Structuralism and the Independence of Mathematics

By Michael D. Resnik

CONSIDERED RELATIVE TO OUR SURFACE IRRITATIONS, WHICH EXHAUST OUR clues to an external world, the molecules and their extraordinary ilk are thus much on a par with the most ordinary physical objects. The positing of these extraordinary things is just a vivid analogue of the positing or acknowledging of ordinary things: vivid in that the physicist audibly posits them for recognized reasons, whereas the hypothesis of ordinary things is shrouded in prehistory....

To call a posit a posit is not to patronize it....Everything to which we concede existence is a posit from the standpoint of a description of the theory-building process, and simultaneously real from the standpoint of the theory that is being built. (W. V. Quine)[1]

Mathematical objects, if they exist at all, exist independently of our proofs, constructions and stipulations. For example, whether inaccessible cardinals exist or not, the very act of our proving or postulating that they do doesn't make it so. This independence thesis is a central claim of mathematical realism. It is also one that many anti-realists acknowledge too. For they agree that we cannot create mathematical truths or objects, though, to be sure, they deny that mathematical objects exist at all. I have defended a mathematical realism of sorts. I interpret the objects of mathematics as positions in patterns (or structures, if you will), and maintain that they exist independently of us, and our stipulations, proofs, and the like.

By taking mathematical objects to be positions in patterns I see all mathematical objects as being like geometrical points in having no identifying features save those arising through the relations they bear to other mathematical objects in the structures to which they belong. Mathematicians talk of numbers, functions, sets and spaces in order to

Michael D. Resnik is University Distinguished Professor of Philosophy at the University of North Carolina at Chapel Hill. He written work is in logic, philosophy of mathematics, and the theory of rational decision making. His most important writings include Frege and the Philosophy of Mathematics, Mathematics as a Science of Patterns, *and* Choices: An Introduction to Decision Theory.

depict structures. Thus they might describe the natural number sequence (0, 1, 2, etc.) as the smallest number structure that has that has exactly one number (position) immediately following each of its numbers (positions) as well as an initial number (position), call it '0', which is preceded by no other numbers (position) in the structure. One of the important features of patterns is that they may occur or be embedded in other patterns. Take for example, simple songs. The pattern of notes exhibited in their initial verses usually recurs in subsequent verses. Furthermore, if we transpose the song into different keys, then the pattern of musical intervals occurs again and again in each new key, with each transposition being a different pattern of notes. I see mathematicians as making observations similar to these, as well as abstracting patterns from practical experience, finding occurrences of patterns in each other and "combining" patterns to arrive at new ones.

Structuralist views of mathematical objects, of which mine is just one, have a reputable history among mathematicians that dates to at least the 1870s.[2] Dedekind expounded a version of structuralism, and we can find kindred themes in Hilbert too. But the recent spate of structuralist writings in the philosophy of mathematics has been in response to two influential papers by Paul Benacerraf, "What numbers could not be" (1965) and "Mathematical truth" (1973).[3] In the first paper Benacerraf reflected on the variety of ways mathematicians have found for defining the natural numbers as sets.[4] Noting that these definitions are equally good from a mathematical point of view, he concluded that there is no fact as to which sets the numbers are, and consequently, that numbers are not sets at all. This was contrary to the teachings of Frege and Russell and many subsequent analytic philosophers, but Benacerraf continued with a more radical thought. Claiming that number theory is just the theory of a certain structure and that numbers have no identifying features except structural ones, he inferred that numbers are not objects at all, or as he put it, "if the truth be known, there are no such things as numbers; which is not to say that there are not at least two prime numbers between 15 and 20".[5]

I found myself unwilling to follow Benacerraf in his last step. His argument that number theory is the science of a certain structure was convincing, but unless something exhibits that structure number theory is vacuous. Moreover, Benacerraf's observations applied to all of mathematics. For throughout mathematics we find alternative (and incompatible) definitions of important mathematical objects. Real numbers may be defined as sets or as infinite sequences or as the sums of infinite series, functions may be defined in terms of sets or sets in terms of functions, and so on. Mathematics may well be the science of structure, but lest it be vacuous, the ontological buck most stop somewhere—things exhibiting the various mathematical structures must exist.[6]

Benacerraf's "Mathematical truth" emphasized a different but much older problem. In the history of the philosophy of mathematics we find views that present a plausible account of mathematical truth by positing a mathematical ontology of abstract entities and views that present a plausible account of mathematical knowledge by emphasizing the role that symbol manipulation and proof play in mathematical practice. Yet nowhere can we find a plausible account of both mathematical truth and mathematical knowledge. If mathematics is about abstract entities that exist timelessly outside of space and time, then it's an utter mystery as to how we can access them and acquire mathematical knowledge. (This is the 'Access Problem'.) On the other hand, if we solve the Access Problem by taking the subject matter of mathematics to be symbols and proofs, we cannot account for the truth of mathematical sentences that purport to make claims about numbers, functions, sets and the like; for we know, thanks to logicians like Frege and Quine, that it is just a confusion to think that these are just symbols. So we are left with a dilemma: we can have a reasonable account of mathematical truth or a reasonable account of mathematical knowledge but not both.

But if we think of mathematical objects as like positions in patterns, then we may be able to solve both of Benacerraf's problems. For just as geometrical points have no identifying features—they all "look alike"—except the ones they have by virtue of their relationships to other geometrical objects, positions in patterns have no identifying features save those which they have in virtue of their relationships to other positions. This would explain why mathematicians don't care whether they define the numbers one way or another so long as the structure of the numbers is preserved, and it would explain why there is no fact of the matter as to whether the numbers are sets. That's just the way positions are. There is a lot more to my interpretation of mathematical objects as positions in patterns than I have presented here, but I want to leave it to discuss my approach to Benacerraf's other problem.[7]

Now one might think that we can access positions through accessing the structures or patterns containing them, and one might also think that something like pattern recognition would be a reliable means for so doing. Some of my earlier papers suggest such an approach, and I know through correspondence and conversation that the idea has found a number of friends. But recently I have not held such a view, and I am not sure that I ever have. Put loosely, I admit patterns that are not concretely instantiated. Now, perhaps, we come to know things about patterns by initially learning things about concretely instantiated ones. But even if this is true, I don't think it will be of much help in accounting for our knowledge of mathematical objects themselves. For mathematical objects, that is, the positions in patterns themselves are very abstract objects, and it is unclear how they could be presented to us by means of the more concrete things occupying them. At least it is unclear how they could be presented to us via an undoubtedly natural process. This is just a structuralist version of the Access Problem.

I don't see how you can avoid foundering on this problem if you face it directly. We have no causal access to mathematical objects or anything that could be taken to be their traces, since they have no traces. For the same reason we don't even have such access to structures though their instances or to types through their tokens. Of course, we often take instances of patterns and tokens of types to reflect features of the patterns and types themselves. For example, speaking of a letter qua type, we might say, "the letter 'A' looks like this," and then inscribe a token 'A'. But letter types are abstract entities, and as such they don't reflect light; so we "see" them only in an extended sense, and can never directly test hypotheses concerning their relation to their tokens.

In my book, Mathematics as Science of Patterns, I approached the Access Problem by applying a postulational epistemology to mathematics. My account had two parts. The first addressed the question of how the first mathematicians could have acquired mathematical beliefs without encountering mathematical objects. This question had puzzled many people influenced by the Access Problem. In answering it I hypothesized that developing, manipulating and studying notations for representing systems of concrete objects eventually led ancient mathematicians to posit mathematical objects qua abstract positions in structures. The second part of my account explained how these beliefs, though initially acquired in way that need not generate knowledge, could indeed count as knowledge, and why standard contemporary mathematics is a body of knowledge too.

The second part of my epistemology is a pragmatic version of confirmational holism—the idea, originating in Duhem and extended by Quine, that hypotheses are confirmed or refuted in bundles rather than individually. The version I favor distinguishes between global (or holistic) conceptions of evidence and pragmatically grounded local conceptions of evidence. The basic idea is that from a logical point of view data will typically bear directly only globally upon relatively large systems of hypotheses, yet we can be pragmatically justified in taking certain data to bear upon a specific hypothesis. Biologists, for example, will be pragmatically justified in appealing to a conception of evidence local to biology to conclude, say, that a certain study refutes a certain biological hypothesis. They need not concern themselves with the fact that from the logical point of view the study also bears upon broader biological, chemical and physical hypotheses and the statistical methods they used. In applying these ideas to mathematics, I take its numerous applications to provide global evidence for mathematics, but I countenance local evidence for mathematical theories too. Indeed, as I see it, a hierarchy of (local) evidence for mathematics parallels the evidential hierarchy of the other sciences. Just as bits of elementary chemistry can support sub-atomic physics, some of the results of arithmetic and geometry can be tested against computations and measurements, analysis can be supported via its arithmetic and geometric consequences, more abstract theories confirmed via their consequences for analysis, and so on. Furthermore, I doubt that the local conception of evidence frees mathematicians from worries about whether the objects they posit exist. The history of the controversies over the negative, imaginary, and infinite numbers, infinitesimals, impredicative sets, and choice functions show that they frequently do concern themselves with the status of newly introduced mathematical entities, and try to find considerations favoring their existence.

I coupled this epistemology with notions of truth and reference that are immanent and disquotational. This means that they apply only to our own language, and serve primarily to permit inferences such as the following:

1) Everything Tess said is true, and she said, "Jones was at home;" so Jones was at home.

2) The term 'the Big Apple' is used to refer to New York City; thus if the Big Apple is hectic, so is New York City.

Even this modest conception of truth and reference allows one to formulate theses committing one to an independent mathematical reality. (One such thesis is that classical mathematical analysis is true whether or not we have proven it to be so.) Moreover, it avoids worries about how our mathematical terms "hook onto" mathematical objects, and explains how initiating mathematical talk can enable us to refer to and describe objects to which we have no causal connection.

Further details of my account need not concern us now. But it is important for me to emphasize that two parts of my account are not tightly connected. It is true that after we have posited positions arranged in various patterns we can refer to them in order to interpret and make better sense of the experiences that led us to posit them. Moreover, I see these experiences as data that give some confirmation to the hypotheses postulating the positions. However, on my view, nothing in the course of positing, including having the experiences that motivated the positing, establishes the existence of the entities posited or the truth of the postulates concerning them. Exactly this feature of my epistemology has been the source of an important objection to it.

The problem is that in an important sense I turned my back on the Access Problem instead of solving it. I did show how we might have arrived at our mathematical beliefs through reasonable means and how they are part of a systematic whole that experience supports. But while this may show that our system of mathematical and scientific theories is internally coherent and squares with experience, it still does not show how mathematics connects to an independent reality.

Here is how Jody Azzouni has expressed his reservation:

> Some philosophers of mathematics marry an ontologically independent mathematical realm to a stipulationist epistemology. The result is unstable if only because such a union still craves explanation for why the stipulations in question correspond to properties of the ontologically independent items

they are stipulations about.[8]

Azzouni is confident that I cannot meet his demand for an explanation of how mathematics is tied to an independent reality, because the practice of science and mathematics "offers no epistemic role for mathematical objects, and so does not respond to the worry that there are no mathematical objects for its theorems to be true of."[9]

Here Azzouni has in mind his 'Epistemic Role Puzzle', that is, the puzzling fact that whether or not mathematical objects exist, they seem to play no role in the things mathematicians do to obtain mathematical knowledge. This is a cousin of the Access Problem; for if mathematical objects are abstract entities, then it's unclear how they could play any role in mathematical practice. Unlike the objects that usually concern science, we cannot interact with them or physically manipulate them.

Now there are two ways one might to respond to the Epistemic Role Puzzle. First, one might explain why mathematical objects, by their very nature, could not and should not have an epistemic role; and then go on to argue that this still does not prevent us from having knowledge that is about them. This is what I tried to do in my book by interpreting mathematical objects as positions in patterns. Since it is the essence of a position that it has no function except to mark a place relative to other places in the pattern containing it, there is no basis for supposing that it has any properties that would allow us to detect it or manipulate it or otherwise involve it in our usual epistemic processes. Mathematics describes structures by telling how objects in them, that is, positions, are related. This is the only reason that it needs objects, and it requires no more of them than that they be related to each other in various ways. Thus in so far as mathematics concerns itself with structure and only structure; it is virtually pointless for its objects to have physically detectable features. Furthermore, if mathematics acknowledged any physical objects as its own proper objects, then it would be obliged to study their physical properties and would sacrifice its focus on structure. Thus, given the goals of mathematics, it makes sense for it to ignore questions of the physical nature of its objects And given that it does, it is impossible for them to have any epistemic role in Azzouni's sense.

The second response would be to argue that in an indirect sense mathematical objects do indeed have an epistemic role. This is the sort of response that formalists who hold that mathematics is about formulas could make. Moreover, in a kind of convoluted way it is open to me too. For, on my view, some mathematical notations mirror the structures they represent. For example, a finite sequence of inscribed unary numerals instantiates an initial segment of an omega sequence; a paper and pencil Turing Machine computation instantiates its abstract counterpart as does a formal derivation or a triangle on a blackboard. So one might argue that here at least structures and their positions do have a role in obtaining mathematical knowledge. I can think of two objections to this response: a) the response appeals to the relationship between types and their tokens, and the former are clearly concrete, so it is they, and not the types which have an epistemic role; b) the response goes through only if we posit structural similarities between the types and tokens in question, and we have no independent way of confirming that these similarities exist.

Later I shall argue that we don't directly access physical objects either but rather only through connections that we posit linking them to sense experience. If this is correct, then neither objection (a) nor objection (b) is compelling.

Let us assume for now that that I can respond successfully to the Epistemic Role Puzzle. I still don't think this sets to rest the general worry about my view. Ultimately, we may have a conflict between what Azzouni calls "coherentist epistemic positions"[10] and more foundational approaches. On the coherentist approach, if our current overall theory of the world is empirically adequate and meets other epistemic virtues, such as simplicity, generality, fecundity and consistency, then we have good reason to believe in the objects that it posits—all of them with no distinction being made between physical and mathematical objects. But according to Azzouni, this is not a true view of science: scientists expect their posits to have an "epistemic role of their own." This may be seen by "noticing how the actual objects under study play an official role in the evidence that epistemic processes are reliable or dependable; in light of this role, scientists are willing to describe such processes as leading to knowledge."[11] For example, suppose that physicists posit a new subatomic particle in order to make a certain group theoretic model apply to their data. Even if their theory very satisfactorily explains their data, typically they will refrain from affirming the existence of the posited particle until they have experimentally detected it. They will require observational evidence that they take to be a reliable indicator of the particle in question. Moreover, in explaining why the evidence reliably indicates the presence of the particle they will ascribe a role to the particle itself in the interactions producing the evidence. It seems then that the holist account of science, at least as expounded by Quine, is inaccurate. And if this is so, then it is reasonable to doubt its application to mathematical knowledge.

Now I think that Azzouni is right that the account of science that he attributes to Quine is not accurate. It is not clear whether this really is Quine's account, since in some of his latter writings Quine retreats from the strong holist theses he advocated in his earlier papers. In any case, if we modify holism, as I have, by distinguishing between local and global conceptions of evidence, then positing in the empirical sciences poses no problems. Empirical scientists are operating with a local conception of evidence which requires them to detect their posits; mathematicians are not.[12]

"Yes," one might object, "but it is exactly because mathematicians are not obligated to detect their posits that mathematical objects are not independent of us." To assess this claim, let us distinguish ontological independence from epistemic independence.[13] An entity is ontologically independent of us if it is not something that we make up, create or construct, etc.; that is, if it could or would exist even if we did not. From physics itself we know that subatomic particles and other unobservable objects are ontologically independent of us, since physics tells us that they (and the universe they inhabit) existed before we did and would have existed even if we had not. However, mathematics proper, being silent about the nature of its objects, simply does not address the question of their ontological independence. Rather it is philosophers, such as myself, who argue for their ontological independence by arguing that only an ontology of abstract entities can verify the existential claims of mathematics. Those offering the objection opening this paragraph think that unless we can show these abstract entities are epistemically independent of us, we should not accept this philosophical argument for the ontological independence of mathematical objects.

Now a major problem with this objection is that it is very difficult to characterize epistemic independence in a reasonably precise way that doesn't beg the question at issue or classify physical objects as epistemically dependent upon us. To illustrate this, I shall examine the following proposal by Azzouni:

> A requirement of our taking an object O to be [epistemically] independent of us is that, given any property attributed to O, we take ourselves as required to explain how we confirm that attribution in a way that non-trivially satisfies (*). Trivial satisfaction (*), or the irrelevance of (*) altogether from knowledge-gathering practices about O indicates that O is [epistemically] dependent on us. [14]

The condition (*) to which Azzouni refers is the following:

> (*) The process by which I come to believe claims about x's is dependable with respect to x's if and only if given that the process has led me to believe S(x) is true, then (under a broad range of circumstances) S(x) must be true, and/or given that the process has led me to believe S(x) is false, then (under a broad range of circumstances) S(x) must be false.[15]

In other words, on this proposal, an object is epistemically independent of us only if: 1) given any property that we attribute to it, we should ordinarily be able to determine by dependable methods whether the property in question applies to that object, and 2) there is a "non-trivial" explanation of why our methods are dependable.[16]

In expounding (*) Azzouni writes that in the empirical sciences "processes which are taken to yield knowledge are seen as doing so precisely because they do (causally) connect us to objects in such a way that what the process gives as an answer covaries with the properties that the objects have."[17] Later he remarks, "Empirical scientific practice routinely worries about when measurements, observations, and instrumental interventions (with objects) can be trusted and when not; when artifacts of our epistemic means of access arise (and how we can recognize them)."[18] Here he is talking about the dependability of quite specific scientific procedures or instruments. Their analogs in mathematics are algorithms, rules of thumb, estimation methods and approximation procedures; mathematicians do worry about the dependability of these things. Of course, they address their worries by taking some body of mathematics for granted and use that to demonstrate that the method in question is sound or sound for a significant number of examples. Accepted mathematics serves both as the source of data by which the methods are assessed and the background theory used to account for their virtues and foibles.

In both the mathematical and empirical cases one probes or checks or justifies a method, instrument, or datum by reference to a supposedly independent standard. Without such a standard it would be pointless to wonder about the reliability of the items in question. Thus we can calibrate a spring scale by weighing objects of known weights, and we can explain how it registers in response to the forces the objects placed upon it generate. But in order to do this we must assume that we have an independent and accurate method for determining the weights of the known objects, and that our theory of the scale is correct. Even when we give an object an epistemic role, doing so is relative to taking some parts of some theory of objects of that type for granted. When we use a telescope to confirm the existence of a planet originally posited to explain perturbations in the orbit of another planet, we presuppose a theory that permits us to conclude that what we are seeing through the telescope is a planet with sufficient mass to do the work. Thus it seems that both mathematicians and empirical scientists are concerned with issues of dependability and use similar means for addressing them.

What happens when we don't have an appropriately independent theory of the objects in question? According to Azzouni, if we simply say, for example, that our theory of the objects states that our methods for investigating them are dependable, then they are not epistemically independent of us—at least not yet—and we are not justified in asserting their ontological independence.[19] This threatens to undercut the epistemic independence of mathematical objects. We can explain the dependability of, say, our algorithms for calculating sums and products of numbers written in decimal notation by appealing to the recursive equations for addition and multiplication and definitions relating decimal numerals to unary numerals. We might explain the dependability of the former by defining numbers in terms of sets, but obviously the process has to end with assumptions that we cannot independently verify. Affirming that these assumptions are simply stipulated to be true will play right into Azzouni's hand, since the only explanation we will have at this point will be the "trivial" one that the methods are dependable simply because they are (according to our theory of them).

Notice that Azzouni writes that it is a "requirement of our taking an object O to be [epistemically] independent of us is that, given any property attributed to O, we take ourselves as required to explain how we confirm that attribution in a way that non-trivially satisfies (*)." [20] The same requirement would hold for those who are realists about subatomic particles. But this seems to be too much to ask even when we consider relatively familiar objects like electrons, whose epistemic role certainly Azzouni acknowledges. The problem is that we sometimes use purely theoretical considerations to attribute properties to electrons that, as a matter of principle, we can't confirm experimentally. For example, electrons have the property of never being in a state in which they have an exact position and an exact momentum. My limited reading in the philosophy of quantum mechanics tells me that a number of theoretical considerations are needed to conclude that this is an objective feature of electrons and not just a limitation of our measuring devices. If so, then it would seem that in principle we cannot confirm this property of electrons by means of a process that satisfies (*). It may well be the case then that the only way we can confirm it, if at all, is by appealing to some well-confirmed scientific theory. Another example that comes to mind is the continuity of space-time, which seems experimentally indistinguishable from its density.[21]

Now if I am right about these examples, the process scientists have used here seems to be this: To confirm claims about physical objects, which cannot be tested directly by experiments, find a well-confirmed theory (in the usual sense) that implies the claim in question. Demanding that we explain why this process is dependable seems to be demanding too much: it is to demand that we explain why a well-confirmed empirical theory asserts the truth. Suppose that in the light of this, we conclude then that sometimes we are not obliged to explain how we can we confirm a property of certain physical objects "in a way that nontrivially satisfies (*)". Isn't this to conclude that (*) is irrelevant in these cases? Now we cannot conclude from this that Azzouni is forced to hold that these objects aren't ontologically or epistemically independent of us. For he only says that "the irrelevance of (*) altogether from our knowledge-gathering practices about O indicates that O is [epistemically] dependent on us ...."[22] But it looks like this amounts to his acknowledging that when as a matter of principle (*) is irrelevant, we don't have to try to explain why our practices satisfy it. At most we need only explain why they fail to satisfy it.

This does not seem so different from the case of mathematics. Sometimes we raise issues of reliability and address them by citing accepted mathematical theories. Sometimes we don't raise considerations of reliability and simply depend upon the theory itself eventually being "confirmed". Moreover, in these cases, we are typically in a position to explain why we cannot apply Azzouni's criterion (*). The difference between mathematics and physics seems more a matter of degree than of kind with independent confirmation of our physical posits being more readily found and more frequently sought.

The difficulties we have found with Azzouni's proposal generalize to the type of position it reflects. This is the type of position that presupposes that we can access reality independently of our conceptual system. The problem is that our only access to any independent reality is through our sensations. Anything else that we access through them is mediated by hypotheses connecting the two. Walking through the woods during the fall I often smell an odor familiar from my medicine cabinet and infer that there must be some witch hazel nearby. My inference is based upon hypotheses linking the smell and the shrub, which I have conjectured but have never independently confirmed. Of course, with enough effort and care, I could test my hypotheses, but only through taking similar hypotheses for granted. Thus one of the first things I would try is to locate a specimen and smell it, but to do that I would need to (assume that I) know what witch hazel looks like. Most everyday physical objects are capable of affecting each of our five senses, and this provides us multiple ways of independently accessing them. And even when something affects only one or two senses—like the sun—we can often access it from multiple locations and at different times. All this confirms our belief that some enduring object is responsible for the sensations we have on these occasions. But each confirmation is relative to taking for granted myriad hypotheses connecting the object we posit and our sensations. Yet even in mathematics we can find independent links to the various structures it studies. Thus, we use numbers to count sheep, measure the length of a field, register the place of competitors in a race, and determine the iterations of an operation. These different empirical routes to the natural numbers give rise to different mathematical models (for example, set theoretic versus. geometric models) of the natural number sequence; and they lend credence to the idea that we are dealing with an independent reality. Again the difference between mathematics and empirical science seems to be a matter of degree.

To quote Quine, "everything to which we concede existence is a posit from the standpoint of a description of the theory-building process." We should add that anything we succeed in accessing we do so only by positing links between them and things whose accessibility we take for granted. Once we realize this, the idea that we can come to know things about patterns through their instances or about types through their tokens becomes much more palatable. As I noted earlier, on the sort of view of mathematical objects I advocate, this does give some mathematical objects an epistemic role.

Clearly, the things we (saints aside) believe in the most are the ones most intimately connected to our senses. We find it harder to doubt that we are standing on firm ground than that the prime numbers go on without end. This may be behind the philosophical intuition that mathematical objects don't exist. Rather than concede to the intuition, I acknowledge that our evidence for mathematical objects is less compelling than it is for every day material bodies, but I deny that we don't have sufficient evidence for the former. I also deny that we have stronger evidence for any physical object to which we have forged some observational connection than we have for any mathematical object. We have "detected" quarks, but I find it a stretch to say that our justification for believing in them is stronger than our justification for believing in numbers.

Where does this leave us? Some philosophers worry that holding that we posit mathematical objects is incompatible with realism. To them mathematical posits smack more of fiction than of empirical science. Perhaps, they came to this view through overlooking my claim that positing mathematical objects does not guarantee their existence and is only an initial step towards obtaining knowledge of the objects posited. In any case, they are likely to argue their point by emphasizing that mathematicians don't even try to detect their posits whereas empirical scientists normally do. Thus empirical scientists meet their obligations towards an independent reality while mathematicians don't. To this I have responded that the role of mathematical objects does not require them to be detectable; the local conception of mathematical evidence does not admit a place for detecting them. The real question is whether we can get "independent" evidence for a set of axioms, and sometimes we can by modeling them in some previously accepted domain. This is something that mathematicians prize.

As Azzouni pointed out, we cannot explain the reliability of mathematical methods in terms of the mathematical objects themselves, whereas in empirical science we regularly account for the reliability of methods by assigning roles to the objects the methods concern. This is evidence of an independent domain. However, we should not overlook the effort mathematicians devote to establishing the soundness of their methods even if in so doing they don't give a role to individual mathematical objects. Moreover, through positing links between structures and their empirical instances, we can bring mathematical objects into the epistemic picture.

In concluding let me note that my defense of the combination of postulationalism and realism turned little upon structuralism or holism. Structuralism played a part in my response to the Epistemic Role Puzzle, but I think it would have been enough for me to say that mathematics concerns itself with only the structural features of its objects whether they are positions in structures or not. Holism occurred in my account of how we might confirm mathematical posits, but the important point that we can support them using the mathematician's (local) conception of evidence should be separable from my more global conception of evidence.[23] φ

Notes

[1] Quine, W.V., Word and Object (Cambridge, MA: MIT Press,1960), p.22.

[2] For a fuller exposition of alternative versions of structuralism as well as a brief account of its history see Shapiro, Stewart, Philosophy of Mathematics: Structure and Ontology (New York: Oxford University Press, 1997).

[3] Benacerraf's papers are reprinted in Philosophy of Mathematics, Second Edition, eds. Paul Benacerraf and Hilary Putnam (Cambridge: Cambridge University Press,1993), pp. 272-294 and pp. 403-420.

[4] From the mathematical point of view any infinite progression of sets will serve for defining the numbers. Thus Zermelo defined them as follows: 0 = Æ (the empty set) 1 = {Æ}, 2 = {{Æ}}, 3 = {{{Æ}}}, and so on. Von Neumann defined them alternatively as follows: 0 = Æ, 1 = {Æ}, 2 = {Æ, {Æ}}= {0, 1}, 3 = {Æ, {Æ}, {Æ, {Æ}}} = {0, 1, 2}, and so on. Another famous definition due to Frege and used by Russell defines 0 as {Æ}, 1 as the class of all unit classes, two as the class of all pairs, and so on. There are infinitely many variations on each of these themes.

[5] "What numbers could not be," Benacerraf and Putnam, p. 294

[6] I believe the thinking I have outlined in this paragraph was quite common 30 years ago. Today ways of resisting it are on the market, but surveying them would take us too far afield.

[7] For more on my structuralist view, see my Mathematics as a Science of Patterns (Oxford: Clarendon Press,1997)

[8] Azzouni, Jody, "Stipulation, logic and ontological independence," in Philosophia Mathematica, Series III, vol.8 (2000): pp. 225-243, p. 232. See also his Deflating Existential Consequence (Oxford: Oxford University Press, 2004), and his "Review of Michael D. Resnik's Mathematics as a Science of Patterns," in Journal of Symbolic Logic 64 (1999): pp.922-3.

[9] "Stipulation, logic, and ontological independence," p. 226.

[10] Deflating Existential Consequence, p.109.

[11] "Stipulation, logic and ontological independence," p.227.

[12] For further discussion, see my Mathematics as a Science of Patterns, chapter 7 and my "Quine and the Web of Belief," in The Oxford Handbook of Philosophy of Mathematics and Philosophy of Logic, ed. Stewart Shapiro (New York: Oxford University Press, forthcoming December 2004).

[13] Azzouni does not draw this distinction. For him, epistemic independence is a necessary condition for ontological independence.

[14] "Stipulation, logic and ontological independence," p.230. Azzouni uses the term "ontological independence" instead of "epistemic independence" in this passage. See the previous note.

[15] "Stipulation, logic and ontological independence," p. 227

[16] According to Azzouni, the trivial explanation is this: A process P is reliable with respect to x's because they have the property that P is reliable with respect to them. Deflating Existential Consequence, p.100.

[17] "Stipulation, logic and ontological independence," p. 228, his emphasis.

[18] "Stipulation, logic and ontological independence," p. 229

[19] This is because we have given the trivial explanation.

[20] "Stipulation, logic and ontological independence," p. 230, my emphasis.

[21] Another exception to Azzouni's requirement may be properties of objects that are true of them "by definition" such as an electron's property of having one unit of negative charge.

[22] "Stipulation, logic and ontological independence," p. 230, my emphasis.

[23] I am grateful to Kenneth Walden for comments and encouragement. This paper grew out of a paper of the same title that I delivered in December 1999 at a symposium with Jerrold Katz and Stewart Shapiro on the philosophy of mathematics at the annual meeting of the Eastern Division of the American Philosophical Association.