Stewart, I.
1987 *The Problems of Mathematics*. Oxford: Oxford University Press.

Tymoczko, T., ed.
1986 *New Directions in the Philosophy of Mathematics*. Stuttgart, Boston: Birkhauser.

Van Bendegem, J. P.
1985 "Dialogue Logic and Complexity." Gent, Belgium Manuscript.
1987 "Fermat's Last Theorem seen as an Exercise in *Evolutionary Epistemology*." Pp. 337–63 in W. Callebaut and R. Pinxten (eds.), *Evolutionary Epistemology: A Multiparadigm Program*. Dordrecht: D. Reidel.
1988 "Non-Formal Properties of Real Mathematical Proofs." Pp. 249–54 in A. Fine and J. Leplin (eds.), *Philosophy of Science Association 1988*, vol. 1. East Lansing: PSA.

Versluys, J.
1914 *96 bewijzen voor het theorema van Pythagoras* (96 Proofs of Pythagoras' Theorem). Amsterdam: A. Versluys.

Welti, E.
1987a *Die Philosophie des strikten Finitismus: Entwicklungstheoretische und mathematische Untersuchungen über Unendlichkeitsbegriffe in Ideengeschichte und heutiger Mathematik*. Bern: Peter Lang.
1987b "The Philosophy of Strict Finitism." *Theoria* 2, 5–6;575–82.

Wright, C.
1980 *Wittgenstein on the Foundations of Mathematics*. London: Duckworth.

MICHAEL D. RESNIK

3

A Naturalized Epistemology for a Platonist Mathematical Ontology

Numbers, sets, functions, and other paradigmatic mathematical objects are, according to the Platonist view, outside spacetime and incapable of interacting with ordinary bodies within it. Taking them thus provides a nicely satisfying metaphysics of mathematics, but it appears to create an immense epistemological gulf between us and the mathematical realm. It is therefore hard to see how we can encompass mathematical objects within our most compelling model of the acquisition of knowledge, a perceptual model, where physical interactions play a central role. For centuries this apparent epistemological contrast between mathematical and physical entities has motivated empiricist critiques of mathematical Platonism.

Paul Benacerraf threw out the empiricist challenge for a generation of philosophers of mathematics in 1973 (Benacerraf, 1973) when he required a satisfactory account of mathematical knowledge to be a species of a general causal epistemology. Since then we have found that causal theories of knowledge stumble over even ordinary material bodies (Maddy, 1982). Yet, Benacerraf's demand was based upon good empiricist intuition: Any satisfactory epistemology should explain our knowledge of mathematical objects without endowing them or ourselves with occult properties or faculties. In today's epistemological circles, this demand often translates as an insistence that the epistemology of mathematics be naturalized.

In this chapter I will take some first steps toward meeting the challenge to naturalize the epistemology of mathematics. I do not do this merely to try to cover myself with a modish mantle, though sticking with fashion at least guarantees one partners in philosophical dialogue. Rather, I do this because meeting the empiricist challenge, in whatever form it currently assumes, is the dialectically strongest position for me as a Platonist to take.

During most of this chapter I will be setting the stage for a postulational account of the genesis of mathematical knowledge. My hypothesis is that our mathematical ancestors brought mathematical objects within our cognizance by positing them. This suggestion raises many questions concerning how positing can generate knowledge about preexisting entities—especially how it can do this when the entities are mathematical ones. At the end of the chapter, I will hint at answers to such questions. The bulk of the chapter will be concerned, however, with addressing the question of how a postulational account of mathematical knowledge could count as a piece of naturalized epistemology—even if it is successful in its own right.

The processes referred to in a naturalized account of knowledge must be natural processes. I will argue later that positing is such a process. Yet, we can posit supernatural objects, for example, spirits, as well as natural ones. Thus positing, no matter how natural a process, can never lead to knowledge of mathematical objects, if they are not natural objects themselves. Now, Quine defines naturalism so that it counts mathematical objects as natural objects, but David Armstrong defines naturalism so that it excludes mathematical objects from the natural universe. To avoid begging the question against Armstrong, I will begin by arguing for a place for mathematical objects within the naturalist's ontology. Then, I will try to specify general parameters for a naturalized epistemology for mathematics, and sketch my postulational account of the origins of mathematical knowledge. Finally, I will prescind some of the problems positing mathematical objects purports to pose.

I will assume the truth of Platonism throughout the chapter, and I shall not argue for it directly. (Of course, the chapter argues for it indirectly, since it tries to disarm one of the major objections to Platonism.) Furthermore, the reasons, described below, that might have led our ancestors to posit mathematical objects probably remain good reasons—although certainly not the only reasons—for us to accept mathematical objects today.

MAKING ROOM FOR MATHEMATICAL OBJECTS

Can a naturalist countenance mathematical objects? Unless the answer to this question is affirmative, our quest for a naturalized epistemology for mathematical objects is bound to fail. So we must address this question before proceeding further.

Yet before we can do that we must answer still another question: What is naturalism? Despite the currency the naturalist philosophy enjoys, I was surprised to find that there is little consensus concerning its content. Quine writes as follows:

> Now how is such robust realism to be reconciled with what we have just been through? The answer is naturalism: the recognition that it is within science itself, and not in some prior philosophy, that reality is properly to be identified and described. (Quine, 1981)

Here Quine throws ontological and epistemological questions into the court of science, leaving open the possibility that Platonic objects are real. And, as is well known, Quine does think that science ontically commits us to such mathematical objects.

David Armstrong's characterization of naturalism is both more definitive and importantly different from Quine's:

> Naturalism I define as the doctrine that reality consists of nothing but a single all-embracing spatio-temporal system. (Armstrong, 1981:149)

Armstrong's position seems to rule out Platonic mathematical objects at the start—a consequence Armstrong is quick to acknowledge.

For Philip Kitcher, writing in the philosophy of mathematics, naturalism goes without a definition, although he declares himself a naturalist and allies himself with empiricism against a priorism (Kitcher, 1988). Evidently, we must decide on a characterization of naturalism before we proceed much further.

Plainly, our target should be Armstrong, since he is completely candid about having no place for Platonic mathematical objects within the "single all-embracing spatio-temporal system" that defines his naturalist universe. Since mathematical entities have no effect in that system, "there is no compelling reason to postulate them" (Armstrong, 1981:154).

It is striking that Armstrong excludes mathematical objects from the naturalist's ontology, while Quine admits them. Armstrong is aware, of course, that Quine and other philosophers argue that we must postulate mathematical objects in order to do science and that our justification for doing so is no different from that used to justify positing electrons and other theoretical entities (Armstrong, 1981:155). But that cuts no ice with him. Consider his forceful response:

> There is this vital difference. [Classes, and so forth] provide objects the existence of which, perhaps, can serve as truth-conditions for the propositions of mathematics. But this semantic function is the only function they perform. They do not bring about anything physical in the way that genes and electrons do. In what way, then, can they help to explain the behavior of physical things? Physics requires mathematics. That is not in dispute. But must it not be possible to give an explanation of the truth-conditions of mathematical statements in terms of the physical phenomena that they apply to? (Armstrong, 1981:155)

The issue here, then, is not that these objects have been introduced as posits but rather that they have no causal powers and play no role in explaining the behavior of physical things. Of course, if it is essential to naturalism that its objects have causal effects on things in the physical world, then Armstrong wins his case hands down. But he also suggests that naturalists can recognize objects that "can help to explain the behavior of physical things." This may be just the crack in his argument Platonists need.

Armstrong concedes that mathematics plays an essential role in physics but denies that its objects play an explanatory role. What function, then, do these objects have? There is the semantic function Armstrong already mentioned: we cannot use mathematics in calculations and deductions unless its terms refer to and its sentences have truth-values.[1] But mathematics is more than a device for calculating and reasoning about physical phenomena; it also helps us describe them. Mathematical ideas infect virtually all of physics. Even at the elementary level, concepts such as instantaneous velocity (change in displacement at an instant, ds/dt) and momentum (mass X velocity) defy reformulation in nonmathematical terms.[2] At the more advanced levels of science we find phase spaces, vector and tensor fields in physics, growth functions and probability distributions in biology, and demand curves and utility spaces in economics. Without such mathematical entities these sciences could not even begin to describe the phenomena they recognize today. Now one cannot explain physical phenomena that one cannot describe. So certainly mathematical objects "can help to explain the behavior of physical things," at least in the sense of being an indispensable tool for the task.

What is more, mathematical facts and properties of mathematical

objects play essential roles in physical explanations themselves. Consider this explanation of why a ball thrown straight up in the air reaches a specific point (rather than another) before it comes back down.

At any instant the velocity of the ball (i.e., the speed and direction with which it travels) is the resultant (vector sum) of its upward and downward velocities. When thrown it has an initial upward velocity of, say, v^* and a zero downward velocity. However, the force of gravity subjects the ball to a positive downward acceleration, a^*, which in turn increases its downward velocity over time until the latter equals and then exceeds its upward velocity. When the two velocities are equal the ball stops its upward course. The downward velocity is identical to a^*t, so the ball stops its upward flight when and only when

$$v^* = a^*t.$$

Thus it will stop when and only when $t = v^*/a^*$. But this value of t determines the exact upward displacement of the ball.

Plainly, this explanation appeals to several mathematical properties of the ball's velocity, for instance, that it is the vector sum of its upward and downward velocities and that when these are identical the velocity equals zero. Moreover, this velocity itself, being a function, is a mathematical object. So the explanation uses mathematical objects and their properties to explain the behavior of a physical thing.

Such considerations do much to undermine Armstrong's argument; but I can easily imagine his defenders protesting that they do not count, because I still have not found any causal role for mathematical objects. This brings us back to Armstrong's primary reason for excluding mathematical objects—their lack of causal powers.

Rather than dealing with this head on, I shall examine a more general assumption implicit not only in Armstrong's thinking but also in many writings in contemporary philosophy of mathematics. This is the assumption that a clear and sharp, causally or spatiotemporally grounded, ontic division obtains between mathematical and physical objects.

The assumption probably comes from thinking about the relative scope of logic, mathematics, and physics: Logic properly includes mathematics; mathematics properly includes physics. Hence it is tempting to see sharp lines between logic and its ontology (or lack thereof), between mathematics and its ontology, and physics and its ontology. It is tempting to think that we can excise the mathematics from physics in order to achieve a fully naturalized ontology.

In finding no mathematical objects within the "all-embracing spatio-temporal system," Armstrong may be presupposing a space-time criterion for differentiating between (abstract) mathematical and (concrete) physical objects: the latter but not the former are within spacetime. But what is it to be in spacetime? To be located in it? To be part of it? To be either? Are spacetime points in spacetime? Is all of spacetime in itself? These are not idle questions. The ontic status of the universal gravitational and electromagnetic fields, prima facie physical entities, as well as that of spacetime points, prima facie mathematical entities, turns on how we answer them. Furthermore, each answer comes with its own set of unresolved controversies (cf. Hale, 1988; Resnik, 1985b).

Moreover, even quantum particles, such as electrons, widely regarded as paradigm physical objects, pose difficulties for a locationally grounded division between the mathematical and the physical. Where are these particles when they are not interacting with each other? On one interpretation of quantum theory, under some circumstances, these particles are not even located within a finite region of spacetime. Then, are they everywhere or nowhere? The answers are controversial, and so is the interpretation of quantum theory, but until the issue is settled we can hardly be satisfied with classifying entities by means of a spacetime criterion.

Quantum particles seem more like mathematical objects than like everyday, commonsense bodies. To see why, read this excerpt from a recent text on particle physics:

> In the most sophisticated form of quantum theory, all entities are described by fields. Just as the photon is most obviously a manifestation of the electromagnetic field, so too is an electron taken to be a manifestation of an electron field and a proton of a proton field. Once we have learned to accept the idea of an electron wavefront extending throughout space . . . it is not too great a leap to the idea of an electron field extending throughout space. Any one individual electron wavefront may be thought of as a particular frequency excitation of the field and may be localized to a greater or lesser extent dependent on its interactions. (Dodd, 1984:27)

But what is a field? The simplest precise description is that it is a function defined on every point in space whose value at that point gives the intensity of the field there. Add to that the reflection that in quantum mechanics talk of intensity at a point (or in a region) is really talk of the probability of an interaction taking place there, and you see how mathematical is the quantum field theoretic conception of particles.

Another commonly accepted way of distinguishing between the physical and the mathematical is to claim that mathematical objects cannot change properties and participate in events. But this way will not stand up to a first set of objections either. Numbers do change some of their properties. Numbering the wives of Henry VIII was only a fleeting property of the number one. Perhaps such properties should not count, on the grounds that they are just accidental properties of numbers. But physical objects can change only their accidental properties too. Perhaps numbering the wives of Henry VIII should not count because it is not a real property. But what makes it unreal? Surely not because it is a relational property, since electrons, say, have properties only by virtue of their relations to other particles. I do not take these quick jabs to be fatal blows to the idea that numbers cannot change, while physical objects can, but they do show that some careful work must be done before we can use the idea of change to distinguish mathematical from physical objects.

Similar difficulties surround the event side of the proposal. The functions used to characterize physical events change their values during the course of an event. We have already seen that in the example of the thrown ball. It is question begging to simply object at this point that the velocity function does not participate in such an event because it cannot. Better to say that the event can be fully and precisely described without referring to the functions. But can such a redescription be carried out? I suppose that in the ball case it might, although the difficulties attending Hartry Field's nominalization should give us pause. However, we have no reason to think that it can be done with events involving subatomic particles, whose basic features, such as charge, spin, and energy level, correspond to no commonsense ideas. This should also remind us that fields participate in events too; they collapse and interact. But fields seem to be hybrid entities hovering between the paradigmatically mathematical and the paradigmatically physical.

The considerations also suggest that it is unwise to exclude objects from the naturalist's universe on the grounds that they have no causal powers or that they play no causal role in explaining the behavior of physical things. For we have seen that it can be unclear whether a given explanatory object (for example, a field) is physical or mathematical or even whether something counts as physical behavior (for example, the collapse of a field).

The tendency for physicists to seek structural explanations of the fundamental features of physical reality also undermines the idea that a fundamental ontic division obtains between the physical and mathematical. The movement began with Einstein's identification of the gravitational field with spacetime itself, which in turn identified masses and their gravitational effects with variations in the geometric structure of spacetime. More recently, physicists have proposed that all of physical reality is an eleven-dimensional space, whose geometrical properties give rise to all of the known physical forces (Freedman and van Nieuwenhuizen, 1985). The trend has been, then, to pass from conceiving of subatomic particles as tiny bodies to conceiving of them as systems of interacting fields spread over spacetime and thence to local variations in the structure of a generalized spacetime. From the point of view of today's science, physical reality is most accurately described as an unchanging structure, whose local variations may be described in less sophisticated terms as bodies and causes, changes and happenings. How then can naturalists recognize just bodies and causes, changes and happenings as real?

Perhaps they will declare that the space—however complicated—embraced by physics is real and then add that other "purely mathematical" spaces are unreal. But what can this mean? It cannot mean that physical reality instantiates the structure in question, for, on the line we have been following, physical reality is nothing but that very structure. A much more appealing answer is to introduce the idea of an observable as a type of local geometric variation and then argue that physical space is real, because in it alone are all and only possible observations also observable events. On this reading, the difference between physical space and other spaces is not that the latter do not contain observable events—for some will, since they are just local structural features—but rather that other spaces differ from physical space by failing to contain all and only the "events" humans could observe. This is an epistemic difference, an important one to be sure, but not enough of one to distinguish physical space ontically from other purely mathematical spaces. None of the difficulties we have encountered with distinguishing physical objects from mathematical objects need arise, if we adopt Quine's characterization of naturalism and let science tell us what exists. Nor need we worry about the existence of mathematical objects, since in asserting that particles have velocities, that reactions reach equilibrium points, and so forth, science commits itself to them again and again.[3] So I will answer in the affirmative the question with which this section began—naturalists can countenance mathematical objects. We are now free to seek a naturalist epistemology for them.

Before we do so, let me deal with the impression one might form that I have been inconsistent in denying the ontic distinction between mathematical objects and physical objects. After all, I began this chapter by emphasizing the abstractness of mathematical objects and the apparent epistemic gap between them and ordinary physical bodies, and lately I have been arguing that the distinction between the abstract and concrete blurs in theoretical science. But I have not denied that there appears to be a striking epistemic gap between ordinary bodies and mathematical objects. Platonists must struggle with the epistemology of mathematical objects, so long as the relative abstractness of mathematical objects seems to prevent the epistemology of ordinary bodies from applying to them. I should also add that Platonism does not need a sharp cleavage between the abstract and concrete for its metaphysics of mathematics to work. Platonism succeeds because, unlike nominalism, materialism, and constructivism, it can supply the vast infinities of objects that mathematics requires—more objects than any mind or minds could construct, more objects than the physical universe contains.[4] Whether these objects be fundamentally different from material or mental ones is not crucial. (For convenience I will continue to refer to mathematical objects as abstract and to ordinary physical bodies as concrete.)

WHAT IS NATURALIZED EPISTEMOLOGY?

As in the case of naturalism qua metaphysical doctrine, there is a perplexing variety of opinions concerning the definition of naturalized epistemology. Hilary Kornblith introduces his anthology on the subject by stating that, in contrast to traditional epistemology, naturalized epistemology holds that the answer to the question "How ought we arrive at our beliefs?" is not independent of the answer to the question "How do we arrive at our beliefs?" and recognizes that psychological investigations can be directly relevant to epistemological ones (Kornblith, 1985). Kornblith's characterization includes normative work, such as Alvin Goldman's, within the scope of naturalized epistemology.

Quine, the source of the term *naturalized epistemology*, formulates his idea of naturalized epistemology in this passage: "Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science." (Quine, 1985:24). The difference between knowledge and mere true belief is usually taken to be normative: knowledge is true belief that passes epistemic muster. To the extent that psychology is not concerned with norms, epistem[illegible] r otherwise, Quine's suggestion that epistemology become a branch of psychology appears to leave behind an important part of the theory of knowledge.[5]

Now I think that Quine's approach is not as far from Kornblith's as the foregoing passages indicate. By letting epistemologists use methods from social sciences other than psychology, we can keep their enterprise within the spirit of Quine's view while also permitting them to describe our epistemic norms and account for the evolution of these norms.[6] Epistemologists can also go beyond natural history and evolutionary theory, if we allow them to systematize our epistemic norms using the method of reflective equilibrium applied in logic, linguistics, and descriptive ethical theory (Resnik, 1985a). As the case of logic shows, organizing our norms thus might yield insights into them more valuable than the system itself.

However, describing, systematizing, and explaining our epistemic practices is one thing; evaluating them is something else—apparently beyond the scope of science, social or physical or formal. Yet even a fair amount of evaluation and criticism could be brought within the purview of epistemology qua science, if such evaluation and criticism were approached from the applied scientific viewpoint of an efficiency engineer. Given a description of our epistemic values and a measure of desirable performance (efficiency), epistemic engineers could determine how close our actual epistemic practices come to the official standards, according to the official yardstick. They could even reform our practices by suggesting methods for bringing our performances closer to the received standards. That is as far as we can go towards normative epistemology without allowing epistemologists to make their own value judgments. That may be far enough. And we may be at the limits of naturalism.

On the other hand, naturalistic epistemologists are members of the scientific community and, as such, free to promote new epistemic values from within that community. In doing so they act no longer as naturalistic epistemologists per se, just as political scientists put their academic roles aside when stepping inside the voting booth. Such actions are consistent with the naturalist's credo, so long as our epistemologists forsake supernatural normative insights.

If the preceding thoughts correctly represent the spirit of Quine's approach, there is no real dispute between him and Kornblith over the limits of naturalized epistemology. Even if it is real, we need not settle it now, since the bulk of the account I will present here is genetic rather than normative.

That is not to say that no work awaits normative epistemologists in the philosophy of mathematics. For all the successes of mathematical logic, we still do not have a good understanding of why we insist on proof in mathematics, why we prize alternative proofs, or of how axioms are justified—to name just a few of the questions that come to mind.[7]

Our concern here, the genetic side of epistemology, is especially pressing for the Platonist. Winning a place for Platonic mathematical objects within the naturalist's universe does not even begin to address the question of how to give a naturalized account of the genesis of our knowledge and beliefs about them. To deal with this problem, we must first develop a firmer conception of what a naturalized account of cognition is.

Cognition is a process, hence a naturalized portrait of it will present it as a natural process. Following Armstrong's definition of naturalism, one could define a natural process as one that takes place wholly within spacetime. This would exclude any theory that involved mathematical objects themselves in the process of cognizing them. As it is, I will not offer such a theory, so I could be at home with an Armstrong-style characterization of natural processes. Yet, on the grounds of consistency and liberality, I suggest the we follow Quine and let a natural process be one that science "has identified and described."

One evident problem with this definition is that the limits of science are vague and unclear. Are psychoanalysis or intensional semantics part of science? If not, then naturalistic epistemologists may not explain the genesis of knowledge in terms of processes hypothesized by those theories. Perhaps they can avoid the problem of demarcating science by not straying into the fuzzy areas between hard-core science and the more controversial disciplines vying for scientific status.

Due to the immature state of the cognitive and social sciences, that may be easier said than done. Many seemingly natural human processes—communicating, learning from experience or from other people, developing preferences or creating theories and works of art—are so complicated that science has only the roughest understanding of them. Undoubtedly hard-core science will eventually bring them under its umbrella in some, perhaps unforeseeable, form. Yet, naturalized epistemologists may need to appeal to these processes now.

Well, let them. But let us restrict naturalistic epistemologists to processes appropriate to the study at hand. Obviously, one should not appeal to communication in giving a naturalized account of communication, but one should be able to appeal to our ability to recognize sound patterns. In a similar way, I think that it will be proper for me to appeal to our ability to acquire knowledge and beliefs about everyday bodies in accounting for the genesis of mathematical knowledge. Of course, someone might point out that we still do not have a satisfactory naturalized account of our knowledge of ordinary bodies. I would be among the first to agree. My idea is to let epistemologists build speculative theories on somewhat less speculative foundations. If we can show how our ability to know and refer to ordinary bodies can by natural means generate an ability to know and refer to mathematical objects, then we will have made the prospects for a naturalized epistemology for the latter all the greater. And we will have shown that if we cannot have natural knowledge of mathematical objects, then it is unlikely that we can have it of ordinary bodies.

I am proposing to give naturalistic epistemologists liberty to hypothesize processes that we might not be able to manipulate experimentally, due to our current lack in knowledge and technology. Furthermore, they should be free to appeal to (hypothetical) events in the remote history of our species. Let us recognize, however, that if they do make such an appeal, then we may never be able to put their theories to a direct experimental test. Such accounts are already common in evolutionary biology and other historical sciences.[8] Evidence for and against them can be sought and is often found. So while requiring naturalized epistemologists to suggest experimental means for testing their theories is too stringent, it is reasonable to ask that they suggest some connection between their speculations and observable evidence.

Testing theories about the cognition of mathematical objects also seems to raise special problems, for we cannot manipulate mathematical objects experimentally. Thus we cannot study human cognition of them as we can human perception of ordinary bodies, where we can alter states of subjects, bodies, and even media in order to see how each affects reports by the subjects of what they perceive. But this difficulty arises only for those who think that some mathematical knowledge is acquired by something akin to perceiving mathematical objects. I disavow this sort of approach.

In these last paragraphs I have urged what I take to be a moderate approach to naturalized epistemology. My approach falls on a spectrum that ranges from the conservative to the speculative. The most conservative approaches recognize only processes that are clearly and uncontroversially described by natural science and that are also subject to experimental control. That would mean describing cognition in terms drawn from physics, chemistry, anatomy, and parts of biology. I know of no epistemologist so rigorous. At the speculative end of the spectrum we find people, for example, Penelope Maddy, willing to posit new cognitive processes or faculties, such as an ability

to perceive certain sets of concrete objects via their members. Clearly this process is not subject to experimental manipulation, since we cannot place an informational screen between the set and its members to allow a subject to see one without seeing the other.

MATHEMATICAL KNOWLEDGE AS KNOWLEDGE ABOUT PATTERNS

In my version of Platonism (Resnik, 1981), mathematical objects are positions in patterns, and mathematical knowledge is knowledge about patterns. On the one hand, this knowledge incompasses much more than the ability to recognize and distinguish simple patterns. Pigeons can do as much. On the other hand, it does not involve some sort of direct perception of patterns qua abstract entities, not even of those patterns that we "see" in perceptible arrangements. Knowing a pattern, in my view, is like knowing a theory. To know a theory is to know what entities and processes it posits and the behavior its laws attribute to both. Similarly, knowing a pattern is a matter of knowing what positions it contains and how they are related to each other.

In discussing knowledge about patterns it is important to distinguish the question of how a research mathematician learns about patterns from the question of how a learner in our society does, and this in turn from the question of how ancient humans learned about them. The mathematician has a gigantic collection of techniques the latter two lack.

I will focus on the question of how ancient peoples might have come to know patterns. One reason for doing this is that I am trying to answer skepticism concerning our ability to acquire knowledge about objects as abstract as my patterns. Focusing on the contemporary mathematician would be to ask instead how we manage to learn about new abstract entities once we already have an abundant fund of knowledge about some abstract entities. So I want to consider people who have no mathematical knowledge and suggest a natural process through which they could acquire it.

I thought about concentrating on the acquisition of mathematics by children. But our children have help from those already in the know. Furthermore, although it is possible that they recapitulate the process through which the human race learned mathematics, the rapidity with which they do so hampers studying it. I will also be concerned with the global question of how humans came to recognize and countenance mathematical objects as a kind of thing rather than with more local questions about how they came to recognize particular mathematical objects such as zero or the square root of two.

Since nobody knows how we developed mathematics, my story is perforce purely hypothetical. Despite this, it is easy to think of the kind of evidence that could bear upon it. Unfortunately, obtaining the evidence itself is much harder. Anthropological studies of primitive peoples could support or correct the initial elements of my narrative, but then a huge evidential gap opens due to the lack of cultures intermediate between us and still existing primitive peoples. History is of little help, too. Elsewhere I argued that the transition from premathematical studies (without ontic commitment to mathematical objects) to full-fledged mathematics (with ontic commitment to them) occurred when Babylonian and Eygptian mathematics developed into Greek mathematics (Resnik, 1982). We know something about the initial stage of this transition from fragments found by archeologists, and we know much about the final stage through Euclid and later commentators. To my knowledge, however, we have no useful evidence about the critical intermediate periods.

With the stage for my account now set, let us try to imagine ourselves in the situation of a primitive people who have no mathematics. Our knowledge of patterns will begin, like our knowledge of everything else, with experience. Experience will also teach us that certain shapes and arrangements work better in certain situations than others. Things having various shapes or arranged in certain patterns will become important to us.

Because of their practical importance, we will find ourselves driven to invent a vocabulary to name some of these patterns. The time will come, however, when we need to instruct a foreigner or novice, and nothing ready to hand is of the right shape. Then we might use a drawing. If so, we will have taken an important step: because we will no longer be restricted to indicating, recognizing, or labelling present things of the same pattern, we will be able to represent how things are arranged or shaped without having those things present.

We need some technical terminology at this point lest we confuse patterns in the concrete sense—in the sense in which we have a drawer full of dress patterns at home—with patterns qua abstract entities, for example, a dress pattern no one has described or drawn. A paper-and-ink dress pattern is an instance of an abstract pattern, for it is a token, a concrete inscription, of a symbol type. However, it is also a concrete representation of a type of dress without being an instance (token) of that type—without being a dress. Right now my concern is with the use of concrete inscriptions to represent other concrete things.

Reserving the term *pattern* for abstract entities, I will use the term *template* to refer to our usual concrete devices for representing how things are shaped, structured, or designed. Concrete drawings, models, blueprints, and musical scores are my paradigm everyday templates. Under the appropriate conventions, templates represent other concrete things, such as buildings, artifacts, or performances, which fit them in the appropriate ways. Templates are thus templates for things of the appropriate kind: blueprints are designs for buildings rather than for sculptures or performances of ballets.

We will go quite some distance in our practical talk about how things are shaped, arranged, or designed without appealing to abstract entities. Concrete templates will do the job perfectly. We will also learn how to design things before we start to manufacture them and how to use modified designs to learn things about modifications in the things themselves. Without introducing abstract entities, we will now be in a position to talk about possibilities, about how things might be arranged or designed or shaped.

Templates have two dimensions. Syntactically they are configurations constructed according to certain conventions. Semantically they represent other concrete things by means of implicit and explicit rules of representation. So far we have only considered templates that successfully fill their representational role. But we can also use the medium for constructing templates of a given kind to construct configurations without any representational role, such as random doodlings on blueprint paper. Surely, we will do that too.

We have thus advanced from the barest recognition of the practical importance of certain shapes and arrangements to a representational system for designing and thence to playful and creative attempts to explore possibilities. Before we move on, let us remember that we use language to construct templates too; for we can often describe an arrangement, shape, or design in words more accurately than we can in a drawing. Some linguistic templates will be sets of directions—instructions on how to build a serviceable lean-to, for instance. Others will describe rather than instruct—a biologist's description of bee dances is an example. Later in mathematics, linguistic templates will be our chief and most reliable methods for representing patterns.

I have not yet touched on the crucial question of how our experience with templates could lead us to knowledge of abstract patterns. The discussion of templates is not in vain, though, because it indicates how we might have begun our initial explorations of patterns and our initial probing of the possible. It also tells us something about the local epistemology of patterns. For although templates now only represent concrete things, they will come to represent the abstract patterns concrete things fit. Looking at the example of a dress, we see that ultimately there will be four entities involved, two concrete ones—the dress and its template, and two abstract ones—the dress pattern and the symbol type of the template. These are related according to figure 3–1.

Thus once we take the step towards countenancing patterns qua abstract entities, it is likely that we will see patterns as associated with

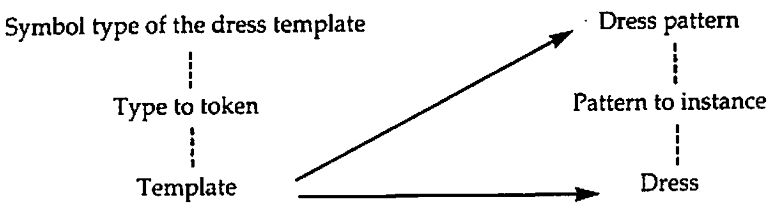

Figure 3–1.
Paradigm for Pattern Representation.

templates according to figure 3–1. This will indicate to us that we can construct and study templates to gain information about patterns. Finally, having bonded patterns and templates, it will become plain that to show that a specific pattern for, say, houses exists, it suffices to exhibit a blueprint for houses of that pattern.

Another important point to notice is that we could not have gotten this far with templates without developing complex syntactic systems (such as place notations for the numbers and geometric diagrams), as well as conventions and rules through which these systems represent concrete entities and criteria for determining whether a given configuration counts as a coherent representation. All of these contribute to easing the step towards full-blown mathematical theories with genuine commitments to mathematical entities.

We might have been led to introduce abstract entities in order to make sense of unending progressions. Many ordinary phenomena might have prompted us to wonder about them, but we can count on our (by then) well-developed numerical notations and geometrical templates to have led us to think of counting without end or subdividing lines into smaller and smaller segments. Perhaps we jumped to abstract mathematical objects at this point. But we might have resisted the move. Good nominalists among us might have shown us how to

account for our intuitions (about counting endlessly and ceaselessly subdividing lines) in terms of the possibility of performing more and more actions—actions which would require, of course, ever more matter, ever increasing lifetimes and attention spans, ever diminishing marks, and the like.

Thinking of mathematical objects as possible concrete ones hits its limit when it comes to limit entities. Perhaps stretching a cord tighter and tighter eventually forces it to be perfectly straight, but subdividing a line into smaller and smaller segments cannot yield an extensionless point; and drawing finer and finer lines cannot produce lines without breadth. Thus in countenancing limit entities it will no longer make sense for us to speak in terms of possible concreta. We will be forced to posit entities, such as points, lines, and circles, as sui generis and existing in their own right or else forego them altogether.

It would have also been simpler and more perspicuous for us to construe limit entities as abstract instead of trying to make do with possibilities. Even today it is unclear how limit entities could be formed from ordinary concrete material by natural processes. The only suggestion that I know of would be to see each limit entity as arising through the completion of an infinite process. One might think, for example, that it is possible to divide a region infinitely many times until nothing but an unextended point remains. But the history of mathematical attempts to understand infinite sequences and sums (as well as Zeno's paradoxes and their variants) shows that talk of completing an infinite process is a metaphor at best, one that breaks down when we ask what counts as finishing the process, what the last steps are like and, what results they produce. Thus it is more intelligible to deny that limit entities are some sort of actual or possible concreta and to posit them as nonmaterial and timeless things to which our concrete objects at most approximate. This move would deftly forestall questions concerning the origins and material properties of limit entities.

Yet such a move would have prompted skeptics to demand an explanation of how one can acquire knowledge about entities so different from the objects of our ordinary experience. I have brought our ancestors to the brink of recognizing abstract entities. Their colleagues' point is that it is not rational to take the plunge unless we (that is, our ancestors) can provide an account of how we can acquire knowledge about the new entities once we have countenanced them.

One way to start would be to posit, in addition to patterns, isomorphisms between certain simple finite templates and the patterns associated with them. Using these we could project some properties of templates onto patterns. We would find that much of our knowledge of the former transferred to the latter.[9] Furthermore, some of the discoveries about templates that led us to limit entities would also help us to discern some of the latter's properties. The considerations that showed us why circles, points, and lines cannot be identified with concrete geometric inscriptions would also indicate where such inscriptions reliably reflect features of abstract entities and where they fail. Thus, we would not be asking our fellows to accept a new mystery that we alone are qualified to interpret. We could point out that we and they have already clearly discerned some of the features of these new entities, that we already have a scheme for representing many of their features, and that we already have some methods for determining their properties and have reason to expect to develop more. Finally, we could also point out that if we do take the plunge and countenance limit entities, then we could also countenance numbers, linguistic types, and a host of other abstract entities our templates represent more adequately than they represent limit entities.

It would have been unfair and premature, however, for our ancestors' skeptical colleagues to demand a complete account of the methods for learning about abstract patterns. Our ancestors would have hardly begun to see the ontological picture. They might have reasonably expected to discover then-undreamed-of methods for exploring the new ontology, just as they could not have conceived of some of the ways we now have for learning about objects as familiar to us as our own bodies.

POSITING MATHEMATICAL OBJECTS

Many abilities developed by humans prior to the onset of mathematics figure in the account of the last section, including the ability to communicate, to use pictures, diagrams, and words to represent things that are absent or merely imagined; the ability to speculate; and, finally, the ability to hypothesize and theorize about new kinds of entities. According to some philosophers, many of these abilities already require interacting with abstract entities. Frege, for instance, maintained that speaking a language, judging, reasoning, indeed, thinking of any kind took place through "grasping" thoughts—his term for abstract entities associated with sentences as their meanings. If he and others like him are right, then my account presupposes an ability to interact with abstract entities at its outset. If they are right, the entire project of naturalizing epistemology seems doomed from

the start. More brightly, if they are right, then, of course, the question of how we know mathematical objects is not fundamentally different from the question of how we know anything about anything at all.

I do not think that Frege was right nor do most contemporary philosophers of cognition. None of them claim to understand the mechanisms involved in thinking, representing, and communicating significantly better than Frege did, but they judge his approach to be a scientific dead end. (One reason is that a system of sentences and brain states can take over the role thoughts play in Frege's account of cognition and communication.) Thus, under the current circumstances, it is fair for us to assume that no abstract entities participate in the premathematical activities that eventuate, on my account, in mathematical knowledge.

Positing mathematical objects first brought them under our ken. Although positing mathematical objects is a variety of verbal behavior, it does not involve a Fregean grasping of abstract entities. To posit mathematical objects is simply to introduce discourse about them and to affirm their existence. It involves nothing more mysterious than the ability to tell fairy stories, invent myths about the gods, or theorize about the forces at work in the observable world. Even Armstrong implicitly recognized as much by questioning our justification for positing mathematical objects without casting aspersions on our ability to do so.

Although I have spoken of positing so far in connection with introducing a whole new category of objects, mathematicians and scientists also use it to introduce single objects within an extant framework. We see physicists positing new particles as additions to extant systems, astronomers affirming new galaxies, and mathematicians postulating new transfinite cardinals. Also, the line between positing and discovering often blurs. Thus people speak of the discovery of the positron, although Dirac posited it long before any one elicited its observable traces. For our purposes, however, it is unimportant that these lines fade.

Yet here is a puzzle we should face, if only briefly. People have posited ghosts, the Ether, and phlogiston with as much ease as they have posited numbers. How can positing lead to knowledge in the one case and not in the others? What distinguishes between them? Primarily, truth and existence. Ghosts, the Ether, and phlogiston do not exist. Hypotheses that they do, are false. Hence no matter how justified people might have been in positing them, doing this could not have led them to knowledge. Of course, truth and existence are no guarantee that positing will lead to knowledge, since positors may lack the appropriate justification for their true beliefs. But our ancestors did not lack an appropriate justification for believing that mathematical objects exist, and they do exist (or so I have assumed throughout this chapter); so our ancestors' positing led them to mathematical knowledge.

As to ourselves, most of us acquire our initial beliefs about mathematical objects from teachers. On the face of it, then, we acquired our mathematical knowledge by having it communicated to us by those in the know. This is no more problematic or unnaturalizable than our ability to learn from our teachers and texts about historical figures or foreign lands and peoples. If a naturalized epistemology can make sense of the transmission of knowledge in the one case, it should be able to make sense of it in the other.

A more subtle problem concerns the aboutness of our mathematical beliefs. What makes them about mathematical objects? And in what sense are they about them? Are they about the same objects that our mathematical ancestors posited? A related problem concerns the apparent lack of "epistemic contact" with mathematical objects which positing does not seem to provide. Some might think that this means that we cannot have knowledge of mathematical objects. This worry is probably due to a mistaken adherence to the mathematics/physics distinction. I think that I can put it to rest along with the other worries I have just canvassed. Unfortunately, I cannot even begin to do so here.[10]

NOTES

I would like to thank Dorit Bar-On, Michael Hand, William Lycan, and Susan Williams for their help in writing the paper this chapter is based on.

1. This point, originally due to Frege (cf. Resnik, 1980:62–63), was emphasized by Benacerraf (1973). Circumventing it was the principal motivation for Field (1980).

2. Field's failed attempt (1980) underscores both the conceptual and technical difficulties of such a project. I discuss the former in Resnik, 1985b and both sorts of problems in Resnik, 1985c.

3. In speaking thus I have not been thinking of science as encompassing so-called pure mathematics. It thus falls short of affirming the existence of many of the entities studied by the far reaches of contemporary mathematics. I do not find this a drawback at this point, since my purposes in this paper will be served if I can naturalize the epistemology of the kind of mathematical entities that figure in science. Furthermore, I think that it is arguable, along the lines I have been pursuing so far, that pure mathematics is part of science.

4. Even the totality of all spacetime points and regions falls short of the number of objects required by mathematics.

5. It is not clear that Quine meant to exclude normative investigations, since he writes that we are still prompted to do epistemology for the traditional reasons, "namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence" (1985:24).

6. Quine may have restricted himself to psychology on the grounds that the other social sciences make such heavy use of the problematic ideas of translation and interpretation that they fail to count as genuine sciences. But even psychology trades heavily in beliefs, which are under Quinean interdiction along with translation. So since none of the social sciences can be taken over intact, properly filtered treatises on (at least) anthropology, history, and sociology should lie alongside of psychology in the epistemologist's library.

7. Important work on these questions is contained in Kitcher, 1983 and Maddy, 1988a, 1988b.

8. For a recent study of such explanations see Resnik (forthcoming a, b).

9. I discuss this in more detail in Resnik, forthcoming a.

10. I address these questions in Resnik, forthcoming b.

REFERENCES

Armstrong, D.
1981 "Naturalism, Materialism and First Philosophy." In D. M. Armstrong, *The Nature of Mind and Other Essays*. Ithaca: Cornell University Press.

Benacerraf, P.
1973 "Mathematical Truth." *Journal of Philosophy* 70:661–80.

Dodd, J. E.
1984 *The Ideas of Particle Physics*. New York: Cambridge University Press.

Field, H.
1980 *Science without Numbers*. Princeton: Princeton University Press.

Freedman, D. Z., and P. van Nieuwenhuizen
1985 "The Hidden Dimensions of Spacetime." *Scientific American* 252:74–81.

Hale, S. C.
1988 "Spacetime and the Abstract/Concrete Distinction." *Philosophical Studies* 53:85–102.

Kitcher, P.
1983 *The Nature of Mathematical Knowledge.* New York: Oxford University Press.
1988 "Mathematical Naturalism." In P. Kitcher and W. Aspray (eds.), *History and Philosophy of Modern Mathematics,* Minnesota Studies in the Philosophy of Science, 11. Minneapolis: University of Minnesota Press.

Kornblith, H.
1985 "Introduction: What is Naturalistic Epistemology?." In H. Kornblith (ed.), *Naturalizing Epistemology.* Cambridge, Mass.: Bradford Books, MIT Press.

Maddy, P.
1982 "Mathematical Epistemology: What is the Question?" *Noûs* 16:106–7.
1988a "Believing the Axioms." Pt. 1. *The Journal of Symbolic Logic* 53:481–511.
1988b "Believing the Axioms." Pt. 2. *The Journal of Symbolic Logic* 53:736–64.

Quine, W. V. O.
1981 "Reply to Stroud." In *Foundations of Analytic Philosophy,* Midwest Studies in Philosophy, 6. Minneapolis: University of Minnesota Press.
1985 "Epistemology Naturalized." In H. Kornblith (1985).

Resnik, D. B.
1989 "Adaptationist Explanations." *Studies in the History and Philosophy of Science* 20, 2:193–213.

Resnik, M. D.
1980 *Frege and the Philosophy of Mathematics.* Ithaca: Cornell University Press, 1980.
1981 "Mathematics as a Science of Patterns: Ontology and Reference." *Noûs* 15:529–50.
1982 "Mathematics as a Science of Patterns: Epistemology." *Noûs* 16:95–105.
1985a "Logic: Normative or Descriptive? The Ethics of Belief or a Branch of Psychology?" *Philosophy of Science* 52:221–38.
1985b "How Nominalist is Hartry Field's Nominalism?" *Philosophical Studies* 47:163–81.
1985c "Ontology and Logic: Remarks on Hartry Field's Anti-platonist Philosophy of Mathematics." *History and Philosophy of Logic* 6:191–209.
1992 "Proof as a Source of Truth." In M. Detlefsen (ed.), *Proof and Knowledge in Mathematics.* London: Routledge and Kegan Paul.
1990 "Beliefs About Mathematical Objects." In A. Irvine (ed.), *Physicalism in Mathematics.* Dordrecht: Kluwer Academic Publishers.

THOMAS TYMOCZKO

4

Mathematical Skepticism: Are We Brains in a Countable Vat?

By the end of this chapter, I will propose a rather extreme sociological view of a branch of mathematics. I will contend that the development of transfinite set theory is guided not by some external domain of sets or by some transcendental logic but by basic features of human nature. However, I develop my suggestion rather slowly. In particular, I will try to develop for mathematics a version of the traditional philosophical problem of skepticism about the external world. My mathematical skepticism will be directed at the distinction that lies at the core of modern set theory, the distinction between countable and uncountable sets. I want to know whether this distinction has any objective reality.

Allow me to begin by setting the issue in a context.

In the anthology *New Directions,* I espoused the cause of quasi-empiricism as an approach in the philosophy of mathematics, borrowing the term from Lakatos and Putnam (Tymoczko, 1986). There are two key aspects to quasi-empiricism: first and foremost, an emphasis on mathematical practice; and second, an openness to scientific methods in mathematics, for example, the use of computer proofs. Although these two features are compatible, sometimes they can pull in opposite directions.

The second feature, which invites us to see mathematical methods as close to scientific methods—or closer than was once thought—goes naturally with a realist view of mathematical objects and with