) and elsewhere claims by assuming that strictly speaking each individual exists only in one world and replacing transworld identity by his counterpart relation; but if we want to allow for transworld identity, the envisaged version of Leibniz' principle is of no help at all, even if the construction to be given should make it true. So, however we go about, it seems to be a mistake to refer Leibniz' principle to the properties objects have actually or in other possible worlds.

The general point is this. We characterize possible objects through their essential properties that cannot vary across the worlds, times, and places at which they exist. The other properties, insofar they are applicable, may vary and open up the objects' range of contingency, and they keep their identity, of course, while changing within that range. Hence, it is inappropriate to characterize their identity by their contingent properties. The properties that fix their identity, insofar it can be fixed at all, can only be their essential properties. Therefore, we should take Leibniz' principle as referring only to the objects' essential properties, i.e., considering the previous conclusion, only to the objects' proper essential properties.

Or can we draw a closer boundary? Within the classifications of the previous section, the only narrower range of properties would be the qualitative properties. But even with respect to the actual qualitative properties, let alone the essential qualitative properties, Leibniz' principle would be simply false. There are so many examples for qualitatively indistinguishable objects (note that spatiotemporal location is always a relational affair). And the knockdown argument is given by symmetric worlds and worlds of eternal recurrence in which ever so many objects with the very same qualitative properties exist. If different objects can be distinguished at all, then mostly in relation to other objects. I am not aware of any further classification of properties bearing on Leibniz' principle. Thus my preliminary conclusion is that when we are looking for a substantial and acceptable version of Leibniz' principle, we have to focus on the objects' proper essential properties explicitly including the relational ones.

In fact, many salient essential properties are relational. For instance, for many bodies their having the origin they have is essential, and this is a relational property. For a natural number its position in the progression of natural numbers is essential. And so on. Hence, the *proper essence* is not only the only remaining, but also at least a plausible candidate for turning Leibniz' principle into a non-trivial truth. The candidate is this:

(L) for all d, $d' \in D$, $d = d'$ if and only if d and d' have the same proper essence (proper Leibniz' principle).

Indeed, I maintain that (L) is true.

The implication from left to right remains uncontestable. What about the converse implication? Is the proper essence only necessary or also sufficient to fix the identity of an object? If it is not sufficient, this would entail in view of the above discussion that there is nothing else to fix the identity of an object. Objecthood would then transcend all conceivable access through properties; the identity of an object would be an inexplicable brute fact. The doctrine of *haecceitism* is explained in the literature in slightly diverging ways, perhaps because of diverging theoretical frameworks. Within our essentialistic framework, I take acceptance versus rejection of (L) to be the exact dividing line between anti-haecceitism and haecceitism.

If haecceitism were true, which kind of situation would we have to conceive as possible? For instance, the following: I assume that my origin is essential for me, i.e., that I have developed from this specific egg of my mother and this specific sperm of my father (where the origin of the sperm and the egg and of my mother and my father is essential for them in turn, and so forth). If this were not sufficient to fix my identity, then it must be possible that the world develops up to this sperm and egg in exactly the same way as the actual world, that they unite in exactly the same way as in this world, and still it is not me, but a phantom twin of mine that thereby comes into existence. Why should this not be a possibility?

Note, by the way, that the same kind of scenario decides about the issue whether actual existence may be among the properties relevant for Leibniz' principle. In this situation the only difference between me and my phantom twin is that I exist and he does not. Thus, the anti-haecceitist would be justified to exclude actual existence from the relevant properties, whereas the haeceitist might hope to save Leibniz' principle by including it. This is why I above deferred the discussion of this issue to the present argument.

I am not sure about this possibility. First, it is clear that if there is one phantom twin, there are many, indeed infinitely many, and there does not seem to be an upper boundary to the number of phantom twins. Secondly, it strikes one as absolutely superfluous to postulate all these possibilities. They do not do any work and appear whimsical. So, they seem to be a case for Occam's razor that would reduce the many possibilities to just one, namely me without phantom twins. Generally, though, I am not a friend of Occam's razor. Being economical is a good maxim in practical life, but not in ontology. Accepting Occam's razor just for the sake of economy leaves me unconvinced. Still, if there is a good application of Occam's razor, this seems to be one.

The crucial point rather seems to be about unknowability. These situations are so defined as to be absolutely indistinguishable. There is no way of telling whether I exist or my phantom twin (except by simply stipulating that it is me who exists and not my phantom twin). We might even consider two scenarios: one which looks like the actual world and in which I have a normal continuous existence and another that also looks like the actual world, in which, however, there are two guys alternating each year, two phantoms, as it were, switching haecceitistically. Again, we could not tell which is the actual scenario. Such cases seem to be absurd.

Is unknowability really the crucial point? If so, we seem to import something alien into our discussion. My intention was to do pure ontology, which must certainly be kept free of epistemological considerations. This is the only worry that makes me feel insecure. Perhaps pure ontology may or must be conceived to be so whimsical. Still, the point is not about the knowability of truth, not about Putnam's internal realism. I am sympathetic to this as well (see Spohn 1991), but it would be an entirely different issue requiring entirely different arguments. No, the point is about the knowability of possibility. No reasonable possibilities are lost by limiting possibility in this way. This is why I accept (L), but presently I don't know how to deepen my argument. It must be clear, though, that by accepting (L) we accept a really strong postulate tightly connecting possible objects in D and possible properties in $\mathbf{P}$.

It was perhaps unhappy that I chose cases of personal identity for the sake of vividness, since these may be special; in any case, we seem especially affected by them. But I am presently not interested in persons (as perhaps opposed to human beings). We could consider the same kind of scenarios with respect to the chair I am sitting on, for which, I take it, the circumstances of its production are essential as well; again we might wonder about its phantom twins.

In my examples I have chosen relational properties concerning origin as an example of relational essences. Let me add that I have done so only because this seems so plausible for many empirical objects in space-time. But nothing depended on this. The argument was intended to be general: Leibniz' principle has to refer to essential properties; it cannot refer to identity and existence; hence we consider the strongest properties that remain; and these form the proper essence of the objects referred to in (L). The question which relational properties are part of that proper essence does not play a role in that argument, even though properties of origin may be most plausible in many cases.

Probably, doubts about (L) are raised by philosophers that have a much weaker notion of an essence of an object. Some think, e.g., that only essential properties (or sortals in Strawson's sense), i.e., properties that apply essentially to all the objects to which they apply at all, belong to the essence of an object. For them, any version of Leibniz' principle referring to essences must be unacceptable. However, they face the old question, what else might fix the identity of an object? If the arguments above hold good, they have to reply: nothing. Thus, the present brief discussion of haecceitism bears on them, too.

Does the point about knowability I was raising depend in any way on the fact that I am always alluding to relational essences because purely qualitative or intrinsic essences are generally insufficient? I don't think that this distinction is relevant with regard to knowability. From my point of view, it is not clear which objects have an intrinsic essence. They seem to be special. Maybe it is only Lewisian possible worlds; everything inside of them is essential to them; nothing could be different; and relations they might have to something external are not relevant. Now, the point about knowability arises for them as well. If (L) holds good, there is only one copy of each possible world. If one rejects (L), one allows for infinitely many duplicates of each world. I am not aware that Lewis explicitly discusses this issue. But if Lewis (1986) says that a possible world is a way a world might have been, he thereby seems to assume that for one way there is one world and not many duplicates. So I conclude that acceptance or rejection of (L) does not depend on relational versus purely intrinsic essences.

This brief section has touched deep issues and has made deep claims that call for a thoroughgoing comparative discussion. I am well aware that there are sharply diverging views that are not refuted in such a superficial manner and many similar views worth inquiring for their subtle differences. However, I cannot afford the required discussion in this brief note. So, let us rush on to the application of (L) in the philosophy of mathematics.

3 Systems, Structures, and the Constitution of Mathematical Objects

Indeed, I believe that with these ideas about the constitution of objects in general in the background we can now provide in a straightforward way what I complained to be missing in current presentations of structuralism. The basic categories of structuralism are systems and structures; they can directly be defined in the framework sketched in section 1:

$S = \langle \mathbf{R}, E\rangle$ is a *system* if and only if $\mathbf{R} \subseteq \mathbf{P}$ and $E \subseteq D$ such that $\mathbf{R}$ and E are applicationally closed and the application operation α is always defined w.r.t. $\mathbf{R}$ and E. In a nutshell, a system is just any closed part of our ontology given by $\mathbf{P}$ and D. One could imagine other closure properties, e.g., closure under ontological dependence, a notion I have not introduced here; but here we may be content with the definition given. Thus, there will be natural and unnatural systems.

It should be clear, though, that systems are not objects in our ontology, at least not so far. They are something we can talk about only in our informal set-theoretic metalanguage. As is well known, the projection of $\mathbf{P}$ into D is always on the verge of paradox and hence something to be done with great care.

What the facts within a system are can be expressed with the help of the actuality operator @. Thereby, systems have certain shapes, and their shapes can be similar or dissimilar. This can be more precisely captured in the following way: Let $S = \langle \mathbf{R}, E\rangle$ and $S' = \langle \mathbf{R}', E'\rangle$ be two systems. Then f is an *isomorphism from S to S'* if and only if f is a bijection from $\mathbf{R}$ onto $\mathbf{R}'$ and from E onto E' such that algebraic operations on $\mathbf{P}$ are preserved under f and for all $P \in \mathbf{R}_0$ $@(P)$ iff $@(f(P))$. S and S' are *isomorphic* if and only if there is an isomorphism from S to S'.

What, then, are structures? Intuitively, they are the forms or shapes of systems; it is systems that have a structure. Hence, a *structure* is just a property of systems (in our metatheoretic sense, i.e., not something belonging $\mathbf{P}$) or, what comes to the same, a relation between a set of properties and a set of objects. However, this is still too wide to be an appropriate characterization of structures. Structures are not any properties of systems, but properties they have in virtue of the facts obtaining in the systems. Thus, isomorphic systems must have the same structure. Finally, a structure must not be empty; there must be systems having that structure. So, to sum up: a property **S** of systems is a *structure* if and only if some system S has **S** and if for any two isomorphic systems S and S' either both have or both lack the property **S**.

The talk of structures is ubiquitous, but of course it is most prominent in mathematics. Indeed, mathematics is in principle interested in structures so generally conceived (of course, uninteresting structures abound). However, if our more specific concern is mathematical ontology, then we have to focus on categorical structures—where **K** is a *categorical structure* if and only if any two systems having the structure **K** are isomorphic. (Since categorical structures are structures, they are consistent.) Let me explain the relation between ontology and categoricity:

There is, for instance, the group structure. It is not categorical, but if we fix the cardinality of a group to be a particular prime number, it is categorical. There is the structure of natural numbers as fixed in the axioms of Peano arithmetic. Gödel's famous incompleteness theorem (together with Gentzen's consistency proof) implied that it is not categorical, but has nonstandard interpretations. However, it is well known that second-order Peano arithmetic *is* categorical. Thus, the examples may be multiplied.

Let us consider natural numbers for a while. What are they? So far, we have no notion of them. We have a lot of objects in D. (That is, we don't even know so far that they are a lot, because I have not stated any existence axiom for objects and no axiom for producing objects out of existing ones.) So, we have a lot of systems. And we may hope that some systems have the categorical structure of the natural numbers, i.e., are progressions. But do we thereby have the natural numbers? It does not seem so. On the other hand, we may say that we have all there is to know about the natural numbers; the progressive structure is, in a way, all there is to them. Structuralists have introduced the metaphor of natural numbers being the *places* of the structure of natural numbers. There are many different systems the objects of which fit into these places, but the places of the structure are unique, and they *are* the numbers. Similar assertions may already be found in Benacerraf (1965).

Note that the categoricity of the natural number structure is essential for speaking of these places. Take groups, by contrast. There are many non-isomorphic groups, and therefore it does not make sense to speak of *the* places of *the* group structure. This is so even when we more specifically consider quaternary groups, for instance, since such groups, although they must have four members, allow for non-isomorphic group operations. Matters become different with, say, the binary or the ternary group, but just because they are categorical structures.

The metaphor of places in a structure (or of offices in an administration as opposed to the occupants of these offices) is clear and powerful, but it remains a metaphor. How then to turn such places into mathematical objects? The classical way was to take some more or less natural distinguished system to represent the structure, for instance, the Frege numbers or the von Neumann numbers. But this does not seem to be correct; what we get thereby always seem to be representatives of numbers, never the numbers themselves. Now, one can try to start from the structure, from the property of systems itself, and to somehow abstract or construct from it the intended objects. But again it seems no construction is involved. The structuralists themselves are evasive concerning this issue; in any case, I do not find a rigorous account of it in Resnik (1991) or Shapiro (1997). So, how might one understand places in a structure as mathematical objects?

I propose to maximally short-circuit the answer. What characterizes the places in a categorical structure is that these places have their locus in the structure essentially and that there is nothing else to be said about them. The number 2 is just that place of the natural number structure which is essentially related to all the other places of that structure as specified in that structure, which in particular is essentially the successor of 1 (which in turn is essentially the successor of 0 which is essentially the only number not being a successor) and which has no property beyond those necessitated by the structure.

In other words: I propose the following powerful existence axiom for objects in D and properties in $\mathbf{P}$:

(M) For each categorical structure $\mathbf{K}$ there is a distinguished minimal system $S^{\mathbf{K}} = \langle \mathbf{R}^{\mathbf{K}}, E^{\mathbf{K}} \rangle$ such that
  (a) $S^{\mathbf{K}}$ has the structure $\mathbf{K}$,
  (b) for each $P \in \mathbf{R}^{\mathbf{K}}_0$, if $@(P)$, then $\Box P$,
  (c) for each $P \in \mathbf{R}^{\mathbf{K}}_0$ $A(P) = (E^{\mathbf{K}})^n$.
  The system $S^{\mathbf{K}}$ is called the *canonical system* having $\mathbf{K}$.

Here, condition (b) says that all the relations obtaining in the system $S^{\mathbf{K}}$ obtain there necessarily; mathematical truth is necessary truth, as everybody agrees. Thus, the essence of all the objects in $E^{\mathbf{K}}$ consists in being related to the other objects in $E^{\mathbf{K}}$ in the way they are. Condition (c) says that the properties in $\mathbf{R}^{\mathbf{K}}$ do not apply beyond the objects in the system $S^{\mathbf{K}}$. This is meant by characterizing the canonical system as minimal. It is often suggested, e.g., in Benacerraf (1965), that such a condition is intuitively correct.

It is not part of (M), and minimality does not mean, that the objects in $E^{\mathbf{K}}$ have no relations to objects outside—correctly so, I think. The number 2, for instance, carries the relation "numbers" to many things, e.g. to binary sets or to plural objects like couples, etc. This should certainly be allowed. May the objects in $E^{\mathbf{K}}$ have more properties and stand in more relations than contained in $\mathbf{R}^{\mathbf{K}}$? Probably yes. What, e.g., about the natural number property "randomly chosen by Merlin"? This is a relational property we may suppose to be applying exclusively to natural numbers, and though it might be extensionally equivalent to some property in Peano arithmetic, it is certainly not a member of the $\mathbf{R}^{\mathbf{K}}$ pertinent to arithmetic.

(M) seems to be a powerful existence axiom. Is it really so? One might think that it is only a conditional existence axiom. Only if $\mathbf{P}$ and D are already so rich as to contain appropriate systems, structures are realized, and only then (M) asserts also the existence of canonical systems in the case of categorical structures. However, this is not the intended reading of (M). (M) is intended as an unconditional existence postulate: Whenever $\mathbf{K}$ *is* a categorical and hence consistent structure, the canonical system $S^{\mathbf{K}}$ and its objects exist and thereby shows the structure $\mathbf{K}$ to be realizable—whether or not we can prove this. In order to also *prove* the existence of these objects, we have to somehow prove the consistency and categoricity of the structure $\mathbf{K}$. But ontology as such is independent of provability.

So, if (M) is as powerful as I say, is it threatened by paradox? I don't see how. If a structure is categorical it must be consistent. Moreover, (M) implies that the canonical systems pertaining to two different categorical structures are disjoint; they cannot generate conflict. So, I don't see any threat. (Indeed, the argument is almost a consistency proof, I find.)

Note, finally, that (M) accurately observe the proper Leibniz' principle (L). (M) specifies the relational essences of mathematical objects, and according to (L) this is all we have to do in order to generate these objects and to fix their identity. These essences are not specified individually. Rather, they are specified for all objects of a canonical system at once; all of them are mutually ontologically dependent. This seems the right thing to say, and the whole point of the previous section was to be able to say this.

All in all, (M) seems to capture exactly what we want. Moreover, it preserves the structuralist answers to the basic questions in the philosophy of mathematics mentioned in the introduction. If (M) tells us what mathematical objects are, it also tells us how we can refer to them; if we can refer to the structure, we can ipso facto refer to the canonical system and its objects. And it tells us how we can know anything about the mathematical objects. This is so because it is no mystery how we can investigate, and know something about, categorical structures. Finally, it explains how mathematics is applicable outside mathematics, since it is no mystery how structures are applicable everywhere and since mathematical objects are nothing but the objects of minimal realizations of such structures, which may thus be rediscovered everywhere. Of course, these are not full answers; but they point in the right direction, I think.

However, to my knowledge, (M), though suggestive, has not been proposed in the literature. Why? One reason is that structuralists have not felt the urgency to introduce objects besides structures. Another reason is certainly that one has to state the ontological framework more or less in the way I did in order to be able to state (M) as above. A more important reason, though, is that there are objections to this procedure. I know of two, a minor and a major objection.

The minor objection is that one would like to say, e.g., that the natural number 2 and the real number 2 are the same number. But according to (M), by necessity, no natural number is a real number and vice versa. However, this appears to me to be a palatable consequence. Strictly speaking, it seems right to insist on this difference. And loosely speaking, natural numbers are, of course, embeddable into the real numbers, and then one might think about defining weaker senses of identity according to which what is embeddable is identical. I shall not pursue here those weaker senses.

The major objection is generated by so-called non-trivial automorphisms. An *automorphism* of a system is an isomorphism from that system to itself. Identity is always such an automorphism, viz. the trivial one. The system of natural numbers has no non-trivial automorphism; that is why it is particularly suited to motivate the structuralist view of mathematics. But Euclidean space, for instance, has uncountable non-trivial automorphisms, rotations, dilations, and translations. Let us look at the simplest example of that kind in order to see what the problem is supposed to be:

A very simple structure is the *dual* structure; any system has the dual structure iff it consists of exactly two different objects. The dual structure is obviously categorical; since identity and difference are the only relations that matter, any two systems consisting of two different objects are isomorphic. So, according to (M) there is the canonical system pertaining to the dual structure which we might call *Duality*. It consists of two objects that we might call the *One* and the *Other*, and the only things we can tell about these objects is that they are self-identical and different from each other. Of course, this is a very shallow system, but it has a non-trivial automorphism mapping the One onto the Other and vice versa. So, why is the One the One and not the Other? The Other could serve just as well as the One and the One as the Other. In short, what distinguishes the two?

I cannot really see a mystery here. First, the One is identical to the One and different from the Other, and hence it behaves in a different way from the Other. This answer may appear to be too trivial, but it is only so trivial because the dual structure is so poor of descriptive means. If it would contain more relations I could state more interesting differences (despite possible automorphisms). Still, the answer seems unsatisfactory. The complaint is that, if one were to put the One and the Other into an urn, so to speak, and to draw one at random, it would be impossible to tell which one was drawn. It seems the One is to be the One simply in virtue of being so called. Yes precisely, and I do not see why this response is not good enough. According to (M), the dual structure is realized in its canonical system, which consists of two objects that we simply stipulate to be the One and the Other. Once we have this canonical structure we may correctly observe that it has a non-trivial automorphism, i.e., that exchanging the roles of the One and the Other generates an isomorphic system, indeed something more similar: an isomorphic system in which the pertinent relations also apply necessarily and solely to the objects of the system. That's all.

Perhaps, though, this example is too simple to be easily assessed. Let us hence look at a more complex, but more familiar example, say, Euclidean space. How does my response carry over to it? Well, there is the categorical axiomatization of three-dimensional Euclidean space. It provides the structure we are considering. It has a canonical system. Here they are, all the points of this system, each self-identical and standing in the various geometrical relations to all the others. Let's give them names. Let us stipulatively and rigidly call one point (0,0,0), and let us stipulatively and rigidly erect a trihedron with (0,0,0) in its center and with (1,0,0), (0,1,0), and (0,0,1) as its stipulated end-points. This helps us to a name for each point in Euclidean space in the familiar way. Starting from this canonical system, we may again observe that each rotation, dilation, and translation generates an isomorphic system (with the stronger similarity already noted w.r.t. duality).

Again, however, you may ask which of the uncountably many points spread out in space is the point (0,0,0)? What are you here asking for? Should I point at (0,0,0)? Of course, I can't. I cannot point at objects that are not in real space and time. You might have asked for deferred ostension. Again, that's either easy or impossible, depending on whether or not you ask for a substantial account of the relevant deference. However, I can tell you what the point (0,0,0) is. In just these words, or with more complicated descriptions ultimately referring to the basic trihedron; there are no other ways of referring to it. If you still miss something I do not know what it is. At least, this is what I have to say from the point of view of (M). It appears good enough to me.

REFERENCES

Benacerraf, P. (1965). What numbers could not be. *Philosophical Review* 74, 47-73.

Fine, K. (1994). Essence and modality. *Philosophical Perspectives* 8, 1-16.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Resnik, M. (1997). *Mathematics as a Science of Patterns*. Oxford: Oxford University Press.

Shapiro, S. (1991). *Foundations without Foundationalism. A Case for Second-Order Logic*. Oxford: Oxford University Press.

Shapiro, S. (1997). *Philosophy of Mathematics. Structure and Ontology*. Oxford: Oxford University Press.

Spohn, W. (1991). A reason for explanation: Explanations provide stable reasons, *Existence and Explanation*, ed. by W. Spohn, B.C. van Fraassen & B. Skyrms, 165-196. Dordrecht: Springer.