Berkeley and the Bridges - What do Empirical Applications of Mathematics Commit us to?

1 Introduction

In his essay 'What is the Problem of Mathematical Knowledge?'[1] Micheal Potter argues that the central epistemological problem about mathematics is the problem of 'accounting for' the consistency of our mathematical axioms. Insofar as the Platonist takes us to have knowledge of mathematical objects, this is supposed to provide him with an initial answer to the consistency question but commit him to saying how we can know about mathematical objects. In contrast Fictionalsits, Formalists and Structuralists face this problem in its pure form, insofar as their use of arithmetic to draw conclusions about e.g. how many fruit two apples and two oranges make commits them too to thinking that the axioms of PA are consistent (and presumably the same holds for applications of set theory via analysis in physics). Thus, everyone is committed to the claim that certain mathematical axioms systems are consistent, and everyone needs to 'account' for this somehow. But what sort of an 'account' do we really owe merely in virtue of being willing to apply mathematics?

In this paper I will argue that our willingness to use a certain mathematical system to form expectations about the empirical world commits us to very little indeed by way of consistency claims - so little that one can get this kind of knowledge empirically. If we imagine Bishop Berkeley as thinking that the calculus was inconsistent but people applying it tended to get correct answers because they made 'compensatory errors' (I gather this is not too far from the historical truth) then it would still be rational for him to walk over bridges built using calculus and even to use calculus himself in building bridges to the extent that he could trust himself to think like the common man in doing so. All that literally using a mathematical practice to form expectations about the empirical world commits one to is a sort of co-variation between two physical systems (yourself doing the math, and bridges or apples or whatnot) both of which you can observe.

A more interesting epistemological question arises if you think that we have reason to believe not only that our use of PA won't lead us to deriving in-

[1]Mathematical Knowledge, Oxford University Press, 2004 , Michael Potter, Mary Leng and Alexander Paseau eds.

compatible answers (and hence perhaps that there is some consistent system which we can be thought of as working in) but that no derivation of inconsistency from PA is possible. I will argue that (contra Potter) the difference between Platonism and Ficitonalism or Structuralism makes no difference here, since the Fictionalist's rational intuition into what proof inscriptions are possible will serve just as well as the Platonist's rational intuition into what proofs exist. I will also briefly suggest how the relevant kind of correct modal intuition can be accounted for naturalistic, quasi-empiricist fashion.

2 Accounting for the intelligibility of consistency claims

So, Potter writes that Platonists Formalists, Fictionalists and Structuralists alike need to 'account' for the consistency of the mathematical axioms systems which they use. What exactly does he think needs accounting for? One possibility is that we need to account for the intelligibility of the claim that mathematics is consistent.

This is an important challenge for a hardcore formalist who says that all mathematical statements are meaningless and that doing mathematics is just a question of producing symbol strings within a certain arbitrary formal system or for a hardcore conventionalist who says that mathematics is a matter of deriving formal logical consequences from purely arbitrary axioms sets. The problem for the formalist is that, intuitively, we understand claims about what can and can't be produced 'in principle' by engaging in formal manipulations in accordance with a certain set of rules e.g. that such manipulations can't yield a string ending in '0=1' or that such and such a position cannot be arrived at by following the rules for chess. Thus, insofar as the formalist believes his system is consistent, or even understands this kind of question about what his system can and cannot prove he must admit the existence of meaningful mathematical questions (namely: can a string with these syntactic properties end in '0=1'?/ is there a proof of 0=1 in such and such formal system?).

Similarly an old-fashioned conventionalist who believes that all mathematics is a matter of first order logic plus arbitrary axioms which determine the meaning of mathematical terms will be sorely tempted to say that there are determinate facts about whether any given thing (such as 0=1) can be proved from their axioms using first order logic. But the claim that contradiction cannot be derived from a given axiom system using first order logic is not itself provable in first order logic. Thus it would seem that there is at least one kind of axiom choice - the choice between adding axioms which let one 'prove' that a given system is consistent vs. adding axioms which let one 'prove' that it isn't consistent - where there is an objectively right answer and we can't just fix the meanings of the terms by adopting whatever convention appeals to you more.

To put the point more generally, many people feel a strong intuition that there are determinate facts about what is and isn't provable in a given formal system or produceable via a given kind of syntactic manipulation and these seem to provide an 'objective' subject matter for at least a portion of mathematics. Thus anyone who wants to claim that all mathematical statements are meaningless, or that we can settle questions about what is provable in a formal system by tossing a coin must indeed 'account' for our apparent understanding of the claim that PA and ZF really don't allow a proof of 0=1 in a way that goes beyond our having adopted a convention of saying so. However, this task of accounting for the apparent meaningfulness and determinate truth or falsity of consistency claims (and other claims about what is provable in principle) doesn't have much to do with epistemology or applications, nor does it pose a problem for most of Potter's intended targets. Clearly the Platonist can make sense objective facts about what recursively enumerable theories are consistent (either because he believes all possible proofs exist or because he believes in some set theoretic surrogate). And the Fictionalist and the Structuralist are explicitly committed (by the very nature of their theories) to thinking that there are definite facts about what would have to be true in certain fictional scenarios, and what structures are and aren't possible, respectively. So they can easily make sense of the claim that the axioms they are working with are consistent e.g. that they describe a possible fictional scenario or that there is a possible structure which provides a model for these axioms. In fact, one might suspect these contemporary theories include appeal to messy facts about modality and truth-in-fiction exactly in order to make sense of the kind of intuitions about first-order-axiom-transcendent mathematical facts which cause problems for older theories like formalism and conventionalism.

3 the problem of justifying our empirical use of mathematics

Instead, I think Potter's epistemological problem of 'accounting' for the consistency of our mathematical axiom systems is better understood as a problem about how to make sense of something like our warrant in assuming that the mathematical system we are working in is consistent, which assumption he takes our empirical applications of arithmetic to commit us to. This is a clever strategy. If we just look at things behaviorally, something every philosopher of math (however anti-realist) will do is form expectations about how many fruit there are in a basket containing 27 apples and 38 oranges by doing sums (that could be reconstructed in PA) and behave - e.g. bet- accordingly. So any philosophy of math owes us an account of why it's reasonable to engage in this behavior. Now if your mathematical doctrine consists of a bunch of inconsistent axioms you can prove both 27+38=66 and 27+38=67. But the apples and oranges can't both constitute 66 fruit and 67 fruit. So if you think your mathematical axioms are likely to be inconsistent there is plausibly something unreasonable about being willing to infer from 'proofs' in that system to conclusions about apples and oranges. Thus we get a general epistemological problem which any adequate philosophy of mathematics must address - explain how we warranted in making the kinds of consistency assumptions about our own mathematical systems which we need to make sense of their every day applications to questions about how many fruit n apples and m oranges constitute?

However, if we want to take this approach seriously we will need to get a good handle on just what kind of consistency claim the relevant mathematical behavior commits us to. For example, in order for someone to be acting rationally in forming beliefs about apples and oranges on the basis of doing calculations in PA, do they have to a) be certain that PA is consistent b) have warrant to assume that PA is probably consistent c) have reason to assume that all the uses of PA which they are disposed to make could be reconstructed in a consistent subsystem of PA d) have reason to assume that all their uses of PA in this kind of situation (though not necessarily all those which they would make in properly mathematical non-predictive contexts) could be reconstructed in a consistent subsystem of PA?

I think the right answer is either d) or something even weaker than that. Firstly, it can certainly be rational to use belief forming methods which one doesn't take to be infallible. Otherwise how could we learn about tables and chairs by looking at them, or about the sun rising tomorrow by doing empirical induction? Thus, I think it's reasonable for someone who thinks it's possible but unlikely that the ZF axioms of set theory are inconsistent to continue using them (and the same goes for PA, though maybe it's harder to imagine what it would feel like to be uncertain about this).

Secondly, I don't think one is even committed to believing that some mathematical system whose axioms you can list and say you are 'working in' is (precisely) consistent. For there have been shown to be inconsistent systems where the shortest proof of a contradiction is very very long, so that even though the axioms are inconsistent there are still definite patterns in what it's possible to prove reasonably briefly and hence (presumably) in what people will ever actually prove in that system. [find this citation] Thus you might keep open the possibility that whatever system you are working in is one of these if you think that the facts about what can be proved reasonably briefly in this system do correctly reflect the facts about how many fruit n apples and m oranges constitute (or whatever your intended application was). In fact, I suspect that one can cook up some recursive axioms which allow short proofs of exactly those arithmetical claims which PA allows, but a long proof of 0=1 and hence all arithmetical claims. In other words, someone could reasonably suspend judgment about whether their mathematical axioms might be subtly inconsistent, while still expecting that they won't ever derive both 27+38=66 and 27+38=67 or other such incompatible results and hence continuing to employ this system. Of course, it would be irrational to think a collection of axioms were true if you thought they could well be inconsistent in this complicated way. But, if we really want to see what kind of consistency claim the mere willingness to employ arithmetic in the ordinary way commits one to, we must admit that someone could rationally draw empirical consequences from the fact that they actually found a proof of something in a given system without ruling out the possibility that everything is provable in principle from that system. You might think there are useful patterns in what statements people can, practically speaking, prove, while being open to the possibility that your mathematical axioms include a subtle inconsistency which renders everything provable in principle.

More generally, we should note that there are a bunch of different ways of describing what we do when we make mathematical calculations and then form expectations about apples and oranges or bridges or concrete spheres. The very same activity of a person writing down a proof can be thought of as writing a proof in ZF, a proof in naive set theory, a proof in various restrictions of PA which forbid you to use comprehension for certain kinds of formulas (formulas which don't happen to occur in the particular argument the person is writing down) etc. And, so far as the rationality of forming a belief in this way goes, it's clearly not required that you think every possible way of describing your actions correspond to a reliable method of forming beliefs.

And in particular, I claim that someone can sufficiently account for the rationality of their using PA to come to conclusions about apples and oranges, by making out that they are warranted in believing that there was a reliable correlation between the kinds of proofs in PA which they can (practically speaking) find and the facts about how many fruit various numbers of apples and oranges make. This is quite compatible with their not being certain that the above correlation exists, and with their suspending judgment about whether there is a proof of inconsistency in PA that is so complex that they would never actually find it. Potter is right that the empirical applications of mathematics pose an epistemic problem for everyone. But the problem is not, strictly speaking, to explain how can be certain that the axioms in our math textbooks are consistent, but how we can be warranted in assuming that they won't let us prove everything.

4 empirical justification for our applications of mathematics

But, now that we have seen concretely what kind of consistency claim any philosophy of math owes us an epistemic 'account' of, and how this demand arises from Potter's applications of arithmetic to statements about apples and oranges, it turns out (as I will now argue) that pretty much any philosophy of mathematics can provide the required account fairly easily. For, as suggested above, I can be justified in forming empirical expectations on the basis of producing proofs in a certain formal system if I merely know there's a correlation between two physical systems: me writing the proofs, and the apples and oranges which I am drawing a conclusion about. I just have to think that there's a negative correlation between there actually being n apples and m oranges which make up p fruit and my 'proving' in the system that n+m=p* for some number p* other than p. But these are a pair of physical systems which I have as much observational access to as anything else. I can observe the number of apples-or-oranges in a basket, and the number of apples, and the facts about what statements I (and others in my community) tend to prove when working in the system just as much as I can observe the positions of the planets or the typical size of a litter of dogs.

Of course, none of these observations are infallible. I might have dreamed that I derived a given sum or miscopied something, or someone might steal one of the apples in the basket before I get to counting it. There are even ways the world could be such that apples always disappeared when people tried to count them, or people constantly dreamed about doing arithmetic and getting sums which they didn't (and couldn't) actually get. But this is nothing new and applies quite as much to claims about observing the location of planets (your telescope could be broken, there could be some basic law of physics which made all telescopes function very differently when turned towards the sky) as to observing the number of apples in a basket, or what sentences people claim to have proved.

So it seems like we have a body of intuitively observable facts which exhibit definite patterns. There are definite patters in what sentences of the form 'm+n=p' people have 'proved' using arithmetic, and there are definite patterns in how many apples-or-oranges there tend to be in baskets containing m apples and n oranges. And, lo and behold, these patterns match up. So, prima facie, any philosopher of math who accepts empirical induction has reason to expect the kind of correlation between what he actually proves using the axioms of arithmetic and the facts about apples and oranges which he forms expectations about on the basis of coming up with some proofs which would justify him in using arithmetic in the way that Potter indicates. One might object that this kind of cheap 'account' of the applications of arithmetic requires applying empirical induction to mathematics, which some people follow Frege in taking to be unwarranted. But this is not the case. For note that the desired conclusion is not one about what proofs are possible, rather it's a conclusion about what proofs in PA people will actually find/write down. If human beings go extinct at some point this will be a finite collection. And in any case it seems intuitively clear that one can make empirical inductions about e.g. the content of American autobiographies, so it's hard to see how 'actual proof inscription' could fail to be an inductible predicate where 'novel inscription' succeeds.

Thus if (as I take Potter to be claiming) the epistemological problem for mathematics is to account for how we have the kind of knowledge of consistency which makes it reasonable for us to come to conclusions about apples and oranges on the basis of mathematical proofs, then it turns out this problem isn't very hard, and doesn't constrain one's overall philosophy of math very much. For, even without even appealing to the Platonist's mathematical intuition or the Fictionalist's modal intuition, everyone can say that we have inductive reason to expect that we won't come up with proofs in PA of both 27+38=66 and 27+38=67, and to expect a correlations between the claims which we will actually find proofs of in PA and facts about how many apples and oranges there are in certain situations. Even the hardcore formalist who doesn't understand consistency can claim to have noticed this correlation between results people actually get in the mathematical game and observations of apples and oranges - and thereby justify his behavior.

The story so far has perhaps been a bit abstract, and maybe complicated by the fact that it's hard for many people to even imagine doubting that PA is consistent and hence hard to evaluate claims about what weaker claims would warrant every day applications of arithmetic, or what evidence we can have for such weaker claims. Thus, I'd like to dramatize the general point that applications of math don't commit you to any stronger consistency claims than empirical evidence can provide by telling this historically inspired story about Berkeley and the Bridges. Berkeley, as you know, was critical of the Calculous, and wrote that it involved shoddy reasoning about infintesimals which he called the 'ghosts of departed quantities'. And indeed some of his objections are considered to be quite reasonable and were only answered by much later efforts to make the calculus formal and rigorous. However, he also had to explain the fact that when people used the Calculous to determine volumes and distances traveled etc. they seemed to get the right answer. He did not (as a logical positivist might have expected) deny the validity of all the empirical observations which seemed to conflict with his mathematical belief that the methods of calculus were wrong. Rater, he proposed that although the methods of the Calculus were faulty - and hence allowed one to derive wrong answers - but when people actually attempted to employ these methods they systematically made "compensatory errors" in such a way as to end up with the right end result. Given this further belief, I think it's pretty clear that Berkeley would (and should) have strongly preferred to walk on bridges that engineers using the calculus predicted would stand up to ones that engineers using the calculus predicted would fall down. He had empirical reason (at least) to think that people nearly always used the calculus in such a way as to make the right kind of compensatory errors and hence get the right answer. Thus, would have had good reason to think that the engineer assessing the bridges would come to the right conclusions when calculating the weight on the spans, the total force of wind and water on the stones etc.

I don't venture to say anything historical about the nature of Berkeley's actual criticisms of the calculus here, but let us imagine that he thought the methods of the calculus were inconsistent in the sense that they allowed one to 'derive' a contradiction which allowed one to derive everything. And, let us suppose, that the most natural formalization of the un-rigorous statements in calculus textbooks of the time really were inconsistent in this way. Thus I think we can imagine him as having justification, indeed knowledge, that the axioms which mathematicians of his time would have given if asked to state what system they were working in were inconsistent. However, suppose that he also had lots of empirical experience in which people employing the calculus seemed to get the right result (or very close to it) e.g. cases where separate measurements confirmed the results of using calculus to estimate the volume of a nearly spherical object with a certain radius, or the distance traveled by something accelerating from standstill at a certain rate. And, suppose he came up with the theory of compensatory errors to explain this kind observation, rather than in some way that would let him formally specify when and how people made the compensatory errors which let them get the right results from using calculus.

In this case, I claim, we have a person who has strong reasons for employing certain mathematical axioms without any commitment to the consistency of these axioms (or any weaker set of axioms which they could explicitly produce). Our imaginary Bishop Berkeley would be quite irrational not to try to use calculus naively to build bridges (or hire someone else to do so if he doubts his capacities to maintain naivite) and yet he knows that the formal system which he or the hired engineer would have to appeal to if asked to axiomatize what they were doing was inconsistent. Thus, it is not the case that using mathematics to form conclusions about empirical objects commits one to more by way of consistency than the expectation that your actual attempts to employ the system are likely to turn out correctly - and this very weak claim can be justified by observation and scientific induction.

Similarly, we can imagine some kind of hardcore formalist or empiricist taking the same attitude toward PA. They might say 'I don't know anything about what PA lets one prove in principle, (indeed the formalist would say he doesn't even know what this expression 'in principle' means) but I do have strong empirical reasons for expecting that my actual employment of the formal game of PA won't involve my proving both 2+2=4 and 2+2=5, and will continue to yield correct results about numbers of apples and oranges as it has always done in the past.

So, to the extent that our willingness to form expectations about apples and oranges and bridges based on mathematical calculations creates a general problem of mathematical knowledge, this problem is easily solved by appeal to empirical observation and employing scientific induction to the effect that our ordinary use of this practice seems to have led to results that cohered with one another and observation in the past. This very weak empirical surrogate for consistency (people don't actually seem to wind up proving everything so there are definite patterns in what a person trying to work in the system is likely to prove) is all that's needed to justify using mathematics.

sectionharder epistemic problems about consistency I have just argued that our actually using of mathematics to draw conclusions about the external world can be justified on merely empirical grounds, so this doesn't constitute any kind of epistemological problem about mathematics. So, perhaps Potter has something more in mind. One possibility is that Potter might take it as a starting point that contradiction couldn't be derived from our mathematical axioms even in principle - and hence require that an adequate account of this knowledge (rather than just a justification for our empirical applications of mathematics). Another is that the intended problem is not one of justification (we can empirically see that currently accepted mathematical doctrines seem to be consistent or close enough to it for our empirical applications of them to be reliable) but of explanation. Perhaps this would make sense of his idea that the Platonist has something special to say. If we get our mathematical doctrines by observing the Forms then this not only gives us some reason to think people won't derive a contradiction from the axioms we currently accept (as the above-mentioned empirical observations might) but explains how we could have gotten a (close to) consistent system in the first place. In contrast, if you think it's random which mathematical axioms we accept, it might seem rather mysterious that we have managed to adopt ones which turn out empirically to be close enough to consistent to make empirical applications a success. I will argue that a similarly empirical (if somewhat more contentious) answer is available to these harder questions about our knowledge of the consistency of mathematical axiom systems.

The first (not too exciting) step is to posit a faculty of rational intuition into whatever you take the ground of mathematical facts to be - mathematical objects, facts about what structures are possible, facts about what following a certain collection of rules will lead you to do etc. That is, you note that people do have the inclination to say that 'the right answer' to 'can contradiction be derived from the PA axioms?' is 'no' (whatever you take this fact to amount to) and say that people are justified in taking this feeling at face value as genuinely telling them about what proofs exist, or could exist, or what following certain formal rules could lead to etc. You can complicate this story by adding in mental pictures (people have a mental picture of e.g. some dots extending off to the right "forever", and then an inclination to think that this mental picture constitutes evidence that the PA axioms are consistent) as desired.

This lets you give the immediate, flat-footed answer to the question of 'how do we know that mathematical axioms systems are consistent?' that everyone probably wants to give: you know that systems are consistent by giving proofs and intuitive arguments which your faculty of mathematical/modal intuition allows you to recognize and be convinced by. And it also lets you give a flat-footed explanation how we have mathematical practices which don't let us infer everything - our rational mathematical/modal intuition allows us to preferentially adopt consistent systems and/or eliminate inconsistent ones. For the Platonist, this works by our adopting axioms which our mathematical rational intuition tells us truly describe abstract objects. For the Fictionalist and moderate Formalist this works by way of our rejecting fictions which our modal intuition tells us don't correspond to a way that things could possibly be/formal games that our modal intuition tells us are boring in the sense that they allow one to arrive at anything. These faculties aren't infallible (as witness naive set theory) but they seem to be pretty good.

But now we get to the interesting questions. How could people have any such faculty of rational intuition that yields accurate verdicts about which systems are truly consistent or even which ones don't actually lead people who try to employ them to prove everything? The first thing is that (as we saw earlier) facts about what people behaving in accordance with certain rules (e.g. creating proofs in a syntactically specified mathematical system) will actually arrive at are as concretely observable as facts about the motion of the planets or the diet of bears. In both cases one needs bridge principles. You can be wrong about whether a person only wrote 'avb' in cases where they wrote 'a' or 'b' on some earlier line, and you can be wrong about whether a certain planet was at a certain location at a certain time.

But in both cases we are talking about concrete facts about the behavior of observable objects (planets or people writing down marks on paper). Thus, intuitions about whether a system that behaves in accordance with rule R can in principle yield result X can be honed and corrected by observation of what actual systems which behave in accordance with rule R actually do yield. These expectations about what a system acting in accordance with rule R can do might be consciously revised, or might arise and be corrected entirely at the level of unconscious habit. That is, you might stop expecting everything which behaves like R to do X after seeing enough examples of things that behave like R not doing X, without ever consciously articulating any principle about systems which behave like R to yourself - we would be hard pressed to articulate the rules behind many of our physical intuitions about how balls and coffee cups will behave, but one can still make sense of these principles being corrected by experience. Maybe even animals have some kind of faculty of forming expectations which gets corrected in this way so that e.g. they form correct expectations about where an insect will end up if it continues to move in accordance with a certain pattern that it has been observed to follow so far.

Thus we can tell a perfectly naturalistic story about how there could be selection against certain kinds of inaccuracy in our rational intuitions about what a system that proceeds in accordance with certain kinds of rules can wind up doing. But remember our distinction between knowledge about what a system can do in principle vs. knowledge about what it won't ever actually do. Strictly speaking, all this kind of naturalistic story gives us is selection in favor of coming to the kind of conclusions about what a system can do in principle which don't conflict with observation of what it actually does do. And presumably there are plenty of contingent restrictions on what one will actually encounter which have nothing to do with what's possible in principle e.g. one only encounters proofs that are fairly short.

This brings us to the next piece of the account, which is the claim that selection for the kind of mathematical/modal faculty which coheres with observations about what actually happens, could easily result in the production of a faculty which is reasonably accurate about what can in principle happen. This is, ultimately, a claim about the basic physics or chemistry of our world. The thought is that the simplest ways of building a calculator which gets sums of numbers less than, say, 140 right is to build one that uses algorithms which would get much larger sums right, so if you knew that a calculator was assembled by a random physical jumbling of parts and that it got the sums of numbers less than 140 right, it would not seem miraculous if it got the larger sums right too. So, if we think that the mechanisms which produce our modal/mathematical intuitions are corrected by experience in such a way as to ensure that they get the right answer about how many fruit a basket which contains n apples and m oranges must contain for the small numbers we actually encounter, it would not be surprising if they got larger sums right too. Similarly it's plausible that in getting the kind of intuitions about what it's possible for a machine/physical system to do so long as it only moves in certain ways which cohere with observations of what such a system actually does, we should get intuitions which are more generally accurate as well.

5 Conclusion

Thus, I claim, we have a good account of how we could have the kind of faculty of rational intuitions which lead us to correct conclusions about what a person employing a certain mathematical system will actually do, (and hence the kind of knowledge of consistency which rationalizes empirical applications of a given axiom system) and a plausible account of how these same faculties of rational intuition could also largely lead us to correct conclusions about what such a person could do in principle - and hence about consistency full stop. Whether Potter's 'problem of mathematical epistemology' is one of justifying our applications of mathematics, or of explaining how we could ever get knowledge of consistency (as we saw these turned out to be slightly different) mere consideration of the way that modal facts -e.g. about what proofs are possible- relate to concrete historical facts -e.g. about what proof-inscriptions will ever be actual- provides an answer. There's no philosophical problem about how we could ever be in a position to get knowledge of consistency facts, or how we could have the kind of rational intuitions about what proofs are possible/what axiom systems correspond to 'a way things genuinely could be' that lead us to form correct expectations/make good inferences about where inconsistency will and won't be found.

This leaves us with a mathematical project and an ontological question. The mathematical project (which, of course, people have already long been working on) is to actually use our faculties of rational intuition to learn about what systems are consistent by creating relative consistency proofs, reviewing the known consequences of axioms which imply the consistency vs. inconsistency of a given system, constructing intuitive pictures of how a certain collection of axioms could all be true etc. It's one thing to account (via optics, anatomy and evolutionary biology) for how we could have a faculty like vision which leads us to largely correct conclusions about the arrangement of physical objects in the world. But positing such a faculty, and giving a naturalized epistemology which fits our possession of such a faculty into our general picture of the world, is just the beginning of the story. Presumably, in addition to being able to give a good philosophical defense for our taking ourselves to be able to see things, one then wants to actually use the faculty of sight to discover things. Similarly, answering this kind of epistemological challenge about how we could ever get knowledge of consistency would be in vain if we didn't also want to actually use our faculties of mathematical/modal intuition to learn things about what is provable in various systems.

The ontological question is whether, given that we can justify our applications of mathematics by positing 'rational intuitions' which embody merely scientific/inductive knowledge of what people will actually prove in a given system, we should take these intuitions to also give us knowledge of what proofs inscriptions are possible (as per the Fictionalist or the Structuralist) or what proofs exist (as per the Platonist). This is a huge question, which I won't try to answer here. However, I'd like to briefly suggest that the issue between those who think there are abstract objects called 'proofs' corresponding to all in principle possibilities of proving something and those who think there are instead primitive modal facts about what it is possible to prove is highly analogous to the question of whether we should take there to be possible worlds corresponding to all possible ways the world could be. That is, on the one hand it can be quite useful and simplifying to quantify over all possible proofs and all possible worlds (and messy to paraphrase them away) so if you are Quinean about ontology you have some reason to accept them. Also if you aren't a Quinean but you want to give a semantic theory which associates the truth conditions of sentences like 'you can prove that there are infinitely many primes in PA' with an object that is supposed to be the truth-maker of this claim, you have reason to accept abstract objects 'witnessing' the possibility of a scenario or of proving something in a formal system. On the other hand, on a more intuitive notion of ontology, claims such as 'there are no unicorns' or 'there could be unicorns' or 'wearing a white tie is more formal than wearing a black tie' can have determinate, objective, truth values without involving any kind of object as a truth maker.

I'm not sure which position on the existence of proofs as abstract objects is right (or whether the dispute turns on some kind of ambiguity). But if you think the possibility of a proof must be witnessed by an abstract object but the possibility of a political system need not, I would like to know why.