

An answer to Hellman’s question: “Does category theory provide a framework for mathematical structuralism?”

Steve Awodey*

June 6, 2003

The straightforward answer is: “yes, obviously”. In my paper [1], to which Hellman [3] is referring, I distinguished between “mathematical” and “philosophical” structuralism; the former refers to a certain, now typical, “abstract” way of practicing mathematics; the later, to its philosophical interpretation. Hellman will surely agree that category theory provides a framework (indeed, *the* currently dominant one) for the practice of modern, abstract mathematics.

In these terms, the intended question really is “Does category theory provide a framework for *philosophical* structuralism?”. In [1], I took a first step toward answering this in the affirmative, by showing how the language and methods of category theory can be used to determine a notion of “structure” that is both precise and yet flexible enough to do some philosophical work. I did not even attempt to lay out a “philosophical position,” some sort of “categorical structuralism,” based on that notion. It is therefore entirely understandable that Hellman’s discussion of what he *infers* such a position would look like is somewhat off the mark.

Indeed, Hellman thinks that my proposal to use category theory to formulate a philosophical structuralism is “naturally viewed in the context of Mac Lane’s repeated claim that category theory provides an autonomous foundation of mathematics as an alternative to set theory” (p. 129; despite my rejection of “categorical foundations”). The reason such a foundation is needed,

*I have benefitted from conversations with André Carus, Henrik Forssell, Geoffrey Hellman, William Lawvere, Colin McLarty, and Michael Warren.

he says, is clear: “if category theory is not autonomous but rather must be seen ultimately as developed within set theory, then Awodey’s suggestion could not be realized... Thus, we cannot hope to assess Awodey’s suggestion without also (re-)examining Mac Lane’s thesis.” (p. 129f.) Like some other discussions of category theory, however, this misses the point of my proposal, which is not to prefer this or that foundation, but to use category theory to avoid the whole business of “foundations”. As for Mac Lane, all that talk about “categorical foundations” was more bridge-building than foundation-building. No one doing category theory thinks we are someday going to find the one “true topos,” in which all mathematics happens. The translations of set theory into topos theory (and other categories) are intended to show that categories like toposes can be used to do a lot of mathematics *for those used to doing mathematics in set theory*; they are not supposed to show that topos theory is the new universal “system of foundations,” intended to replace set theory. Indeed, the idea of “doing mathematics categorically” involves a different point of view than the customary foundational one, as I shall try to explain in this note.

A related misunderstanding is Hellman’s “problem of the home address.” The question “where do categories come from and where do they live?” (p. 136) asks for something that only seems reasonable from a foundational perspective. This requirement exemplifies a general pattern in the discussion: Hellman sympathizes with the structural viewpoint, and even appreciates and accepts many of the details of the categorical approach to it, but he thinks that there is still something missing in the overall category theoretical position. The something is to be provided by his modal-structuralism. But I contend that what is missing is only a correct understanding of the categorical approach.¹ That would hardly be surprising; after all, no such categorical position has yet been articulated.² So let me now have a go at it—at least then we’ll be talking about the same position.

¹This is not to say that the modal structural approach is intended as “foundational” in the sense to be defined below.

²Actually, Mac Lane began trying to develop such a position under the heading of “the protean character of Mathematics” in a couple of late papers [4, 5]. My notion of the “schematic” character of mathematics, below, derives from that source.

Categorical structuralism

As a first, very rough, approximation, we may say that the point of view that we are going to describe emphasizes form over content; descriptions over constructions; specification of assumptions over deductive foundations; characterization of essential properties over constitution of objects having those properties.

With only slightly more precision, we can say that the “foundational perspective,” to which we are proposing an alternative, is based on the idea of building up specific “mathematical objects” within a particular “foundational system,” in such a way that:

1. there are enough such objects to represent the various kinds of numbers, as well as the spaces, groups, manifolds etc. of everyday mathematics, and
2. there are enough laws, rules, and axioms to warrant all of the usual inferences and arguments made in mathematics about these things, as well as at least some of the most obvious “rounding off” statements dealing with features of the system itself (like the well-foundedness of all sets, a question that does not arise in non-set-theoretic mathematics).

As opposed to this one-universe, “global foundational” view, the “categorical structural” one we advocate is based instead on the idea of specifying, for a given theorem or theory³ only the required or relevant degree of information or structure, the essential features of a given situation, for the purpose at hand, without assuming some ultimate knowledge, specification, or determination of the “objects” involved. The laws, rules, and axioms involved in a particular piece of reasoning, or a field of mathematics, may vary from one to the next, or even from one mathematician or epoch to another. The statement of the inferential machinery involved thus becomes a (tacit) part of the mathematics; functional analysis makes heavy use of abstract functions and the axiom of choice, some theorems in algebra rely on the continuum hypothesis; many arguments in homology theory are purely algebraic, once given the non-algebraic objects that they deal with; theorems in constructive analysis avoid impredicative constructions; 19th century analysis employed

³As in the topologist’s sense of “homotopy theory,” not the logician’s sense of “first-order theory”.

other methods than modern-day analysis, and so on. The methods of reasoning involved in different parts of mathematics are not “global” and uniform across fields or even between different theorems, but are themselves “local” or relative.

Thus according to our view, there is neither a once-and-for-all universe of all mathematical objects, nor a once-and-for-all system of all mathematical inferences. Are there, then, various and changing universes and systems? How are they determined, and how are they related? Here I would rather say that there are *no* such universes or systems; or rather, that the question itself is still based on a “foundationalist” preconception about the nature of mathematical statements.

Top-down versus bottom-up

Let us try a different tack. To understand (describe) a piece of mathematics (say, that in the complex numbers $i^5 = i$) the foundationalist must “construct” the terms involved (the complex numbers and their multiplication operation, and perhaps even the identity relation) and then prove that the specific entities so constructed do indeed have the stated property. The structuralist can simply observe that

- (i) in any ring, if $x^2 = -1$ then $x^5 = x$, and
- (ii) the complex numbers are by definition a ring with an element i such that $i^2 = -1$, and having a couple of other distinctive properties.

The foundationalist may now object that he, too, can show by the same simple proof that:

- (i') any ring in his universe with an element such that $x^2 = -1$ also has $x^5 = x$, and
- (ii') his complex numbers are a ring in which i is a root of -1 .

But this statement is vastly different from the one the structuralist makes; it involves consideration of a possibly huge but fixed range of specific rings, as well as of a particular ring consisting of equivalence classes of pairs of Dedekind cuts of ... The very simple statement (i) made by the structuralist involves the relevant features of the situation, described from the top down, as it were; and it applies in any other situation sharing those features: two

associative operations, one distributing over the other, etc. Isn't this just the statement (i') made by the foundationalist?

No. Statement (i') is a universal quantification over a specific range of specific "objects", presumed or constructed, but somehow fixed and given. Statement (i), although it talks about "any ring", is not about all fixed rings in a fixed universe, or even all "possible" rings in a range of "possible universes". It is a *schematic* statement about a *structure*—the structure of rings—which can have various *instances*. We next need to make these notions more precise.

An example

Suppose instead of $x^2 = -1$ we consider the condition,

$$x^2 + x + 1 = x , \tag{1}$$

which is equivalent to the former one in rings, but also makes sense in similar structures having no inverses, like the natural numbers ("semi-rings with unit"). Now, one can still show that any element satisfying (1) also satisfies $x^5 = x$, by the following clever argument.⁴ Multiplying (1) by $x^{(n-1)}$, we see that for any n :

$$x^{(n+1)} + x^n + x^{(n-1)} = x^n . \tag{2}$$

⁴Due to Marcelo Fiore.

Thus, expanding and contracting by (2), we can calculate as follows:

$$\begin{aligned}
x^5 &= x^6 + x^5 + x^4 \\
&= x^6 + x^5 + x^5 + x^4 + x^3 \\
&= x^6 + x^5 + x^5 + x^4 + x^4 + x^3 + x^2 \\
&= x^6 + x^5 + x^5 + x^4 + x^4 + x^3 + x^3 + x^2 + x \\
&= x^6 + x^5 + x^5 + x^4 + x^4 + x^3 + x^3 + x^2 + x^2 + x + 1 \\
&= x^5 + x^5 + x^4 + x^3 + x^3 + x^2 + x^2 + x + 1 \\
&= x^5 + x^4 + x^3 + x^2 + x^2 + x + 1 \\
&= x^4 + x^2 + x^2 + x + 1 \\
&= x^4 + x^3 + x^2 + x^2 + x + x + 1 \\
&= x^3 + x^2 + x + x + 1 \\
&= x^2 + x + 1 \\
&= x
\end{aligned}$$

Now one and the same simple equational proof shows that $i^5 = i$ in the complex numbers and that any element satisfying (1) in a semi-ring with unit also satisfies $x^5 = x$. Moreover, it also shows that e.g. any *set* X such that $X^2 + X + 1 \cong X$ also has a canonical isomorphism $X^5 \cong X$, obtained by essentially the same calculation.⁵ And the same holds for any object X in any category with finite products and coproducts that are distributive, say, the category $\text{Sets}^{\mathbb{C}}$ of all set-valued functors on a small category \mathbb{C} , as well as for countless other situations that are “formally similar”.

This is not a matter of universally quantifying over a fixed range of particular instances of, say, specific sets with specific binary operations. Rather it is a single, simple proof at a certain level of abstraction, that applies to various different cases, each at perhaps other levels of abstraction, and none of which must be assumed to be any more or less general or specific for the calculation to apply.

Mathematical theorems are “schematic”

Every mathematical theorem is of the form “if such-and-such is the case, then so-and-so holds”. That is, the “things” referred to are assumed to have

⁵Consider e.g. the set of freely generated terms in a language with one constant symbol, a unary operation $\neg t$, and a binary operation $s * t$.

certain properties, and then it is shown, using the tacitly assumed methods of reasoning, that they also have some other properties. Any finitely generated, abelian group has a certain decomposition; the fundamental group of any topological space acts in a certain way on the universal covering space; and so on. Of course, many theorems don't literally have this form, but every theorem has some conditions under which it obtains.

Theorems state connections, relations, and properties of the structures involved: group, topological, continuous actions, etc. The proof of a theorem involves the structures mentioned, and perhaps many others along the way, together with some general principles of reasoning like those collected up in logic, set theory, category theory, etc. But it does not involve the specific nature of the structures, or their components, in an absolute sense. That is, there is a certain degree of "analysis" or specificity required for the proof, and beyond that, it doesn't matter what the structures are supposed to be or to "consist of"—the elements of the group, the points of the space, are simply *undetermined*.

This lack of specificity or determination is not an accidental feature of mathematics, to be described as universal quantification over all particular instances in a specific foundational system as the foundationalist would have it — a contrived and fantastic interpretation of actual mathematical practice (and even more so of historical mathematics!). Rather it is characteristic of mathematical statements that the particular nature of the entities involved plays no role, but rather their relations, operations, etc.—the "structures" that they bear—are related, connected, and described in the statements and proofs of theorems. It is a theorem in topology that the first homology group of an arcwise-connected space is naturally isomorphic to the abelianization of the fundamental group of the space. This statement doesn't depend on the specific points of the space, or even on the specific space; it is about a connection between homology and homotopy. In this sense, mathematical statements (theorems, proofs, etc., even definitions) are about connections, operations, relations, properties of connections, operations on relations, connections between relations on properties, and so on.

The "schematic" element in mathematical theorems, definitions, and even proofs is not captured by treating the indeterminate objects involved as universally quantified variables, as quantification requires a fixed domain over which the range of the variable is restricted. This schematic character is more akin, rather, to the phenomenon Russell's "typical ambiguity" was intended

to capture.⁶ A somewhat analogous distinction is that between a general statement about real numbers and a statement about the indeterminate x in the ring $\mathbb{R}[x]$ of polynomials with real coefficients. For every real number x , the number $x^2 + 1$ has a square root; but there is of course no polynomial $p(x)$ such that $p(x)^2 = x^2 + 1$. In the case of the quantified statement, we consider what is true for an arbitrary real number, while in the indeterminate case we, in effect, consider what holds for an arbitrary number x , in an arbitrary ring over \mathbb{R} . One could say, rather speculatively, that the difference in both cases seems to be related to that between what is *true* for all of a fixed range of values, as opposed to what can be *proved* for an indeterminate value. Presumably, Hellman would use a modal logical approach here. I think a more straightforward analysis is desirable. But first let us consider a possible objection.

That sounds like Russell’s if-then-ism!

Of course, the idea that mathematics is in some sense hypothetical, and that it is about relations of relations of relations of ... is not new—a version of it was stated by Russell in the *Principles of Mathematics*. We distinguish the position proposed here from old-fashioned “if-then-ism” in two respects:

First, if-then-ism is sometimes understood to pertain to the underlying methods of reasoning—the laws, axioms, and rules of the foundational system. The theorems of mathematics, even when stated in the form “if A , then B ”, are supposed to “really” mean “if the laws, axioms, and rules of the system are true and correct, then if A , then B ”. That’s the way, e.g., one sometimes hears set theoretical foundations described. The argument against if-then-ism in this form is that it makes all theorems hypothetical; they can never really be known, because the antecedent conditions will always remain in doubt. We may never know whether the axioms of ZFC are true, or they may even be inconsistent, and so it will not do to carry them along as conditions on every theorem. But in our case, the conditions are rather of the kind “if G is a finitely generated abelian group”, not “if the axioms of ZFC are true”. There is usually no question about *whether* such conditions are ever satisfied; rather, like axiomatic definitions, they serve to specify the range of application of the subsequent statement. A theorem of the form

⁶In [2], Feferman recognizes the similarity between Russell’s typical ambiguity and category theory’s relative use of the concept of “smallness”.

“If G is a normal subgroup of the fundamental group of the space X , then ...” is not hypothetical, in the sense that the critic of if-then-ism objects to. The truth of the consequent statement doesn’t depend on some unknown or unknowable antecedent conditions; rather it *applies* only to those cases specified by the antecedent description. In cases where we are not sure whether the conditions at issue are ever satisfied, i.e. whether they are consistent, we have no recourse but to investigate their consequences in order to gain more information. Of course, establishing any “if ..., then ...” implication requires some tacitly assumed methods of reasoning, from simple chains of equations, to, say, ZFC. And where these methods are not conventionally assumed or obviously inferred, the statement of the theorem will generally include them: “assuming the axiom of choice, ...” or “given a measurable cardinal, ...” or “intuitionistically, ...”.⁷

But the essential difference between the position being sketched here and old-fashioned, relational structuralism is the idea of a top-down description, which presupposes no bottom-up hierarchy of things. For Russell, every relation had to be a relation on *some things* which, even if they were themselves analyzable into relations, had to be among some other *things*, ..., and this process had to either stop somewhere (atoms), or an account had to be given of infinite analysis.

The difficulty arises in the preoccupation with *relations* as the fundamental notion of “structure”; for a relation presupposes its relata, and off we go into the descent of Russellean analysis. If we take instead the perfectly autonomous notion of a morphism in a category, we can build structures out of them to our heart’s content, without ever having to ask what might be in them.

Why category theory?

No one claims that category theory is the *only* way to talk about structures of structures of ... Or even that it is the *best* way (although I know of no better one). The only claim being made in this connection is that it is a *very good* way. There are reasons why this is so, of which some are historical and accidental. Thus, while things could have gone differently, with Eilenberg

⁷There is also a kind of hypothetico-deductivism sometimes asserted by logicians: theorems are proofs from premisses in deductive systems. As this is not the kind of relational structuralism the early Russell had in mind, nor a position the one advocated here is in any danger of being confused with, I will not address it.

and Mac Lane instead inventing “schmategories”, given the same 50 years of study, by the same people, for the same ends, we would now have a very good theory of schmategories. But there are also theoretical reasons why that the work was done on *categories* and not in graph theory, universal algebra, first-order logic, or descriptive set theory. Category theory was developed so extensively because the notion of a category, and the related notions of functoriality, naturality, and adjointness, proved to be so effective in modern, abstract mathematics. And the reason for this broad applicability has a lot to do precisely with their effectiveness at specifying and manipulating *structures*. I’ll try very briefly to indicate why this is so.

In category theory, such notions as relation, connection, property, and operation are all subsumed under the primitive notion of a morphism. It is general and flexible enough that it can be made to do the work of all those other notions and more. Need relations? use products and monomorphisms; operations? morphisms on products; homomorphisms? consider the category of structures; connections between structures? use functors between categories; connections among connections? categories of functors; and so on it goes. Many apparently different phenomena can be described in a uniform way, and thus easily related to each other, by the language of category theory.

But where do all these categories live? What is their home address? When you consider, for example, the notion (involving just a few objects and arrows and a couple of diagrams) of a group G in an arbitrary category \mathbb{C} , and pass to the category $\text{Group}(\mathbb{C})$ of all groups in that category, and then to the category $\mathbb{C}^{\text{Group}(\mathbb{C})}$ of all functors from that category of groups back down to \mathbb{C} , and so on, then these constructions have to take place somewhere! They require some collection principles!

No; the idea that one is “going up” in a hierarchy, and that this requires stronger and stronger collection principles and existence assumptions rests on the “foundationalist” conception that the “objects” involved are fixed and determinate. From a categorical perspective, one is rather “going down”, by specifying more of the ambient structure to be taken into account: here, say, a cartesian closed category \mathcal{E} in which the original category \mathbb{C} is small. Where does *that* category \mathcal{E} come from? We describe it, by the handful of axioms for cartesian closed categories, and then assume further that there is in it a category \mathbb{C} with whatever properties we were interested in—in particular, it has a group G . Is that *the same* category \mathbb{C} we started with? The question makes no sense. Neither G nor \mathbb{C} nor \mathcal{E} are specific things here, they are

schematic structures, as it were, specified or determined by configurations of objects and arrows and conditions on them, which can be assumed or found to hold in various different situations, which in turn may also be schematic.

Thus rather than saying, for example, “now suppose this particular solar system is an atom in some huge piece of matter in an enormous solar system”, one is instead saying “now suppose this particular configuration of bodies occurs, not as a solar system, but as an atom in some piece of matter in a solar system”. The former assumption indeed requires additional (outrageous) existence assumptions, while the latter requires none. A configuration is only assumed as a structure from the start, and so it can be specialized by assuming it to occur in more special situations. The schematic character of statements about structures is clearly essential to this approach; whatever we were saying about \mathbb{C} originally (e.g. that there is a group in it), can still be said about \mathbb{C} after we have put it into some ambient category \mathcal{E} , because we weren’t assuming anything particular about it in the first place.

How, then, do we make precise the notions of *schematic* statements about *structures* that have different *instances*? Simply by using the usual language and methods of category theory; they automatically treat mathematical objects as “structures”, and categorical statements about them are inherently “schematic”, in the required sense. This is what makes category theory a good language for structuralism. It is also what gives it an essentially different perspective from the foundational one.

The universal structures

Certain structures are canonical, in the sense that they are distinguished by universal properties: the natural numbers, the free algebras, the finite sets, and other such structures (and their duals). These universal structures reappear in many different categories; they are characterized, up to isomorphism, by universal properties similar to those characterizing other canonical structures like cartesian products and exponentials.

Such structures are in a sense more specific than the general or arbitrary ones like general categories, groups, spaces, etc. As a particular structure, the natural numbers, for instance, have a certain autonomy; statements about them stand alone, in a way that statements about arbitrary categories or groups do not. A statement about the natural numbers is not conditional or hypothetical, but rather a specific statement about a particular structure.

On the other hand, even such universal structures as the natural numbers have various different “instances” or “occurrences”, in the sense that they can be found in various different categories. The natural number objects in a particular category of sets and in an arbitrary elementary topos will both satisfy the Peano postulates, for instance, but may differ with respect to some other properties—even logically definable ones.⁸ Thus, although we can recognize different “instances” or “occurrences” of the natural numbers in different categories by their universal property, it is not clear whether we should call them “the same” numbers, or “the same” structure.

One interesting, if somewhat speculative, possible way to pursue this question is to restrict the language that we use in formulating the properties of these structures, i.e. to be more specific about what counts as a “property of the structure”. Consider the case of the finite sets, for instance. They can be characterized as the initial boolean topos, or as the free category with initial object 0 and binary coproducts $A + B$ on one object $*$, or in several other ways. These descriptions make sense in many different categories (specifically, in any category with finite limits). And a category of finite sets can be shown to exist, for instance, in any topos with a natural numbers object. The properties of the finite sets will thus depend to some extent also on the ambient category. In particular, there are certain core statements about them that are “absolute”, in the sense that they hold in all cases, and certain other statements that may hold or fail from instance to instance. The absolute statements include all those in the language of toposes, for instance.⁹ In that sense, the description of the finite sets as the free boolean topos fixes all of their structural properties, independently of the ambient category.

The study of such absolute and structural properties of universal structures has not yet been developed. In fact, it is not even known whether e.g. the real numbers or larger categories of sets than just the finite ones are universal, in the sense of being determined by universal properties. Only then could one sensibly ask which of their properties are not just typical, but absolute. This seems to me like the sort of question philosophers of mathematics might fruitfully pursue.

⁸This follows from the incompleteness of arithmetic and the topos completeness of higher-order logic. Note that it makes no sense to ask whether the natural numbers objects in *different* topoi are isomorphic.

⁹This requires a proof, omitted here.

References

- [1] Awodey, S.: 1996, “Structure in mathematics and logic: a categorical perspective,” *Philosophia Mathematica* (3), vol. 4, pp. 209–237.
- [2] Feferman, S.: forthcoming, “Typical ambiguity: trying to have your cake and eat it too,” to appear in the proceedings of the conference Russell 2001, Munich, June 2–5, 2001.
- [3] Hellman, G.: 2003, “Does category theory provide a framework for mathematical structuralism?,” *Philosophia Mathematica* (3), vol. 11, pp. XXX.
- [4] Mac Lane, S.: 1992, “The protean character of Mathematics,” *The space of mathematics*, Berlin: de Gruyter, pp. 3–13.
- [5] Mac Lane, S.: 1993, “Categorical foundations of the protean character of Mathematics,” *Philosophy of mathematics today*, Dordrecht: Kluwer, pp. 117–122,

Philosophy Department
Carnegie Mellon University
Pittsburgh, PA 15213
USA
awodey@cmu.edu

Abstract: An affirmative answer is given to the question quoted in the title.