Comments on Maddy and Tymoczko

Paul Benacerraf

PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Vol. 1984, Volume Two: Symposia and Invited Papers. (1984), pp. 476-485.

Stable URL:
http://links.jstor.org/sici?sici=0270-8647%281984%291984%3C476%3ACOMAT%3E2.0.CO%3B2-9

PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association is currently published by The University of Chicago Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/ucpress.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org
Sat May 12 00:54:44 2007

Comments on Maddy and Tymoczko[1]

Paul Benacerraf

Princeton University

My remarks are largely exploratory, and not at all faithful to the title of the symposium; indeed, they skirt most of the fascinating issues discussed by Maddy and Tymoczko. Furthermore, I have no settled position on the issues that I raise. In fact, in a recent paper (Benacerraf 1985) I advance a platonist line much like the one so effectively derided by Tymoczko (and which he has confessed he believes as well). Although I will not argue the point, I do believe this much: they are interesting issues that deserve to be raised and given careful consideration, no matter what their ultimate disposition. I will let the questions speak for themselves.

In the first part I raise a skeptical problem that is intended to illustrate a perplexity I imagine most of us feel when we contemplate the following two questions side by side (they are evidently prompted by Maddy's fascinating account of recent work in set theory):

(1) What is the structure (content) of our mathematical concepts?

and

(2) What about us makes it so that we possess concepts having such a structure (content)?

[Not: How did we come by them? That is a further question, interesting in its own right, but further. More like: What "facts" about those that have them make it the case that they have them?]

In the second part I reiterate what I take to be obvious: that the natural direction from which to expect an answer is in terms of the notion of meaning which, in these cases, intimately involves the concept of a rule. This leads to some brief remarks on how recent work by Kripke on Wittgenstein (Kripke 1982) might bear on the issue -- a matter to which Tymoczko devotes a significant portion of his own paper.

I end in perplexity.

1. A Skeptical Suggestion

I begin by inviting you to think with me in some detail about a question that forces itself upon us when we listen to Maddy, Moschovakis,

PSA 1984, Volume 2, pp. 476-485
Copyright © 1985 by the Philosophy of Science Association

et al., -- when we read Gödel and others describing work in set theory (much like the work Maddy has related to us) as the "unfolding" or "exploration" of certain features of our concept of set:

> "...the axioms of set theory by no means form a system closed in itself, but, quite on the contrary, the very concept of set on which they are based suggests their extension by new axioms which assert the existence of still further iterations of the operation "set of"....These [strong axioms of infinity] show clearly, not only that the axiomatic system of set theory as used today is incomplete, but also that it can be supplemented without arbitrariness by new axioms which only unfold the content of the concept of set explained above." (Gödel 1947/64, pp. 476-77)

The concept is there -- it is ours -- and one part of what we do in mathematics is to explore and unpack what is in some sense "in" it. So, there is immediately a presumed contrast among (1) unfolding and uncovering what is there, or isn't to be found there; (2) seeing that something or other is incompatible with what is there; and (3) deciding to modify or extend the concept, perhaps by adopting some axiom that is either incompatible with or indeterminate in our present conception (and making other suitable adjustments).

What sense does it make to say in a given case that we are exploring *vs.* expanding *vs.* modifying, etc. "our concept" of set? Given the purported complexity of the concept, what must be the structure of our minds or brains for them to contain a representation or an encoding of the structure that we attribute to the cumulative hierarchy, *including* in the encoding enough of that structure that transcends what we consciously or explicitly believe, as that might be represented by the axioms we accept,[2] to make sense of the notion of our working at actualizing it in terms of further articulating our beliefs -- say as *new* axioms?

Note that, although in the end it may come down to that, this question does not concern in any *immediate* way whether we have knowledge of anything independent of us. It concerns what it is to entertain (contain) even ever so vaguely, a representation of something so complex as, say, the cumulative hierarchy[3] in all its delicious intricate detail -- a representation itself sufficiently detailed to be adequate to the task, *whether or not such a representation represents matters as they are*, i.e., whether or not there exists a structure that corresponds to it in the desired way.

For those -- call them conceptualists -- who see mathematics as the externalization and articulation in mathematical language of such conceptual structures, there is a fact of the matter (just as there was for Gödel, say, though no conceptualist he) whether the conceptual structure we are in the process of elaborating (and have, probably at least partially, expressed as the Zermelo-Fraenkel axioms (ZFC)) is one in which the Generalized Continuum Hypothesis (GCH) holds, to choose a tired example. Notice that this is *not* the question of the *independence* of GCH from ZFC. It concerns rather whether GCH holds in, contradicts, or is underdetermined by "our concept of set" -- the concept we are struggling to make explicit in our set-theoretical investigations. The model is something like this:

(1) We have a "concept of set" -- perhaps an iterative conception;

(2) We make explicit some of the properties of this concept by writing down principles. Some get to be axioms, others theorems. Think of it on the model of the witness who is describing the suspect to the police artist -- only we are both witness and artist. Clearly (a) there is more to our representation of the suspect than we articulate; (b) we may at any point "articulate" something that isn't there -- i.e., we may be mistaken in saying she had a moustache -- and that mistake could be a mistake not only about the suspect (which is not what concerns us here), but also about our own representation of the suspect, something we could later come to describe simply as having misremembered. Our representation is not something we can read like a book. (c) We could have such a "representation" whether or not we had in fact seen anyone. It suffices that we merely think we have. [I just put 'representation' in shudder quotes simply to sidestep the issue of whether, when it did not arise from a "veridical" (more shudder quotes) context, what we have should properly be called a representation, and its content, its representational content.]

What would make someone a "conceptualist" is the supposition that all we have to go on is our concepts, and that in doing mathematics we are simply elaborating these, as well as the consequences of the axioms we put down that purport at least partly to express them. Mistakes can come by seeing only dimly what is there, in our minds, in some form or other; or they can come by the very "structures" we are trying to express themselves being confused and incoherent. The distinction between these is impossible to pin down absent a reasonable psychological theory of what it is to have certain concepts [sentences-in-the-head vs. other "representational" structures, for example]. But in either case, the conceptualist must and does allow for error -- transcriptional error, so to speak. Even though it is our own concepts that we are expressing, we can mess up.

Such a view of what we are doing in mathematics (at various stages of development) leaves out much. It is also laced with the very real problems that Wittgenstein, Kripke, and now Tymoczko have called to our attention. To be sustained to any degree at all it must answer, for example, the obvious questions that arise concerning the possibility of communication among mathematicians. But what I am concerned to bring out here is that in writing or assessing the adequacy or plausibility of axioms, the conceptualist is somehow comparing the axioms and their consequences with some (internal?) "structure" they are meant to express.

So far, the platonist and conceptualist are boatfellows; for both are articulating or assessing propositions which are meant to express aspects of some conception each claims to have. The platonist feels in addition that this conception is of some reality [external to and][4] independent of the working mathematician. The representational structures he is struggling to articulate either correspond or fail to correspond to this independent reality. And his conception is, presumably, somehow responsive to this external reality (though this last is a more narrowly epistemic question). He is subject to error in both of the ways typical of our witness-cum-artist: he could be misrepresenting his own conception -- seen but dimly; or his conception could itself fail accurately to represent the reality to which it is meant to correspond; or both.

Mysteries abound concerning both what that reality might be and how the mathematician's evolving conceptions can be responsive to it -- how they can be seen to require adjustment because they are seen as misrepresenting that reality. Inconsistency, of course, is the easy case -- that form of "misrepresentation", although not always easy to spot, once spotted, is easy to label as a misrepresentation. There are no inconsistent structures or objects, only descriptions (*pace* Cantor and "inconsistent totalities" -- treating *them* as decoy ducks would not do violence to his view).[5]

The platonist's opponent, be it a nominalist, neo-intuitionist, or Skolemite, always seems to end up either denying that we know what it seems obvious that we do know, or reinterpreting what we obviously know as some unsuspected, epistemically sanitary, thesis. The platonist doggedly persists in maintaining that our sentences mean just what they seem to mean, and that we possess the knowledge we would possess if we knew the relevant propositions that we in fact seem to know. That this presents a problem *for the philosopher*, who must present us with an adequate epistemology is meant to be the principal thrust of Benacerraf (1973). But whatever the inadequacies of the view, at least this much is *right* in the platonists' persistence in it despite their singular lack of success in producing a satisfactory explanatory epistemological account: it is plainly a foolish *policy* simply to adjust our beliefs to fit our (admittedly inadequate) theory of what we can come to know and how we can come to know it. Our cognitive processes function. They should be assumed to be functioning, and perfectly adequately, despite our own inability, on reflection, to grasp and explain *how* they function. (This is what I think Kreisel (1965) means as he frequently adverts to remarks about how skeptical, or restrictive accounts fail to do justice to our "mathematical" or "intellectual experience".) Still, all this notwithstanding, philosophical reflection on these processes must continue, and it must be taken seriously. It is in this vein that I pursue here lines that are almost bound to lead to skeptical conclusions.

What is hard is to strike the proper balance.

So my aim here is to raise a question concerning whether (or how) a picture such as that offered above of the common ground between the platonist and conceptualist could be right at all if it is also to imply that the host of propositions with which Professor Maddy has entertained and instructed us have determinate truth values *in the conception(s) of set of the people she discusses* -- never mind whether they have determinate truth values *tout court*.

Given the extreme finitude of our minds and representational systems [both of which I take as given], how determinate a picture (representation, conception) *can we possibly have, now or ever*, of structures as complex as the cumulative hierarchy is alleged to be? Cardinality arguments suggest themselves. 'Determinate' is not the right word, of course. 'Detailed' is best. Or, since that way of putting the question smacks of platonism, we can ask instead about a hypothetical structure, leaving open the question of whether the hypothetical structure is even being conjectured to exist. What seems obvious is that the degree of complexity of these hypothetical structures far outstrips any possible degree of complexity of the *representing* structures, be they our internal, conceptual systems of representation or the linguistic apparatus that we use to express these. (I am convinced that some such story is what lurks behind Skolemite arguments for the relativity of set theoretic concepts. Or "justifies them" if you prefer to think of them as justified.) Call this "the complexity problem".

[WARNING: There is a danger lurking here, at least *prima facie*: For the question being posed even to make sense, as posed by (presumably) finite *me*, it would appear that I had to have compared *us* on the one hand, and *it* on the other, and seen that there was indeed a representational gap. As photographer I am aware of the mismatch between the coarse-grained film I am using and the detail I might like to capture in this picture: that silhouette will of necessity appear only as an almost amorphous blur, not as the svelte figure that it is. I see the figure and know from experience (or theory, no matter) the limits of this old-fashioned fast film. But in the mathematical case, what corresponds to seeing the figure? It would seem that in order to be in a position to claim the disparity I must, self-defeatingly, compare the conception of set at issue with the representational capacity of my powers of representation, *and this I can only do through my representations*. So, on this model, in order even to pose this skeptical concern specifically about the conception of set under discussion, it would seem that I would have to be in full possession of the concept and come to realize, *on that basis*, that *such* a conception cannot be adequately represented by my (our) conceptual apparatus, owing presumably to some shortcomings of that apparatus (of which I would also have to have some minimal conception). I don't *think* the problem is intractable. But I do think it is a problem.

Reflection is a source of insights *and* problems, particularly reflection on reflection.]

Note that the complexity problem, if it exists at all, exists also for the Gödelian platonist whose conceptual structures are constantly being shaped by some form of "interaction" with the reality that exhibits the structures we speak of. *Apprehending* is something *we* can do only with the representational equipment at our disposal. So even if we have some form of epistemic contact with ("objectual apprehension" of?) this independent reality, we are limited in the conceptual representations we make of it by our representational equipment. You can't take color pictures with black and white film. Nor can you represent on film of very large grain aspects of the pictured object that are below the threshold of representation of the film, to revert to a picture that is undoubtedly somewhat misleading, but that still *seems* quite apt.

In thinking through what it is that we do when we investigate the outer reaches and foundations of set theory in the way that Professor Maddy has illustrated for us, we must form *some* model of the investigating mathematician (or mathematical community). One set of questions that arise are an expression of my own bewilderment at how *whatever representational devices we have* could possibly be sufficiently fine-grained to represent the wealth of detail that our *theory* of sets *purports to express* -- measurable and inaccessible cardinals, constructible *vs.* unconstructible sets of reals, to mention only the tip of the iceberg.

A related question, on which I will pause only briefly, concerns how our linguistic structures themselves can represent -- i.e., purport to express, this level of detail. As I suggested above, Skolem's arguments for the relativity of the cardinality concepts (but in truth of all the set-theoretic concepts) seem to me simply to fall out of the basic query I am addressing now. But they do so with greater precision because he bases his views on an explicit _analysis_ of our representational structures -- in this case any first order set-theory -- and argues that such a theory, in consequence of its first-order structure _alone_ cannot forestall interpretation over denumerable domains. So one aspect of the expressive complexity of which I spoke cannot reside in the first-order structure of the language.

We are tempted to reply, as I have done (Benacerraf 1985): Who expected it to reside there in the first place? It resides perhaps in the meaning we intend for '$\in$', which intentions are not honored in interpretations which assign 'TRUE' to the sentence that _purports_ to express Cantor's theorem, but on which it is in fact _false_, when the sentence is interpreted as meaning what Cantor's theorem means. The denumerable interpretation fails to be a model of _Cantorian_ set theory, although it is one of the models of the first-order formalism in which we sometimes express the theory -- a kind of model the theory as expressed in that formalism _must_ have if it is consistent.

That strikes me as the right reply. But what entitles us to make it? Yet, if we _can't_ justifiably make it, much of the thrust of the question I am trying to raise here -- concerning the apparent gap between the potential representational power of our conceptual apparatus and the structures being represented -- simply disappears. We can _have_ no relevantly adequate conception of anything more complex than our conceptual apparatus can mirror in some obvious way. Our _present_ theories express propositions about the universe of sets that are determinate _at most_ up to the first- order expressive power of our language. So we can rest easy, all that those "large cardinal axioms" are capable of expressing are features of denumerable collections.

But of course, their ability to express _even these_ may depend in some essential way on their being imbedded in infinite languages -- ones with more than finitely many (finite) formulas. Otherwise, quantifiers could perhaps be regarded as a mere convenience, abbreviating conjunction and disjunction. _These_ are certainly idealizations relative to the expressive and representational powers of the mere finite humans that manipulate them. So whatever is in our _heads_ is certainly outstripped by even the most rudimentary theory with _only_ infinite models (and, for some not-very-large n, for models with cardinality n or greater).

These are some of the intuitions that I see behind many of the skeptical arguments of Skolem and others, perhaps even Kripke's Wittgenstein. If we can mean what we think we mean by our mathematical talk, we must do so because in some sense we transcend these finite (denumerable, inaccessible, etc...) limitations. Otherwise, 'plus' _does_ mean 'quus', for some not-very-large n. Of course, if we can't even _think_ we mean what we think we think we mean,...

2. Rules

How *do* we transcend these limitations -- by *rules* is the customary answer, indeed rules with more than finite scope. We don't need, in order to grasp what it means to square a number, to store all ordered pairs of the form n, n^2 in our brains, minds, or what-have-you. We need only learn how to *apply the rule*: multiply n by n. The *rule-as-written* is a finite object which is meant to determine (represent? encapsulate?) an operation with infinite scope. But looking more closely -- in what could my *understanding* of this rule reside? How could I have gotten it *right*? What mental (or other) items or ability could constitute my having gotten it right? Note that this doesn't mean: how could I have incorporated an inscription of the right abstract rule into my mental life? [Unless we understand that in a way that begs the whole question: E.g., an inscription is of the "right" abstract rule if(f) we use it in the right way.] Rather it means: how could I have understood aright what I was (presumably meant to have been) taught when I was taught the rule? What "understanding" was I given of the rule-as-written that transcended my finite limitations: either I was given *examples* (from which I generalize), or referred to some other rules, themselves with potentially infinite outputs but themselves also given only on a finite basis. In the end, my general "understanding", which means my understanding *tout court*, (for *all* understanding is general) rests on a finite basis and depends on my making "correct" projections from that finite basis.

The argument is meant to undermine the claim that I can *know* that one projected continuation is better than another because *the very judgment* that one is better is itself a judgment based on an interpretation of the rule-as-written [of a mental or other token of the rule] -- an interpretation that is *presumed* right, but given the nature of which it remains a total mystery how I could have gotten right (or indeed even what constitutes getting it right). With Tymoczko, let us call some version of the foregoing the "Private Language Argument" (PLA).

It is not entirely obvious what the argument shows. For Kripke's Wittgenstein, what is uppermost is the *normative* element -- the normative role of the rule in defining the correct continuation. This leads naturally to a generalization beyond mathematics to all of meaning. In my opinion the argument, if it has any bite at all, certainly chews up Tymoczko's formalist in the starting gate. He cannot even *present* his formal system without running headlong into the very problem we have been discussing -- to begin with, in terms of what constitutes a *formula*. The rules defining wffhood would be rules every bit as much in need of interpretation as the laws of addition or quaddition.

On some accounts, including Tymoczko's, both of Wittgenstein and of the issues in dispute, it is not possible to resolve the matter at the individual level. What *makes* a given continuation or application "correct" is that it is [the] one sanctioned by the linguistic community of the speaker -- one that agrees with the practice of the community. The *community* provides the standard of correctness that the individual's intentions were powerless to supply and without which the crucial notion of "correct application" was argued not to make sense. It anchors the use of speakers in their shared practice.

Some ancient questions come naturally to mind: Is the Private Language Argument one to which the mathematician who works alone, anti-social and cut off from the rest of the community, or on a desert isle must succumb? If so, shouldn't he take everyone else with him? For how is the community in better shape than the lonely mathematician (except perhaps for not being lonely)? Suppose there are two. Isn't it every bit as much a matter of interpretation -- open to differing interpretations -- whether Sam and Cecily agree whether Sam today is acting in accordance with yesterday's intentions? Surely it is not memory that is at stake here. And if there is, to that extent, no more fact of the matter as to whether mathematicians taken pairwise agree or disagree than there is when they are taken onewise, how can there come to be when they are aggregated into even larger clumps?

The reply might be: Yes, Mathilda, there is no fact of the matter in that sense about clumps either. There is just the ur-phenomenon of signifying agreement, and by extension, of participation in, and therefore conformity with, a practice.

But one is tempted to reply: Can one make the extension? Or does there remain a gap? Even granting some "irreducible", perhaps even syncategorematic, notion of signifying agreement, does that carry us all the way to the notion of a practice? For a practice would appear to require not just signified agreement (and/or disagreement), but at the very least, conformity to a rule; whether or not such conformity is acknowledged by acts overtly signifying agreement with the practice. (If dispositions to signify agreement would help here, we could have helped ourselves to them much earlier in the argument and the skeptical problem would not have gotten off the ground.)

On the other side, in defense of closet mathematics, if I go into my closet finally to work out my ideas for a proof of Goldbach's conjecture, and if they jell, have I not solved the problem, even if I die of joy on the spot? Or what is worse, even if acid rain has warped the minds of those outside my closet in such a way that they can't see my proof as a proof when I emerge from my closet and circulate it?? I proved it -- whether anyone else ever knows I did; even if everyone else denies I did. Which is not to say that just anything I might have scribbled in the closet would have counted as a proof.

As to whether we are stuck with the skeptical conclusion -- first in the one-person case and then in the community case -- I think that depends in part on whether we can make something of the crucial remark in *Philosophical Investigations*-201 "What this shows is that there is a way of grasping a rule which is not an interpretation, but which is exhibited in what we call 'obeying the rule' and 'going against it' in actual cases." But this can be of help only if it can be shown that PLA as employed for the skeptical conclusion about mathematics was dependent on repeated reliance on the interpretation of rules. And even then, it will undermine the skeptical conclusion only to the extent that it rested on PLA. If the argument from complexity with which I began these remarks can be sustained independently of PLA, a significant skeptical conclusion would stand even if *PI*-201 does indeed undermine the Kripke-Wittgenstein argument.

Notes

[1]Comments on (Maddy 1985) and (Tymoczko 1985), constituting, along with comments by George Boolos, a Symposium entitled "New Directions in the Philosophy of Mathematics", Philosophy of Science Association Meetings, Chicago, October 1984. I am grateful to Paul Boghossian for a number of discussions of Kripke's Wittgenstein. He is blameless for the remarks I make here.

[2]Given the intimate connection that exists between a portion of set theory and second-order axioms in "standard interpretation", it would be appropriate here to distinguish first- and second-order versions of the axioms we accept. If -- and this is only an example -- as Kreisel maintains (cf. Kreisel 1965) the Continuum Hypothesis is decided in full second-order Zermelo-Fraenkel Set Theory (ZFC^2), although it is independent of second-order formalizations of ZFC, then much of the mystery I am urging you to contemplate concerns what constitutes "accepting" ZFC^2 as opposed to just ZFC (or at least has a counterpart in that problem). There is much to be said here, but here is not the place to say it. For now, we should content ourselves with noting the complexity of what lies dormant under the phrase "the axioms we accept".

[3]I pick the cumulative hierarchy of pure sets because that seems to be everyone's favorite candidate for the clearest and simplest to elucidate -- and hence for the one we are most likely to conceive clearly. Very few claim to have a clear grasp of the other conception(s). Some of the other candidates: the ramified type hierarchy with reducibility, Quine's systems of New Foundations or Mathematical Logic. See discussions of the cumulative hierarchy in the articles by Boolos, Gödel, Parsons, and Wang in Part IV of (Benacerraf and Putnam 1983), as well as in (Hallett 1984), (Scott 1974), and (Van Aken forthcoming). Opinion appears to be divided concerning whether "capping" the cumulative hierarchy with a layer of "proper classes" [super-collections that are neither sets nor elements], as in Gödel-vonNeumann, yields a natural conception when the capping is "taken seriously" -- i.e., not simply as a way of bounding models of the hierarchy which, viewed from another perspective, are seen as only partial. There is, of course, some question as to whether even this tentative explanation of what it is to "take seriously" a theory of proper classes even points in any direction, much less expresses a definite proposition.

[4]The internal-external metaphor is not only overworked, it is misleading to boot. Presumably, the cumulative hierarchy is not in space-time. So, although it is not internal to us, neither is it external. "Independent" seems better. At least it leaves the door open for the platonist to find an account in terms of which our theories can be objective (though, of course, it doesn't guarantee their objectivity).

[5]The salience of inconsistency as an undesirable trait, coupled with the mystery surrounding what else we can use to decide if our mathematical theories accord with the "mathematical facts", may suffice to explain the formalist's insistence on consistency and sometimes, "fruitfulness", as the sole touchstones of mathematical correctness.

References

Asquith, Peter D. and Kitcher, Philip. (eds.). (1985). PSA 1984. Volume 2. East Lansing: Philosophy of Science Association.

Benacerraf, Paul. (1973). "Mathematical Truth." Journal of Philosophy 70: 661-680. (As reprinted in Benacerraf and Putnam (1983). Pages 403-420.)

---------------- and Putnam, Hilary. (eds.). (1983). Philosophy of Mathematics, Selected Readings. 2nd ed. New York: Cambridge University Press.

----------------. (1985). "Skolem and the Skeptic." Proceedings of the Aristotelian Society, Supplementary Volume LIX: 85-115.

Gödel, Kurt. (1947/64). "What is Cantor's Continuum Problem." American Mathematical Monthly 54: 515-525. (As reprinted in Benacerraf and Putnam (1983). Pages 470-485.)

Hallett, Michael. (1984). Cantorian Set Theory and Limitation of Size. Oxford: Clarendon Press.

Kreisel, Georg. (1965). "Informal Rigour and Completeness Proofs." In Problems in the Philosophy of Mathematics. (Proceedings of the International Colloquium in the Philosophy of Science, London, 1965, Volume 1.) Edited by Imre Lakatos. Amsterdam: North Holland. Pages 138-171.

Kripke, Saul A. (1982). Wittgenstein on Rules and the Private Language Argument. Cambridge, MA: Harvard University Press.

Maddy, Penelope. (1985). "New Directions in the Philosophy of Mathematics." In Asquith and Kitcher (1985). Pages 427-448.

Scott, Dana. (1974). "Axiomatizing Set Theory." In Axiomatic Set Theory. (Proceedings of Symposia in Pure Mathematics, Volume 13, Part II.) Edited by T.J. Jech. Providence: American Mathematical Society. Pages 207-214.

Tymoczko, Thomas. (1985). "Gödel, Wittgenstein, and the Nature of Mathematical Knowledge." In Asquith and Kitcher (1985). Pages 449-468.

van Aken, James. (Forthcoming). "Axioms for the Set Theoretic Hierarchy." Journal of Philosophical Logic.

Wittgenstein, Ludwig. (1953). Philosophical Investigations. (trans.) G.E.M. Anscombe. New York: MacMillan.