NOÛS 46:1 (2012) 127–158

Rational Imagination and Modal Knowledge

Jonathan Ichikawa
The University of St Andrews

Benjamin Jarvis
Queen's University Belfast

Abstract

How do we know what's (metaphysically) possible and impossible? Arguments from Kripke and Putnam suggest that possibility is not merely a matter of (coherent) conceivability/imaginability. For example, we can coherently imagine that Hesperus and Phosphorus are distinct objects even though they are not possibly distinct. Despite this apparent problem, we suggest, nevertheless, that imagination plays an important role in an adequate modal epistemology. When we discover what is possible or what is impossible, we generally exploit important connections between what is possible and what we can coherently imagine. We can often come to knowledge of metaphysical modality a priori.

The Red Sox won last night, but they could have lost. Knowledge of the first conjunct is secured via perception; what of the second? We cannot *see* that the Red Sox could have lost in the same way we saw that they won. Of course, we could learn the modal fact by testimony; one can learn pretty much anything by testimony. Still, not everybody could have learned it by testimony; someone has to have learned it in some other way. But how? If not through testimony, how can we know actually-false propositions to be possible? This is a central question of modal epistemology.

We also know facts about impossibility: although it's possible that the Sox lost, it is impossible that they both won and lost—and we know this, too. An attractive explanation for this knowledge is that we can perform a *reductio* on the alleged possibility. But not everything we know to be impossible conceptually entails an obvious absurdity in this fashion. (More on conceptual entailment in §4 below. For now, think of conceptual entailments as

entailments that are transparent.) It is impossible, for instance, for Hesperus to be closer to the Earth than Phosphorus, or for some water sample to contain no hydrogen. Call these latter kinds of propositions—ones that do not conceptually entail absurdities—*conceptually possible* propositions. We know many conceptually possible propositions to be impossible; another central question of modal epistemology is to explain how we can have this knowledge.

In tackling these central questions throughout the history of modal epistemology, two ideas have emerged. The first is that our capacity for modal knowledge is closely tied to our faculty of *imagination*. The second is that our capacity for modal knowledge is partly explained by our *rational* capacities broadly construed.[1] Our project in this paper is to develop these ideas with (relatively) minimal commitments to show how a priori modal knowledge is possible. We call the resulting position "moderate modal rationalism."

Our moderate modal rationalism departs from other versions of modal rationalism developed recently. David Chalmers (2002) and Frank Jackson (1998), for instance, have developed forms of modal rationalism on the basis of a two-dimensionalist semantic framework. This framework is controversial; we mean to do without it. George Bealer (2002) and (2004) has developed another approach on the basis of his theory of intuitions. While our position may well be compatible with his account, we are wary of many of his commitments concerning intuitions. (The word 'intuition' does not appear further in our paper.)[2] One of the primary goals of this paper is to show that a version of modal rationalism can be developed without the substantive commitments of either of these two frameworks.[3]

Throughout our paper, we will rely on the notion of a *proposition*. For the purposes of this paper, we understand propositions as truth evaluable entities that are differentiated by the way that concepts are put together to grasp them—and hence more finely than the objects and properties the propositions are about. It should be possible, for those whose commitments require it, to reinterpret our comments in light of opposing (neo-Russellian and nominalist) views; we leave such exercise to the reader.

In the next section, we will tackle some preliminaries, at which point will we give an overview of the paper.

§1. Preliminaries

We contend that our imaginative capacities underlie our capacity for modal knowledge. But what sort of imaginative capacities do we have in mind? In discovering what is possible or necessary, surely we deploy imagination in a variety of ways.

For our purposes, we will begin by focusing on the sort of propositional attitudes we principally come to have when engaging with fictions. Even on first glance it is clear that these imaginings are relatively unfettered when it comes to content. While there are notable instances of imaginative resistance—we typically resist imagining, for instance, that killing someone for our own pleasure is morally right, even if the story we're reading says it is—these instances are the exception, rather than the rule, even in cases of fictions with impossible contents. People regularly engage imaginatively with fictions in which there are incidents of time travel that, upon some reflection, prove to be impossible. Moreover, philosophers have no difficulty engaging imaginatively with a thought experiment story in which *per impossibile* it is discovered that Hesperus is not Phosphorus after all.[4]

We contend that the propositional attitude we principally come to have when engaging with fictions is *supposing*. In our view, to say that someone imagines that *p* in response to fiction is roughly to commit oneself to the person's supposing that *p*, and, furthermore, drawing what are, from the person's point of view, immediately good inferences from this supposition without landing in an apparent absurdity. (For our purposes, "landing in an apparent absurdity" need not be inferring to outright contradictions. For certain propositions—to borrow an example from Stephen Yablo (2002), the proposition that someone found a five-fingered maple leaf that is also (simultaneously) shaped like a regular egg—one might infer that one should be able to visualize events sufficient for the truth of them, and, for the purposes of our rough account, this might *ipso facto* land one in an apparent absurdity if one finds oneself unable to carry out the corresponding offline visualization.) Imaginative resistance might occur when inferences that land the would-be imaginer in absurdities are altogether too immediate. It might also occur when, for one reason or another, the would-be imaginer disapproves (morally, aesthetically, etc.) of an author's invitation (via fictionalizing) to imagine the proposition in question.

For our purposes, then, when we talk about imaginative capacities underlying our capacity for modal knowledge, we are effectively proposing that engaging in *supposition* can be a guide to possibility and necessity. Henceforth, we will use 'imagine' and 'suppose' interchangeably, unless we state otherwise explicitly.[5] No doubt, there will be some who do not share our view about the connection between propositional imaginings in response to fiction and supposing; indeed, some may reject that supposition is a kind of imagination of any sort. We invite such individuals to treat our language use here as stipulative.

Supposing is, of course, wholly unfettered when it comes to content. So long as a supposition is merely for the sake of argument, we generally find ourselves capable of supposing any proposition that we can entertain.[6] Indeed, for the purposes of a *reductio*, we frequently suppose propositions that we believe to be absurd.

In light of these considerations, the Naïve Modal-Imagination Hypothesis (NMIH) is obviously problematic:

> NMIH: For any proposition <*p*>, <*p*> is possible if you can imagine/conceive <*p*>, and impossible if you cannot.[7]

Some philosophers have advocated something that sounds like NMIH, but if we understand 'imagining'/'conceiving' as supposing, NMIH is clearly false. (No doubt, these philosophers understand 'imagining' or 'conceiving' in some other way.) Thinkers can suppose any proposition that can entertain, but it is clearly not the case that every proposition they can entertain is possible.

Matters do not improve much when we alter NMIH to restrict it to "imagining" not in the sense of merely supposing, but in the previously discussed sense of "imaginative engagement with some fictional story"—supposition along with the drawing of obvious inferences. Consider NMIH*:

> NMIH*: For any proposition <*p*>, <*p*> is possible if you can (propositionally) imagine <*p*> in engagement with some fictional story in which <*p*> is true, and impossible if you cannot.

On our suggested view, (propositional) imagining in engagement with some fictional story requires not only supposing, but also drawing what are, from the thinker's point of view, immediately good inferences from this supposition without landing in an apparent absurdity. Unlike NMIH, NMIH* does not have the implication that every entertainable proposition is possible; for some suppositions, *e.g.* the supposition that the Red Sox both won and lost, drawing immediate inferences does quite quickly land us in apparent absurdities. According to NMIH*, these supposed propositions are impossible.

Nonetheless, NMIH* is also obviously inadequate. Someone could read a fictional story whose plot depends on spacetime's being curved and resist imagining that spacetime is curved. Indeed, he may find he cannot so imagine. Perhaps he does not have the relevant concepts to do so, or perhaps he thinks that supposing that spacetime is curved leads to absurd consequences. Obviously, this is not an indication that the proposition that spacetime is curved is impossible—for him or anyone else.

A philosophical egoist might imaginatively engage with a fictional story in which the protagonist *rightly* kills someone merely for a large sum of money without resistance. He finds drawing those inferences he takes to be immediately good from the supposition that someone rightly kills someone merely for a large sum of money does not land him in an absurdity apparent to him. Obviously, this is not sufficient to show that it is possible to rightly kill someone merely for a large sum of money.

It is almost certain that thinkers vary in what they can imagine in engagement with fictional stories, precisely because they vary in what they take to be immediately good inferences or when they take themselves to have landed in an absurdity. Possibility and necessity, however, do not vary according to the thinker.

These obstacles for NMIH and NMIH* suggest another proposal. Perhaps when it comes to possibility and necessity, what matters is not whether the proposition might be supposed, but rather whether it might *coherently* be supposed. Of course, a proposition might be coherently supposable whether or not any particular thinker is capable of supposing it. Some thinkers might not have the relevant concepts necessary for entertaining the proposition in question. More importantly, coherent supposition is an objective matter that is not relative to a thinker. We coherently suppose <*p*> when we suppose <*p*> and supposing <*p*> and drawing good inferences from this supposition could never *in fact* lead one to conclude a genuine absurdity, e.g. a blatant contradiction. It makes no difference whether any particular thinker would find these (good) inferences good or whether upon drawing these (good) inferences he would take himself to have concluded a genuine absurdity or not.

We will call the proposal under consideration the Strong Modal-Imagination Hypothesis (SMIH):

> SMIH: For any proposition <*p*>, <*p*> is possible if one might coherently imagine/conceive <*p*> and impossible if she could not.

(Again, 'imagining'/'conceiving' stands for supposing.) While SMIH has an air of plausibility, we contend it is also false. There is nothing incoherent about imagining that

> h: Hesperus is closer to the earth than is Phosphorus.

An agent who, in imagining this proposition, carried out all of the inferences to which he was thereby rationally committed could imagine <*h*> without concluding a genuine absurdity; nevertheless, <*h*> is impossible, because

> i: Hesperus is Phosphorus.

Since <*i*> is knowable only a posteriori, we cannot tell, merely from our ability to suppose and rationally infer from supposition, whether <*h*> is possible.

Is a connection between possibility and (coherent) imaginability therefore untenable, or would another conservative revision yield a plausible option? Our project is to suggest a conservative revision. The result will be a moderate modal rationalism that explains knowledge of possibility and necessity,

including the necessary a posteriori. It will also yield a framework in which a priori knowledge of metaphysical modality is frequently possible.

Here is an overview of the remainder of the paper: In §2, we defend the rejection of SMIH. In §§3–5, we introduce and defend a notion of conceptual possibility and necessity. In §§6–7, we relate conceptual modality to metaphysical modality, and show how the former can be useful in coming to knowledge of the latter. In §8, we argue that in many cases, the connection between conceptual modality and metaphysical modality can yield a priori knowledge of metaphysical modality.

§2. Imaginable Impossibilities

At least some philosophers have not been convinced by the argument from instances of the necessary a posteriori that SMIH must be false. They dispute that it is after all possible to imagine propositions like <*h*>; the best we can do, they may think, is to imagine some other, possible state of affairs, and to *confuse* that for a state in which *h*. In this section, we argue that this misidentification response (MR) is untenable.[8]

Comments Kripke makes himself suggest the MR to alleged counterexamples to SMIH:[9]

> . . .though we can imagine making a table out of another block of wood or even from ice, identical in appearance with this one, and though we could have put it in this very position in the room, it seems to me that this is *not* to imagine *this* table as made of wood or ice, but rather it is to imagine another table, *resembling* this one in all external details, made of another block of wood, or even of ice. (1980, p. 114)

Over the last thirty years, many have at least *seemed* to embrace the MR. A brief sampling follows.

> In this sort of case, one might misinterpret the imagined situation as a situation in which S; here, the situation is merely one in which one has evidence for S. (Chalmers, 2002, p. 153)

> To imagine myself truly believing that Hesperus and Phosphorus were distinct, I would have to imagine them being distinct; and that I cannot do, no more than I can imagine Venus's being distinct from Venus. (Yablo, 1993, p. 23)

> In response to questions, she replies that she is imagining a world in which there is a colorless, tasteless liquid that comes out of taps and fills lakes but that is not H_2O. Now we have a possible defeater . . . it is not unreasonable to suppose that she is really just imagining that something superficially resembling water is not H_2O rather than water itself is not H_2O. (Tye, 1995, p. 186)

It can be difficult to discern the extent to which any of these authors actually embrace the MR—whether they do or not depends on how they understand 'imagining'. In the case of Yablo (1993), for instance, "imagining" is not coherent supposition or propositional imagining in response to a fictional story, but rather what he calls "objectual imagining." Exploiting the ambiguity of the word 'imagining', perhaps one could find an interpretation of the sentence we used to express SMIH that is necessarily true. Doing so, however, does not necessarily constitute progress in modal epistemology, nor does it involve accepting SMIH, as we understand it.[10] Indeed, we might well wonder what "objectual imagining" really is. (We will say a bit more about objectual imagining in §8.) That an interpretation of the sentence we used to express SMIH is necessarily true matters very little unless one can give an independent characterization of the relevant type of "imagining" according to which it is possible for thinkers regularly to discern whether they are "imagining" as opposed to merely seeming to "imagine."[11]

Whatever its prevalence, there are at least some philosophers who have been attracted to the MR as a way of retaining SMIH (as we have explained it, where 'imagining' stands for supposing). However, to insist that it is impossible to imagine metaphysical impossibilities is at odds with platitudinous views about propositional imagination (understood as supposition). A highly plausible thesis connects propositional imagination to belief:

> HPT1: For any proposition <*p*>, if it is possible for someone coherently to believe <*p*>, it is possible for someone coherently to imagine <*p*>.

(Most people make a distinction between beliefs that are merely false and/or unwarranted, and beliefs that not only are false and/or unwarranted, but fundamentally just don't make any sense at all, e.g. the belief that you're being intentionally persecuted by the number two. A coherent belief is one that even if false and unwarranted, can be made sense of.) A second highly plausible thesis maintains that it is possible to believe metaphysical impossibilities:

> HPT2: It is possible for someone coherently to believe metaphysical impossibilities like <*Water does not contain hydrogen*> or <*Hesperus is closer than Phosphorus*>.

Our two highly plausible theses together undercut the misidentification response, for they entail that it is possible to imagine the impossibilities the subjects represent themselves as imagining, and therefore that SMIH is false. Both highly plausible theses are true.

Of the two theses, HPT2 is perhaps the more obvious. Suppose someone, at the advent of modern chemistry, performs an experiment that misleadingly indicates that a water sample contains no hydrogen. He sincerely reports: "I believe that water contains no hydrogen"—he speaks truly in so reporting,

and therefore believes the metaphysical impossibility that water contains no hydrogen.[12] The belief this person refers to is certainly not incoherent—it makes good sense for him to believe what he does. Moreover, it would be absurd to offer an analogue of the misidentification response, thus: "you *think* you have the coherent belief that water contains no hydrogen, but *actually*, you have the coherent belief that some non-water but watery *stuff* contains no hydrogen, and confusing that state with a state in which *water* contains no hydrogen." His is a paradigmatic case of a false but coherent belief about water. Even if it is necessarily false, it cannot be criticized for being false merely on the basis of rational reflection.

Examples are easily multiplied. Lois believes the necessary falsehood that Superman is stronger than Clark, not the contingent falsehood that some *other* guy who *looks like* Superman is stronger than Clark. Or, for any a posteriori false <*p*>, someone may coherently believe the necessary falsehood <*Actually, p*>—necessarily false because facts about the actual world do not change when evaluated at other worlds. Many pre-Kripke philosophers believed that *Hesperus could have been distinct from Phosphorus*; since this possibility claim is false, it is necessarily false (assuming the correct modal logic is at least as strong as S4) and therefore also an instance of HPT2.

There are also compelling theoretical reasons to accept HPT1. The best philosophical and psychological theorizing about the propositional imagination relates propositional imagination closely to belief. On one widespread approach to imagination, for instance, the propositional imaginings under consideration are *simulations* of beliefs; when one imagines that *p*, one enters into a state that is in some senses similar to the belief that *p*.[13] It has been widely recognized that propositional imagination plays many of the functional roles of belief—when we imagine something sad, for instance, we feel sad, not unlike how we'd feel if we believed the sad content. If imaginings are simulations of beliefs, then, it would be very odd indeed if some beliefs—the ones with metaphysically impossible contents—could not be simulated. If someone can believe that water contains hydrogen, someone else can simulate that belief, and thereby imagine a metaphysical impossibility.

According to a leading alternative approach, propositional attitudes like beliefs, desires, and imaginings involve having bits of syntax represented in cognitive 'boxes'.[14] To believe <*p*> is to have in one's belief box a sentence expressing that *p*; to imagine it is to have such a sentence in the imagination box. The denial of HPT1 on this model would amount to the claim that the imagination box admits different sentences than does the belief box. But this does not seem to be true; the mechanisms that regulate the contents of our belief boxes seem to be just the same mechanisms that regulate the contents of our imagination boxes. Certain incoherent sentences are automatically removed from both boxes by a particular cognitive mechanism (the 'Updater', in Nichols and Stich's terminology)—and this mechanism operates without regard to which box houses the relevant sentences. Indeed, the parallels

between belief and imagination vis-à-vis patterns of inference prompt Nichols (2004) to suggest that belief and imagination are "in the same code"—by which he means that a wide variety of cognitive mechanisms process beliefs and imaginings in the same ways.[15] A belief is treated in a very similar way to an imagining with the same content. It is implausible, then, that a metaphysically impossible proposition could be represented in the belief box, but not in the imagination box; no appropriate mechanism is sensitive to the difference.

So there are compelling reasons to accept both HPT1 and HPT2. So far as we can see, the only way to retain SMIH in light of our discussion is to adopt a two-dimensionalist semantics so that in (coherently) supposing, for instance, that water does not contain hydrogen, one is not supposing the secondary intension of the sentence 'Water does not contain hydrogen', which is necessarily false, but rather the primary intension, which is true at some possibilities, even thought it is not true of the actual world.[16]

While we are not prepared to enter a full critique of two-dimensionalist semantics here, we are inclined to reject it. Given our own commitments, we see no way to accept SMIH. It is possible coherently to imagine metaphysical impossibilities.

§3. Conceptual Possibility

We must admit, then, that coherently imagining some proposition does not ensure that the proposition is metaphysically possible. But it does seem as though there is some interesting status in the neighborhood of possibility being picked out by coherent imagining. An alternative to SMIH involves considering *conceptual* possibility, to be contrasted with metaphysical possibility, as a status important to modal epistemology. Conceptual possibility, on this line, could be established by recognizing one's imagining as coherent; the counterexamples to SMIH, though metaphysically impossible, can be counted as *conceptually* possible.

(Some philosophers object to the term 'conceptual possibility' on the grounds that propositions like *Hesperus is not Phosphorus* ought not to be judged possible *in any sense*.[17] As far as we can see, this is a mere terminological disagreement; those with insuperable aversion to the idea of something weaker than metaphysical possibility traveling under the name 'conceptual possibility' are invited to substitute their own preferred term.)

To employ this strategy, we face two challenges: first, to explain the notion of conceptual possibility at work, i.e. the notion of coherently imagining, and second, to defend against the charge of having changed the subject. Modal epistemology is, at least traditionally, ultimately thought to be about metaphysical possibility. We begin now with the first challenge.

What is conceptual possibility? A conceptual possibility can be coherently imagined to obtain; it is a situation that the constraints of rationality make

room for. Metaphorically, it is a point in the conceptual space of an agent. More precisely, a proposition is conceptually possible just in case it does not conceptually entail an absurdity. <*Some green things have no color*> is conceptually impossible, because <*x is green*> conceptually entails <*x has a color*>. <*Hesperus is closer to the Earth than Phosphorus*> is not conceptually impossible, because the entailment from <*x is Hesperus*> to <*x is Phosphorus*> is not a conceptual entailment.

It is controversial that there is a coherent and useful notion of conceptual entailment. The challenge for our gloss on conceptual possibility, then, is to articulate an adequate characterization of it.[18] As a precursor to our pursuit of an adequate characterization, we note our unequivocal position that conceptual entailment is not a matter of "truth (preservation) by convention." We consider any like view to be thoroughly refuted.[19] An adequate characterization of conceptual entailment can be given, instead, in terms of rational commitment and inference—or so we will suggest.

§4. Conceptual Entailment

Let us begin with an example. Suppose while reading a story, we encounter LV:

LV: It was raining in Las Vegas.

Suppose that we do not encounter LVW:

LVW: The streets in Las Vegas were wet.

Unsurprisingly, in engaging with this fiction, we will imagine that it was raining in Las Vegas. More remarkably, we will, in typical cases, also imagine that the streets in Las Vegas were wet (at least if the issue becomes relevant) even though the story does not say as much explicitly. Indeed, we are likely to *infer* from our imagining that it was raining in Las Vegas to an imagining that the streets in Las Vegas were wet. We suggest this tendency to infer reflects compliance with a rational commitment to so infer at least should the question of the wetness of the streets of Las Vegas arise.

This inference in imagination parallels the inference we would make in our beliefs were we to learn that it was raining in Las Vegas. Given our (perhaps tacit) background knowledge, a belief that it was raining in Las Vegas rationally commits us to believing that the streets in Las Vegas were wet (should the question arise); indeed, very typically, an inference from the former belief to the latter belief will preserve knowledge. Thus, if someone knows that it was raining in Las Vegas and infers on that basis to the belief that the streets in Las Vegas were wet the latter belief typically constitutes knowledge as well.

Of course, this inference might be defeated. While one might know that it was raining in Las Vegas, one might also know that the streets in Las Vegas were covered. This latter knowledge would undercut the inference to the belief that the streets in Las Vegas were wet. In exactly parallel fashion, if it is part of the story we are reading that the streets in Las Vegas were covered, we will not typically imagine that the streets in Las Vegas were wet even if we are imagining that it was raining in Las Vegas.

Our rational commitment to infer in imagination also has the same sort of basis as our rational commitment to infer in belief. Analogous to the case of belief, our rational commitment to infer from imagining that it was raining in Las Vegas to imagining that the streets in Las Vegas were wet rests, for instance, on our (perhaps tacit) empirical knowledge that there is a correlation between rain in a city area and wet streets. The warrant for this knowledge is earned via experience. In general, rational commitments to infer in belief coincide with evidentiary relations between propositions. Some of these evidentiary relations, e.g., between <*The lights are on*> and <*Someone's home*>, are established a posteriori, and hence hold only contingently and relative to the subject's past experiences. Other evidentiary relations obtain in a way independent of past experiential warrant. To embrace skepticism about rational commitments to infer in belief would effectively commit one to skepticism about propositional justification—to think we can't make sense of a rational commitment to infer to a belief is effectively to think we can't make sense of having justification for a belief we don't yet have. To have justification for a belief just is to be rationally committed to having that belief should the question of the belief's truth arise.[20]

Our LV-LVW example illustrates the point that rational commitments in imagining largely parallel our rational commitments in believing. If believing that p rationally commits you to believing that q (should the question arise), imagining that p will also typically rationally commit you to imagining that q. Frequently what we are rationally committed to imagining (at least should the question arise) extends well beyond those explicit and tacit invitations we choose to accept. It often makes perfect sense to ask of a scenario imagined to obtain whether it is the case that p even if <p> was not explicitly or tacitly used to specify the scenario.[21] In general, when we question of a scenario imagined to obtain whether it is a scenario in which p, we have not two but three choices:

(1) We answer that it is a scenario in which p because we are rationally committed, in imagining that the scenario obtains, to imagining that p.
(2) We answer that it is not a scenario in which p because we are rationally committed, in imagining that the scenario obtains, to not imagining that p.
(3) We reply that it is indeterminate, or we need more information, because imagining that the scenario obtains does not rationally commit us to

> imagining one way or the other whether or not *p* is true. (Suppose that it is raining in Vegas. Is this a scenario in which it is cold?)

Notice that these choices parallel the three choices we face when, in doxastic deliberation, we ask whether the actual world is such that *p* by consulting what our current beliefs dictate:

> (1) We answer that the actual world is such that *p* because we are rationally committed by our current beliefs to believing that *p*.
> (2) We answer that the actual world is not such that *p* because we are rationally committed by our current beliefs to not believing that *p*.
> (3) We withhold judgment because our current beliefs are not strong enough to commit us one way or the other.

For those like us who understand imagining as at least partly constituted by supposing, the strong parallels between belief and imagining should not be surprising. When we suppose some proposition for the sake of argument, we (rightly) draw the same sorts of conclusions from that supposition that we would were we to believe the proposition in question. It is precisely for this reason that we can generally test the coherence of a supposition as proxy for testing the coherence of a potential belief of the same proposition we supposed.

Our intention is to use rational commitments to infer in imagination in order to characterize conceptual entailment. Of course, not all rational commitments to infer in imagination correspond to conceptual entailment. The rational commitment to infer from imagining that it was raining in Las Vegas to imagining that the streets in Las Vegas were wet does not coincide with an entailment relation between the proposition that it was raining in Las Vegas and the proposition that the streets in Las Vegas were wet. The former proposition obviously does not entail the latter in any sense. We can see as much by remembering how rational commitment can be defeated by further imaginings. As we already indicated, one might, for instance, imagine that the streets of Las Vegas are covered. If the streets of Las Vegas are covered, then they may not be wet even if it rains there.

Perhaps though, *indefeasible* rational commitments to infer in imagination coincide with conceptual entailments? If someone has an indefeasible rational commitment to infer from imagining that *p* to imagining that *q*, they're rationally committed to imagining that *q* by imagining that *p*, and that rational commitment persists in *any possible case* in which the subject continues to imagine that *p*—including cases in which one imagines various other things (like that the streets are covered). Note that if the proposition that *p* does not at least metaphysically entail the proposition that *q*, then there is some metaphysical possibility such that *p and not-q*. If it is imagined that this possibility obtains, this imagining ought to defeat the inference

from imagining that *p* to imagining that *q*—but then, that inference cannot have been indefeasible. From this argument we can see that an indefeasible rational commitment to infer does at least entail metaphysical entailment.

Unfortunately, almost no rational commitments to infer in imagination are in this sense indefeasible.[22] Remember that rational commitments to infer in imagination share their basis with rational commitments to infer in belief. Our rational commitment to infer from imagining that *p* to imagining that *q* (if the question is raised) is explained by whatever explains our rational commitment to infer to believing that *q* (if the question is raised) so long as we might continue to believe that *p*. Consequently, rational commitments to infer in imagination are indefeasible only if corresponding rational commitments to infer in belief are indefeasible. However, it is quite plausible that no rational commitments to infer in belief are indefeasible.

Indeed, almost any rational commitment to infer can be defeated *due to a subject's rational limitations*. Rational limitations may arise from limitations in one's conceptual repertoire, limitations in computational capacity (e.g. in the time it takes to draw an inference), and tendencies to make performance errors in drawing inferences, or, for that matter, any other sorts of proclivities the subject has to make or exhibit confusions in attempting to execute in accordance with what he has reason to think. (This list may not be exhaustive, but it indicates that the term 'rational limitations' is not a catchall; rational limitations contrasts with limitations in experience that result specifically in a paucity of *evidence*, which might limit the subject's ability to competently infer in a wholly different way.[23] Rational limitations are limitations concerning the processing of evidence.) Testimony from a panel of expert logicians can defeat John's rational commitment to infer in accordance with *modus ponens*, but only because of John's rational limitations vis-à-vis logic. If John were an acknowledged über-logician, he would have no reason to kowtow to the panel of "expert" logicians any more than we have reason to defer to elementary school children on matters of basic arithmetic.

Similarly, almost any rational commitment to infer can be defeated *due to evidence regarding a subject's (current) rational limitations*. Evidence to the effect that Jane has taken a pill that inhibits rational capacities can defeat Jane's rational commitment to infer in accordance with *modus ponens* as can evidence to the effect that Jane is crazy even if in fact Jane has not taken such a pill and is not crazy.

Nevertheless, despite the defeasibility of almost all rational commitments to infer, we maintain that conceptual entailment can still be characterized in terms of rational commitments to infer in imagination. We suggest that what is distinctive of rational commitments to infer in imagination that coincide with conceptual entailment is that they are defeasible *only* due to a subject's rational limitations or due to evidence regarding a subject's (current) rational limitations. More precisely:

> CE: A set of propositions $\{<p_1>, <p_2>, \ldots, <p_n>\}$ conceptually entails $<q>$ just in case any defeat of the rational commitment to infer to imagining $<q>$ (at least should the question arise) when one imagines $<p_1>$, imagines $<p_2>$, ..., and imagines $<p_n>$ must be due at least partly to either (a) the subject's rational limitations or (b) the subject's having evidence concerning his (current) rational limitations.

One can prove that, by CE, conceptual entailment implies metaphysical entailment, and hence is necessarily truth-preserving, as any respectable entailment relation ought to be. Proof: Suppose that the set, S, of propositions $\{<p_1>, <p_2>, \ldots, <p_n>\}$ conceptually entails $<q>$, but it is metaphysically possible that the members of S are true while $<q>$ is not true. Imagining this metaphysical possibility (that nevertheless $<q>$ is not true) *in addition to* imagining $<p_1>$, $<p_2>$, ..., and $<p_n>$ ought to defeat the rational commitment to infer in imagination from $<p_1>$, $<p_2>$, ..., $<p_n>$ to $<q>$ *in and of itself* unless this metaphysical possibility is a genuine absurdity. (Imagining a genuine absurdity rationally commits one to imagining anything at all—it leads to an imaginative explosion, so to speak. A scenario in which an absurdity is true is one in which anything goes.) Presumably, metaphysical possibilities are not absurdities—because they are genuine possibilities, a rational agent can make sense of them. But then, by CE, S does not conceptually entail $<q>$, for there is a way to defeat the rational commitment to infer that does not essentially involve (a) or (b).

So far as we can tell, CE is equivalent to CE*:

> CE*: A set of propositions $\{<p_1>, <p_2>, \ldots, <p_n>\}$ conceptually entails $<q>$ just in case any defeat of the rational commitment to infer to imagining $<q>$ (at least should the question arise) when one imagines $<p_1>$, imagines $<p_2>$, ..., and imagines $<p_n>$ must never be wholly due to further imaginings.

The implication from CE to CE* is pretty clear. If defeat of a rational commitment must be due to either conditions (a) or (b), then it cannot be wholly due to further imaginings. The implication from CE* to CE is less clear, but it is difficult to see what else could defeat a rational commitment if further imaginings cannot do so by themselves.

It's worth highlighting that CE and CE* are merely characterizations of conceptual entailment, not explanations of it. For the purposes of showing that conceptual entailment is a legitimate relation, nothing more than a characterization is required. It is entirely possible, for instance, that conceptual entailment relations between propositions explain our rational commitments to infer, rather than the reverse. What's more, conceptual entailment could be characterized just as well without invoking imaginings. An alternative characterization might be offered using belief and knowledge. We outline this characterization in footnote 25.

CE is perhaps best understood through considering examples:

- If one imagines that *Stephen knows that* p, what could possibly defeat one's rational commitment to infer (should the question arise) to imagining that *p*? Our contention is that defeaters arise only due to rational limitations or evidence concerning rational limitations. If so, then by CE, <*Stephen knows that* p> conceptually entails <*p*>.
- Suppose that in imagining the events of a story, the question arises as to whether <*p or not-p*> is true. Presumably, you have a rational commitment to infer (on the basis of no particular previous imagining) to the imagining that *p or not-p*. Under what circumstances can this rational commitment be defeated? We might suppose it is only due to your rational limitations or evidence concerning your rational limitations. If so, then by CE the null set conceptually entails <*p or not-p*>.

Examples of conceptual entailment are bound to be controversial, but we must take care to make challenges on the appropriate grounds. For instance, the proposed conceptual entailment from <*Stephen knows that* p> to <*p*> is not threatened by the mere existence of Stu, a sophisticated thinker with complex theoretical reasons for thinking the rational commitment to make that inference might be defeated in other ways. Of course, if Stu is correct, then <*Stephen knows that* p> does not conceptually entail <*p*>. Nevertheless, so long as it is the case that Stu's rational capacities might be further improved upon so as to see through these misleading complex reasons, the fact that Stu is, in some sense, *rational* for rejecting that inference *does not* imply that <*Stephen knows that* p> does not conceptually entail <*p*>.[24] Stu's rationality in *this* sense and his cognitive sophistication are compatible with his rational limitations helping explain the defeat of his rational commitment. Likewise, we may have reasons to doubt an instance of excluded middle, but as long as these reasons are misleading in that an über-rational logician would see through them, it makes no difference when it comes to conceptual entailment because the defeat of this inference is explained by our own rational limitations.

Why can't the rational commitment to infer from <*Stephen knows that* p> to <*p*> be defeated by further imagining that *Stephen knows that* p *even though not-p*? If <*Stephen knows that* p> in fact conceptually entails <*p*>, the proposition <*Stephen knows that* p *even though not-p*> is a genuine absurdity that *ipso facto* will conceptually entail any proposition, including <*p*>. We parenthetically made note of this phenomenon earlier; further imagining an absurdity never defeats a rational commitment to infer.

The following examples emphasize that, plausibly, not all a priori inferences with belief coincide with conceptual entailment.

- Some have argued that an inference on the basis of no previous premise to the belief that *one is not a brain in a vat* is a priori warranted. Suppose so.

Even still, does this imply that the null set conceptually entails that one is not a brain in a vat? Obviously not. To be sure, one probably is rationally committed to imagining that one is not a brain in a vat if one is imagining anything about oneself at all. After all, when one imagines that one is having certain sensory experiences, it is natural also to imagine that these sensory experiences are veridical. Nevertheless, the rational commitment to infer in this way is clearly defeated by further imagining that people have kidnapped you, extracted your brain, and put it in a vat that keeps it alive while they feed it electrical signals that mimic incoming sensory signals. This further imagining wholly explains the defeat of our rational commitiment to infer, so by CE, the null set does not conceptually entail that one is not a brain in a vat.

- Very plausibly, an inference on the basis of no previous premise to the belief that one is here now (if one exists, now exists, and one is anywhere at all now) is a priori warranted. Does this imply that the null set conceptually entails that one is here now? Obviously not. One might just as well imagine that one is somewhere else—in Fiji, say. Indeed, it is not obvious in this particular case that there is a corresponding rational commitment to infer in imagination to the conclusion that one is here now. (Patterns of rational commitment to infer in imagination evidently *largely* but do not *entirely* parallel those in belief.)

This example shows that conceptual entailment does not coincide with metaphysical entailment.

- By CE, conceptual entailment presupposes a rational commitment to infer in imagination. In the event that one does not know that Phosphorus is Hesperus, there may well be *no* rational commitment to infer to the imagining that *Hesperus is a planet* from the imagining that *Phosphorus is a planet*. This would lead us to conclude that <*Phosphorus is a planet*> does not conceptually entail <*Hesperus is a planet*>. Indeed, even for those of us who know that Phosphorus is Hesperus and hence are rationally committed to the inference, we might still imagine that it is not the case that Phosphorus is Hesperus; for instance, one might imagine that the putative discovery that Phosphorus is Hesperus is a mistake. On its own, this further imagining defeats this rational commitment to infer in imagination from <*Phosphorus is a planet*> to <*Hesperus is a planet*>. Hence the former proposition does not conceptually entail the latter (even though it metaphysically entails the latter).[25]

§5. Conceptual Possibility Revisited

CE distinguishes conceptual entailment from metaphysical entailment and the a priori broadly construed. Notably, it does so without assuming, problematically, a 'molecular analysis' structure of concepts. It picks out an interesting relation, which we can use clarify conceptual necessity: a proposition is conceptually necessary if and only if it is conceptually entailed by any

proposition. We will also use the following equivalence: <p> conceptually entails <q> just in case <$p \supset q$> is conceptually necessary.

We have already suggested that a proposition is conceptually possible just in case the proposition does not conceptually entail an absurdity. But what is an absurdity? The paradigm instances of absurdities are contradictions, but we can use conceptual entailment to clarify what an absurdity is as well. Absurdities conceptually entail any proposition.[26] This understanding of absurdity allows us to simply our definition of conceptual possibility: a proposition is conceptually possible just in case the proposition is not an absurdity.

Because we explain conceptual possibility in terms of conceptual entailment and conceptual entailment is characterized in terms of rational commitments to infer, our grasp on conceptual possibility comes through our apprehension of these rational commitments, i.e. through exercising our rational capacities.

Indeed, our grasp on conceptual possibility is through rational capacities that track wholly a priori rational commitments, i.e. rational commitments with no empirical basis. Rational commitments to infer that coincide with conceptual entailments must not depend upon any empirical basis. Any rational commitment to infer in imagination that depends upon an empirical basis can be wholly defeated *merely by imagining that the relevant empirical discoveries away*. This is effectively illustrated by recalling the rational commitment to infer in imagination from <*It was raining in Las Vegas*> to <*The streets in Las Vegas were wet*>. This rational commitment is defeated by imagining the empirical basis for this rational commitment away. One can simply imagine (for one reason or another) that there is no correlation between raining and wet streets—one might imagine, for instance, that rain always evaporated before it hit the ground.

Learning about conceptual possibility involves recognizing that rational capacities with no empirical basis do not lead one to infer an absurdity. It is wholly plausible this process can be carried out a priori. We will return to this point in §8 when we discuss further how conceptual possibilities might be recognized.

§6. Conceptual and Metaphysical Modality

We have established a notion of possibility—conceptual possibility—that is co-extensive with coherent imaginability. But why should that be of any interest to modal epistemology? Modal epistemology is about objective modal facts, not about rational commitments relating to propositions and concepts.

Broadly speaking, there are two general strategies one might employ to answer this worry. One might maintain that, once the distinction between conceptual possibility and metaphysical possibility is made clear, it is the former that modal epistemology ought, after all, to be concerned with. This

suggestion would be supplemented with either the claim that modal epistemology has actually been about conceptual possibility all along, or an argument that we ought to change the subject in this way. We will not take this strategy; on our proposal, the relevant questions in modal epistemology have always been, and should continue to be, questions about metaphysical possibility and necessity.

On our approach, then, conceptual modality is useful as an intermediary step. Knowledge of conceptual possibility and necessity is useful, from the point of view of modal epistemology, because it can help us get to knowledge of metaphysical possibility and necessity.

The main challenge, of course, is the necessary a posteriori. It is impossible that *Hesperus is closer to the Earth than Phosphorus*—but it is coherently imaginable and hence conceptually possible. As a result, it may be tempting to conclude that only facts about conceptual modality, and never about metaphysical modality, can be known from the armchair.[27] But this is an overreaction. That some propositions are necessary and a posteriori entails only that conceptual possibility does not entail metaphysical possibility. But we have more armchair resources than single applications of tests of conceptual possibility.

Although some propositions cannot be known a priori to be metaphysically possible, others look like plausible candidates:

- *The Red Sox lost*.
- *Some green squares have 13-inch sides*.
- *Most tigers are bred in captivity*.
- *Someone has a justified true belief that is not knowledge*.

All of these propositions are conceptually possible; we can know this by recognizing that we can coherently imagine them. (More on the epistemology of conceptual modality in §8.) We mean to defend the intuitive suggestion that we can know them to be metaphysically possible, too—but how can this be so, given the observation that conceptual possibility doesn't entail metaphysical possibility? As in the case above, it is helpful to examine features of the counterexamples to the entailment.

How did we come to realize that necessarily, water comprises H_2O?

Putnam told us a story that effectively invited us to imagine that in some far-off world, some non-H_2O substance XYZ has many of the surface properties of water. He then asked us to judge whether that faraway substance was water. We replied that it was not. Why? He did not specify whether or not it was. Nevertheless, we took ourselves to be rationally committed to imagining that XYZ was not water. Something in the Putnam story rationally committed us so to imagine. What could so commit us? Presumably we *are* rationally committed to infer in imagination from *<x is not composed of*

H_2O> to <*x is not water*>. But what does this rational commitment have to do with our ultimate judgment that *necessarily* water is composed of H_2O?

The answer cannot be mere recognition of the rational commitment to infer. We are rationally committed to infer in imagination from <*It was raining in Las Vegas*> to <*The streets in Las Vegas were wet*>, but recognizing this rational commitment in no way warrants a judgment of necessity. We are not thereby warranted in judging that necessarily if it was raining in Las Vegas, the streets in Las Vegas were wet. As in the case of the rational commitments to infer in imagination from <*x is not composed of H_2O*> to <*x is not water*>, this inference does not correspond to conceptual entailment. These rational commitments are defeated by imagining that contrary to present scientific consensus, water is not H_2O—the scientists have just been confused by misleading evidence. In fact, the rational commitment to infer in imagination from <*x is not composed of H_2O*> to <*x is not water*> has an empirical basis much as the rationally commitment to infer in imagination from <*It was raining in Las Vegas*> to <*The streets in Las Vegas were wet*> does. In the former case, the empirical basis is the discovery that the samples of water we interact with are composed of H_2O; this empirical basis explains why the rational commitments to infer do not correspond to a conceptual entailment, and hence, why water's being some non-H_2O substance XYZ is a conceptual possibility.

Why then should we judge on the basis of Putnam's thought experiment that necessarily water is composed of H_2O when we do not judge on the basis of imagining that it was raining in Las Vegas that necessarily if it was raining in Las Vegas, the streets in Las Vegas were wet? What is different about these two cases?

We suggest the basis for a judgment of necessity in response to the Putnam thought experiment comes in recognizing *other* conceptual entailments. To make these conceptual entailments salient, we invite you first to imagine not only that in some far-off world, some non-H_2O substance XYZ has many of the same surface properties of water, but also that in the actual world, contrary to present scientific consensus, the samples of water we interact with are not uniformly composed of H_2O but that same underlying substance XYZ. Now we ask, "Is that faraway substance water?"

In this case, the clear verdict is "yes". Indeed, this verdict strongly indicates rational commitments to infer in imagination that <*x is water*> from <*In the actual world, the samples of water we interact with are uniformly composed of y*> and <*x is composed of y*>.

Now consider a third case flipping Putnam's completely. Imagine that in some far-off world some substance, comprising H_2O, has many of the same surface properties of water, while continuing to imagine, as before, that actually, contrary to present scientific consensus, water is not uniformly composed of H_2O but some non-H_2O substance XYZ. Now we ask, "Is *that* faraway substance water?"

In this case, the clear verdict is "no". Indeed, this verdict strongly indicates rational commitments to infer in imagination that <*x is not water*> from <*In the actual world, the samples of water we interact with are uniformly composed of y*> and <*x is not composed of y*>.

Unlike the rational commitments to infer in imagination from <*x is not composed of H_2O*> to <*x is not water*>, *these* rational commitments to infer *do* seem to correspond to conceptual entailment. Facts about what sorts of substances qualify as water *are* conceptually entailed but only by a posteriori propositions about what chemical composition the samples of water we interact with actually have.

Indeed, the fact that in the actual world, the samples of water we interact with are uniformly composed of H_2O is a likely to be a crucial tacit background imagining when engaging with the Putnam thought experiment. We suggest that it is through tacitly recognizing that this tacit background imagining conceptually entails what sorts of substances qualify as water that we come to judge (and are warranted in judging) that necessarily, water is composed of H_2O. Far from undercutting the importance of the imagination and conceptual modality in modal epistemology, close attention to how Kripke-Putnam thought experiments result in judgments of necessity motivates drawing a close connection; recognition of conceptual entailments appears to play an important role in explaining our judgments about the necessary a posteriori.

We claim that in light of the fact that it is false that in the actual world, the samples of water we interact with are not uniformly composed of H_2O, the fact that <*Water is not composed of H_2O*> conceptually entails <*In the actual world, the samples of water we interact with are not uniformly composed of H_2O*> implies it is metaphysically impossible for water not to be composed of H_2O. Generalizing on this example, we suggest if <*p*> conceptually entails something about the actual world that is in fact not the case, then <*p*> is metaphysically impossible. Using symbols we get FAMI.

FAMI: $\exists(<q>)[\sim A(<q>) \,\&\, \Box_c(p \supset A(<q>)] \supset \Box_m(<\sim p>)$

Here, '$\Box_c$' means "it is conceptually necessary that", '$\Box_m$' means "it is metaphysically necessary that," and '$A(<q>)$' means "the actual world is such that *q*." The contrapositive of FAMI expresses a necessary condition for metaphysical possibility:

FAMI*: $\Diamond_m(<p>) \supset \sim\exists(<q>)[\sim A(<q>) \,\&\, \Box_c(p \supset A(<q>))]$

FAMI* says that the metaphysical possibility of <*p*> implies that there isn't a <*q*> such that <*p*> conceptually entails that in the actual world, *q* even though it is not the case that in the actual world, *q*. This necessary condition for metaphysical possibility is not met by conceptual possibility. This is

unsurprising, since we've asserted all along that conceptual possibility is insufficient for metaphysical possibility—SMIH is false. However, we see no reason to think that the right-hand side of FAMI* isn't, in addition to being a necessary condition for <*p*>'s metaphysical possibility, also a *sufficient* condition, in conjunction with <*p*>'s conceptual possibility. Call this claim moderate modal rationalism (MMR):

MMR: $\Diamond_m(<p>) \equiv (\Diamond_c(<p>) \ \& \sim\exists(<q>)[\sim A(<q>) \ \& \ \Box_c(p \supset A(<q>))])$

To deny MMR right-to-left is to assert that there are some propositions that are conceptually possible, but metaphysically impossible, and for which imagining them true does not necessitate imagining anything false about the actual world. This is the claim that metaphysical possibility requires something more than either conceptual possibility or the condition exploited by Kripke-Putnam thought experiments. What could this mystery ingredient to metaphysical possibility be? The fact that a proposition meets the necessary conditions expressed by the right-hand side of MMR is at least generally thought to settle the question as to whether the proposition is metaphysically possible.

Anyone defending the mystery ingredient view must show either that we do make an additional distinction that figures into our conclusions about metaphysical possibility, or defend the view that metaphysical possibility has necessary conditions about which we're entirely in the dark. Neither alternative looks particularly plausible. Certainly the standard examples of the necessary a posteriori do not motivate any such mystery ingredient. We conclude that MMR is true.

§7. Examples, Moral and Mathematical

In order to substantiate our conclusion, it helps to see how MMR accommodates the metaphysical necessity of basic moral principles and mathematical truths.

Basic Moral Principles

Plausibly, (NoPain) is metaphysically necessary:

> NoPain: We ought not cause others to feel excruciating pain merely for the purposes of superficial entertainment.

If NoPain is also conceptually necessary, then MMR easily explains why it is metaphysically necessary as well: conceptual necessity entails metaphysical necessity. However, even if NoPain is metaphysically necessary, some version of G. E. Moore's open question argument might make us doubt whether it is conceptually necessary. Perhaps it is not incoherent to suppose this conclusion false. Suppose NoPain is conceptually contingent.

According to MMR, NoPain can still be necessarily true if there is some <*q*> that is actually false even though NoPain's being false in some scenario (real or imagined) conceptually entails <*q*> is actually true. (To see that this follows from MMR, negate both sides of the biconditional, and substitute <*not p*> for <*p*>.) Is there some such <*q*>? Yes.

Before we point to the <*q*>, let us try to explain why there should be one. It can be settled merely by rational reflection that the normative strongly supervenes on the non-normative. This, even if it cannot be settled by such means what the bridge principles of supervenience specifically are. The specifics of the bridge principles, of course, constitute the correct moral theory (perhaps it will be a general theory, or perhaps, as the particularist thinks, it will be only a collection of truths). In supposing that moral principles like NoPain are not conceptually necessary, we are supposing that our discovery of the correct moral theory does not proceed via purely rational reflection; perhaps there is a quasi-perceptual or even constructivist aspect to this discovery. Whatever the basis of this discovery, though, it is clearly a matter of conceptually necessity that these discoveries project to non-actual scenarios in the same way that discoveries of say, the material composition of natural kinds, project to non-actual scenarios. This is enough to establish the metaphysical necessity of NoPain.

How? Let *basic moral principles* be moral principles that follow from the correct moral theory. Let <*q*> be the proposition that NoPain is not a basic moral principle. That NoPain is false (in some scenario) conceptually entails that actually, NoPain is not a basic moral principle. But actually, NoPain is a basic moral principle, so it follows by MMR that NoPain is metaphysically necessary.

In other words, MMR allows that the metaphysical necessity of NoPain may be conclusively settled by the fact that 1) it is conceptually necessary that if it is false (in some scenario), then NoPain does not actually follow from the correct moral theory and 2) whatever actually makes it the case that some moral theory is correct and NoPain follows from it. Obviously, this result generalizes for any basic moral principles. The upshot is we can recognize that basic moral principles are metaphysically necessary simply by combining our knowledge of conceptual necessity with the knowledge we gain through moral inquiry (however that works).

Mathematical Truths

There is a longstanding debate in the philosophy of mathematics over whether our warrant for believing mathematical truths has a purely rational basis, or whether it is gained via exercising, for instance, quasi-perceptual or constructivist faculties. This debate is largely analogous to the debate just considered over whether our warrant for believing basic moral principles has a purely rational basis, or whether it is gained in some other way. Our treatment here will be the same as in the previous case.

If our warrant for believing mathematical truths has a purely rational basis—if it is just incoherent to suppose them false—then they are generally conceptually necessary, and hence metaphysically necessary. We could, of course, learn that these mathematical truths are metaphysically necessary by learning that they can be established on a purely rational basis.

If it is not generally incoherent to suppose mathematical falsehoods (because of their existential commitments, for instance), then they are not conceptually necessary. Still, what is conceptually necessary is that whatever settles mathematical truth in actuality projects to non-actual scenarios. More specifically, that <*not p*> is a mathematical truth (in some scenario) conceptually entails that <*not p*> is actually true. Thus, by MMR, if <*not p*> is not actually true, then <*p*> is a metaphysically necessary mathematical truth. In this way, we can explain the metaphysical necessity of mathematical truths even if they are not conceptually necessary.

An example that often raises concern in the modality of mathematics is the continuum hypothesis (CH), which is provably independent from the ZFC axioms of set theory. There is a mistaken tendency to take the independence results for CH as a clear indication that CH and its negation are both conceptually possible. No such consequence follows; myriad conceptual necessities (e.g., <*If something is known, it is true*>) and impossibilities (e.g., <*Julius Caesar is identical to the number two*>) fail to be settled by the axioms of set theory.[28] Mathematicians might yet establish CH if they can show that certain models of ZFC have particularly good properties, and show that any one of these good models is one in which CH is true. (Whether this would show that CH is conceptually necessary depends on whether the existential commitments of certain mathematical truths are compatible with the status of conceptual necessity.) We are not able to foresee whether mathematicians might adopt new axioms that settle the matter.[29]

Even if it turns out that there is no basis for accepting CH or accepting its negation, MMR can accommodate that one of CH or its negation might nevertheless be true as a matter of metaphysical necessity. So long as CH (being true in some scenario) conceptually entails that CH is actually true and the analogue is true for its negation, we can explain the metaphysical necessity of whichever one is true via MMR. Given this conceptual entailment, MMR tells us the truth or falsehood of CH will be projected onto non-actual metaphysical possibilities even if establishing the truth or falsehood of CH is impossible.

Reflection on the application of MMR in the cases of morality and mathematics reveals both its power and flexibility. According to MMR, the facts about conceptual necessity and actuality determine the facts about metaphysical necessity.[30] Even if powerful, this is a very plausible claim. MMR is flexible because it allows metaphysical necessity in two different ways. This flexibility makes it difficult to refute; refuting it requires showing not only that there are metaphysical necessities that are not conceptual necessities,

but also that there are metaphysical necessities that do not arise in connection with what must, as a matter of conceptual necessity, be true actually given that it is true in some non-actual supposed scenario. That metaphysical necessity could so thoroughly divorced from the rationally constrained imagination is difficult to fathom.

§8. Conceptual Modality and Apriority

We have, in MMR, a straightforward methodology for deciding facts of metaphysical possibility: first, check and see that the proposition in question is conceptually possible, then check and see that imagining it doesn't commit one to imagining a falsehood about the actual world. In what sense does our MMR deserve the name 'moderate modal rationalism'? Obviously, our view doesn't have the result that in general, knowledge of metaphysical modality is a priori.

Nevertheless, we count our view as a species of moderate rationalism because, according to it, grasp of (metaphysical) possibility and necessity is mediated through sensitivity to conceptual possibility and necessity. This sensitivity *does* result solely from exercising cognitive capacities that have a purely rational basis.[31]

Furthermore, our view does allow for a considerable amount of knowledge of metaphysical possibility and necessity that results solely from exercising our purely rational capacities. Conceptual necessity entails metaphysical necessity (and this is a matter of conceptual necessity), so to the extent the former is a priori, the latter will be also. It is also quite plausible that there are many propositions where as a matter of conceptual necessity, conceptual possibility is sufficient for metaphysical possibility. Assuming that one can recognize conceptual necessity a priori, these are cases of conceptually possible propositions, for which we can recognize a priori that the second MMR condition drops out—cases in which we can recognize a priori that imagining $<p>$ doesn't conceptually entail anything at all, much less anything false, about the actual world. If we can recognize a priori that imagining that a green square has six-inch sides does not conceptually entail anything about the actual world, then the second MMR condition—the usually-a-posteriori one—is trivially met. Insofar, then, as we can recognize the relevant facts about conceptual possibility and necessity a priori, we can know a priori that it is metaphysically possible that there is a green square with six-inch sides.[32]

To what extent, then, are facts about conceptual necessity, possibility, and entailment—facts about what can and can't be coherently imagined and what we are rationally committed to infer in imagination—knowable a priori?[33]

It is not at all difficult to see that we can often recognize conceptual impossibilities a priori. We can know a priori that some imagining is incoherent by finding an obvious *reductio ad absurdum* using inferences that

correspond to conceptual entailments. Doing so merely requires exercising rational capacities whose basis survives further imaginings. (As we saw in §5, such rational capacities cannot depend upon an empirical basis.) Whether we are capable of this depends on how smart we are—but we are smart enough often to find them. Imaginings of blatant contradictions are incoherent; likewise with imaginings that there are green apples with no color. (A one line *reductio* is sufficient for these propositions; a world where contradictions are true, or where green apples have no color, is a world where anything goes!) Other conceptual impossibilities are less transparent, but still recognizable in this fashion. Some propositions, such as <*There is a largest prime number*>, or <*Some set contains all sets that don't contain themselves*>, may appear to us mortals at first glance to be perfectly coherent to imagine, but further a priori investigation reveals them not to be. Conceptual impossibilities can be known a priori to be conceptually impossible.

Prima facie, there is more difficulty in supporting the claim that we can know a priori that imaginings of some propositions are coherent. To know as much a priori requires knowing a priori that there *isn't* a *reductio ad absurdum* for some proposition. How could we know this? One way to acquire at least some warrant for this negative claim is the straightforward one: we examine the proposition, searching for a *reductio*, and find none. Since we are generally pretty good at recognizing a *reductio* when there's one there, a negative result in our search provides a *pro tanto* reason to think that there's no *reductio* to be found, and the proposition is coherent.

This *pro tanto* reason can be strengthened by filling in the specifics of a scenario in which the proposition is true. Considering a fleshed out scenario raises the likelihood of finding absurdities that might otherwise have escaped one's notice. The most straightforward way to flesh out a scenario is mentally to construct a structural facsimile of it. Structural mental representation (as opposed to linguistic or conceptual representation) is most familiar from the process of visualization. Offline perceptual simulations can be used to structurally represent spatial relations—but they might also be used to represent other sorts of relations as well, just as a visual graph can model the change in population of India over the last fifty years. Indeed, people may well have mental states besides offline perceptual simulations that also can be used to for the purposes of structural representation.

The validity of conceptual entailment guarantees that there is no *reductio* for any imagining we can succeed in structurally representing. In this way, structural representation can positively establish conceptual possibility. Very plausibly, it is through this method that we know that it's coherent to imagine, for instance, that there are blue swan-like creatures. Because it's obviously true that there are blue swan-like creatures in a scenario we can easily visualize, we conclude that there couldn't be a *reductio* on the imagining that there are blue swan-like creatures. We draw this conclusion a priori; no genuine perceptual experience need play a role in warranting the conclusion. Consequently,

there's every reason to think that we know a priori that the proposition is coherently imagined. The same could be said of many propositions we take to be coherent.

How do we know that we have succeeded in structurally representing a scenario in which the relevant proposition is true? The answer is, in one sense, very straightforward: we have stipulative authority over what the fundamental elements of our mental facsimile stand for. Once we set (either implicitly or explicitly) what the fundamental elements of the facsimile stand for, verifying whether the proposition is true in the represented scenario is merely a matter of recognizing whether the world's being as it is simulated conceptually entails the proposition in question. Quite plausibly, this recognition operates on the basis of purely rational capacities.[34]

Of course, there is a complication here—there is an obvious limit to our stipulative authority. If we're looking to establish conceptual possibility: we'd better not stipulate so as to represent incoherent scenarios. One cannot just stipulate that a visualized person is bachelor and then go on to stipulate that the person is a woman. We must not *overstipulate*.

We can avoid overstipulation by being sensitive to potential conceptual entailment relations. If we take care to stipulate only matters conceptually independent from what is already stipulated, then our mental facsimile will represent coherent scenarios. Taking such care is not unduly difficult, nor does it require antecedent knowledge of what exactly is conceptually impossible and what isn't. One can take care not to overstipulate a represented person's gender when one has already stipulated he's a bachelor just by being aware that there is some conceptual link between something having to do with gender and bachelorhood. One can take care not to overstipulate that a represented object is orange when one has already stipulated it is green just by being aware that there are conceptual necessities concerning color exclusion. This sort of care can plausibly be taken a priori. Being sensitive to which features of a scenario are conceptually independent so as to restrict stipulation to those features that are thus independent is the sort of ability we have in virtue of rational competences.[35]

This is not to say, of course, that we're infallible at avoiding pitfalls. Sometimes, we overstipulate without realizing it. Accidental overstipulation, however, is the exception rather than the rule. More to the point, it can be recognized and abandoned by a priori means; we can derive a *reductio* by inferring in accordance with conceptual entailments. (Fleshing a scenario out even further can help us do so.)

Earlier, we briefly encountered the suggestion of Yablo (1993) that a proposition is possible if and only if one can "objectually imagine" it. It seems likely that objectual imagination involves scenario construction in roughly the way we have suggested. The difference between our suggestion and Yablo's is that Yablo seems to be suggesting that in objectually imagining we construct *metaphysical* rather than conceptual possibilities. When it

comes to representing metaphysical possibilities, the question as to how one can be sure that you've avoided overstipulating becomes *significantly* more pressing.[36] Here, overstipulation is stipulating the *metaphysically* impossible. To avoid such overstipulation, one must be sensitive to dependence relations that are metaphysically necessary. How is this sensitivity to be achieved in light of the necessary a posteriori? We do not see how it could be except by exploiting the connections between metaphysical modality and conceptual modality stated by MMR. And yet, given MMR, knowing what is conceptually possible and necessary obviates the need for representing metaphysical possibilities directly. We conclude "objectual imagining" is not central to modal epistemology.

We do not expect completely to mollify the skeptic as to how knowledge of conceptual possibility is possible. But whatever difficulties remain, discovering that imaginings are coherent can't be any more difficult than discovering that beliefs are coherent, and most of us are pretty confident that most of the beliefs that we have, even if they are false, are not incoherent.

Finally, we note that thinking that we frequently can have a priori knowledge of coherence or incoherence is not thinking that for every proposition, we know whether it is coherent to imagine it or not, much less that we always in fact know so a priori. The status of some propositions is vigorously contested. An über-rational agent would know whether it is coherent to imagine that some creature has a brain exactly like the brain of David Chalmers but has no phenomenal experience. At most one of the present authors is sufficiently rationally capacitated to attain that knowledge.

The Kripke-Putnam thought experiments show that possibility can't be so straightforwardly tied to imagination and rationality as many have suggested. Nevertheless, the methodology those thought experiments presuppose accepts that there is still a very tight relationship between possibility on the one hand and imagination and rationality on the other. We have attempted to codify this relationship in MMR. By exploiting this relationship, we can achieve knowledge of metaphysical possibility, and we can sometimes do so a priori.[37]

Notes

[1] See Gendler and Hawthorne (2002) and Yablo (1993) for some discussion on the importance of these two ideas in the history of modal epistemology.

[2] In fact, we do not even think that intuitions are often used as evidence, nor that intuitions are much more than beliefs or inclinations to believe. (See Williamson (2007), Chapter 7.) Indeed, we suspect that talk of "exercising rational capacities" should replace any talk about the role of "intuitions" in epistemology. As a result, we're also not sympathetic to the details of Sosa (2007), Lecture 3 where intuitions also appear to play a central justificatory role. This, even though we are very sympathetic to the general idea that a priori knowledge—including a priori modal knowledge—results from exercising epistemic competencies. Our approach fits with Peacocke (1999), Chapter 4 although we won't try to draw comparisons explicitly.

[3] One caveat: our paper will focus primarily on the question of how we *can* have modal knowledge, rather than on the related question how we *do* come to have it. How we do have modal knowledge is as much a question of psychology as it is epistemology—we won't pretend to know from the armchair how people do know except to say that they are somehow sensitive to the method by which they can come to have it.

[4] Tamar Szabó Gendler (2000) has even argued that one can even engage imaginatively with fictions that contain arithmetical impossibilities.

[5] We are thus in terminological agreement with the approach of Currie & Ravenscroft (2002).

[6] Some philosophers engaging with the 'puzzle of imaginative' have made much of the apparent fact that some propositions seem to defy imaginability—in prototypical examples, if we are reading a fiction that enjoins us to imagine a proposition that strikes us as morally repugnant, we are, according to at least some authors, *unable* to comply. Even in these cases, it should be clear that we're at least able to imagine/suppose the deviant propositions *for the sake of argument*. This is presumably what is involved, for instance, when we draw inferences from such moral falsehoods. ('If gratuitous suffering is good, then the Holocaust was good.')

[7] When we write '<*p*>' we are indicating a mention of some proposition; when we write '*p*' we are indicating a use.

[8] We are not first to have pointed out that MR is untenable. See, for instance, Bealer (2004), §5.

[9] Byrne (2007) argues that that interpreting Kripke as embracing the MR is a mistake.

[10] Peacocke (1999), Chapter 4. Gendler and Hawthorne (2002).

[11] Gendler and Hawthorne (2002).

[12] Stalnaker (1984) thinks sets of possible worlds are the objects of belief, and that it is therefore impossible to believe the impossible. If it is impossible to believe the impossible on a possible-worlds approach, so much the worse for possible-worlds approaches. (Suppose that Stalnaker is right, contrary to our professed belief: it is impossible to believe impossibilities. Then when we believe that someone believes an impossibility, we believe the impossible.) (See Sorensen (1996).) King (2007) argues that Stalnaker can and should admit that we can believe the impossible.

[13] See Currie and Ravenscroft (2002), Chapters 1–2. Goldman (2006) is also a clear proponent of this approach. It's worth noting that at least one of us thinks that propositional imagining/supposing does not simulate belief; instead, propositional imagining/supposing that one believes simulates beliefs. Even so, HPT1 is overwhelmingly plausible.

[14] Nichols & Stich (2000).

[15] Apparently, Nichols intends to draw a contrast between beliefs and imaginings on one hand, and, for instance, desires on the other when he says that beliefs and imaginings are "in the same code." It is unclear what this contrast could be other than merely that beliefs and imaginings have similar functional roles. After all, accepting the cognitive box picture implies that desires and beliefs are in "the same code" as well at least in the sense that both beliefs and desires involve tokens of Mentalese.

[16] Cf. the distinction Chalmers (2002) makes between primary and secondary conceivability.

[17] Bealer (2002); van Inwagen (1998).

[18] As we point out later, a characterization of conceptual entailment is not the same as an analysis (or definition) of it. In offering a characterization, one merely gives necessary and sufficient conditions. In giving an analysis, one is typically suggesting that the analysans is conceptually or explanatorily prior to the analysandum. In putting forward a characterization, we intend to make no such suggestion. We leave open the possibility, for instance, that conceptual entailment should be understood as a primitive relation between propositions.

[19] Quine (1936) and Quine (1962). See also Peacocke (2004), especially p. 27.

[20] Two different anonymous referees have expressed doubts about whether rational commitment to infer can really be made sense of. We are not moved by these doubts largely because the

notion of rational commitment to infer in belief is so closely tied to the notion of propositional justification, a notion in epistemology that is not likely to be dislodged. The notion of rational commitment to infer in imagination is just the analogue of rational commitment to infer in belief. That there must be some such analogue is mandated by the uses to which we put our imaginative faculties. Cf. Currie and Ravenscroft (2002).

These points also make it clear why rational commitment to infer is not likely to be supplanted by counterfactual reasoning as one of these anonymous referees suggested might be the case. Counterfactual reasoning antecedently requires a grasp on what is evidence for what; establishing any counterfactual requires having a grasp on whether we would be entitled to infer that counterfactual given the experiences we have had.

[21] We distinguish scenarios from possibilities in that scenarios need not be genuinely possible in any sense.

[22] We intend to use 'rational commitment', so that if someone is rationally committed to infer and they so infer because they are rationally committed, then the resulting mental state is rational. In other words, if someone is rationally committed to infer, they are not irresponsible for so doing. In this sense of 'rational commitment', it will be possible to defeat someone's rationally commitment to infer by raising concerns about the inference, even if these concerns ultimately prove to be the result of confusion, which through more capable reasoning, might be alleviated. Defeasibility of "rational commitment" in this sense coincides with what Peacocke (2004), p. 30 terms 'defeasibility of identification'.

One might alternatively use 'rational commitment' so that a person's rational commitment to infer can never be defeated by concerns about the inference that, upon further and more capacitated rational scrutiny, ultimately prove to be the result of confusion. Defeasibility of "rational commitment" in this sense coincides with what Peacocke (2004), p. 30 terms 'defeasibility of grounds'.

[23] Here we effectively rely on a principled distinction between experience in its role as enabler and in its role as the provider of propositional warrant. We acknowledge that Williamson (2007) challenges whether this distinction can be made out in a principled way. Although we engage with that challenge in ongoing work, we will ignore it here.

[24] We therefore avoid the critique of conceptual modality in Williamson (2007), Chapter 4.

[25] It is also possible to characterize conceptual entailment in epistemic terms, relating inferences and knowledge, instead of as commitments to imagine. We focused on imagination in the main text above because our direct interest, for purposes of modal epistemology, was in commitments of imaginings. A characterization in epistemic terms could be given by reference to the property of a *knowledge-preserving inference*. If there is a knowledge-preserving inference from $<p>$ to $<q>$, then someone, if she knows that p, and infers that q on that basis, while continuing to know that p, will come to know that q. (Harman (1986) emphasizes the importance of the 'continuing to know' clause.) This feature of inference is roughly analogical with the rational commitments discussed above. As in the case with rational commitment in imagination, that an inference is knowledge-preserving does not imply that it corresponds to a conceptual entailment. The same counterexample suffices; the inference from <LV> to <LVW> is knowledge-transmitting but no conceptual entailment. As before, *defeasibility* of the inference seems to be playing an important role; if the streets are covered, then the inference is defeated, and does not preserve knowledge. Total indefeasibility of knowledge-preservation for beliefs is also problematic, for the same reason indefeasible rational commitments for imaginings was. So, under what circumstances, when $<p>$ conceptually entails that $<q>$, could someone know that p, and infer on this basis that q, continuing to know that p throughout the process, and yet fail to know that q? Only, as above, by virtue of rational limitations, or by having evidence concerning one's rational limitations. In generality:

> CE_k: A set of propositions $\{<p_1>, <p_2>, \ldots, <p_n>\}$ conceptually entails $<q>$ just in case, for any subject, there is a knowledge-preserving inference from $\{<p_1>, <p_2>, \ldots, <p_n>\}$ to $<q>$ that can be defeated only due at least partly to (a) the subject's

rational limitations, or (b) the subject's having evidence concerning his (current) rational limitations.

Our new CE_k delivers the same verdicts about conceptual entailment as did CE. For example, if one knows that *Stephen knows that p*, and infers to the belief that *p*, this latter can only fail to constitute knowledge by virtue of the (a) or (b) conditions above, so *<Stephen knows that p>* conceptually entails that *<p>*.

That there are such parallel characterizations of conceptual entailment, one in terms of rational commitments of imaginings, and the other in terms of preservation of knowledge in belief, should not be surprising. As we emphasized above, beliefs and imaginings are similar in a number of respects. And no one should be surprised to find tight connections between rationality and knowledge.

[26] A referee worries that this explanation of absurdities cannot be reconciled with many (putatively) paradigm cases of absurdities, for instance, the proposition that something is simultaneously pink and turquoise all over. Why think that imagining this proposition rationally commits one to imagining anything? Of course, an analogous skeptical question could be raised for any proposed absurdities including an outright contradiction that *p and not p*. Generally, there can't be any demonstration that any putative absurdity is an absurdity that does not rely on rules of inference that the skeptic is contending. A dialetheist isn't going to be satisfied by a classical logician's deployment of classical rules, but that hardly shows that the classical logician isn't right. (In the pink-turquoise case the relevant rules, of course, concern the "logic" of color. Add the right axioms of color (or equivalent rules of inference) with a strong enough logic and the desired result will follow.)

We're not particularly concerned to defend any particular proposition as an absurdity; if it turns out that the proposition that something is simultaneously pink and turquoise all over is not an absurdity as we've explained it, that's fine by us. We can show how these sorts of propositions are impossible on other grounds. See §7.

[27] Cf. the view discussed, with citation, in Jackson (1998), p. 69.

[28] 'Settled' obviously here means something stronger than 'conceptually entailed'. So likewise is the provable 'independence' of CH from ZFC weaker than conceptual independence.

[29] Even if we had conclusive reason for thinking that the matter could not be settled, this may constitute conclusive reason for thinking that although one of CH and its negation is true in every model of ZFC, neither is true *simpliciter*. Far from showing that CH and its negation are both conceptually possible, this might show that neither is conceptually possible, for both might be indeterminate as a matter of conceptual necessity.

[30] This is an epistemic claim rather than a metaphysical claim about fundamentality. For the purposes of this paper, we need not decide whether conceptual necessity is more fundamental than metaphysical necessity—it's enough that they are related in the way MMR claims they are.

[31] Compare with Thesis (II) and Thesis (III) of Peacocke (1999), pp. 168–171.

[32] Williamson (2007) argues that there is no principled distinction between the a priori and the a posteriori. We are here presupposing him to be incorrect.

[33] Again, we do not assume that a priori knowledge is empirically indefeasible. We might say that a priori knowledge is weakly a priori. Kitcher (2000); Field (2000).

[34] Obviously, we are rejecting the insistence of Yablo (2002), p. 457–61 that "peeking" is by nature a method of inquiry that is a posteriori. He rejects "peeking" as a priori in the first place because discerning features of imagined situations is too close to introspecting that one has a headache. We do not see the connection. To be sure, learning something from a mental simulation requires, in some sense, being aware of what is going on in the simulated scenario just as learning something from supposing requires being aware, in some sense, of what is supposed. We do not see, though, that this sort of "awareness" should constitute any serious form of introspection that is incompatible with learning a priori in either case. After all, learning from a simulation need not depend on one's consciously realizing that one is simulating

just as learning from supposing need not depend on one's consciously realizing that one is supposing.

Yablo rejects "peeking" as a priori in the second place because it requires exercising a "perceptual faculty rather than a cognitive one." To the extent that modeling requires exercising a perceptual faculty rather than a cognitive one, we do not see that it is done in a way that is incompatible with a priority. There is no incoherence in the idea of a priori knowledge being gained via the supposition of some empirical proposition, so long as the conclusion of the argument does not ultimately depend on the empirical proposition's truth, or on one's having supposed that the empirical proposition was true. Likewise, judging about some particular scenario using perceptual faculties can lead to knowledge a priori so long as the truth of the conclusion does not ultimately depend on the scenario's obtaining or on one's having exercised perceptual faculties. What matters for a priority is not whether a perceptual or cognitive faculty is exercised, but whether experience plays a justificatory role, and if it does, whether the role is merely hypothetical or not.

Yablo rejects "peeking" as a priori in the third place because the recognitional capacities one uses to determine what is true in a simulation have an empirical basis. The rational relations between perceptual experience and beliefs show that there can be (and frequently are) rational commitments to infer that need not have any empirical basis. One does not need to do any empirical research to know what propositions are true if one's perceptual experiences are veridical. Obviously, some recognitional capacities will have an empirical basis, but not all.

[35] For one (rather extreme) way of guaranteeing coherence, we could stipulate only microphysical facts; that which those microphysical facts conceptually entail would certainly be coherent. Cf. Chalmers (1996), pp. 76–77; Jackson (1998), pp. 81–84.

[36] For this reason, we take "objectual imagining" or simulation of metaphysical possibility to be subject to the criticisms of Byrne (2007), §§6–7 in a way that simulation of conceptual possibility is not.

[37] For helpful comments and discussion, we are grateful to Björn Brodowski, Melissa Ebbers, Richard Heck, Carrie Jenkins, Jason Stanley, Ernest Sosa, Brian Weatherson, Timothy Williamson, and three anonymous *Noûs* referees.

References

Bealer, G. (2002). Modal epistemology and the rationalist renaissance. In T. Gendler, & J. Hawthorne (Eds.), *Conceivability and possibility* New York: Oxford University Press. (pp. 71–126).

Bealer, G. (2004). The origins of modal error. *Dialectica*, 58(1), 11–42.

Byrne, A. (2007). Possibility and imagination. *Philosophical Perspectives*, 21, 125–144.

Chalmers, D. (1996). *The Conscious mind*. New York: Oxford University Press.

Chalmers, D. (2002). Does conceivability entail possibility? In T. Gendler, & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 145–200). New York: Oxford University Press.

Currie, G. & Ravenscroft, I. (2002). *Recreative minds*. New York: Oxford University Press.

Field, H. (2000). A priority as an evaluative notion. In P. Boghossian, & C. Peacocke (Eds.), *New essays on the a priori* (pp. 117–49). New York: Oxford University Press.

Gendler, T. (2000). The puzzle of imaginative resistance. *Journal of Philosophy,* 97(2), 55–81.

Gendler, T. & Hawthorne J. (2002). Introduction: conceivability and possibility. In T. Gendler, & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 1–70). New York: Oxford University Press.

Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. New York: Oxford University Press.

Harman, G. (1986). *Change in view.* Cambridge, MA: MIT Press.

Jackson, F. (1998). *From metaphysics to ethics*. New York: Oxford University Press.

King, J. (2007). What in the world are the ways things might have been? *Philosophical Studies*, 133(3), 443–453.
Kitcher, P. (2000). A priori knowledge revisited. In P. Boghossian, & P. Benacerraf (Eds.), *New essays on the a priori* (pp. 65–91). New York: Clarenden Press.
Kripke, S. A. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
Nichols, S. (2004). Imagining and believing: The promise of a single code. *Journal of Aesthetics and Art Criticism,* 62(Special issue on Art, Mind, and Cognitive Science), 129–139.
Nichols, S., & Stich, S. (2000). A cognitive theory of pretense. *Cognition,* 74(115–47)
Peacocke, C. (1999). *Being known*. New York: Oxford University Press.
Peacocke, C. (2004). *The realm of reason*. New York: Oxford University Press.
Quine, W. V. O. (1936). Truth by convention. Reprinted in P. Benacerraf and H. Putnam (Eds.), *Philosophical of Mathematics* (pp. 329–354). New York: Cambridge University Press.
Quine, W. V. O. (1962). Carnap and logical truth. Reprinted in P. Benacerraf and H. Putnam (Eds.), *Philosophical of Mathematics* (pp. 355–376). New York: Cambridge University Press.
Sorensen, R. (1996). Modal bloopers: Why believable impossibilities are necessary. *American Philosophical Quarterly*, 33(1): 247–61.
Sosa, E. (2007). *A virtue epistemology*. Oxford: Oxford University Press.
Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
Tye, M. (1995). *The problems of consciousness*. Cambridge, MA: The MIT Press.
van Inwagen, P. (1998). Modal Epistemology. *Philosophical Studies*, 92, 67–84.
Williamson, T. (2007). *The Philosophy of Philosophy*. Malden, MA: Blackwell Publishing.
Yablo, S. (1993). Is conceivability a guide to possibility? *Philosophy and Phenomenological Research, 53*(1), 1–42.
Yablo, S. (2002). Coulda, woulda, shoulda. In T. Gendler, & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 441–492). New York: Oxford University Press.