

Ways a World Might Be

Robert C. Stalnaker

Print publication date: 2003

Print ISBN-13: 9780199251483

Published to Oxford Scholarship Online: Apr-05

DOI: 10.1093/0199251487.001.0001

On what Possible Worlds could not be (1996)

Abstract and Keywords

This paper explores the analogy between mathematical Platonism and modal realism, and between Benacerraf's dilemma and the epistemological objection. It is argued that the parallels and contrasts may clarify both modal realism and the general problem of model epistemology. The paper begins with a sketch of Benacerraf's reasons for thinking that there is a prima facie conflict between a straightforward account of mathematical truth and a reasonable account of mathematical knowledge. It then develops presents a strategy to respond to the dilemma, and argues that this does not suggest a parallel response to the epistemological objection to modal realism. Finally, a more general problem for an epistemology of necessary truth is examined.

Keywords: mathematical Platonism, modal realism, Paul Benacerraf's dilemma, mathematical truth, mathematical knowledge, epistemological objection, truth

DOI:10.1093/0199251487.003.0003

Abstract and Keywords

This paper explores the analogy between mathematical Platonism and modal realism, and between Benacerraf's dilemma and the epistemological objection. It is argued that the parallels and contrasts may clarify both modal realism and the general problem of model epistemology. The paper begins with a sketch of Benacerraf's reasons for thinking that there is a prima facie conflict between a straightforward account of mathematical truth and a reasonable account of mathematical knowledge. It then develops presents a strategy to respond to the dilemma, and argues that this does not suggest a parallel response to the epistemological objection to modal realism. Finally, a more general problem for an epistemology of necessary truth is examined.

Keywords: mathematical Platonism, modal realism, Paul Benacerraf's dilemma, mathematical truth, mathematical knowledge, epistemological objection, truth

1. Introduction

According to David Lewis's modal realism,¹ the actual world is only a proper part of a reality that consists of many parallel universes that are spatially and temporally disconnected from each other. The actual world is defined indexically as that part of reality that stands in some spatio-temporal relation to ourselves. Talk about what is merely possible, according to this thesis, is not talk about what might have existed but does not; instead, it is talk about what does exist, but only in some different part of reality than the one we inhabit. Merely possible cabbages, kings, and croquet games are things of the same kind as the actual cabbages, kings, and croquet games that inhabit our world; their nonactuality consists simply in their being located somewhere else.

Modal realism is really two theses, one semantic and one metaphysical. The semantic thesis is that claims about what is possible or necessary should be analyzed as claims about what is true in some or all of the appropriately individuated parts of reality. The metaphysical thesis is that there exists a plurality of parts of reality individuated in this way—a plurality that is full enough to make true many of our ordinary beliefs about what is possible and necessary when those beliefs are interpreted in accordance with the semantic thesis.

It is a familiar objection to the semantic part of the thesis of modal realism that if it were true, then it would not be possible to know any of the facts about what is merely possible, or to have any reason to believe any modal claims about what is merely possible. If the semantic part of the thesis were true, it is argued, then the metaphysical part would be an unsupported and insupportable speculation. For example, it seems that we can know, or at least have reason to believe, such claims as that the Federal Reserve Board might not have voted to raise interest rates on a particular occasion on which they in fact did vote to raise interest rates. But on Lewis's analysis, this claim entails that there is a parallel universe, spatio-temporally disconnected from ours but similar to it, in which there is a Federal Reserve Board populated with counterparts of the members of the actual Board who (p. 41) met as the actual members did, but who voted not to raise interest rates. How could one know that, and why should the reasons we have for believing that the Board might not have raised interest rates be reasons for believing in the existence of such people and events?

This epistemological objection to modal realism may seem to parallel closely Paul Benacerraf's dilemma about mathematical truth and knowledge.² Benacerraf saw a tension between the demand for a plausible account of what mathematical statements say and a defensible account of the way we know that what such statements say is true. If we take the statements of, say, number theory at face value, then we will interpret them as statements about a domain of entities—the numbers—which lack spatio-

temporal location, and with which no one can be causally related. But if that is what mathematical statements are about, then it is not clear how it is possible to know that any of them are true. Lewis points out the parallel between the epistemological objection to modal realism and Benacerraf's dilemma, and exploits it in his response to that objection. We can agree, I think, that if modal realism is no worse off than Platonism about such things as sets and numbers, then it will find itself in better company than many critics suppose.

My aim in this paper is to explore the analogy between mathematical Platonism and modal realism, and between Benacerraf's dilemma and the epistemological objection. I think the parallels, as well as the contrasts, are illuminating, and may help to clarify both modal realism and the general problem of modal epistemology. I will first sketch Benacerraf's reasons, as I understand them, for thinking that there is a *prima facie* conflict between a straightforward account of mathematical truth and a reasonable account of mathematical knowledge, and Lewis's response to it. Second, I will sketch a strategy for responding to the dilemma, and argue that this strategy does not suggest a parallel response to the epistemological objection to modal realism. Finally, I will look at a more general problem for an epistemology of necessary truth.

2. The dilemma, and Lewis's response

It seems to be a reasonable constraint on an adequate philosophy of mathematics that it should explain mathematical statements as genuine statements making claims that are true or false. And we should expect a semantics for mathematical statements to be continuous with a semantics for statements in general. In particular, one should expect existential statements about sets, functions and numbers to be interpreted with the same referential truth#conditional semantics that we use to interpret existential statements about tables and chairs, quarks and photons, symphonies and (p. 42) rock bands. But on the other hand, we should also expect our account of mathematical knowledge to be continuous with our account of knowledge in general. The procedures we use to evaluate and justify mathematical statements should be explained by a general account of knowledge, together with our account of mathematical truth; we should be able to explain why, if mathematical statements say what our semantics says they say, those methods and procedures are appropriate ways of telling whether they are true. Benacerraf argues that some philosophical accounts of mathematics—Platonistic accounts—give a natural semantics but allow for no plausible epistemology; others—what he calls combinatorial accounts, such as formalism and conventionalism—may account for the mathematicians' methods and procedures, but fail to give a plausible account of what is said by the statements justified by those methods and procedures. Benacerraf poses the problem, but does not offer a solution. We are not invited just to opt for one or the other horn of the dilemma. Instead, it is suggested that we need

to find some kind of reconciliation—some account that respects both the prima facie account of mathematical truth and the need to explain mathematical knowledge. 'Ultimately I will argue that each kind of account has its merits and defects: each addresses itself to an important component of a coherent overall philosophic account of truth and knowledge.'³

Why, according to Benacerraf, is there a conflict between a straightforward Platonistic account of mathematical statements and the facts about the way we evaluate and justify those statements? The central reason is that the straightforward semantics explains truth in terms of reference or satisfaction, and (Benacerraf argues) genuine reference requires a causal connection between an act of reference and the object of reference. But the objects hypothesized by a Platonistic account of mathematical statements are not the kinds of things with which the speakers in a mathematical discourse, or anyone else, can be causally related. Lewis's response to Benacerraf's dilemma is to accept the Platonistic semantics, but to reject the assumption that reference requires a causal transaction with the referent. Then once the precedent is set, we can move on from numbers and functions to possible cabbages, kings, and croquet games. If reference and knowledge are sometimes possible in the absence of causal connection, one can no longer refute a thesis simply on the ground that it implies we can have knowledge of a certain domain without interacting with it.

The proponents of the epistemological argument against modal realism reject the parallel between mathematical objects and possibilities, realistically construed. They insist that reference to and knowledge of *concrete* things requires causal interaction, even if reference to and knowledge of *abstract* things does not. But, Lewis responds, why should the line between what (p. 43) requires causal interaction and what does not be drawn in terms of this distinction (a distinction, Lewis adds, that is none too clear in any case)? Instead, we should say that reference to and knowledge of *contingent* facts requires causal interaction, while knowledge about modal reality does not. In the context of modal realism, this is to say that indexical knowledge, or knowledge of one's own place in reality, requires causal interaction with that part of reality we are in, but impersonal knowledge does not. According to this line, a causal requirement for knowledge of an existential proposition depends, not on the kind of thing that is said to exist but only on whether it is claimed to exist in the knower's own part of reality. One might have knowledge, without causal connection, of the existence not only of numbers, sets, and functions, but also of objects, such as sidewalks, that are concrete.

3. Liberal Platonism

I am inclined to agree with Lewis that Platonism about mathematical objects does not imply that knowledge of such objects requires causal interaction with them, and so

that Benacerraf's dilemma, in the form in which it is presented, can be dissolved. But I still feel the force of the epistemological objection to modal realism. The challenge is to explain the disanalogy between reference to and knowledge of numbers and reference to and knowledge of cabbages, kings, and croquet games that are alleged to exist in spatio-temporally disconnected parts of reality. I once tried to meet this challenge by sketching a cheap or liberal version of Platonism, a way of reconciling a straightforward semantics with mathematical epistemology that could not be extended to a defense of modal realism. According to my liberal Platonist, one begins with the facts of mathematical discourse—the existence of a practice involving activities that seem to resemble assertion, inference, argument. One notes that if a Tarskian semantics is given for the products of the discourse (hypothesizing a domain of things of the kind that the discourse seems to be talking about), then that semantics helps to explain facts about the discourse, for example that people tend to try to assert and accept what is true, according to the semantics, and to make and endorse inferences that preserve truth. Now, the liberal says:

if we agree that this is a legitimate practice, that its (perhaps evolving) standards for evaluating the products of its discourse are relatively clear, then isn't this sufficient reason to accept that the practice really is a practice of making assertions, and its semantics really tells us what the statements say? Can't we, with a clear conscience, engage in this practice, and endorse some of the existential statements that meet its standards, without making some further supernatural assumptions about causal interaction with a realm of immaterial entities?

(p. 44) Then the liberal (incautiously) said, 'the existence of numbers *is just constituted by* the fact that there is a legitimate practice involving discourse with a certain structure, and that certain of the products of this discourse [statements that have an existential form] meet the standards of correctness that it sets.'⁴ Hartry Field rightly took my liberal Platonist to task for this last statement, which suggests that he has mislabeled his position, a position that Field suggested would be more accurately described as a kind of linguistic idealism.⁵ But my liberal's last statement was careless. The position that he was trying to espouse was in the spirit of Carnap's 'Empiricism, Semantics and Ontology'.⁶ In trying to defend the thesis by saying, from an external point of view, what numbers are, he distorted the view, and violated his Carnapian principles. I think the thesis, more carefully stated, really is a version of Platonism, and not a kind of linguistic idealism—at least it accepts the two theses that Field takes to define Platonism: that numbers, functions, and sets exist, and that they are 'mind-independent and language-independent'.⁷ Speaking from within the practice, one might ask counterfactual questions about the entities whose existence one has endorsed, and if one does, one will surely affirm that those entities would have existed even if no one had ever thought about them, or engaged in a practice of talking about them. The Carnapian, like all of us,

will say that a statement would not have been meaningful if the language of which it is a part had not existed, but that is compatible with agreeing that what it says might still have been true. ⁸

Our liberal Platonist need not follow Carnap in making a sharp distinction between internal and external questions—questions asked from within a framework, and purely practical questions about whether to adopt or endorse the framework. Let us suppose that our post-Carnapian liberal Platonist, having read Quine, recognizes that all questions are asked from within some linguistic context, and that the questions, ‘is it reasonable to adopt the framework of numbers?’ and ‘are there numbers?’ are not so easily separated. But we can still distinguish questions that are external in the sense that they are questions about the practice, or the framework, that is at issue. The kind of external perspective I have in mind is like the one (p. 45) Quine adopts at the beginning of *Word and Object* when he says: ‘I propose...to ponder our talk of physical phenomena as a physical phenomenon, and our scientific imaginings as activities within the world that we imagine.’ ⁹ Without pretending that we are stepping outside of the world, or setting aside our commitments, we can still describe ourselves, our discourse, and the fact that we make the commitments we make as features of the kind of world we believe ourselves to be in. It is in this kind of inquiry that the causal theory of reference has a place. When we ask: ‘how must a speaker’s use of a name be related to a person or physical object in order for the object to be the referent of the name?’ we are pondering our acts of referring as phenomena within the world to which we refer. Now when our liberal Platonist ponders his and our talk of mathematical objects as an activity within the world that we imagine, he finds that it is an activity with reasonably well-defined rules and procedures. He further finds that it is an activity that is usefully described and systematized as a practice that involves making assertions and inferences using expressions that have a Tarskian semantics. Since the practice in question is one that the liberal Platonist engages in, and since he is a good Quinean who believes that to be is to be the value of a bound variable, he feels free to refer to the members of the domain—the numbers—in giving the semantics for the discourse. But now we ask, still viewing the practice as a phenomenon that is a subject of empirical inquiry, ‘how must a speaker be related to the numbers in order to be in a position to refer to them?’ It is with this question that the disanalogy between mathematical and empirical discourse emerges. The liberal Platonist says that the legitimacy of the practice does not presuppose that there is any story to be told, in our empirical theory of the practice, about an external relation between acts of reference and the referents. The liberal was tempted to put the point by saying that the existence of numbers is just constituted by the fact that the practice is a legitimate one with a certain structure and well-defined standards, but this had the false consequence that numbers were contingent objects whose existence was dependent on the existence of the activity of talking about numbers. The point is rather

that the commitment to numbers is just constituted by the endorsement of the practice with this structure and standards.

One might be inclined to think that the liberal Platonist is trying to occupy a middle ground that is not really there—that his position is simply an equivocation between a thesis that is not really realism, and a version of Platonism that is not distinguishable from Lewis's—a thesis about mathematical objects that can provide a precedent for a defense of modal realism against the epistemological objection. Let me address the two sides of this objection in turn.

(p. 46) First, is the liberal Platonist's thesis really realism, rather than some kind of fictionalism, instrumentalism, or quasi#realism? I don't think it would be appropriately assimilated to fictionalism, since a fictionalist position requires a commitment to the intelligibility of a thesis that, if true, would make the fiction fact. The liberal might grant the intelligibility of something like Gödelian Platonism—the thesis that there is a realm of entities with which we have causal contact through a quasi#perceptual faculty of mathematical intuition—but it would seem gratuitous to interpret the liberal's thesis as the thesis that we are pretending that some such theory is correct. The liberal denies not only that such a story is true, he denies that it is needed to make sense of the content of mathematical claims.

More plausibly, one might argue that the liberal is really an instrumentalist who does not take mathematical truth seriously. Benacerraf notes, in his criticism of conventionalism, that to give a real truth#conditional semantics for a language, it does not suffice simply to sprinkle the word 'true' onto some well#defined subset of the sentences. 'Suppose now that the language is set theory, in some first order formalization....Let those [sentences] with an even number of horseshoes be "true".' Such a trivial sprinkling might, Benacerraf notes, be 'parlayed into a truth theory', at least a theory that had the form of a truth theory. But to be a theory that is really about truth, 'the explanation must proceed through reference and satisfaction and, furthermore, must be supplemented with an account of reference itself'.¹⁰ What is required, Benacerraf suggests, is the kind of external explanation of the relation between acts of reference and referents that the liberal Platonist denies is required.

Now I agree that the matching of the word 'true' with expressions, and even the possibility of formulating an account of the matching in a recursive theory that has the form of a Tarskian truth definition, is not sufficient to justify saying that we are really talking about truth. But it is not clear that it is a causal account of reference that is what is required to distinguish genuine from counterfeit semantics. One might instead focus on the practice of assertion. One needs, for a division of expressions into two classes to count as an account of truth, to connect that assignment with the way the expressions are used in some practice that involves something that looks like assertion

and reasoning—some explanation of the role of the distinction between the sentences matched with 'true' and those that are not in some kind of cognitive and communicative enterprise. The account of the epistemology of mathematics answers to this demand, and suffices, the liberal Platonist argues, to justify the straightforward truth#conditional semantics.

One might argue that simply because the liberal Platonist rejects the demand for any answer to the external question about the relation between his acts of 'reference' and the objects of reference, his position is best called (p. 47) by another name—'quasi#realism' rather than 'realism'.¹¹ To this suggestion, the liberal Platonist might reply, with some exasperation: 'This is just the kind of issue that Carnap taught us was meaningless. We agree about what to do: about what practices are legitimate, about what to say within those practices, and about what to say in answer to the external questions about the practice as an activity within the world. All we disagree about is whether to call it "realism". Have it your way: if you want to deny that I am making an ontological commitment, I won't object. I have no stake in any metaphysical label; I simply want to defend the legitimacy of the practice, and of the theoretical account of it provided by the straightforward semantics.' Whether it is called 'realism' or 'quasi#realism' the liberal Platonist's thesis does seem to meet Benacerraf's requirement that sentences such as 'There are at least three large cities older than New York' be seen to have the same semantic structure as sentences such as 'There are at least three perfect numbers greater than 17',¹² or more generally that 'the semantical apparatus of mathematics be seen as part and parcel of that of the natural language in which it is done'.¹³

If we succeed in defending the liberal Platonist against the charge that he isn't really a Platonist, we will face the charge that he is not really a liberal. We have not yet succeeded in distinguishing our Platonist's way of rejecting a causal requirement for reference to and knowledge of mathematical objects from Lewis's way of rejecting a causal requirement for reference to and knowledge of possible things and worlds. 'Your liberal responds to Benacerraf's dilemma exactly as I do,' Lewis might say:

by rejecting the general demand for a story about how *this* kind of evidence can be evidence for *that* kind of claim. My view [Lewis might continue] is that epistemology should be naturalized epistemology. We should do it by describing the methods that we in fact use to form and justify our opinions of various kinds, by systematizing these methods and bringing them into reflective equilibrium. If our most plausible semantic hypothesis tells us that a certain bit of discourse is about things that speakers are not causally connected to, and if epistemological principles that seem on reflection to be reasonable tell us that we can know about such things, then there is no reason that we should not accept both the semantics and the epistemology. Isn't that

just what your liberal Platonist thinks? Why, then, should he change his tune when he turns from mathematics to modal realism? ¹⁴

(p. 48) Of course Lewis will agree that the liberal Platonist might, in assessing the costs and benefits—in bringing his epistemological and semantic beliefs into reflective equilibrium—decide that numbers pass the test—are worthy of belief—while possibilities, realistically construed, do not. But, I think Lewis would argue, he cannot use the general epistemological objection to justify this assessment without making some principled distinction between the two kinds of non-contingent objects.

I think there is a difference between the liberal Platonist's defense of existential commitment without a causal connection and Lewis's, but it is difficult to pin down. The source of the difference is that my liberal Platonist, unlike Lewis, is still, in his heart, a verificationist. He has given up attempts to articulate an empiricist criterion of meaningfulness, since he recognizes that there is no neutral, non-question-begging language in which to describe the procedures by which we verify and confirm statements. But he continues to hold onto the idea that answers to questions about how we justify our beliefs, and about how we must be related to things in order to have beliefs about them, tell us something about the contents of the beliefs, and about the natures of the things. ¹⁵ The liberal Platonist cannot say what he is tempted to say—that the existence of numbers is just constituted by the fact that there is a legitimate practice involving discourse with a certain structure, since that statement distorts the commitments of the practice by implying that numbers are contingent objects that would not exist under certain counterfactual circumstances. But he can say that commitment to numbers is just constituted by an endorsement of the legitimacy of the practice, and that the account of the epistemology of number theory—of how we are related to the numbers we know about—tells us something, not only about the conditions under which it is reasonable to believe in those things, but also about the extent and nature of the commitment that one makes when one endorses the practice, and affirms the belief. The answer to the Quinean external question about how we must be related to objects in order to refer to and know about them should be seen as part of an answer to a question about what those objects are like. One is therefore not free to combine the epistemological account of the way we must be related to numbers in order to know about them with just any account of the nature of the numbers themselves.

Lewis will agree, of course, that the story told by a theory about the nature of the objects to which the theory is committed must be compatible (p. 49) with the story told about the way the speakers affirming the theory are related to those objects. But just what constraints does the ontology put on the epistemology? I am inclined to say that what distinguishes the liberal Platonist from the modal realist is that the latter, but not the former, makes specific claims about the nature of the entities whose existence is

affirmed that are incompatible with the account of the way knowers are related to the objects. The modal realist insists that possible persons and physical objects are persons and physical objects, things of the same kind as actual persons and physical objects, things that exist in space and time (though not, in all cases, our space and time) and that interact causally with other persons and physical objects. But those are the sorts of things knowledge of which requires a kind of explanation—an answer to the Quinean external question—that the modal realist's theory does not allow for. It is the liberal Platonist's verificationist principles that lend support to this last claim. He holds that we understand what it is to be a physical object or a person partly by understanding how one must be related to one in order to know things about it. That, I think, is why there seems to be a conflict between the insistence that *possibilia* are things of the same kind as contingent physical objects, and the claim that they are things whose existence we can know about in the way we know about things of a very different kind, such as numbers, functions, and sets.

Of course Lewis will reject the assumption that there is a conceptual connection between the way we know about things and our conception of the nature of those things. In the end, I have to grant that I have no general principle that the modal realist must accept that will permit liberal Platonism about numbers, but exclude commitment to real but merely possible cabbages, kings, and croquet games. The modal realist will insist that we know about the existence of *all* persons and physical objects (actual or merely possible) in the same way, a way that requires no causal connection with them. Knowledge of the existence of *actual* things is distinctive, not because of the nature of the things, but because it is indexical knowledge—knowledge not just of the existence of the things, but of the fact that we stand in a particular relation to them. According to this picture, *all* impersonal ontological beliefs are to be justified on a priori grounds. Some may find this picture hard to swallow, and for those who do, I think that even if the liberal Platonist line I have sketched does not give us a decisive refutation of modal realism on its own terms, it does at least offer a plausible strategy for defending a kind of realism about mathematical objects that does not provide a precedent for an analogous defense of modal realism. But there remains an epistemological puzzle about how we know about possible worlds, whether one is a modal realist or not, and whether or not one is able to swallow Lewis's epistemological picture. I will conclude by sketching this more general problem about knowledge of necessary truth.

(p. 50) 4. Content and necessary truth

I suggested above that to distinguish a genuine truth#conditional semantics from a mere division of sentences into two classes, one labeled with the word 'true', one should focus not on the explanation of reference, but on the practice of assertion. It is because assertion is not simply the game of trying to name the true that an assignment of truth#

values to the sentences is not in itself sufficient to provide a basis for an account of this and other speech acts. An explanation of assertion requires a notion of the *contents* of the sentences used to make assertions—of the information that the sentences can be used to communicate. One thing wrong with Benacerraf's example of an arbitrary assignment of truth#values to the sentences of the language of set theory (based on the number of horseshoes in some formalization of those sentences) is that the 'semantics' assigns truth#values without telling us what the sentences are supposed to say. The 'truth' assignment tells us nothing about why one might 'assert' one or another of the sentences that one judges to be 'true', or what effect it might have if one did. But whatever the details of one's semantics, if all the sentences in the part of the language in question are either necessarily true or necessarily false, then it is not clear how to give an account of content or information for those sentences. A semantics for sentences that express contingent propositions will give rules that determine the truth#values for those sentences as a function of the facts, and so will say how those sentences might convey information about the facts on which their truth depends. But if the semantic rules determine truth#values unconditionally, then it is harder to see how the sentences might be used to say something about what they are about, and how the semantics might be distinguished from an arbitrary assignment of truth#values. Behind Benacerraf's dilemma is a problem more general than a problem about Platonistic ontology, a problem that could not be resolved either by simply accepting a deviant semantics, or by hypothesizing a Gödelian faculty of mathematical intuition, unless one of these moves was a way of explaining mathematical propositions as propositions that were in some way contingent. The problem is to explain how necessary truths can have content.

According to Lewis, one of the virtues of modal realism, one of the benefits to weigh against the costs, is that it provides us with an account of the contents of propositional attitudes and speech acts. Among all the possible worlds, there will be a certain subset—the doxastically or epistemically accessible possible worlds—that are the ones compatible with some agent's beliefs or knowledge. The set of doxastically accessible worlds should not be (p. 51) thought of as defined by the propositions that the agent believes. Rather, Lewis suggests, we should 'characterise the content of knowledge or belief from the outset in terms of something like the epistemically or doxastically accessible worlds'.¹⁶ An explanation of belief should begin with an account of how the facts about a person's internal functional organization and causal relations with his or her external environment determine a set of doxastically accessible worlds. Then a proposition is believed (at least implicitly) if it is true in all of those possible worlds.¹⁷ How, on this conception of belief, is one to explain the possibility of inconsistent belief, or of failures to believe all the necessary consequences of one's beliefs? We must recognize, Lewis says (and here I am in agreement with him), that belief attributions may describe a person's state of belief in some kind of indirect way. 'The connection of

belief sentences [sentences attributing belief] with belief as characterized by doxastic alternatives is complicated and multifarious.'¹⁸ Part of the answer to this question, according to Lewis, is that beliefs may be fragmented and compartmentalized: one may believe several propositions separately without having integrated them into a single belief state. But while this might explain how it is possible to believe several different propositions without believing some consequence of them, it cannot, by itself, explain how one might fail to believe a necessary truth. There is just one necessary truth, on this conception of content, and it is true in all possible worlds. Even if our beliefs are characterized by several nonintegrated sets of doxastic possibilities, a necessarily true proposition still must be true in all of them, so to account for the possibility of ignorance and error of necessary truths, we need to recognize a different way in which belief attribution is 'complicated and multifarious'. Failure to believe a necessary truth, Lewis suggests, is not failure to believe the proposition it expresses, but rather failure to believe a sentence expressing it. Doubt and error about a necessary truth expressed by some sentence depends on the believer not fully understanding what the sentence says.¹⁹

(p. 52) This notion of content, and this kind of explanation of attitude attributions in the case where the content clause expresses a necessary truth, suggest a perspective on propositions about modality itself that is different from the perspective suggested by the possible worlds story. This story seems to suggest a domain of objects that makes our modal propositions true or false in the way that the domain of animals makes our zoological statements true or false. There are zoological facts about the domain of animals. We might ask whether things meeting certain conditions are in this domain. (Are there any unicorns—or at least horse-like animals with one long horn growing out of their foreheads? No, nothing like that.) It might seem that the modal realist believes that in a similar way there are modal facts about the domain of possible worlds. Can't we ask about whether there are possible worlds meeting certain conditions? (Are there any worlds with water that has a chemical composition different from H₂O? No, nothing like that.) But the correct answers to all such questions will be necessary truths. The content of any such answer will be the same as the content of a sentence of the form *P* or not *P*. Sentences of this form obviously tell us nothing about the character of the domain of all possible worlds: no matter what that domain is like, the union of any subset of it with its complement will be the whole set, and this will suffice to ensure that provided that the sentence that goes in for *P* expresses *some* proposition, then the corresponding sentence of the form *P* or not *P* will characterize the whole domain of all possible worlds. If there were a doubt about whether a sentence of that form expressed a true proposition, it would be a doubt about what the sentence said—about whether it expressed the proposition that it in fact expressed—and not a doubt based on any ignorance of what the possible worlds are like. The same will be true for any sentence expressing a necessary truth. If it is true that there might have been talking

donkeys, or that it is necessary that water is H₂O, or that there are inaccessible cardinal numbers, then each of these claims expresses this same proposition, and so what these sentences say also tell us nothing about the character of the possible worlds. Any ignorance or mistake about the truth of such statements will be ignorance or mistake about the relation between the statement and the proposition, and not about the proposition itself. The idea of the possible worlds analysis of content is that the content of a statement should be understood in terms of the possibilities the statement rules out. On this kind of account, there can be no substantive characterization of the whole domain of possibilities, since a substantive characterization would be a characterization that distinguished the way the domain of all possible worlds is from the way it might instead have been. The Tractarian Wittgenstein had this conception of content, and he clearly recognized this problem posed by it for an account of the content of necessary truth and for a characterization of the space of possibilities. His talk of what can be shown, but not said, is an attempt to address the problem, but it doesn't really help (p. 53) since showings must have content as well as sayings. As Frank Ramsey said, 'what we can't say we can't say, and we can't whistle it either'.²⁰

According to Lewis's way of reconciling this account of content with the fact of ignorance and error about modal truths and deductive relations, we can fail to know a necessary truth, or falsely believe an impossible proposition, only by failing fully to understand some sentence that expresses it. Knowledge and ignorance of modal truths (truths that seem to be about the character of the possible worlds, as well truths of mathematics) are to be explained as semantic knowledge and ignorance—knowledge and ignorance of facts about the relation between the expressions we use to state modal truths and the content of those truths. If this is right, then there is no *information* about the character of modal reality itself, considered independently of the way we think and talk about it. The picture that the thesis of modal realism suggests of a domain of possible worlds, divided up into worlds, about which we can have substantive opinions, is thus difficult to reconcile with the conception of content that, according to Lewis, provides one of the benefits of the possible worlds theory.

Empiricists used to say that necessary truths were true by convention, or true in virtue of meaning, or true because of the relations between our ideas—not because of the nature of anything in reality. Lewis's modal realism seems to be a conception of modal truth that is at the opposite extreme. 'Why', Lewis once asked,

should it be human conventions that create and destroy facts about what is possible? All that human conventions can do is to select one verbal expression rather than another to enjoy the privilege of truth by virtue of the facts about the possibility of worlds. In just the same way, human convention selects one verbal expression rather than another to

enjoy the privilege of being truth by virtue of the facts about tomorrow's weather. ²¹

Conventions match sentences with propositions, but the propositions, whether necessary or contingent, are true independently of the conventions. But on Lewis's account of content, there is only one necessary proposition—the set of all possible worlds—and knowing that it is true—knowing that the actual world is in this set—does not require knowing any facts about modal reality. On Lewis's account of what we are doing when we attribute knowledge and ignorance of what appears to be modal fact, we are talking about knowledge and ignorance of the relations between sentences and the propositions they express. The conventionalist account of necessary truth and Lewis's are not so far apart as they at first appear to be.

I think this account of content helps explain both why mathematical Platonism is appealing, and why it has seemed problematic. Platonism seems to be a way of characterizing the content of mathematical statements—of giving them a subject matter—since it suggests a picture according to which (p. 54) they are contingent. If one thinks of mathematics as being about a domain of entities, then it is hard to resist the temptation to believe that the domain might not have existed. The Platonist's ontological commitment, one is tempted to think, is constituted by his exclusion of the possibility that there is no such domain of objects. But then the Platonist adds that the domain is one that exists necessarily, and so no such possibility has been excluded after all—no such commitment has been made. Perhaps it is better to say, as the liberal Platonist does, that the ontological commitment is not really a commitment to the exclusion of a possibility, and so is not a commitment that requires a defense in terms of some external relation between acts and objects of reference.

If true mathematical claims are necessarily true, then there are no facts, about a domain of entities, or about anything else, that make the proposition they express true. They would be true no matter what the facts were. But one may still be ignorant of the answers to questions about the consequences of the rules and procedures by which certain sentences are matched with the necessarily true and the necessarily false propositions (the proposition that is true if and only if the actual world is located somewhere in the space of all possible worlds, and the proposition that is false if and only if it is). Proof, calculation, construction, and other kinds of mathematical inquiry and argument are ways of answering this kind of question.

Lewis's account of content is separable from his modal realism. An actualist about possible worlds (one who believes that what are called possible worlds are not really worlds, but are possible states of the world—properties that the world might have had) can understand content in the same way, and can use the same strategies that Lewis uses to reconcile this conception of content with the fact of ignorance and error

about modal truths and deductive relations. I have, elsewhere, defended this account of content and the strategies Lewis endorses, and I think it suggests a plausible way to address Benacerraf's problem about how to reconcile a straightforward semantics for mathematical discourse with the facts about way the sentences interpreted by such a semantics are established and refuted. But I don't think this account of content, and these strategies, are compatible with modal realism. Recall that according to Lewis's theory, all purely objective, impersonal beliefs—all beliefs about what there is in reality and not about the believer's place in that reality—are beliefs whose contents are necessarily true or necessarily false. But the necessarily true proposition does not really say anything. The information conveyed by sentences that express it is information about something else—something not impersonal and objective, but about the relation between the believer's language and the world. It seems that according to a theory that includes both modal realism and the possible worlds account of content, reality as it is in itself is something one cannot talk or think about. I don't think this is a consequence that a realist should be happy with.

Notes:

- (1) See Lewis ([1986a](#)) for a systematic exposition and defense of modal realism.
- (2) Benacerraf ([1973](#)).
- (3) Benacerraf ([1973](#) : 666).
- (4) The liberal's incautious statement is taken from Stalnaker ([1988](#) : 119).
- (5) Field ([1989](#)).
- (6) Carnap ([1950](#)).
- (7) Field ([1989](#) : 1).
- (8) Compare: Barry Stroud ([1984](#) : 192) asks 'whether it is Carnap's view that the statement we can now make and understand about mountains in Africa would no longer be true if we abandoned the thing language, or would not have been true if we had never adopted it'. If Carnap were to allow himself to take such counterfactuals seriously, he surely would deny that the existence of characteristics of mountains in Africa depends on anyone's linguistic practices, and there is no reason why he could not say the same thing in the formal mode: that true statements that we in fact make would still be true even if we did not make them, or even have a language in which they could be made. Questions about the truth of those counterfactuals should be regarded by Carnap as

internal questions asked in the context of the thing language, and answered according to its rules.

(9) Quine (1960 : 5).

(10) Benacerraf (1973 : 677).

(11) The 'quasi#realist' is Blackburn's name for a person who, 'starting from a recognizably anti#realist position, finds himself progressively able to mimic the intellectual practices supposedly definitive of realism' (Blackburn 1993 : 15).

(12) Benacerraf (1973 : 663).

(13) Ibid. 666.

(14) Cf. Lewis (1986a : 113–15).

(15) Lewis makes clear, in a remark in a footnote, his lack of sympathy with the idea that there is some conceptual connection between our conception of the nature of a thing and our beliefs about how one may know about it. Commenting on a proposal of Michael Dummett to distinguish between abstract and concrete things in terms of 'how we could understand their names', Lewis says that even if such a distinction can be drawn, 'it tells us nothing directly about how the entities on opposite sides of that border differ in their nature. It is like saying that snakes are the animals we instinctively most fear—maybe so, but it tells us nothing about the nature of snakes' (Ibid. 82, n. 56).

(16) Ibid., 28.

(17) This is not quite right. In order to account for *de se*, or indexical beliefs—beliefs about one's time and place within a particular possible world—Lewis argues that contents of belief should be properties, or sets of centered worlds, rather than sets of possible worlds. But I think we can ignore this complication in this context without distorting the issues at hand.

(18) Ibid. 34.

(19) As Lewis makes clear, one can't use that#clauses to deny belief, or attribute ignorance, in a case where the subject has no idea what the sentence means. But when the subject has a 'pretty good idea' what a sentence means, but has failed to see that it expresses a necessary truth, then one may use the sentence, in indirect discourse, to describe the situation.

(20) Ramsey (1929/1990 : 146).

(21) Lewis (1969 : 207).

