

Conceivability and Metaphysical Possibility

Eden Lin

A thesis submitted for the degree of Bachelor of Philosophy at the
University of Oxford

Date of award: 2008

Table of Contents

Abstract	2
I. Introduction	3
II. Metaphysical Possibility	5
III. Chalmers's Modal Rationalism	13
Prima Facie Conceivability and Ideal Conceivability	
Negative Conceivability and Positive Conceivability	
Primary Conceivability and Secondary Conceivability	
Two-Dimensionalism: Primary Intensions and Secondary Intensions	
Modal Rationalism	
The Modal Status of Modal Rationalism	
The Objection from Kripkean <i>A Posteriori</i> Necessities	
Aside: Complications Due to Externalism	
IV. Chalmers's Case for Modal Rationalism	57
The Case Against Strong Necessities	
The Argument for the <i>A Priority</i> of Modal Rationalism	
V. Modal Rationalism Is Not <i>A Priori</i>	65
VI. Two Arguments Against Modal Rationalism	74
First Argument	
Second Argument	
Assumptions	
VII. The Conceivability Premises	81
General Considerations	
The First Argument's Conceivability Premises	
The Second Argument's Conceivability Premises	
VIII. Conclusion	97
Works Cited	98

Abstract

David Chalmers maintains that certain kinds of conceivability are reliable guides to metaphysical possibility. In this essay, I argue that Chalmers is wrong about the relationship between conceivability and metaphysical possibility. I begin by explaining the notion of metaphysical possibility and by distinguishing it from those of other kinds of so-called possibility (including epistemic, logical, conceptual, physical, and natural possibility). I then explain Chalmers's view on the relationship between conceivability and metaphysical possibility—a view that he calls *modal rationalism*. Next, I examine his positive case for his view, and I argue that it is unpersuasive. In particular, I argue, his attempt to show that modal rationalism is *a priori* is clearly a failure: given a proper understanding of the notion of metaphysical possibility, the balance of reasons overwhelming favors the hypothesis that modal rationalism is not *a priori*, and we are entitled to assume that it is not *a priori* unless compelling reasons to think otherwise are revealed. Finally, I produce two arguments that show that because modal rationalism is not *a priori*, it is false. I conclude that since there is every reason to think that modal rationalism is not *a priori* and no plausible reason to think that it is *a priori*, and since the falsity of modal rationalism follows from its not being *a priori*, there is every reason to think that modal rationalism is false.

I. Introduction

It is plausible and relatively uncontroversial that *a priori* reasoning can, by itself, give us knowledge of metaphysical possibility and necessity in the following way: if such reasoning reveals a contradiction or incoherence in a proposition, then that proposition is metaphysically impossible.¹ Some philosophers assert a much stronger link between *a priori* reasoning and metaphysical modality, however. In their view, the fact that *a priori* reasoning of a particular sort reveals a proposition to be consistent shows that it is metaphysically possible. In effect, they take what is widely and plausibly regarded as a necessary condition on metaphysical possibility to be a sufficient condition. Let us call the view that *a priori* reasoning and metaphysical modality are connected in this strong manner ‘modal rationalism’.

In this essay, I examine and criticize the modal epistemology of the most prominent advocate of modal rationalism, David Chalmers. I begin by explaining the notion of metaphysical possibility. Then, I present Chalmers’s modal rationalism, which consists of two theses about the relationship between *conceivability* and metaphysical possibility. I then lay out Chalmers’s positive case for his position—namely, that his modal rationalism is *a priori* because there is no notion of metaphysical possibility on which one can coherently

¹ If there are not only contingent *a priori* sentences, but contingent *a priori* propositions, then this principle must be qualified to accommodate those propositions. For in that case, the negations of contingent *a priori* propositions will be metaphysically possible even though *a priori* reasoning reveals them to have a certain kind of incoherence.

suppose that it is false. Finally, I argue that there *is* such a notion of metaphysical possibility, that Chalmers's modal rationalism is not *a priori*, and that because it is not *a priori*, it is false.

II. Metaphysical Possibility

The easiest way to explain the notion of metaphysical possibility is to distinguish it from those of other kinds of so-called possibility and to present various simple hypotheses about the nature and extent of reality, explaining, for each of these hypotheses, what would be metaphysically possible if reality were as it describes.

Let us begin with the former task. The 'must' of metaphysical necessity is obviously neither the 'must' of moral obligation (e.g., 'you must feed the poor') nor the 'must' of epistemic necessity (e.g., 'he must have gone to the store'). Accordingly, the concept of metaphysical possibility is neither that of moral permissibility nor that of epistemic possibility.

The notion of metaphysical possibility ought also to be distinguished from those of the various modalities that some philosophers call 'logical possibility' (of the 'narrow' and 'broad' kinds) and 'conceptual possibility'. The truths of logic (i.e., the logically valid sentences) are said to be logically necessary in the narrow sense. A broader set of truths, which includes (but is not limited to) the truths of mathematics as well as the analytic sentences (e.g., 'red is a color', 'all bachelors are unmarried'), are said to be logically necessary in the broad sense.² Apparently overlapping with the broadly logical necessities are 'conceptual truths', which are said to be 'conceptually necessary' because the concepts they

² This term seems to originate with Plantinga (1974).

involve, together with the laws of logic, ensure their truth.³ Corresponding to these three kinds of necessity are three kinds of possibility: narrowly logical, broadly logical, and conceptual. The notion of metaphysical possibility does not rule out the hypothesis that metaphysical possibility is coextensive⁴ with one of these three sorts of possibility, just as the notion of the tallest man does not rule out the hypothesis that the tallest man is the fattest man. But, as in the case of the notions of the tallest man and the fattest man, the concept of metaphysical possibility is not that of narrowly logical, broadly logical, or conceptual possibility.⁵

The notion of metaphysical possibility is not that of physical possibility either. Nor is it that of natural possibility. A sentence or proposition is physically possible just if its truth is compatible with the laws of physics and it could really be true. A sentence or proposition is naturally possible just if its truth is compatible with the laws of nature and it could really be true.⁶ It is difficult to give illuminating accounts of what it is to be a law of physics and what it is to be a law of nature, but we have at least a rough idea, based on a sense of what

³ I remain neutral on the question of whether broadly logical necessity and conceptual necessity are the same thing. I am reporting other philosophers' talk of different kinds of necessity without endorsing it.

⁴ Two kinds of possibility, F and G, are coextensive just if, for all sentences S, S is F-possible if and only if S is G-possible.

⁵ Bealer (2002, pp. 78-79) says that 'metaphysical possibility' is a synonym for 'logical possibility', but he does not thereby disagree with me, for he thinks that 'logical possibility' just means 'possibility' and has nothing essential to do with logical validity. What I am calling 'logical possibility' is different from the thing that he calls 'logical possibility' and that he identifies with metaphysical possibility.

⁶ I choose to focus on sentences and propositions, but it is fairly easy to see how my account could be modified for other types of entity. For example, a state of affairs is physically possible just if its obtaining is compatible with the laws of physics and it could really obtain, and it is naturally possible just if its obtaining is compatible with the laws of nature and it could really obtain.

counts as natural and what counts as physical. Something is natural just if it is not supernatural—just if it does not involve God, miracles, and other such things. What it is to be physical can be defined with reference to paradigmatic cases of non-sentient physical objects (e.g., rocks, tables, chairs) and paradigmatic cases of physical properties and relations (e.g., mass, spin, charge): the physical properties and relations are those properties and relations that are broadly of *that* kind and that are required to give a complete account of objects like *that*.⁷ As was the case with the logical and conceptual modalities, nothing about the notion of metaphysical possibility rules out its being coextensive with natural possibility (and, if natural possibility and physical possibility are coextensive, with physical possibility). Nevertheless, the concept of metaphysical possibility is not that of physical possibility or that of natural possibility.

What is metaphysical possibility, then? It is best to approach this question by contrasting natural and physical possibility with the logical and conceptual modalities mentioned earlier. Whether a sentence⁸ is logically or conceptually possible is, at bottom, a question of whether it is *consistent*. ‘Not all bachelors are unmarried’ is logically possible in the narrow sense but not in the broad sense.

Its being narrowly logically possible just is its being consistent in the following

⁷ This way of defining the physical is proposed by Jackson (1998, p. 7). It is plausible to suppose that there is no conceptual room for physical things that are not natural, but that there is conceptual room for natural things that are not physical (e.g., the fundamental phenomenal properties of Chalmers (1996), which are neither physical nor supernatural).

⁸ I focus on sentences because on some views of propositions, logical possibility and necessity are properties of sentences but not of propositions. If ‘Hesperus is not Hesperus’ and ‘Hesperus is not Phosphorus’ express the same proposition, then although the former is logically impossible but the latter is not, logical possibility and necessity do not seem to be properties that could be possessed by the proposition they both express.

way: it is not a logically invalid sentence. Its being broadly logically impossible just is its being inconsistent in a different way: given the synonymy of 'bachelor' and 'unmarried man', it expresses a contradiction. Whether a sentence is physically or naturally possible is not a question of whether it is consistent, however. All physically or naturally possible sentences are, in fact, consistent in the two ways just mentioned, but their being consistent is not what it is for them to be physically or naturally possible. To say that a sentence is physically or naturally possible is to say, among other things, that it could really be true. To say that a sentence is physically or naturally necessary is to say that reality must be the way it says it is. However tight the correlation between a sentence's being consistent and its being genuinely capable of being true in reality, and however plausible it is that the former is a necessary condition of the latter, the two are separate issues: the question of whether reality could genuinely be a certain way is a separate question from that of whether its being that way is consistent.⁹ Physical and natural possibility belong to a fundamentally different genus of possibility than the logical and conceptual modalities, the genus of possibility that directly concerns what could be true in reality. Let us call any type of possibility in this genus a type of *real* possibility.

Strictly speaking, real possibility is the only genuine kind of possibility—

⁹ Consistency is sometimes defined in terms of possibility: a set of sentences, it is said, is consistent just if it is possible that the conjunction of its members is true. (Here, the kind of possibility in question is evidently logical.) This gets things backwards: logical possibility is understood in terms of consistency, not the other way around. In any case, even if one defines consistency in terms of logical possibility, my point still holds: the question of whether reality could genuinely be as a certain sentence represents it to be is a separate question from that of whether that sentence is logically possible.

and this is what gives it its name. There is no single genus, possibility, of which real possibility, epistemic possibility, and logical/conceptual possibility are all types: the logical and conceptual types of “possibility” are types of consistency, not possibility, and epistemic “possibility” is just a certain kind of ignorance.¹⁰ (To see that no kind of consistency is a genuine kind of possibility, note the absurdity of saying that consistency with classical sentential logic is a kind of possibility: ‘Everything is F and not F’ is consistent in this way, but surely it is not possible.)¹¹ This drives home the point that whether something is really possible (i.e., whether reality could genuinely be a certain way) is a separate question from whether it is “possible” in some other sense (e.g., not ruled out by what one knows, or consistent according to a certain standard of consistency).

The notion of metaphysical possibility is just that of the broadest type of real possibility: if something is possible in any real modality, it is metaphysically possible. The notion of metaphysical necessity is just that of the strictest type of real necessity: if something is metaphysically necessary, then it is necessary in every real modality. Metaphysical possibility is absolute, unrestricted real possibility. Metaphysical necessity is absolute, unrestricted real necessity. In light of the fact that, strictly speaking, real possibility is the only kind of possibility and real necessity is the only kind of necessity, metaphysical

¹⁰ A number of philosophers agree, though they do not use my terminology of ‘real’ possibility and impossibility. van Inwagen writes that “there is no such thing as logical possibility—not, at least, if it is really supposed to be a species of possibility” (2001, pp. 247-48). And although Bealer claims that the term ‘logical possibility’ has historically been used to mean what I call ‘real possibility’, he argues that the sort of thing that I refer to as ‘logical possibility’ is not a kind of possibility, and he rejects the practice of calling every kind of consistency a kind of possibility (2002, pp. 78-79).

¹¹ Bealer (2002), p. 79.

possibility is ‘possibility *tout court*’ and metaphysical necessity is ‘necessity *tout court*’.¹²

The best way to solidify one's grasp of the notion of metaphysical possibility is to entertain a succession of simple hypotheses about the nature and extent of reality and to determine what is metaphysically possible according to each hypothesis. Imagine a philosopher who believes that there absolutely could not exist anything supernatural, but who nevertheless admits that the existence of supernatural beings is consistent in the broadly logical sense. For her, the sentence ‘a supernatural being exists’ is broadly logically possible but not metaphysically possible: not only are broadly logical possibility and metaphysical possibility different concepts, but they are not coextensive. Moreover, in her view, metaphysical possibility is coextensive with natural possibility, since nothing supernatural could ever exist. Now imagine a philosopher who thinks that the existence of supernatural beings is consistent in the broadly logical sense and who advocates the following thesis, which he considers a highly non-trivial result of extended philosophical reflection: any sentence that is broadly logically possible could be true in reality, and no sentence that is not broadly logically possible could be true in reality. In his view, metaphysical possibility is broader than natural possibility and coextensive with broadly logical possibility. But even for him, it is conceptually distinct from broadly logical possibility, for the claim that they are coextensive is a substantive thesis, not a trivial one. In spite of their disagreement about the breadth of metaphysical possibility, these philosophers share the correct view that whether

¹² See, e.g., van Inwagen (2001), pp. 248-49.

something is metaphysically possible is a separate matter from whether it is broadly logically possible.

The common practice of informally characterizing metaphysical possibility as “what God could bring about, if he existed”¹³ is not only misleading, but it commits those who engage in it to theism in a way of which they seem to be unaware. Those who speak of metaphysical possibility in this way evidently mean thereby to convey an impression of what they take to be the breadth of metaphysical possibility: since God could bring about violations of the laws of nature and physics but could not violate the laws of logic, metaphysical possibility is broader than natural and physical possibility but not broader than logical possibility. But this way of making that claim about the breadth of metaphysical possibility implies that it is metaphysically necessary that God exists: for God, if he existed, could not bring about his never having existed.¹⁴ More importantly, that claim about the breadth of metaphysical possibility is highly controversial, and because (as I have explained) it is not a consequence of the very notion of metaphysical possibility, it is misleading to use it to introduce or explain that notion. Contemporary philosophers are in the habit of giving examples of metaphysically possible things that are physically impossible—they seem particularly fond of massed objects that travel faster than the speed of light—but even if such things are metaphysically possible, this is not guaranteed by the very notion of metaphysical possibility. Such examples are useful if they are meant to flesh out certain substantive views about what is and is not

¹³ For just one example of this widespread practice, see Chalmers (1996), p. 138.

¹⁴ This was pointed out to me by Tim Williamson.

metaphysically possible, but they are highly misleading if they are meant to introduce or elucidate the very concept of metaphysical possibility. The notion of metaphysical possibility is just that of the broadest kind of real possibility—nothing more, nothing less.¹⁵

¹⁵ Gideon Rosen understands metaphysical possibility more or less as I do. He suggests that we “call any modality that is alethic, non-epistemic, and sometimes substantive or synthetic a *real* modality,” and he writes: “among the real modalities, the metaphysical modalities are *absolute* or *unrestricted*. Metaphysical necessity is the strictest real necessity and metaphysical possibility is the least restrictive sort of real possibility in the following sense: If *P* is metaphysically necessary, it is necessary in every real sense: If *P* is really possible in any sense, then it’s possible in the metaphysical sense” (Rosen, forthcoming).

III. Chalmers's Modal Rationalism

Chalmers's modal rationalism is couched in the language of conceivability: the sorts of *a priori* reasoning which, in his view, give us knowledge of metaphysical possibility are varieties of conceiving. It would be misleading to say that Chalmers thinks that anything that is conceivable is metaphysically possible, however. To see just what his modal rationalism consists in, one must see which kinds of conceivability he has in mind, and one must grasp two important distinctions in his framework: that between primary and secondary intensions, and that between primary and secondary possibility.

Chalmers takes conceivability to be a property of sentence tokens, or as he calls them, *statements*.¹⁶ Although he does not explicitly say what kind of sentence type he thinks statements are tokens of, it is plausible that he regards statements as tokens of *linguistic* sentence types¹⁷—where two expression tokens are tokens of the same linguistic sentence type if and only if they are

¹⁶ Remarkably, although the term 'statement' is ubiquitous in Chalmers's corpus, he does not say what he means by it. However, he says the following two things: (i) statements are the entities whose primary and secondary intensions return truth values at the worlds where they are evaluated [Chalmers (2002a), pp. 163-64], and (ii) sentence tokens are the entities whose primary and secondary intensions return truth values at the worlds where they are evaluated [Chalmers (2006), p. 78]. (As I shall explain later, primary and secondary intensions are certain semantic values which are, in Chalmers's view, possessed by expression tokens.) It is clear, then, that he thinks that statements are sentence tokens.

¹⁷ I conjecture this because when Chalmers explains why the bearers of primary and secondary intensions are expression tokens rather than expression types, he writes: "It is often the case that two tokens of the same linguistic type can have *different* epistemic [i.e., primary] intensions" [Chalmers (2006), p. 95]. If the expression tokens that possess primary and secondary intensions are tokens of linguistic expression types, then surely the sentence tokens that possess primary and secondary intensions are tokens of linguistic sentence types.

tokens of the same sentence type in the same language. (Thus, two tokens of the English sentence 'water is H₂O' are tokens of the same linguistic sentence type, but no expression tokens in other languages that are orthographically identical to these two tokens would be of the same linguistic sentence type as them: if there were a language different from English in which the orthographic string 'water is H₂O' was a sentence, tokens of that non-English sentence would not be tokens of the same linguistic sentence type as the two aforementioned tokens of the English sentence 'water is H₂O'.)¹⁸ According to Chalmers, there are three dimensions along which the claim that a statement is conceivable can be disambiguated: *prima facie* versus ideal conceivability, negative versus positive conceivability, and primary versus secondary conceivability.¹⁹

PRIMA FACIE CONCEIVABILITY AND IDEAL CONCEIVABILITY

A statement *S* is *prima facie* conceivable for a given subject just if, after some consideration, that subject judges that *S* is conceivable, on one of the substantive accounts of conceivability to be discussed below (primary positive, primary negative, etc.). A statement is *ideally* conceivable just if the judgment that it is conceivable cannot be defeated by better reasoning. The notions of better reasoning and of undefeatability by better reasoning are supposed to be intuitive: the reasoning that yields the judgment that $32 \times 43 = 1,376$ is better

¹⁸ My definition of a linguistic sentence type is based on Chalmers's discussion of linguistic expression types in Chalmers (2006), pp. 66-67.

¹⁹ Chalmers (2002a), pp. 146-47.

than the one that yields the judgment that $32 \times 43 = 1,375$, and while the latter judgment is defeatable by better reasoning, the former is not.²⁰ Prima facie conceivability is subject-relative: a statement S may be prima facie conceivable for a relatively poor reasoner without being prima facie conceivable for a better reasoner. Ideal conceivability is not subject-relative, since the idea of undefeatability by better reasoning idealizes away from the cognitive limitations of particular subjects.²¹ Although Chalmers considers prima facie conceivability a variety of conceivability, it would be consonant with the spirit of his view to think of it as merely *apparent* conceivability, and to reserve the title of conceivability proper for ideal conceivability: if S is prima facie conceivable for me but not ideally conceivable, S appears to me to be conceivable, but it is not really conceivable.²² Since Chalmers does not think that any kind of merely prima facie conceivability is a guide to metaphysical possibility, I will ignore prima facie conceivability in what follows, and I will use ‘conceivable’ and ‘conceivability’ to stand for ‘ideally conceivable’ and ‘ideal conceivability’.

²⁰ My example.

²¹ At one point, Chalmers seems to say that because different subjects can associate different primary intensions with the same sentence, even ideal conceivability is subject-relative [Chalmers, (2002a), p. 172]. Presumably, what he means is that the ideal conceivability of a sentence *type* is subject-relative, since different tokens of the same sentence type, uttered by different subjects, might have different primary intensions. But since we are concerned with the ideal conceivability of statements (i.e., sentence *tokens*), and since each statement presumably only has one primary intension, we have no reason to suppose that the ideal conceivability of statements is subject-relative.

²² Chalmers (2002a), pp. 147-49.

NEGATIVE CONCEIVABILITY AND POSITIVE CONCEIVABILITY

A statement *S* is *negatively* conceivable just if *S* cannot be ruled out *a priori*—where the specific kind of ruling out will vary between the primary and secondary versions of negative conceivability, which I shall explain in the next section.²³ A statement *S* is *positively* conceivable just if one can coherently imagine a situation that verifies *S*—where “a situation is (roughly) a configuration of objects and properties within a world.”²⁴ Here, imagination can be perceptual (as when one imagines a flying pig), but it need not be: one can imagine situations “beyond the scale of perception” (e.g., molecules of H₂O), situations that are “unperceivable in principle” (e.g., “the existence of an invisible being that leaves no trace on perception”), and distinct but perceptually indistinguishable situations (e.g., “the existence of a conscious being and its zombie twin”).²⁵ A situation is *coherently* imaginable just if “it is possible to fill in arbitrary details in the imagined situation such that no contradiction reveals itself.”²⁶ Thus, although one can arguably imagine the geometrically impossible situations depicted in the paintings of M. C. Escher, one cannot coherently imagine them: when one tries to fill in the details of how the surfaces are put together to make three-dimensional objects, one finds a contradiction.²⁷ A situation *verifies* *S* just if it is incoherent to suppose that the situation obtains

²³ Chalmers (2002a), p. 149.

²⁴ Chalmers (2002a), p. 151.

²⁵ Chalmers (2002a), p. 151.

²⁶ Chalmers (2002a), p. 153.

²⁷ My example.

without S's being true—or, equivalently, just if ideal rational reflection reveals that the situation is one in which S is true. Verification, in this sense, is not just an evidential relation, but is rather “broadly analogous to an entailment of one statement by another.”²⁸ An imagined situation in which a typically honest philosopher claims that his favorite number is 2 does not verify the statement ‘2 is the favorite number of at least one philosopher’ even though it supports it evidentially: after all, it is coherent to suppose that, in spite of his usual honesty, he is not telling the truth. By contrast, an imagined situation in which 2 is the favorite number of a certain philosopher does verify ‘2 is the favorite number of at least one philosopher’, since it is incoherent to suppose that the situation obtains even though that statement is false.

In describing positive conceivability as he does, Chalmers takes himself to be giving an account of a type of conceivability that philosophers are already familiar with (even if not under the name ‘positive conceivability’) rather than introducing and defining a new kind of conceivability of his own discovery or invention. “A typical philosophical thought experiment starts with *prima facie* positive conceivability,” he writes: the philosopher imagines a situation “with certain important features specified” but not in “fine detail,” and he judges, after some consideration, that this situation verifies a certain statement S. (For example, one imagines a situation in which Bill happens by chance to judge correctly that a particular object in his field of vision is a barn while driving in an area replete with fake barns that are perceptually indistinguishable from the object he takes to be a barn, and one judges that this imagined situation verifies

²⁸ Chalmers (2002a), p. 152.

the statement ‘Bill does not know that the object he sees is a barn’.) “For the thought experiment to yield the intended conclusion, this *prima facie* judgment must be correct, so that *S* is ideally positively conceivable”: the philosopher must be capable of filling in arbitrary details in the situation without any contradiction revealing itself, and it must be truly incoherent (not merely apparently so) to suppose that the situation obtains without *S*’s being true.²⁹ Since philosophers are already familiar with the process of conducting a typical philosophical thought experiment, they are already familiar with what Chalmers calls ‘positive conceivability’.³⁰ And since Chalmers’s account of positive conceivability is meant to describe something of which philosophers already have an understanding (however implicit and imprecise), it is in principle possible for that account to be incorrect: just as a particular philosophical account of, say, knowledge can be wrong because it contradicts our understanding of knowledge, so can Chalmers’s account of positive conceivability be incorrect because it contradicts philosophers’ understanding of the typical method of philosophical thought experiments. Moreover, it is legitimate for us to use our preexisting understanding of that method to guide our interpretation of Chalmers’s account of positive conceivability: the principle of charity enjoins us to interpret his account in a way that accords with what we know to be true about that method. This will be important later, when I argue against Chalmers’s modal rationalism.

²⁹ Chalmers (2002a), pp. 153-54.

³⁰ Chalmers also writes: “The characterization of positive conceivability that I have given here... cannot be considered a reductive definition. At best, it is something of a clarification. Nevertheless, there seems to be a reasonably clear intuitive notion in the vicinity, of which most people seem to have a grasp. It may be that the notion can be given a more rigorous definition, or it may be that it should be taken as primitive; this is one of the central open questions in the area.” (2002a, p. 156)

PRIMARY CONCEIVABILITY AND SECONDARY CONCEIVABILITY

There is a sense in which the negations of Kripkean *a posteriori* necessities (e.g., 'Water is not H₂O' and 'Hesperus is not Phosphorus') are conceivable, and there is a sense in which they are not. A statement S is *primarily* conceivable just if it is conceivable in the sense in which the negations of Kripkean *a posteriori* necessities are conceivable. A statement S is *secondarily* conceivable just if it is conceivable in the sense in which the negations of Kripkean *a posteriori* necessities are not conceivable.³¹

For all we can know *a priori*, water is not H₂O in the actual world; and for all we can know *a priori*, Hesperus is not Phosphorus in the actual world.³² Indeed, we can imagine many situations in which water is not H₂O and which actually obtain for all we can know *a priori* (e.g., a situation in which the dominant, clear liquid found in the oceans and the lakes around here is a chemical kind, XYZ, distinct from H₂O). We can also imagine many situations in

³¹ Chalmers (2002a), pp. 156-57; (2005), section 2.

³² Scott Soames and Nathan Salmon think that the proposition that water is H₂O is identical to the proposition that water is water, and that the proposition that Hesperus is Phosphorus is identical to the proposition that Hesperus is Hesperus, and they reason that since we know *a priori* that water is water and that Hesperus is Hesperus, we know *a priori* that water is H₂O and that Hesperus is Phosphorus. In their view, our inclination to think that we do not know the latter propositions *a priori* is a consequence of the fact that we do not know *a priori* that the sentences 'water is H₂O' and 'Hesperus is Phosphorus' express the propositions that they do. (For Salmon's version of this view, see Salmon (1986).) Chalmers rejects Soames's and Salmon's view, and for ease of exposition, I shall assume in this essay that he is right to do so. If Soames and Salmon are right, one will have to present Chalmers's view wholly in terms of epistemic relations to statements rather than in terms of propositional attitudes. Such a re-writing would not be trivial, and it might reveal additional complexities, but I will not undertake it in this essay.

which Hesperus is not Phosphorus and which actually obtain for all we can know *a priori* (e.g., a situation in which Mars and Venus are distinct, as they actually are, and in which the heavenly body that appears where Venus actually appears in the morning is Mars while the heavenly body that appears where Venus actually appears in the evening is Venus). In this sense, 'water is not H₂O' and 'Hesperus is not Phosphorus' are conceivable.³³ It is this kind of conceivability that philosophers have in mind when they say that the negations of Kripkean *a posteriori* necessities are counterexamples to the claim that conceivability entails metaphysical possibility: those sentences are conceivable in this sense, but they are metaphysically impossible. It is also this kind of conceivability that one is likely to have in mind when one first encounters Chalmers's discussion of the difference between negative and positive conceivability: when Chalmers explains that a statement is negatively conceivable just if it cannot be ruled out *a priori*, and when he says that a statement is positively conceivable just if one can coherently imagine a situation which verifies it, one tends to interpret him as talking about the negative and positive versions of the sense of conceivability in which 'water is not H₂O' and 'Hesperus is not Phosphorus' are conceivable.

There is also a sense in which 'water is not H₂O' and 'Hesperus is not Phosphorus' are inconceivable, however. It is inconceivable in this sense to

³³ Given that Chalmers takes conceivability to be a property of statements rather than of sentence types, this should be put more precisely as: "in this sense, typical tokens of the English sentence types 'water is not H₂O' and 'Hesperus is not Phosphorus' are conceivable." I define the notion of a typical token of an expression type using the notion of a primary intension, which I shall explain later: a token T of an expression type K is a *typical* token of K just if the primary intension of T is the primary intension typically associated with tokens of K. In what follows, when I speak loosely of expressions by using quote-names, I should be read as speaking of typical tokens of the English expression types named by those quote-names.

anyone who knows that water is H₂O that water might counterfactually have not been H₂O: even if there were some non-H₂O stuff that was outwardly indistinguishable from water and that played the role that water actually plays in our environment, given the fact that water is H₂O, this stuff could not conceivably be water. Likewise, it is inconceivable in this sense to anyone who knows that Hesperus is Phosphorus that Hesperus might counterfactually have failed to be Phosphorus: even if the heavenly body that appears where Venus actually appears in the morning were distinct from the heavenly body that appears where Venus actually appears in the evening, given the fact that Hesperus is Phosphorus, Hesperus could not conceivably have failed to be Phosphorus. As far as this kind of conceivability is concerned, 'water is not H₂O' and 'Hesperus is not Phosphorus' are inconceivable. Indeed, it is plausible that philosophers come to agree with Kripke's claims that 'water is H₂O' and 'Hesperus is Phosphorus' are metaphysically necessary precisely because they assume that conceivability in this sense is a necessary condition of metaphysical possibility and they judge that 'water is not H₂O' and 'Hesperus is not Phosphorus' are inconceivable in this sense.

The first of these two senses of conceivability, on which 'water is not H₂O' and 'Hesperus is not Phosphorus' are conceivable, is primary conceivability. The second of these two senses of conceivability, on which 'water is not H₂O' and 'Hesperus is not Phosphorus' are inconceivable, is secondary conceivability. Although our initial fix on the distinction between these two kinds of conceivability is given by examples, Chalmers provides us with explicit, general accounts of primary and secondary conceivability (of both the negative and

positive varieties).

A statement *S* is *primarily negatively* conceivable just if its negation, $\sim S$, is not *a priori*.³⁴ Thus, since ‘water is H₂O’ and ‘Hesperus is Phosphorus’ are not *a priori*, their negations, ‘water is not H₂O’ and ‘Hesperus is not Phosphorus’, are primarily negatively conceivable. As I remarked earlier, this is the most natural reading of Chalmers’s claim that a statement is negatively conceivable just if it cannot be ruled out *a priori*: when one encounters Chalmers’s definition of negative conceivability, one thinks by default of primary negative conceivability. Indeed, even before he distinguishes primary and secondary conceivability, Chalmers writes: “*S* is ideally negatively conceivable when it is not *a priori* that $\sim S$.”³⁵

A statement *S* is *secondarily negatively* conceivable just if *a priori* reflection and complete empirical nonmodal knowledge cannot rule out the hypothesis that *S* might (counterfactually) have been true.³⁶ To see why it would not be enough to say that a statement *S* is secondarily negatively conceivable just if *a priori* reflection cannot rule out the hypothesis that *S* might have been true (or, equivalently, just if it is *a priori* that *S* could not have been true), consider an ideal reasoner who is ignorant of the fact that Hesperus is Phosphorus. Even though ‘Hesperus is not Phosphorus’ is not secondarily negatively conceivable,

³⁴ Chalmers (2002a), p. 158: “We can say that *S* is primarily negatively conceivable when it is not ruled out *a priori* that *S* is actually the case, or, more briefly, if *S* is not ruled out *a priori*.” For simplicity’s sake, I assume in this essay that bivalence is true and I ignore Chalmers’s brief discussion of indeterminate statements at (2002a), pp. 149-50. Given this assumption, to say that *S* is not ruled out *a priori* is just to say that $\sim S$ is not *a priori*.

³⁵ Chalmers, (2002a), p. 149.

³⁶ Chalmers (2002a), p. 159.

for this ideal reasoner, no amount of *a priori* reflection can rule out the hypothesis that Hesperus might not have been Phosphorus. ‘Hesperus is not Phosphorus’ is inconceivable in the sense that we have in mind only in light of the knowledge that Hesperus is Phosphorus: secondary conceivability is conceivability in light of empirical knowledge. Not just any empirical knowledge is sufficient, however. We can suppose that our ideal reasoner, who cannot rule out *a priori* that Hesperus might not have been Phosphorus, knows all empirical nonmodal facts other than the fact that Hesperus is Phosphorus (and any facts that imply it). *A priori* reflection and knowledge of all empirical nonmodal facts other than this one (and those that imply it) cannot rule out the hypothesis that Hesperus might not have been Phosphorus, but ‘Hesperus is not Phosphorus’ is secondarily negatively inconceivable. Secondary negative conceivability must therefore be conceivability in light of complete empirical nonmodal knowledge.³⁷

Chalmers’s accounts of the positive versions of primary and secondary conceivability rely on the distinction between two ways one can consider a situation when one imagines it in the way relevant to positive conceivability. To *consider a situation as actual* is to consider it as a way the world might actually be, for all that we can know *a priori*. To *consider a situation as counterfactual* is to consider it as a way the world might have been, given all the empirical nonmodal facts about how it actually is. Suppose that R is a situation in which the dominant clear liquid that fills the oceans and the lakes around here is XYZ. When R is

³⁷ The empirical knowledge in question is restricted to *nonmodal* knowledge to avoid trivializing the link between secondary conceivability and possibility. See Chalmers (2002a), p. 159.

considered as actual, it is (according to Chalmers) correct to describe it as a situation in which water is XYZ. When R is considered as counterfactual, though, it is incorrect to describe it as a situation in which water is XYZ. For to consider a situation as counterfactual is to consider it as a way the world might have been, given all the empirical nonmodal facts about how it actually is; and given the empirical fact that water actually is H₂O, water could not have been XYZ. Thus, no situation considered as counterfactual can correctly be described as one in which water is XYZ—though, of course, a subject who is ignorant of some of the empirical nonmodal facts about the actual world (including the fact that water is actually H₂O) might, on account of his ignorance, mistakenly judge that when R is considered as counterfactual, it can correctly be described as a situation in which water is XYZ.³⁸

In general, when one describes a situation, one must specify whether one is describing it when it is considered as actual or when it is considered as counterfactual. For example, one cannot simply say, 'U is a situation in which water is H₂O', for this utterance is ambiguous between two different thoughts: the thought that U is a situation in which the dominant clear liquid that fills the oceans and the lakes around here is H₂O (or so, at least, Chalmers thinks), or the thought that U is a situation in which H₂O is H₂O. Instead, one must either say 'when considered as actual, U is a situation in which water is H₂O' (in which case one expresses the former thought) or 'when considered as counterfactual, U is a situation in which water is H₂O' (in which case one expresses the latter). There are exceptions to this rule, however. Intuitively, if a description D is such that, for

³⁸ Chalmers (2002a), p. 157.

any situation U, D says the same thing about U when U is considered as actual as it does about U when U is considered as counterfactual, then it is acceptable to use D to describe a situation without specifying whether one is considering the situation as actual or as counterfactual. Arguably, the descriptions ‘a situation in which the dominant clear liquid that fills the oceans and the lakes around here is H₂O’ and ‘a situation in which H₂O is H₂O’ meet this condition—or so, at least, Chalmers thinks. That is why it is (arguably) acceptable for one to use these descriptions to distinguish the two thoughts that might be expressed by ‘situation U is one in which water is H₂O’. Likewise, the description ‘a situation in which the dominant clear liquid that fills the oceans and the lakes around here is XYZ’ arguably meets this condition as well. That is why one can arguably speak of a situation R in which the dominant clear liquid that fills the oceans and the lakes around here is XYZ, without specifying whether one’s description of R is meant to apply to it when it is considered as actual or when it is considered as counterfactual.³⁹

The clauses ‘for all that we can know *a priori*’ and ‘given all the empirical nonmodal facts about how it actually is’ are crucial to the definitions of these two ways of considering a situation, and it is somewhat misleading for Chalmers to give these two ways of considering a situation names that do not remind us of the inclusion of these clauses in their definitions. It is crucial that we ignore

³⁹ It is beyond the scope of this essay to characterize more precisely those descriptions which, intuitively, say the same thing about a situation whether it is considered as actual or considered as counterfactual. They must, at minimum, be semantically neutral (to use terminology that I shall define in the next section), but semantic neutrality of the basic kind that I shall define is not enough to make them behave as they must if we are to use them to describe situations without specifying whether we are considering them as actual or as counterfactual.

everything except what we can know *a priori* when we consider a situation as actual—that we consider it not just as a way the world might be for all we know, but as a way the world might be for all we can know *a priori*. For if ignoring everything but what we can know *a priori* were not required for considering a situation as actual, then given that we know *a posteriori* that water is H₂O, R could not correctly be described as a situation in which water is XYZ even when considered as actual. Likewise, it is crucial that all the empirical nonmodal facts about how the world actually is be taken into account when we consider a situation as counterfactual. For if we were not required to take all these facts into account when considering a situation as counterfactual, then we could ignore the fact on account of which it is incorrect to describe R, when it is considered as counterfactual, as a situation in which water is XYZ—namely, the fact that water is actually H₂O—and it would be acceptable so to describe R when considering it as counterfactual. It is therefore crucial to remember that part of the distinction between considering a situation as actual and considering a situation as counterfactual is that, in the case of the former, one ignores everything but what we can know *a priori*, while in the case of the latter, one takes into account all the empirical nonmodal facts about the actual world.

A statement *S* is *primarily positively* conceivable just if one can coherently imagine a situation which, when considered as actual, verifies *S*.⁴⁰ A coherently imaginable situation verifies a statement *S* when considered as actual just if it is incoherent to suppose that the situation actually obtains without *S*'s being true.⁴¹

⁴⁰ Chalmers (2002a), p. 157.

⁴¹ Chalmers (2002a), pp. 157-58.

For example, according to Chalmers, a situation R in which the dominant clear liquid that fills the oceans and the lakes around here is XYZ verifies the statement 'water is XYZ' when considered as actual: for when it is considered as actual, R is correctly described (in Chalmers's view) as a situation in which water is XYZ, and it is incoherent to suppose that R actually obtains without its being the case that 'water is XYZ' is true. As I indicated earlier, primary positive conceivability is what one naturally thinks of when one first encounters Chalmers's account of positive conceivability: when he says that a statement is positively conceivable just if one can coherently imagine a situation that verifies it, one naturally interprets him as talking about situations considered as actual.

A statement S is *secondarily positively* conceivable just if one can coherently imagine a situation which, when considered as counterfactual, verifies S.⁴² A coherently imaginable situation verifies a statement S when considered as counterfactual just if ideal rational reflection and complete empirical nonmodal knowledge reveal that S would have been true if the situation had obtained.⁴³ For example, when considered as counterfactual, a situation R in which the dominant, clear liquid that fills the oceans and the lakes around here is XYZ verifies 'water is H₂O' but fails to verify 'water is XYZ': for when considered as counterfactual (i.e., as a way the world might have been, given all the empirical nonmodal facts about how it actually is), R is correctly described as a situation in which water is H₂O but is not XYZ, and although it is incoherent to suppose that 'water is H₂O' would not have been true had R

⁴² Chalmers (2002a), p. 157.

⁴³ Chalmers (2002a), pp. 158-59.

obtained, it is coherent to suppose that 'water is XYZ' would not have been true had R obtained. An ideally rational subject who is ignorant of certain empirical nonmodal facts (including the fact that water is H₂O) might incorrectly judge that one can coherently imagine a situation which, when considered as counterfactual, verifies 'water is XYZ'. Thus 'water is XYZ' might be *prima facie* secondarily positively conceivable for such a subject. In reality, however, one cannot coherently imagine a situation which, when considered as counterfactual, verifies 'water is XYZ', so this statement is not secondarily positively conceivable.

To get the hang of Chalmers's four-fold distinction of conceivability into the primary positive, primary negative, secondary positive, and secondary negative varieties, it will be useful to look at a few statements and to determine in which, if any, of these four senses of conceivability each of them is conceivable according to Chalmers.

Let S_1 be a typical token of the English sentence 'water is not water'. S_1 is not primarily negatively conceivable, because its negation, 'water is water', is *a priori*. S_1 is not secondarily negatively conceivable, because *a priori* reflection and complete empirical nonmodal knowledge can rule out the hypothesis that it might have been true. (Indeed, since it is *a priori* that water could not have failed to be water, *a priori* reflection alone can rule out the hypothesis that S_1 might have been true.) S_1 is not primarily positively conceivable because one cannot coherently imagine a situation that verifies it when considered as actual: every coherently imaginable situation, when considered as actual, is correctly described as one in which water is water. S_1 is not secondarily positively conceivable, because one cannot coherently imagine a situation that verifies it

when considered as counterfactual: every coherently imaginable situation, when considered as counterfactual, is correctly described as one in which water is water.

Let S_2 be a typical token of the English sentence 'water is not H_2O '. S_2 is primarily negatively conceivable, because its negation, 'water is H_2O ' is not *a priori*. S_2 is not secondarily negatively conceivable, however, because *a priori* reflection and complete empirical nonmodal knowledge can rule out the hypothesis that it might have been true: given that water is H_2O , one can rule out *a priori* that water might not have been H_2O .⁴⁴ S_2 is primarily positively conceivable, because one can coherently imagine a situation that verifies it when considered as actual—namely, the situation R, in which the dominant, clear liquid that fills the oceans and the lakes around here is XYZ. When considered as actual, R is correctly described as a situation in which water is XYZ and hence not H_2O ; thus, it is incoherent to suppose that R actually obtains without S_2 's being true. S_2 is not secondarily positively conceivable, though, because one cannot coherently imagine a situation which verifies it when considered as counterfactual. Given all the empirical nonmodal facts, including the fact that water is H_2O , every coherently imaginable situation (including R) will, when considered as counterfactual, correctly be described as one in which water is H_2O . Thus, every coherently imaginable situation will fail to verify S_2 when considered as counterfactual: every such situation will be such that ideal rational reflection and complete empirical nonmodal knowledge do not reveal that S_2

⁴⁴ I assume (as most philosophers do) that the necessity of identity is *a priori*—i.e., that it is *a priori* that for any x and any y , if $x = y$, then it is metaphysically necessary that $x = y$.

would have been true if it had obtained.

Let S_3 be a typical token of the English sentence 'water does not exist'. S_3 is primarily negatively conceivable because its negation, 'water exists', is not *a priori*. S_3 is secondarily negatively conceivable because *a priori* reflection and complete empirical nonmodal knowledge cannot rule out the hypothesis that it might have been true: from the totality of empirical nonmodal facts, it does not follow by *a priori* reflection that S_3 could not have been true. S_3 is primarily positively conceivable, because one can coherently imagine a situation which verifies it when considered as actual—e.g., a situation in which nothing clear and nothing liquid exists. S_3 is secondarily positively conceivable too, because one can coherently imagine a situation that verifies it when considered as counterfactual—e.g., a situation which, when considered as counterfactual, is correctly described as one in which H_2O does not exist.

It is worth noting that the following relationship holds between primary positive conceivability and primary negative conceivability:

(PN) For any statement S , if S is primarily positively conceivable, then S is primarily negatively conceivable.

Unpacking the definitions of primary positive conceivability and primary negative conceivability, this principle—which we might call the *posneg* principle—amounts to:

(PN) For any statement S , if one can coherently imagine a situation which, when considered as actual, verifies S , then $\sim S$ is not *a priori*.

The proof of this principle is as follows. A statement S is verified by a coherently

imaginable situation considered as actual just if it is incoherent to suppose that that situation actually obtains without S's being true. For any statement S, if S is *a priori*, then it is incoherent to suppose that S is not true. If it is incoherent to suppose that a statement S is not true, then for any coherently imaginable situation, it is incoherent to suppose both that that situation actually obtains and that S is not true. Thus, for any statement S, if S is *a priori*, then for any coherently imaginable situation, that situation verifies S when considered as actual. Now, let T be any statement that is primarily positively conceivable. If T is primarily positively conceivable, then there is a particular coherently imaginable situation, D, which verifies it when considered as actual. If $\sim T$ is *a priori*, then $\sim T$ is verified by D when D is considered as actual. Thus, if $\sim T$ is *a priori*, then when D is considered as actual, D verifies both T and $\sim T$. But no coherently imaginable situation can verify two contradictory statements when considered as actual, so D does not verify both T and $\sim T$ when considered as actual. Since, by assumption, T is verified by D when D is considered as actual, $\sim T$ is not verified by D when D is considered as actual. Thus, $\sim T$ is not *a priori*. By the definition of primary negative conceivability, it follows that T is primarily negatively conceivable. Since T stands for any primarily positively conceivable statement, the posneg principle is true.⁴⁵

One cannot know *a priori* whether a certain statement is secondarily conceivable. For in order to know whether it is secondarily conceivable, one must know whether, in light of all the empirical nonmodal facts about the actual

⁴⁵ Chalmers affirms the posneg principle, but to my knowledge, he does not give an argument for it. (See Chalmers (2002a), p. 155.) The argument above is my own.

world, it is conceivable that it might have been true—and this can only be known *a posteriori*.⁴⁶ By contrast, one can know *a priori* whether a certain statement is primarily conceivable. For when we attempt to determine whether it is primarily conceivable, we need not rely on any *a posteriori* knowledge: we need only determine whether its negation is *a priori* (in the negative case) or whether a coherently imaginable situation that might actually obtain for all we know *a priori* verifies it (in the positive case)—and all this can be determined *a priori*.⁴⁷

Because all knowledge of secondary conceivability is *a posteriori*, any knowledge of metaphysical modality inferred from knowledge of secondary conceivability will be *a posteriori*. Thus, of the four kinds of conceivability distinguished by Chalmers, only primary positive conceivability and primary negative conceivability have the potential to give us *a priori* knowledge of metaphysical modality. Chalmers does not think that the primary conceivability of a statement (either positive or negative) entails its metaphysical possibility, though. “[P]rimary conceivability does not entail metaphysical possibility: ‘water is not H₂O’ is primarily conceivable, but it is not metaphysically possible.”⁴⁸ To see why he says this and how he can still maintain that knowledge of primary conceivability can give us *a priori* knowledge of metaphysical modality, one must understand two crucial distinctions in Chalmers's framework: that between primary and secondary intensions, and that between primary and secondary possibility.

⁴⁶ Chalmers (2002a), p. 159.

⁴⁷ Chalmers (2002a), p. 158.

⁴⁸ Chalmers (2005), section 2.

TWO-DIMENSIONALISM: PRIMARY INTENSIONS AND SECONDARY INTENSIONS

It is plausible that part of the meaning of an expression token is an intension—a function from possible states of the world to extensions that reflects the way in which the token’s extension depends on how the world is.⁴⁹ According to Chalmers, there are *two* ways in which the extension of an expression token depends on the state of the world, and hence two intensions associated with each expression token: the token’s primary intension and its secondary intension.

To see why it is plausible that there are no less than two ways in which the extension of an expression token depends on the state of the world, consider a typical token of the English expression ‘water’. As Kripke argued, given that water is H₂O, there is no possible state of the world such that, had it obtained, water would not have been H₂O: even if the dominant clear liquid in the oceans and lakes around here were filled with XYZ, for example, water would be H₂O. Thus, there is a sense in which the extension of that token does not vary from one possible state of the world to another: given that its actual extension is H₂O, no matter what the world were like, its extension would still be H₂O. One might think, then, that insofar as the dependence (or independence) of that token’s extension on the state of the world can be represented by an intension, that token’s intension should be a constant function that returns H₂O at every world. This intension is what Chalmers calls the secondary intension of a typical token

⁴⁹ Like Chalmers, I regard the truth value of a statement at a world as its extension there.

of 'water'.

According to Chalmers, however, there is another pattern of dependence of the extension of our token of 'water' on the state of the world. We do not know *a priori* that water is H_2O : for all we know *a priori*, the scientists are wrong, and water is something other than H_2O . Suppose it turns out that the dominant clear liquid that fills the oceans and the lakes around here is actually XYZ. Then it will turn out (according to Chalmers) that the extension of that token is not H_2O but XYZ. But if it turns out that the scientists are right, and the dominant clear liquid that fills the oceans and the lakes around here is H_2O , then it will turn out that that token's extension is H_2O . It seems, then, that there is a different pattern of dependence of the extension of that token on the state of the world, one that can be represented by a function that does *not* return H_2O at every world at which it is evaluated—a function that returns XYZ at worlds where (roughly) the dominant clear liquid that fills the oceans and the lakes around here is XYZ and H_2O where the dominant clear liquid that fills the oceans and the lakes around here is H_2O . This is what Chalmers calls the primary intension of a typical token of 'water'.

When one thinks about it, it makes sense that a typical token of 'water' would have a primary intension in addition to a secondary intension. The secondary intension of such a token returns the same extension at every possible state of the world. Thus, it captures a way in which the extension is such a token is independent of the state of the world and invariant. However, there clearly seems to be a way in which the extension of a typical token of 'water' depends on the state of the world and varies in accordance with variations in it: it turns out

to be different things depending on which possible state of the world turns out to be actual. The primary intension of typical tokens of 'water' captures this pattern of dependence.

The *primary* intension of an expression token is the function from possible worlds to extensions that reflects the way in which the token's extension depends on which possible world turns out to be actual: at each world where it is evaluated, it returns the extension that the token will turn out to have if that world turns out to be the actual world. To borrow terminology from our discussion of primary and secondary conceivability, the primary intension of an expression token T is the function such that, at any world W where it is evaluated, it returns the extension of T at W when W is considered as actual (i.e., considered as a way the world might actually be, for all we can know *a priori*).⁵⁰

51

⁵⁰ Chalmers's official account of primary intensions, which is best presented in (2006), requires a highly elaborate conceptual framework that is beyond the scope of this essay to explore more fully. Certain aspects of his account, including the role of what he calls 'canonical descriptions', are not described in this essay, since they do not bear on my argument.

⁵¹ Chalmers is adamant that the evaluation of the primary intension of a token T at a world W does not require the existence of T at W [Chalmers (1996), p. 366 n26; (2002a), p. 168; (2006), p. 81]. He insists, for example, that the primary intension of a particular typical token of 'water' returns XYZ at a world where the dominant clear liquid that fills the oceans and the lakes around here is XYZ, even if that token does not exist at that world. One might think, however, that this is ruled out by the account of primary intensions just given. For one might think that a token cannot have an extension unless it exists, and thus that if T does not exist at W, then if W turns out to be actual, T will turn out not to have an extension. If this line of reasoning is correct, then one will have to modify the account of primary intensions given above. Perhaps one could say the following (where the first universal quantifier is interpreted substitutionally): for any S, and for any world W, if 'S' is a statement, then the primary intension of 'S' is true at W just in case if W turns out to be actual, it turns out that S. ("For any world W, if 'water is XYZ' is a statement, then the primary intension of 'water is XYZ' is true at W just in case if W turns out to be actual, it turns out that water is XYZ" is a substitution instance of that sentence, and it seems to

The *secondary* intension of an expression token is the function from possible worlds to extensions that reflects the way in which the token's extension at counterfactual worlds is determined, given that its actual extension and those of its parts (if it is made up of other expression tokens) are already fixed: at each world where it is evaluated, it will return the extension that the token would have had if that world had obtained, given that the token and its parts have the extensions that they actually have.⁵² The secondary intension of an expression token T is the function such that, at any world W where it is evaluated, it returns the extension of T at W when W is considered as counterfactual (i.e., considered as a way the world might have been, given all the empirical nonmodal facts about how it actually is).⁵³

capture what Chalmers has in mind as far as the primary intension of 'water is XYZ' goes.) More generally, one might say the following (where, again, the first universal quantifier is interpreted substitutionally): for any T, and for any world W, and for any x, if 'T' is an expression token, then the primary intension of 'T' is x at W just in case if W turns out to be actual, then T turns out to be x. ("For any world W, and for any x, if 'water' is an expression token, then the primary intension of 'water' is x at W just in case if W turns out to be actual, then water turns out to be x" is a substitution instance of this sentence, and it seems to capture what Chalmers has in mind as far as the primary intension of 'water' is concerned.) I doubt that these proposals will ultimately withstand scrutiny, however. Chalmers does not give us a rigorous, general account of primary intensions on which it is perfectly clear how the primary intension of a token can return an extension at a world where that token does not exist, and it is beyond the scope of this essay to produce one on his behalf. Nothing in my argument turns on this issue.

⁵² Chalmers assumes that what primary intension a complex expression token has is determined by how the expression tokens that constitute it are arranged and what primary intensions these tokens have, and that what secondary intension a complex expression token has is determined by how the expression tokens that constitute it are arranged and what secondary intensions these tokens have.

⁵³ As is the case with primary intensions, Chalmers thinks that the secondary intension of a token can return an extension at a world where that token does not exist. As before, it is beyond the scope of this essay to give a rigorous, general account of secondary intensions on which it is perfectly clear how this is possible. I

Here are some examples of how Chalmers's theory of primary and secondary intensions plays out in particular cases. Consider a typical token of the English expression 'the funniest conscious being'. What the actual extension of this token is depends on which possible state of the world turns out to be actual. If it turns out that the actual world is one in which John is the funniest conscious being, then the extension of that token is John; if it turns out that the actual world is one in which Bill is the funniest conscious being, then the extension of that token is Bill. Thus, the primary intension of that token is a function from possible worlds to extensions that returns John at worlds (considered as actual) in which John is the funniest conscious being, and that returns Bill at worlds (considered as actual) in which Bill is the funniest conscious being. More generally, the primary intension of that token is such that at every world where it is evaluated, it returns whoever is the funniest conscious being at that world when that world is considered as actual. Now, suppose that John is actually the funniest conscious being. We can ask what the extension of our token would have been if things were different from how they actually are, given that its actual extension is fixed. Given that John is the token's actual extension, if Bill were the funniest conscious being, then Bill would be its extension, and if George were the funniest conscious being, then George would be its extension. Thus, the secondary intension of our token returns Bill at worlds (considered as counterfactual) in which Bill is the funniest conscious being, and it returns George at worlds (considered as counterfactual) in which George is the funniest conscious being. More generally, the secondary intension of that token is such

shall stick with the account in the text.

that at every world where it is evaluated, it returns whoever is the funniest conscious being at that world when that world is considered as counterfactual.

Consider a particular token of mine of the English expression type 'I', and suppose that it is a typical token of that type. For all I know *a priori*, I am John Locke. If it turns out that the actual world is one in which I am John Locke, then the extension of my token of 'I' turns out to be John Locke. For all I know *a priori*, I am Aristotle. If it turns out that the actual world is one in which I am Aristotle, then the extension of my token of 'I' turns out to be Aristotle. Thus, the primary intension of my token of 'I' returns John Locke at worlds (considered as actual) where I am John Locke, and it returns Aristotle at worlds (considered as actual) where I am Aristotle. More generally, the primary intension of my token of 'I' is such that at every world where it is evaluated, it returns whoever I am at that world when that world is considered as actual. We can also ask what the extension of our token would be if things were otherwise than they actually are, given that its actual extension is fixed. Given that the person I actually am is the extension of my token of 'I' in the actual world, the person I actually am is the extension of that token in all counterfactual worlds: if even if I had very different character traits and a very different life history than I actually have, as long as I existed, I would still be the person I actually am. Thus, the secondary intension of my token of 'I' returns the person I actually am everywhere it is evaluated.

Finally, consider a particular token of mine of the English sentence type 'I am a philosopher', and suppose that it is a typical token of that type. If it turns out that the actual world is one in which I am a philosopher (e.g., one in which I am John Locke and John Locke is a philosopher) then the extension of that token

of mine is the true. If it turns out that the actual world is one in which I am not a philosopher (e.g., one in which I am Aristotle but Aristotle is not a philosopher), then the extension of that token of mine is the false. Thus, the primary intension of my token of 'I am a philosopher' returns the true at worlds (considered as actual) in which I am a philosopher, and it returns the false at worlds (considered as actual) in which I am not a philosopher. We can also ask what the extension of my token would be if things were different from how they actually are, given that the extension of my token and those of its parts are fixed. Given that the token of 'I' in my token of 'I am a philosopher' actually has as its extension the person I actually am, if the person I actually am were not a philosopher, the extension of 'I am a philosopher' would be the false, whereas if the person I actually am were a philosopher, the extension of 'I am a philosopher' would be the true. Thus, the secondary intension of my token of 'I am a philosopher' returns the true at any world where the person I actually am is a philosopher, and it returns the false at any world where the person I actually am is not a philosopher.

As the last two examples suggest, the arguments taken by primary intensions are, strictly speaking, not possible worlds but *centered* possible worlds—ordered pairs consisting of a world and a center representing the viewpoint in space and time of a being at that world.⁵⁴ This is required not only because some expression tokens are or contain tokens of what Kaplan calls pure indexical expressions⁵⁵ (e.g., 'I', 'here', 'now'), but also because some expression

⁵⁴ Chalmers (1996), p.60; (2002a), p. 166; (2006), p. 76.

⁵⁵ Kaplan (1989), p. 491.

tokens, such as typical tokens of 'water', have a hidden indexical element. (According to Chalmers, roughly speaking, the primary intension of a typical token of 'water' returns at each centered world the dominant clear liquid that fills the oceans and the lakes around the center of that world.) Unlike primary intensions, however, secondary intensions take as arguments normal, uncentered worlds. This is because expression tokens with indexical elements behave rigidly when evaluated at worlds considered as counterfactual, and hence do not depend on centers for the determination of their extensions at worlds so considered.

For any expression token T, the primary and secondary intensions of T will return the same extension at the actual world. We can make this claim more precise by using the notion of a centered possible world's being *based on* a possible world, which I define as follows:

For any centered world W' and any world W , W' is *based on* $W =_{\text{def}} W'$ is an ordered pair $\langle a, b \rangle$ of a world and a center such that $a = W$ and b is the center.

For any expression token T and any subject S, if T is thought or uttered by S, then the extension returned by the secondary intension of T at $W_{@}$ is the same as the extension returned by the primary intension of T at $W_{@}'$ (where $W_{@}$ is the actual world and $W_{@}'$ is the centered world based on $W_{@}$ that is centered on S). Thus, although clarity generally demands that we speak not of the extension of a token at a world, but of the extension returned by its primary intension at a centered world or of the extension returned by its secondary intension at a world, there is

no ambiguity in speaking simply of the extension of a token at the actual world.⁵⁶

Tokens of pure indexical expression types (such as 'I', 'here', and 'now') have different secondary intensions depending on the where, when, and by whom they are uttered. It is plausible to suppose, however, that for any pure indexical expression type K, and any two tokens T₁ and T₂ of K, T₁ and T₂ have the same primary intension. For example, it is plausible that the primary intension of every token of the English expression type 'I' is the function such that, at every centered world, it returns the being at the center of that world; and it is plausible that the primary intension of every token of the English expression type 'now' is the function such that, at every centered world, it returns the time at the center of that world. By contrast, Chalmers argues, some expression types are such that different tokens of them have different primary intensions but the same secondary intension. Suppose that Leverrier introduced the term 'Neptune' to stand for whatever perturbs the orbit of Uranus, and suppose that his wife acquires that term from him but does not know that it is stipulated to refer to whatever perturbs the orbit of Uranus. Then Leverrier's tokens of 'Neptune' will have as their primary intension a function that returns, at every centered world considered as actual, whatever perturbs the orbit of Uranus at that world. His wife's tokens of 'Neptune' will not have this function as their primary intension, however. Instead, Chalmers argues, it is plausible that the primary intension of her tokens of 'Neptune' will be the function that returns, at every centered world considered as actual, whatever Leverrier refers to with his tokens of 'Neptune' at that world. However, since the thing that actually perturbs the orbit of Uranus is

⁵⁶ Chalmers (1996), p. 61.

the thing that Leverrier actually refers to with his tokens of ‘Neptune’, the primary intensions of Leverrier’s and his wife’s tokens of ‘Neptune’ return the same extension at the actual world—namely, Neptune. And since the term ‘Neptune’ behaves rigidly when evaluated at worlds considered as counterfactual, Leverrier’s tokens of ‘Neptune’ and his wife’s have the same secondary intension—namely, the function which picks out Neptune at every world where it is evaluated.⁵⁷ Cases like ‘I’ and ‘here’, where tokens of the same expression type have different secondary intensions, and cases like ‘Neptune’, where tokens of the same expression type have different primary intensions, make it clear why primary and secondary intensions are possessed by expression tokens rather than by expression types.

When a primary or secondary intension is not shared by all tokens of a linguistic expression type, it cannot be part of the “standing meaning” of that type: for the standing meaning of a linguistic expression type is common to all tokens of that type. Thus, the secondary intension of a particular token of mine of ‘I’ (viz., the function that returns the person I actually am at all worlds) cannot be part of the standing meaning of the English expression type ‘I’, and the primary intension of Leverrier’s tokens of ‘Neptune’ cannot be part of the standing meaning of the French expression type ‘Neptune’. By contrast, when a primary or secondary intension is shared by all tokens of a linguistic expression type, it might plausibly be regarded as part of the standing meaning of that type. Thus, the primary intension of a token of ‘I’ (viz., the function that returns, at any centered world, the being at the center of that world) might be regarded as part

⁵⁷ Chalmers (2002a), pp. 167-68; (2006), p. 95.

of the standing meaning of the English expression type 'I'.⁵⁸

Some expression tokens have primary and secondary intensions that behave the same way as each other, except that the former is evaluated at centered worlds and the latter is evaluated at uncentered worlds. We can define the property of *semantic neutrality* as follows:

For any primary intension x and any secondary intension y , x is *equivalent* to $y =_{\text{def}}$ for any world W in the domain of y , and for any centered world W' that is based on W and that is in the domain of x , the extension returned by y at W is the extension returned by x at W' .

For any expression token T , T is *semantically neutral* $=_{\text{def}}$ the primary intension of T is equivalent to the secondary intension of T .⁵⁹

Typical tokens of 'the smallest prime number' are arguably semantically neutral, and so (arguably) are typical tokens of 'the funniest conscious being'. By contrast, typical tokens of 'water', 'Hesperus', 'Phosphorus', 'I', 'here', and 'now' are not semantically neutral. Indeed, it follows from the definition of semantic neutrality that any expression token whose extension at worlds considered as actual varies with variations in centering will fail to be semantically neutral: a token T is semantically neutral only if, for any world W in the domain of its secondary intension, and any two centered worlds W' and W'' that are based on W and are in the domain of its primary intension, the extension returned by T 's

⁵⁸ Chalmers (2006), pp. 97-98.

⁵⁹ This is my attempt to make precise Chalmers's claim that "[i]ntuitively, a semantically neutral expression is one that behaves the same whether one considers a world as actual or as counterfactual" in (2006), p. 86. Chalmers would not *define* semantic neutrality as I do, however, since in his official account of primary and secondary intensions, semantic neutrality is conceptually prior to primary and secondary intensions. In his view, the best one can do is to take the notion of semantic neutrality as primitive.

primary intension at *W*' is the same as the extension returned by *T*'s primary intension at *W*". Since typical tokens of pure indexical terms fail to meet this condition, no such token is semantically neutral. According to Chalmers, however, typical tokens of a large range of English expressions—including 'and', 'philosopher', 'friend', 'consciousness', 'cause', 'doctor', and 'square', to name a few—are semantically neutral. Intuitively, he claims, typical tokens of "descriptive" expressions are semantically neutral.⁶⁰

There are two ways one can think about the possible worlds and centered possible worlds that primary and secondary intensions take as arguments: one can regard them as metaphysically possible centered worlds and worlds, or one can regard them as what Chalmers calls *scenarios*—maximally specific ways the world might actually be, for all that we can know *a priori*. Let us call something *a priori consistent* just if it cannot be ruled out *a priori*.⁶¹ Scenarios are related to *a priori* consistency in the same way that metaphysically possible worlds are related to metaphysical possibility: just as everything that is metaphysically possible can be represented by a space of metaphysically possible worlds, each of which is a maximally specific way the world could really be, so can everything that is *a priori* consistent be represented by a space of scenarios, each of which is a maximally specific way the actual world might be, for all we can know *a priori*.⁶²

⁶⁰ Chalmers (1996), p. 62; (2006), p. 86. One can define stronger kinds of semantic neutrality than the basic kind I define in the text, but since none of these is relevant to my argument, I will stick with the basic kind.

⁶¹ Thus, negative conceivability is a type of *a priori* consistency—namely, the type that applies to statements.

⁶² Chalmers (2006), pp. 76, 81-86; (2002a), p. 166.

On the face of it, although it is highly plausible that every metaphysically possible world corresponds to a scenario, it is not at all obvious that every scenario corresponds to a metaphysically possible world. (Indeed, the latter correspondence is just the sort of thing that Chalmers is trying to establish in arguing for modal rationalism.) Even if there are scenarios that do not correspond to any metaphysically possible worlds, however, it is plausible that the behavior of an expression token's primary and secondary intensions at those scenarios will be partially constitutive of that token's meaning. After all, if two expression tokens behaved in exactly the same way at all metaphysically possible worlds but did not behave in exactly the same way at some scenario, their difference of behavior at that scenario would plausibly constitute a difference between their meanings. Thus, in the absence of a compelling reason to think that every maximally specific way the actual world might be for all we can know *a priori* is a maximally specific way the world really could be, understanding primary and secondary intensions as taking scenarios as arguments seems to be the safer option, provided that one's main goal is to capture a certain aspect of the *meanings* of expression tokens with these intensions. In Chalmers's view, understanding primary and secondary intensions in this way also allows one to identify *a priority* for statements with primary necessity: a statement is *a priori* just if its primary intension is true at every scenario.⁶³

⁶³ Chalmers (1996), p. 69; (2006), pp. 64, 81, 84-85. There are potential problems for this identification. Given that every scenario is centered on a being, and given that a typical token of 'I' picks out the being at the center of every scenario at which it is evaluated, a typical token of 'I exist' will be primarily necessary. But it is not

Thinking of primary and secondary intensions as taking metaphysically possible centered worlds and worlds as arguments has its advantages too, however: it allows one to make claims about the (metaphysical) modal status of statements using the terminology of primary and secondary intensions. For the purposes of this paper, it is more important to make such claims than it is to use Chalmers's two-dimensional framework to explain meaning. Thus, when I speak of primary and secondary intensions in what follows, I will be referring to primary and secondary intensions conceived of as taking metaphysically possible centered worlds and worlds as arguments.

A statement *S* is *primarily* possible just if there is a centered metaphysically possible world at which its primary intension is true. A statement *S* is *secondarily* possible just if there is a metaphysically possible world at which its secondary intension is true.⁶⁴ Thus, if there is a centered metaphysically possible world at which the primary intensions of 'water' and 'H₂O' return XYZ and H₂O, respectively, and at which the primary intensions of 'is' and 'not' return what they actually return, then 'water is not H₂O' is primarily possible, because its primary intension is true at that world. But 'water is not H₂O' is not secondarily possible even if there is such a world, for its secondary intension is false at every metaphysically possible world.⁶⁵

plausible to think that typical tokens of 'I exist' are *a priori*. It is beyond the scope of this essay to investigate these issues any further.

⁶⁴ Primary necessity, contingency, and impossibility are defined from primary possibility in the usual way; and secondary necessity, contingency, and impossibility are defined from secondary possibility in the usual way.

⁶⁵ Chalmers (2002a), pp. 164-65. For simplicity's sake, I shall assume that even at worlds where H₂O does not exist, the secondary intension of 'water is not H₂O' is false and the secondary intension of 'water is H₂O' is true. More generally, I shall

The account I have given of Chalmers's two-dimensional framework is incomplete, but it will do for the purposes of this essay. The view that every token of every linguistic expression type has both a primary intension and a secondary intension is highly controversial, and it is a worthy subject of debate in its own right.⁶⁶ I will not contest it in this essay, however. Chalmers formulates his modal rationalism in a way that presupposes his framework of primary and secondary intensions. My goal in this essay is not to refute the presuppositions of his formulation of modal rationalism, but to argue that, as he formulates it, modal rationalism is false. Thus, I will presuppose the truth of Chalmers's two-dimensional framework. Moreover, I will accept the various claims that he makes about specific cases in the course of presenting that framework (e.g., that the primary intension of a typical token of 'water' returns, at every centered world, the dominant clear liquid that fills the oceans and the lakes around the center; that typical tokens of 'consciousness' are semantically neutral; etc.).

MODAL RATIONALISM

With the distinctions between primary and secondary intensions and between primary and secondary possibility in mind, we are finally in a position to present

assume that every thing is self-identical at every metaphysically possible world, including worlds at which it does not exist. This accords with the general practice of treating statements asserting the self-identity of metaphysically contingent entities (e.g., 'Hesperus is Phosphorus' and 'water is H₂O') as metaphysically necessary.

⁶⁶ See the discussions in Soames (2005) and in García-Carpintero and Macià (2006).

the two theses that constitute Chalmers's modal rationalism:

(CP+) Primary positive conceivability entails primary possibility.

(CP-) Primary negative conceivability entails primary possibility.

In other words:

(CP+) For any statement *S*, if one can coherently imagine a situation which, when considered as actual, verifies *S*, then the primary intension of *S* is such that there is a centered metaphysically possible world at which it is true.

(CP-) For any statement *S*, if $\sim S$ is not *a priori*, then the primary intension of *S* is such that there is a centered metaphysically possible world at which it is true.

Chalmers does not claim that the primary conceivability of a statement entails its metaphysical possibility because he takes it as a datum about the notion of metaphysical possibility (as it is applied to statements) that the negations of *a posteriori* necessities such as the ones presented by Kripke are not metaphysically possible, whether or not there are centered metaphysically possible worlds at which the primary intensions of these negations are true. To say that a statement is metaphysically possible is not just to say that one of its intensions is true at some centered or uncentered metaphysically possible world; it is to say that its *secondary* intension is true at a metaphysically possible world. For a statement to be metaphysically possible is for it to be secondarily possible. Thus, whether or not (CP+) and (CP-) are true, statements such as 'water is not H₂O' and 'Hesperus is not Phosphorus' are metaphysically

impossible.⁶⁷

How, then, does conceivability give us *a priori* access to metaphysical modality, in Chalmers's view? Although the primary conceivability of a statement *S* does not entail its metaphysical possibility, it does entail the existence of a centered metaphysically possible world at which the primary intension of *S* is true. Thus, even though primary conceivability does not give us direct access to the metaphysical possibility of *statements*, it gives us direct access to the metaphysical possibility of centered *worlds*. Moreover, if a statement *S* is primarily conceivable, then any statement whose secondary intension is equivalent to the primary intension of *S* is metaphysically possible: for if *S* is primarily conceivable, its primary intension will (according to (CP+) and (CP-)) return the true at some centered metaphysically possible world *W'*, and any secondary intension equivalent to its primary intension will therefore return the true at the metaphysically possible world *W* on which *W'* is based. If *S* is semantically neutral, then *S* itself will be a statement whose secondary intension is equivalent to its primary intension. Thus, if Chalmers is correct, the primary conceivability of a semantically neutral statement entails its metaphysical possibility.

The posneg principle says that every statement that is primarily positively conceivable is primarily negatively conceivable. It therefore implies that if every primarily negatively conceivable statement is primarily possible, then every primarily positively conceivable statement is primarily possible—i.e., that if (CP-) is true, then (CP+) is true. The posneg principle is true, so if (CP-) is

⁶⁷ Chalmers (1996), pp. 67-68; (2002a), p. 164.

true, then (CP+) is true as well. (CP+) entails (CP-) if and only if the following principle, which Chalmers calls the *negpos* principle, is true:

(NP) For any statement S, if S is primarily negatively conceivable, then S is primarily positively conceivable.

Equivalently:

(NP) For any statement S, if $\sim S$ is not *a priori*, then one can coherently imagine a situation which, when considered as actual, verifies S.

Unlike the *posneg* principle, it is not *a priori* obvious that the *negpos* principle is true. Chalmers finds it plausible, but he does not claim to have proved it conclusively.⁶⁸ It is beyond the scope of this essay to determine whether it is true. I shall not assume it, so I shall not assume that (CP+) entails (CP-).

I shall use 'Chalmers's modal rationalism' (or 'modal rationalism' for short) to refer to the conjunction of (CP+) and (CP-). Of course, each of these two theses asserts a very strong link between metaphysical possibility and the *a priori* domain. If it turns out that (CP+) is true but that (CP-) is not, that will still be a philosophically important outcome and a victory for Chalmers.

THE MODAL STATUS OF MODAL RATIONALISM

Chalmers thinks that (CP+) and (CP-) are metaphysically necessary if they are

⁶⁸ Chalmers (2002a), pp. 173-84, 186-89, 194-95.

true at all.⁶⁹ He is almost certainly correct, since it is bizarre to think that although (CP+) and (CP-) are actually true, they could have been false if things had turned out otherwise: surely, if there actually are deep connections between conceivability and possibility of the kind asserted by (CP+) and (CP-), those connections could not have failed to obtain. Thus, Chalmers believes the following two nearly indubitable theses:

(M+) If (CP+) is true, then (CP+) is metaphysically necessary—i.e., its secondary intension is true at every metaphysically possible world.

(M-) If (CP-) is true, then (CP-) is metaphysically necessary—i.e., its secondary intension is true at every metaphysically possible world.

Now, it is highly plausible to suppose that (CP+) and (CP-) are semantically neutral. By all appearances, they are not like ‘water is H₂O’ or ‘Hesperus is Phosphorus’, but like ‘the number of conscious beings exceeds the smallest prime number’. If (CP+) and (CP-) are semantically neutral, however, the following two theses follow from (M+) and (M-):

(N+) If (CP+) is true, then (CP+) is primarily necessary—i.e., its primary intension is true at every centered metaphysically possible world.

(N-) If (CP-) is true, then (CP-) is primarily necessary—i.e., its primary intension is true at every centered metaphysically possible world.

For (CP+) and (CP-) are semantically neutral just if their primary intensions are equivalent to their secondary intensions. According to (M+), if (CP+) is true, then

⁶⁹ Chalmers writes that modal rationalism “is surely necessarily true, if true,” and he accepts the inference from the possibility of the falsity of modal rationalism to its falsity (1999, p. 485).

the function that returns the true at every metaphysically possible world is the secondary intension of (CP+); and according to (M-), if (CP-) is true, then that function is the secondary intension of (CP-). Thus, if (M+) is true, then if (CP+) is both true and semantically neutral, the function that returns the true at every centered metaphysically possible world is the primary intension of (CP+); and if (M-) is true, then if (CP-) is both true and semantically neutral, the same function is the primary intension of (CP-).

Although none of (M+), (M-), (N+), and (N-) is part of modal rationalism proper, any defender of modal rationalism will probably affirm them as well. If modal rationalism is true, then in all likelihood, it is both primarily and secondarily (i.e., metaphysically) necessary.

THE OBJECTION FROM KRIPKEAN *A POSTERIORI* NECESSITIES

A common and superficially plausible objection to the idea that conceivability is a guide to metaphysical possibility is that the negations of Kripkean *a posteriori* necessities are conceivable but metaphysically impossible. For example, it is said that although 'water is not H₂O' and 'Hesperus is not Phosphorus' are conceivable, neither is metaphysically possible. It is clear, though, that typical tokens of English sentences such as these are not counterexamples to Chalmers's modal rationalism. Typical tokens of 'water is not H₂O' and 'Hesperus is not Phosphorus' are indeed secondarily (i.e., metaphysically) impossible, but they are not secondarily conceivable either: given that water is H₂O and that

Hesperus is Phosphorus, it is not conceivable that water might not have been H₂O or that Hesperus might not have been Phosphorus. Such tokens are primarily conceivable, however: their negations are not *a priori*, and there are coherently imaginable situations which verify them when considered as actual. And it is plausible, Chalmers says, that they are primarily possible as well: it is plausible that there is a centered metaphysically possible world at which the primary intension of typical tokens of 'water is not H₂O' is true, and one at which the primary intension of typical tokens of 'Hesperus is not Phosphorus' is true. At the very least—and this is the important point—nothing that Kripke says rules out the existence of such worlds. Thus, it is plausible that typical tokens of the negations of Kripkean *a posteriori* necessities are not counterexamples to (CP-) or (CP+). Indeed, as long as we line up the right kind of conceivability with the right kind of possibility, they are not obviously counterexamples to the more general thesis that conceivability entails possibility: the primary conceivability of a statement entails its primary possibility, and the secondary conceivability of a statement entails its secondary (i.e., metaphysical) possibility.⁷⁰

It is beyond the scope of this essay to rehearse the many other purported counterexamples to Chalmers's modal rationalism. Chalmers argues that no particular statement is a known, clear counterexample to (CP+) or (CP-).⁷¹ Let us assume that he is correct. Even so, he must give a positive argument in favor of those theses. Modal rationalism is not so intuitive a view that merely defending it from supposed counterexamples suffices to justify it. Why should we believe that

⁷⁰ Chalmers (1996), pp. 67-68; (2002a), p. 165.

⁷¹ Chalmers (2005).

it is true?

ASIDE: COMPLICATIONS DUE TO EXTERNALISM

Before we examine Chalmers's positive case for modal rationalism, we must consider a potential threat to modal rationalism that might arise from externalist views of meaning. (CP+) and (CP-) are interesting theses only if not every statement is primarily possible. After all, if it turns out that every statement is primarily possible, then every statement will satisfy their consequents, and instead of expressing an interesting relationship between conceivability and possibility, they will turn out to be trivial. On the face of it, not every statement is primarily possible. For example, it certainly seems that a typical token of the English sentence type 'Everything is alive and everything is not alive' will be primarily impossible, since it will be logically invalid. However, one might think that one can apply externalist views about meaning to Chalmers's framework of primary and secondary intensions in such a way that it turns out that every statement is primarily possible, and hence that (CP+) and (CP-) are trivial.

To see just one way one might follow this train of thought, consider this definition:

For any expression token T: the primary intension of T is *social* =_{def} the primary intension of T and the linguistic expression type K of which T is a token are such that, for any centered world W, the extension returned by the primary intension of T at W is whatever the social environment at the center of W determines the extension of tokens of type K to be.

Thus, if the primary intension of a typical token of the English expression type 'arthritis' is social, then it is such that, at a centered world where the social environment at the center determines that certain ailments of the thigh are in the extension of tokens of that type, it returns (among other things) those ailments of the thigh. Likewise, if the primary intension of a typical token of the English sentence type 'arthritis can occur in the thigh' is social, then it is such that, at any centered world, it returns the true if and only if the social environment at the center determines that the true is the extension of tokens of the English sentence type 'arthritis can occur in the thigh'.

Now consider the following thesis about primary intensions:

(SOC) For any statement *S*, the primary intension of *S* is social.

Someone might argue that this thesis implies that every statement is primarily possible. For one might argue that for any linguistic sentence type *K*, there is a centered metaphysically possible world *W* where the social environment at the center of *W* determines that tokens of type *K* are true (i.e., have the true as their extension). Since every statement is a token of a linguistic sentence type, it would follow from this and from (SOC) that every statement has a primary intension that is such that, at some centered metaphysically possible world, it returns the true. It would follow, in other words, that every statement is primarily possible.

It is beyond the scope of this essay to determine whether (SOC) is true, or whether it implies that every statement is primarily possible. As I mentioned earlier, my goal in this essay is not to contest Chalmers's formulation of modal

rationalism, but to show that modal rationalism, as he formulates it, is false. Thus, I shall assume either that (SOC) fails to imply that every statement is primarily possible or that (SOC) is false. More generally, I shall assume that some statements are primarily impossible, and hence that (CP+) and (CP-) are not trivial in the way described earlier. Having noted and set aside the threat from externalism, I turn now to Chalmers's positive arguments for modal rationalism.

IV. Chalmers's Case for Modal Rationalism

Chalmers makes his positive case for modal rationalism in two ways. First, he argues indirectly for the truth of modal rationalism by arguing against the existence of a certain kind of *a posteriori* metaphysical necessity that he calls 'strong necessities'. Second, he gives an argument that purports to show that modal rationalism is *a priori*.

THE CASE AGAINST STRONG NECESSITIES

Chalmers defines strong necessities (of the negative and positive kinds) and weak necessities as follows.

For any statement *S*: *S* is a *strong a posteriori metaphysical necessity* (or a '*strong necessity*' for short) =_{def} *S* is not *a priori*, *S*'s primary intension is true at all centered metaphysically possible worlds, and *S*'s secondary intension is true at all metaphysically possible worlds.

For any statement *S*: *S* is a *negative strong a posteriori metaphysical necessity* (or a '*negative strong necessity*' for short) =_{def} *S* is not *a priori*, *S*'s primary intension is true at all centered metaphysically possible worlds, *S*'s secondary intension is true at all metaphysically possible worlds, and $\sim S$ is primarily negatively conceivable.

For any statement *S*: *S* is a *positive strong a posteriori metaphysical necessity* (or a '*positive strong necessity*' for short) =_{def} *S* is not *a priori*, *S*'s primary intension is true at all centered metaphysically possible worlds, *S*'s secondary intension is true at all metaphysically possible worlds, and $\sim S$ is primarily positively conceivable.

For any statement *S*: *S* is a *weak a posteriori metaphysical necessity* (or '*weak necessity*' for short) =_{def} *S* is not *a priori*, *S*'s primary intension is not true at all centered metaphysically possible worlds, and *S*'s

secondary intension is true at all metaphysically possible worlds.⁷²

Given the definition of primary negative conceivability, the definition of a negative strong necessity is equivalent to that of a strong necessity: because a strong necessity is, by definition, not *a priori*, it is true by definition that the negation of a strong necessity is primarily negatively conceivable. Every positive strong necessity is, by definition, a strong necessity (and thus a negative strong necessity), but the reverse is not true by definition: every negative strong necessity (and thus every strong necessity) is a positive strong necessity only if the negpos principle is true. Since I do not assume that principle in this essay, I shall not assume that every negative strong necessity (and thus every strong necessity) is a positive strong necessity.

The negations of weak necessities are not counterexamples to (CP+) or (CP-): for if a statement *S* is a weak necessity, then $\sim S$ is primarily possible and is therefore not both primarily conceivable (positively or negatively) and primarily impossible. The point about the negations of Kripkean *a posteriori* necessities that I made earlier can be put in terms of weak necessities. According to Chalmers, Kripkean *a posteriori* necessities are most plausibly regarded as weak necessities: although the secondary intensions of typical tokens of 'water is not H₂O' and 'Hesperus is not Phosphorus' are not true at any metaphysically possible world, their primary intensions are most plausibly regarded as true in some centered metaphysically possible worlds. That is why it is plausible that

⁷² Chalmers (1996), pp. 136-37; (1999), p. 477; (2002a), p. 189; (2005), section 7. It is not clear that Chalmers's talk of strong and weak necessities in (1996) comes to the same thing as it does in his later works. My exposition relies primarily on his (2005).

the negations of these necessities are not counterexamples to (CP+) or (CP-).

If there are any strong necessities, then (CP-) is false: for if a statement S is a strong necessity, then its negation, $\sim S$, is a statement that is both primarily negatively conceivable and primarily impossible. If there are any positive strong necessities, then (CP+) and (CP-) are false: for if a statement S is a positive strong necessity, then its negation, $\sim S$, is a statement that is primarily positively conceivable, primarily negatively conceivable, and primarily impossible. Every strong necessity is such that its negation is a counterexample either to (CP-) alone (if it is not a positive strong necessity) or to both (CP-) and (CP+) (if it is a positive strong necessity). Thus, showing that there are no strong necessities (either positive or negative) is a necessary condition of showing that (CP-) is true, and showing that there are no positive strong necessities is a necessary condition of showing that (CP+) is true.

Chalmers writes that (CP-) and (CP+) are “equivalent to the theses that there are no negative strong necessities and that there are no positive strong necessities respectively,” and he therefore thinks that any argument against the existence of strong necessities is *ipso facto* an argument against the existence of counterexamples to (CP-) and (CP+).⁷³ He raises two general considerations against the supposition that there are strong necessities. First, he says, there is no reason to believe that strong necessities exist. Second, he says, the existence of strong necessities would put “brute and inexplicable” constraints on the space of possible worlds.⁷⁴ Since he takes (CP-) to be equivalent to the thesis that there

⁷³ Chalmers (2005), section 7.

⁷⁴ Chalmers (1996), p. 137.

are no strong necessities, and since he takes (CP+) to be equivalent to the thesis that there are no positive strong necessities, he thinks that these two general considerations count in favor of (CP-) and (CP+) exactly to the extent that they count against the existence of strong necessities. In this way, he argues indirectly for the truth of (CP-) and (CP+) by arguing against the existence of strong necessities.

Chalmers's two general considerations are feeble and ineffectual, however. Even if he is right about there being no reason to believe in strong necessities, for all that he has said so far, there is no reason to believe in modal rationalism either. Moreover, it seems that the truth of modal rationalism would be just as brute and inexplicable as the existence of strong necessities. Why should it be true that every coherently imaginable situation corresponds to part of a metaphysically possible world? Why should it be true that every statement that cannot be ruled out *a priori* really could be true? What could explain such a neat and convenient alignment of the *a priori* domain and metaphysical modality? Chalmers must give us an argument that proves that there cannot be any strong necessities, or at least one that establishes that the truth of modal rationalism would not be brute and inexplicable.

Moreover, even if Chalmers's general considerations ruled out the existence of strong necessities, it is not obvious that they would thereby establish (CP-) and (CP+). For it is not obvious that Chalmers is right when he says that these two principles are equivalent to the theses that there are no negative strong necessities and that there are no positive strong necessities, respectively. Accordingly, it is not obvious that showing that there are no strong

necessities is sufficient for showing that (CP-) and (CP+) are true, or that showing that there are no positive strong necessities is sufficient for showing that (CP+) is true.

(CP-) is equivalent to the claim that there are no strong necessities if and only if the following principle is true:

(-PS) For any statement S: if S is not *a priori*, then if S is primarily necessary, then S is secondarily (i.e., metaphysically) necessary.

If (-PS) is true, then the negation of any counterexample to (CP-) will be a strong necessity, since from the fact that it is primarily necessary and not *a priori*, it will follow that it is also secondarily necessary. Thus, if (-PS) is true, then showing that there are no strong necessities will be sufficient for showing that (CP-) is true. But if (-PS) is false, then the negation of some counterexample to (CP-) will fail to be a strong necessity, since it will be primarily necessary, not *a priori*, but not secondarily necessary. Thus, if (-PS) is false, then in order to show that (CP-) is true, one must show both that there are no strong necessities and that there are no statements which are primarily negatively conceivable, primarily impossible, but secondarily possible.

(CP+) is equivalent to the claim that there are no positive strong necessities if and only if the following principle is true:

(+PS) For any statement S: if $\sim S$ is primarily positively conceivable, then if S is primarily necessary, then S is secondarily (i.e., metaphysically) necessary.

If (+PS) is true, then the negation of any counterexample to (CP+) will be a

positive strong necessity, since from the facts that it is primarily necessary and that its negation is primarily positively conceivable, it will follow that it is also secondarily necessary. Thus, if (+PS) is true, then showing that there are no positive strong necessities will be sufficient for showing that (CP+) is true. If (+PS) is false, however, then the negation of some counterexample to (CP+) will fail to be a strong necessity, since it will be primarily necessary, its negation (viz., the counterexample) will be primarily positively conceivable, but it will not be secondarily necessary. Thus, if (+PS) is false, then in order to show that (CP+) is true, one must show both that there are no positive strong necessities and that there are no statements which are primarily positively conceivable, primarily impossible, but secondarily possible.

Since Chalmers thinks that (CP-) and (CP+) are equivalent to the claims that there are no negative strong necessities and no positive strong necessities, respectively, he evidently assumes that (-PS) and (+PS) are true. He does not give us reasons to think that they are true, however, and it is not obvious that they are. Thus, the general considerations that Chalmers raises against strong necessities fail to establish (CP-) or (CP+) not only because they fail to rule out strong necessities, but also because they would not obviously be enough to establish (CP-) or (CP+) even if they ruled out strong necessities. Chalmers needs to give us a different argument in support of the truth of (CP-) and (CP+).

THE ARGUMENT FOR THE *A PRIORITY* OF MODAL RATIONALISM

Chalmers only has one other positive argument for (CP+) and (CP-), and it is by his own admission a mere “sketch” of an argument. Our modal concepts are used to analyze psychological and rational phenomena, he observes: we use them when “analyzing the contents of thoughts and the semantics of language, giving an account of counterfactual thought, [and] analyzing rational inference.” Thus, we need a “*rational* modal concept, tied constitutively to consistency, rational inference, or conceivability,” even if we accept strong necessities. For example, even if we think that the laws of nature are strong necessities, we must use a space of “rationally coherent possibilities involving different laws”—a space of worlds that we take to be broader than that of metaphysically possible worlds—in order to explain what is going on when a scientist wonders what would have been the case if the laws of nature had been different. But once we admit such a rational concept of possibility, he says, there is no need for a distinct notion of metaphysical possibility: the rational modal concept explains “all modal phenomena that we have reason to believe in.”⁷⁵ To posit a second, metaphysical, modality that is not reducible to the first, rational, modality is to posit “a primitive that answers to no-one and does no work.”⁷⁶

Indeed, Chalmers asserts, the hypothesis that (CP+) or (CP-) is false is incoherent, because there is no notion of metaphysical possibility on which it is coherent to suppose that either of them is false: “we do not even have a distinct

⁷⁵ Chalmers (2002a), pp. 193-94.

⁷⁶ Chalmers (1999), p. 491.

concept of metaphysical necessity to which the second primitive can answer....

The only concept of a 'metaphysically possible world' that we have is that of a logically possible [i.e., ideally conceivable] world."⁷⁷ Thus, according to Chalmers, (CP+) and (CP-) are *a priori*, and their conjunction, modal rationalism, is therefore *a priori* as well.⁷⁸ Since (CP+) implies that there are no positive strong necessities, and since (CP-) implies that there are no strong necessities, it is *a priori* that there are no strong necessities. It turns out, then, that modal rationalism and the thesis that there are strong necessities are not on a par. Not only would the existence of strong necessities be brute and inexplicable, but it is *a priori* that there are no such necessities. Modal rationalism, on the other hand, is far from brute or inexplicable: it is *a priori*.

⁷⁷ Chalmers (1999), p. 491. Elsewhere, Chalmers describes logical possibility as "corresponding roughly to conceivability" (1996, p. 35). He also describes the space of logically possible worlds as one that contains "one world for every ideally conceivable scenario" (1999, p. 490).

⁷⁸ Chalmers (1999), pp. 489, 485.

V. Modal Rationalism Is Not *A Priori*

In its current form, Chalmers's case for the *a priori* of modal rationalism is not really an argument, but rather a pair of claims: (i) there is no theoretical use for a concept of metaphysical possibility that is not constitutively tied to consistency and conceivability, and (ii) we do not have a concept of metaphysical possibility that is not constitutively tied to consistency and conceivability. Accordingly, one cannot really oppose it in any other way than by showing that one or more of these claims is incorrect. It is obvious, however, that both of Chalmers's claims are incorrect. It is manifestly true that we have a concept of metaphysical possibility that is not constitutively tied to those of logical possibility, conceivability, or consistency: as I argued earlier, the concept of metaphysical possibility is just that of the broadest kind of real possibility—a notion that is distinct from the rational notions to which Chalmers refers and that lacks any conceptual connections to those notions. Furthermore, it is not true that there is no theoretical need for such a concept: we need it to express certain claims, such as that a certain philosopher believes that although it is logically and conceptually consistent to suppose that there are violations of the laws of physics, there absolutely could not be any such violations in reality. Chalmers's sketch of an argument manifestly fails to establish—or even make it plausible to think—that modal rationalism is *a priori*.

Indeed, a clear understanding of the concept of metaphysical possibility and of how the notion of real possibility differs from those of consistency and

other kinds of so-called “possibility” entitles us strongly but defeasibly to believe that neither (CP+) nor (CP-) is *a priori*. For anyone who has such an understanding of these notions will see that whereas no reason to suppose that either (CP+) or (CP-) is *a priori* can be found even after extended reflection, it seems obvious, both initially and after extended reflection, that for all we can know *a priori*, (CP+) and (CP-) are false.

Consider a particular typical token T_1 of the English sentence type ‘every conscious being has parts’. T_1 is semantically neutral, so it is primarily necessary just in case it is metaphysically necessary. It seems obvious, even after extended reflection, that for all we can know *a priori*, T_1 is metaphysically necessary: no matter how hard we look, we are unable to find any way to rule out *a priori* the hypothesis that although it is consistent to suppose that there are conscious beings without parts, there absolutely could not be any conscious beings without parts. Thus, it seems obvious, even after extended reflection, that for all we can know *a priori*, T_1 is primarily necessary. It seems equally obvious after extended reflection that T_1 is not *a priori*, however: no matter how hard we look, we cannot find any way to reason *a priori*, from the proper analyses of the concepts of consciousness and parthood or from anything else, to the conclusion that T_1 is true. Thus, it seems obvious that the negation of T_1 , ‘not every conscious being has parts’, is primarily negatively conceivable. Moreover, it seems obvious after extended reflection that ‘not every conscious being has parts’ is primarily positively conceivable—i.e., that one can coherently imagine a situation which verifies it when considered as actual (e.g., a situation which, when considered as actual, is correctly described as one in which a conscious monad exists). Thus,

even after extended reflection, it seems obvious that for all we can know *a priori*, T₁ is primarily necessary and secondarily necessary even though it is not *a priori* and its negation is both primarily negatively conceivable and primarily positively conceivable. In other words, it seems obvious even after extended reflection that for all we can know *a priori*, T₁ is both a negative strong necessity and a positive strong necessity.

The same pattern of reasoning can be repeated with a particular typical token T₂ of the English sentence type 'every conscious being is extended'. It seems obvious even after extended reflection that T₂ is metaphysically necessary for all we can know *a priori*: no matter how hard we try, we cannot rule out *a priori* the hypothesis that although it is consistent to suppose that there are non-extended conscious beings, there absolutely could not be any such beings in reality. Since T₂ is semantically neutral, it seems obvious after extended reflection that T₂ is primarily necessary for all we can know *a priori*. But it seems equally obvious after extended reflection that T₂ is not *a priori*, and that its negation 'not every conscious being is extended', is therefore primarily negatively conceivable. Moreover, it seems obvious after extended reflection that 'not every conscious being is extended' is primarily positively conceivable—i.e., that one can coherently imagine a situation which verifies it when considered as actual (e.g., a situation which, when considered as actual, is correctly described as one in which Cartesian souls exist). Thus, it seems obvious after extended reflection that for all we can know *a priori*, T₂ is both a negative strong necessity and a positive strong necessity.

T₁ and T₂ are just two among many statements about which it seems

obvious, even after extended reflection, that for all we can know *a priori*, they are both negative strong necessities and positive strong necessities. But if there are any negative strong necessities, then (CP-) is false; and if there are any positive strong necessities, then (CP-) and (CP+) are both false. Thus, if some statements are both negative strong necessities and positive strong necessities for all that we can know *a priori*, then neither (CP+) nor (CP-) is *a priori*. It therefore seems obvious, even after extended reflection, that (CP+) is not *a priori*, that (CP-) is not *a priori*, and thus that their conjunction, modal rationalism, is not *a priori* either.

By contrast, it does not appear to be the case, either at first glance or after extended reflection, that either (CP+) or (CP-) is *a priori*. The apparent obviousness of their not being *a priori* is not counterbalanced to any degree by any plausible reasons to suppose that they are *a priori*, or even by any appearance of their being *a priori*. Just as one cannot come up with any plausible reasons to suppose that 'Gordon Brown is the Prime Minister' or 'there are bicycles in Oxford' are *a priori*, so is one unable to come up with any plausible reasons to suppose that either (CP+) or (CP-) is *a priori*. (Chalmers's case for the *a priority* of modal rationalism purports to give us such reasons, but the total implausibility of these reasons is evident upon reflection.) And just as neither of the two sentences just mentioned even *appears* to be *a priori*, neither (CP+) nor (CP-) has even the appearance of being *a priori*.

After extended reflection, the balance of reasons and intuitions overwhelmingly favors the hypothesis that neither (CP+) nor (CP-) is *a priori*. This balance of reasons does not conclusively rule out the possibility that there are very deeply hidden contradictions in the supposition that (CP+) or (CP-) is

false—contradictions that are not apparent even after extended reflection; thus, it does not show conclusively that neither (CP+) nor (CP-) is *a priori*. It does, however, entitle us to believe that neither (CP+) nor (CP-) is *a priori* until a compelling reason to believe otherwise is produced, just as we are entitled to believe with a high degree of confidence that ‘Gordon Brown is the Prime Minister’ is not *a priori* until a compelling reason to think otherwise emerges. Moreover, absent a compelling reason to believe that (CP+) or (CP-) is *a priori*, the balance of reasons entitles us to use the propositions that (CP+) is not *a priori* and that (CP-) is not *a priori* as premises for further reasoning. The burden of proof is on the philosopher who claims that (CP+) and (CP-) are *a priori*: until he gives us reasons in favor of his view that outweigh the considerable reasons we already have against it, we are entitled to assume that his view is wrong.

This stance on what we are entitled to assume about the *a priori* of (CP+) and (CP-) given the balance of reasons after extended reflection is a consequence of the following general principle:

(AP1) For any statement S, if it seems obvious, after extended reflection, that S is not *a priori*, and extended reflection does not produce plausible reasons to think that S is *a priori*, then one is entitled to believe that S is not *a priori* (and to assume so for further reasoning) until a compelling reason to think otherwise is produced.

Given the definition of primary negative conceivability, (AP1) is equivalent to:

(AP2) For any statement S, if it seems obvious, after extended reflection, that S is primarily negatively conceivable, and extended reflection does not produce plausible reasons to think that S is not primarily negatively conceivable, then one is entitled to believe that S is primarily negatively conceivable (and to assume so for further reasoning) until a compelling reason to think otherwise is produced.

And although it is not implied by (AP1) or (AP2) unless the negpos principle (which I do not assume in this essay) is true, the following principle seems as plausible as they are:

(AP3) For any statement S, if it seems obvious, after extended reflection, that S is primarily positively conceivable, and extended reflection does not produce plausible reasons to think that S is not primarily positively conceivable, then one is entitled to believe that S is primarily positively conceivable (and to assume so for further reasoning) until a compelling reason to think otherwise is produced.

These three principles are surely true. It would surely be unreasonable to demand that one meet a higher evidential standard than the one required by these principles in order to be entitled to a defeasible belief in the non-*a priori*, primary negative conceivability, or primary positive conceivability of a statement. It would be plainly unreasonable to maintain, for example, that unless one gives a sound demonstrative proof showing that a certain statement S is not *a priori*, then even if extended reflection produces no plausible reasons to think that S is *a priori*, and even if it seems obvious after extended reflection that S is not *a priori*, one is not entitled to believe—even defeasibly—that S is not *a priori*. The evidential bar mentioned in (AP1), (AP2), and (AP3) is the evidential bar the passing of which entitles us to believe defeasibly that ‘Gordon Brown is the Prime Minister’ and ‘there are bicycles in Oxford’ are not *a priori*. Since we are plainly entitled to believe defeasibly that these sentences are not *a priori*, (AP1), (AP2), and (AP3) are surely true.

Chalmers would agree—or at least, he should. For in his book, he

endorses a principle very similar to the three just mentioned when he defends the claim that zombies—beings that are physically identical to beings with phenomenal consciousness, but which lack phenomenal consciousness—are logically possible. He writes:

Arguing for a logical possibility is not entirely straightforward. How, for example, would one argue that a mile-high unicycle is logically possible? It just seems obvious. Although no such thing exists in the real world, the description certainly appears to be coherent. If someone objects that it is not logically possible—it merely seems that way—there is little we can say, except to repeat the description and assert its obvious coherence. It seems quite clear that there is no hidden contradiction lurking in the description.... In general, a certain burden of proof lies on those who claim that a given description is logically *impossible*. If someone truly believes that a mile-high unicycle is logically impossible, she must give us some idea of where a contradiction lies, whether explicit or implicit. If she cannot point out something about the intensions of the concepts “mile-high” and “unicycle” that might lead to a contradiction, then her case will not be convincing... If no reasonable analysis of the terms in question points toward a contradiction, or even makes the existence of a contradiction plausible, then there is a natural assumption in favor of logical possibility.⁷⁹

In this passage, Chalmers essentially says that if something seems obviously logically possible upon reflection, and extended reflection does not produce any plausible reasons to think it logically impossible, one is entitled to believe that it is logically possible until a compelling reason to believe otherwise is produced. There is no principled reason for him to endorse this principle but to reject (AP1), (AP2), or (AP3).

More importantly, (AP1), (AP2), and (AP3) must be accepted by anyone who wants to use (CP+) and (CP-) to come to know that some statement that he

⁷⁹ Chalmers (1996), p. 96.

does not know to be *a priori* is primarily possible. One can use (CP+) to come to know that some statement S is primarily possible only if one knows (CP+) and one knows that S is primarily positively conceivable; and one can use (CP-) to come to know that some statement S is primarily possible only if one knows (CP-) and one knows that S is primarily negatively conceivable. Now, from the fact that a statement S is *a priori*, it follows that its negation, $\sim S$, is not *a priori*, and hence that S is primarily negatively conceivable and (via the posneg principle) primarily positively conceivable. Thus, if one knows that a certain statement is *a priori*, one can know via a deductive argument that it is primarily negatively conceivable and primarily positively conceivable. If one does not know that a statement S is *a priori*, however, one has no plausible deductive path to the knowledge that it is primarily negatively conceivable or primarily positively conceivable, and one can only have this knowledge by weighing the evidence for and against the hypothesis that it is primarily negatively conceivable or primarily positively conceivable: provided that this hypothesis is true, one can know it by correctly judging that the evidence in favor of it sufficiently outweighs the evidence against it. But if this is a legitimate way to come to know that a certain statement S is primarily negatively conceivable or primarily positively conceivable, then surely (AP1), (AP2), and (AP3) are true: for belief is necessary for knowledge, and these principles merely assert that if the evidence overwhelming favors the conceivability of a statement, one is entitled to the defeasible belief that that statement is conceivable. The paradigmatic statements whose primary possibility modal rationalists want to use (CP+) and (CP-) to come to know are not known by them to be *a priori*. Chalmers, for example,

wants to use (CP+) and (CP-) to come to know that 'there exists a zombie' is primarily possible, even though he does not know that it is *a priori*. Thus, any proponent of modal rationalism will have to accept (AP1), (AP2), and (AP3). Rejecting these principles might save the modal rationalist from the conclusion that we should believe that (CP+) and (CP-) are not *a priori* until we are given compelling reasons to think otherwise, but it would do so only at the cost of making it impossible to use (CP+) or (CP-) to come to know the primary possibility of those statements whose primary possibility modal rationalists are most intent on using (CP+) or (CP-) to come to know.

It seems obvious, both initially and after extended reflection, that neither (CP+) nor (CP-) is *a priori*, and extended reflection produces no plausible reasons to think that either of them is *a priori*. (Chalmers's argument for the *a priori* of (CP+) and (CP-) purports to give us such reasons, but it takes only a bit of reflection to see that it does not.) By (AP1), this entitles us to assume that neither (CP+) nor (CP-) is *a priori* until compelling reasons to believe otherwise are produced. Not only is (AP1) surely true, but it must be accepted by any modal rationalist. Thus, we are entitled to assume—and I shall assume—that neither (CP+) nor (CP-) is *a priori* until compelling reasons to believe otherwise are produced, and any modal rationalist must concede that we are entitled to this assumption.

VI. Two Arguments Against Modal Rationalism

If neither (CP+) nor (CP-) is *a priori*, then they are not merely put back on a par in terms of bruteness and inexplicability with the claims that there are negative strong necessities and positive strong necessities: they are demonstrably false.

There are at least two ways to show this.

FIRST ARGUMENT

Let M be a typical token of the English sentence type 'there exists a metaphysically necessary conscious being', and let \sim M be a typical token of the English sentence type 'there does not exist a metaphysically necessary conscious being'. Typical tokens of these two English sentence types are semantically neutral, so M and \sim M are semantically neutral. Now, since each of these statements is the negation of the other, and since neither of them is *a priori*, both are primarily negatively conceivable. Both statements are primarily positively conceivable as well: one can coherently imagine a situation which, when considered as actual, is correctly described as one in which there exists a metaphysically necessary conscious being, and one can coherently imagine a situation which, when considered as actual, is correctly described as one in which there does not exist a metaphysically necessary conscious being.⁸⁰ Thus, according to each of (CP+) and (CP-), both statements are primarily possible.

⁸⁰ I will justify these claims about conceivability in the next section. For the moment, I am just presenting the arguments.

This cannot be the case, however, because only one of these statements is primarily possible. On the standard S5 picture, according to which exactly the same worlds are metaphysically possible from the point of view of every metaphysically possible world, something is metaphysically necessary at the actual world just in case it is metaphysically necessary at every metaphysically possible world. M is true just if there is a conscious being that exists at every metaphysically possible world, and given the S5 picture, this is so just if every metaphysically possible world is such that, at every world metaphysically possible from its point of view, the being in question exists. $\sim M$ is true just if no actually conscious being exists at every metaphysically possible world, and given the S5 picture, this is so just if every metaphysically possible world is such that there is no actually conscious being that exists at every world that is metaphysically possible from its point of view.⁸¹ Thus, if M is true, then it is metaphysically (i.e., secondarily) necessary, and if M is false, then it is metaphysically (i.e., secondarily) impossible. Likewise, if $\sim M$ is true, then it is metaphysically necessary, and if $\sim M$ is false, then it is metaphysically impossible. Since one of M and $\sim M$ is true and the other is false, one of M and $\sim M$ is metaphysically necessary and the other is metaphysically impossible. Now, any statement that is both metaphysically impossible and semantically neutral is primarily impossible. After all, a statement is metaphysically impossible just if its secondary intension is the function that returns the false at every metaphysically

⁸¹ Note that M does not imply that at every metaphysically possible world, there exists a conscious being. It says that there is a being that is conscious at the actual world and that exists at every metaphysically possible world, but it does not imply that this being is conscious at every world at which it exists.

possible world, and it is semantically neutral just if its secondary intension is equivalent to its primary intension. Thus, if a statement is both semantically neutral and metaphysically impossible, its primary intension is the function that returns the false at every centered metaphysically possible world—precisely the primary intension the possession of which is constitutive of primary impossibility. M and $\sim M$ are both semantically neutral, and one of them is metaphysically impossible. Thus, one of them is primarily impossible.

Either M is primarily impossible or $\sim M$ is primarily impossible.

Nevertheless, each of (CP+) and (CP-) implies that both of these statements are primarily possible. Even though we do not know which of these two statements is primarily impossible, we know that one of them is, and this is enough to refute both (CP+) and (CP-).⁸²

SECOND ARGUMENT

Consider the negations of (CP+) and (CP-):

\sim (CP+) It is not the case that, for any statement S , if one can coherently imagine a situation which, when considered as actual, verifies S , then the primary intension of S is such that there is a centered metaphysically possible world at which it is true.

\sim (CP-) It is not the case that, for any statement S , if $\sim S$ is not *a priori*, then the primary intension of S is such that there is a centered metaphysically possible world at which it is true.

⁸² Yablo (1999) gives an argument like this one in a very abbreviated form. His version is unnecessarily weak, however, because it involves God—a being of which Chalmers (1999, p. 484) claims to have difficulty conceiving for reasons other than its necessary existence.

(N-) says that if (CP-) is true, then it is primarily necessary. (CP-) is primarily necessary just if \sim (CP-) is primarily impossible, so (N-) implies that if (CP-) is true, then \sim (CP-) is primarily impossible. By (CP-), if \sim (CP-) is primarily impossible, then \sim (CP-) is not primarily negatively conceivable. Thus, given (N-), if (CP-) is true, then \sim (CP-) is not primarily negatively conceivable. To say that \sim (CP-) is not primarily negatively conceivable is just to say that its negation, (CP-), is *a priori*. Thus, given (N-), if (CP-) is true, then (CP-) is *a priori*. But as I have argued in section V, (CP-) is not *a priori*. Thus, given that (N-) is true, (CP-) is false.

A similar line of reasoning shows that (CP+) is false. (N+) says that if (CP+) is true, then it is primarily necessary. (CP+) is primarily necessary just if \sim (CP+) is primarily impossible, so (N+) implies that if (CP+) is true, then \sim (CP+) is primarily impossible. By (CP+), if \sim (CP+) is primarily impossible, then \sim (CP+) is not primarily positively conceivable. Thus, given (N+), if (CP+) is true, then \sim (CP+) is not primarily positively conceivable. But \sim (CP+) is primarily positively conceivable: one can coherently imagine a situation which verifies it when considered as actual. Thus, given that (N+) is true, (CP+) is false.⁸³

⁸³ If the negpos principle is true, then there is an alternate version of this half of the argument that relies not on the primary positive conceivability of \sim (CP+), but on the premise (defended in section V) that (CP+) is not *a priori*. Having established that, given (N+), if (CP+) is true then \sim (CP+) is not primarily positively conceivable, we infer via the negpos principle that, given (N+), if (CP+) is true then \sim (CP+) is not primarily negatively conceivable. We then reason that since to say that \sim (CP+) is not primarily negatively conceivable is just to say that (CP+) is *a priori*, and since (CP+) is not *a priori*, it follows that that if (N+) is true, (CP+) is false. But (N+) is indeed true, so (CP+) is false. Since I am neutral in this essay on the question of whether the negpos principle is true, I will ignore this alternate version of the argument and stick

ASSUMPTIONS

The first argument, but not the second, assumes the standard S5 picture of the space of worlds, according to which exactly the same worlds are metaphysically possible from the point of view of every metaphysically possible world. It also assumes that M and $\sim M$ are semantically neutral.

The second argument, but not the first, assumes that $(N+)$ and $(N-)$ are true. It makes no difference whether one considers it an argument against modal rationalism that relies on these two extra assumptions or an argument against the conjunction of modal rationalism with these two assumptions. I shall think of it in the former way.

Both arguments make assumptions about what is conceivable. The first assumes that M and $\sim M$ are both primarily positively conceivable and primarily negatively conceivable. The second assumes that $\sim(CP-)$ is primarily negatively conceivable (since it assumes that $(CP-)$ is not *a priori*), and that $\sim(CP+)$ is primarily positively conceivable. Let us call these assumptions the *conceivability premises*.

There are only four ways to challenge the two arguments just presented. The first is to reject the standard S5 view of the space of worlds. The second is to deny that M and $\sim M$ are semantically neutral. The third is to deny $(N+)$ or $(N-)$. The fourth is to reject these arguments' conceivability premises. Even if only one

with the version presented in the text.

of the two arguments presented above is sound, (CP+) and (CP-) are both false. Thus, in order to save modal rationalism from refutation, one must show that neither of these arguments is sound. Denying the S5 view or the semantic neutrality of M and \sim M has no impact on the soundness of the second argument, and denying (N+) or (N-) has no impact on the soundness of the first. There are only two ways to disarm both arguments: (i) deny the S5 picture or the semantic neutrality of M or \sim M *and* deny (N+) or (N-), or (ii) deny at least one of the conceivability premises for each argument.

I have no more to say in defense of (N+) and (N-) than I did when I first introduced them: it seems obvious that (CP+) and (CP-) are metaphysically necessary if true, and it is also highly plausible that they are semantically neutral, so it is highly plausible that (N+) and (N-) are true. Nor do I have much to say in defense of the assumption that M and \sim M are semantically neutral: it seems obvious that as far as the relationship between their primary and secondary intensions go, M and \sim M are not like typical tokens of 'I am a philosopher' and 'water is H₂O', but like typical tokens 'there are even numbers', whose primary and secondary intensions are equivalent. Nor, finally, shall I attempt to justify the S5 picture, which is accepted by the vast majority of contemporary metaphysicians, including Chalmers⁸⁴ (though there are notable dissenters⁸⁵). Only the conceivability premises remain. In the next section, I will show that they are overwhelmingly likely to be true. Given the plausibility of my other assumptions, the overall result will be that the modal rationalist cannot plausibly

⁸⁴ When Chalmers considers Yablo's abbreviated version of my first argument, he says that "denying S5" would be "unpromising" [Chalmers (2005), section 8].

⁸⁵ For example, Salmon (2005).

evade both of my arguments.

VII. The Conceivability Premises

The first argument assumes that both M and $\sim M$ are primarily positively conceivable and primarily negatively conceivable. The second argument assumes that $\sim(\text{CP-})$ is primarily negatively conceivable and that $\sim(\text{CP+})$ is primarily positively conceivable. If all of these assumptions are false, then none of the arguments is sound.

There are two ways these conceivability premises could turn out to be false: they could turn out to be false for the same reason—that is, because of a feature shared by all of the statements whose primary conceivability they assert—or they could turn out to be false for different reasons. Since it would be unnecessary to consider the truth of each of these premises individually if one already knew of a reason on account of which they are all false, I begin by considering whether there is any feature, held in common by all of the statements whose conceivability is asserted in these premises, that renders these premises false.

GENERAL CONSIDERATIONS

There is no particular subject matter (e.g., conscious beings, metaphysically necessary existence, the relationship between conceivability and possibility) shared by all of the statements mentioned in the conceivability premises that would account for these premises' all being false. Those premises do share a

salient feature, however: they all assert the primary conceivability of a statement that involves modality. Perhaps the modal rationalist could reject all of the conceivability premises in one blow by denying that any statements involving modality are positively or negatively primarily conceivable.

This would not be a promising strategy, though. As far as primary *negative* conceivability is concerned, it is obvious that some statements involving modality are primarily conceivable. To say that a statement is primarily negatively conceivable is just to say that its negation is not *a priori*, and there are plenty of statements involving modality whose negations fail to be *a priori*. For example, 'John can finish the exam in less than three hours' is a modal statement whose negation, 'John cannot finish the exam in less than three hours' is not *a priori*, so 'John can finish the exam in less than three hours' is primarily negatively conceivable.

It is also obvious that some statements involving modality are primarily *positively* conceivable. A statement S is primarily positively conceivable just in case one can coherently imagine a situation U which, when considered as actual, verifies it; and a situation U verifies a statement S when considered as actual just if it is incoherent to suppose that U actually obtains without S's being true. Let us say that a statement S_a is an *obvious a priori consequence* of a statement S just if it is obvious that S_a can be inferred from S using *a priori* reasoning alone. Surely, if a coherently imaginable situation U verifies a statement S when considered as actual, then any statement S_a that is an obvious *a priori* consequence of S is also verified by U when it is considered as actual. After all, if S_a is an obvious *a priori* consequence of S, then it is obviously incoherent to suppose that S is true but

that S_a is false. Thus, if it is incoherent to suppose that U actually obtains without S 's being true, then it is also incoherent to suppose that U actually obtains without S_a 's being true. Now, for every primarily positively conceivable non-modal statement S , we can construct a statement S' that results from affixing 'it is possible that' to the beginning of S . Given the fact that it is an obvious *a priori* truth about possibility that anything that is actual is possible, every such statement S' will be an obvious *a priori* consequence of its corresponding statement S . Thus, every such statement S' will be primarily positively conceivable. If 'there is a panda in Brazil' is primarily positively conceivable, then so is 'it is possible that there is a panda in Brazil'; if 'a zombie exists' is primarily positively conceivable, then so is 'it is possible that a zombie exists'. Since some non-modal statements are primarily positively conceivable, some modal statements are primarily positively conceivable as well.⁸⁶

It is clear, then, that there are no general grounds for denying the primary conceivability, negative or positive, of all statements involving modality: many such statements are primarily negatively conceivable, and many of them are primarily positively conceivable. The modal rationalist cannot reject all of the conceivability premises at once simply by denying that statements involving modality are primarily conceivable.

There is another salient feature shared by the statements mentioned in the conceivability premises, however—a feature more specific than simply

⁸⁶ This is a modification of a point I got from Tim Williamson, who also observes that, if a modal rationalist thinks that every primarily possible statement is primarily positively conceivable, she will have to maintain that statements involving modality are primarily positively conceivable in order to avoid the absurd conclusion that all such statements are primarily impossible.

involving modality: they are all statements of necessity or impossibility. The first argument assumes that M and $\sim M$ are both primarily conceivable, and thus that statements asserting or denying the *necessary* existence of a particular being are primarily conceivable. The second argument assumes that $\sim(\text{CP-})$ is primarily negatively conceivable, and thus that statements asserting the primary *impossibility* of some primarily negatively conceivable statement are primarily negatively conceivable. It also assumes that $\sim(\text{CP+})$ is primarily positively conceivable, and thus that statements asserting the primary *impossibility* of some primarily positively conceivable statement are primarily positively conceivable. The modal rationalist might argue that, even if some statements involving modality are primarily conceivable, no statement of necessity or impossibility is primarily conceivable.

Once again, as far as primary *negative* conceivability is concerned, this is obviously not a promising line of argument. Consider any statement of necessity or impossibility, S , and consider its negation, $\sim S$. For example, S could be (CP+) and $\sim S$ could be $\sim(\text{CP+})$; or S could be M and $\sim S$ could be $\sim M$. Either S is *a priori* and $\sim S$ is not, or $\sim S$ is *a priori* and S is not, or neither S nor $\sim S$ is *a priori*. In the first case, S will be primarily negatively conceivable, since its negation, $\sim S$, is not *a priori*. In the second, $\sim S$ will be primarily negatively conceivable, since its negation, S , is not *a priori*. In the third, both S and $\sim S$ will be primarily negatively conceivable, since they are each other's negations and neither of them is *a priori*. In every possible situation, at least one of S and $\sim S$ will be primarily negatively conceivable. Since S and $\sim S$ stand for any two contradictory statements of necessity or impossibility, some statements of necessity or impossibility are

primarily negatively conceivable. There is no general reason to deny the primary negative conceivability of such statements.

Nor is this line of argument promising when it comes to primary *positive* conceivability. Let S be any statement of necessity or impossibility that is knowable *a priori* (e.g., 'necessarily, if P then P'). For reasons I already explained in my proof of the posneg principle, given that S is *a priori*, any coherently imaginable situation will verify it when considered as actual. Thus, any statement of necessity or impossibility that is knowable *a priori* will be primarily positively conceivable. There is therefore no general obstacle to the primary positive conceivability of statements of necessity or impossibility.⁸⁷

What about statements of necessity or impossibility that are not *a priori*? Could there be a general reason to deny their primary conceivability? If so, a modal rationalist could say that, because none of the statements mentioned in the conceivability premises is *a priori*, none of them is primarily conceivable.

This approach obviously does not work as far as primary negative conceivability is concerned, because there obviously are statements of necessity or impossibility that are not *a priori* and whose negations are also not *a priori* (e.g., 'it is metaphysically necessary that there exists an odd number of conscious beings'). To say that a statement is primarily negatively conceivable is just to say that its negation is not *a priori*, so these statements are all primarily negatively conceivable.

When it comes to primary positive conceivability, however, things get a little tricky. As I mentioned when I explained Chalmers's account of positive

⁸⁷ I owe this point to Tim Williamson.

conceivability, Chalmers writes that a situation is “a configuration of objects and properties within a world.”⁸⁸ To say that a statement is primarily positively conceivable is to say that one can coherently imagine a situation that verifies it when considered as actual. Thus, given Chalmers’s account of situations, to say that a statement is primarily positively conceivable is to say that it is verified by a coherently imaginable configuration of objects and properties within a world when this configuration is considered as actual. Now, as I have just argued, there is nothing problematic about the verification of an *a priori* statement of necessity or impossibility by an imaginable configuration of objects and properties within a single world: the mere fact that such a statement is *a priori* guarantees its verification by any coherently imaginable situation whatsoever, so we need not even look at any particular situation to see if it verifies that statement. If a statement of necessity or impossibility is not *a priori*, however, one does need to look at a particular coherently imaginable situation to see whether ideal rational reflection reveals it to be a situation in which the statement is true. Here, one might worry that there is a problem. For one might think that certain modal properties (e.g., *necessarily existing*, *being necessarily prime*) are properties that one cannot imagine unless one imagines the entire space of possible worlds, and one might therefore conclude that Chalmers is referring to *non-modal* properties when he writes that a situation is “a configuration of objects and properties within a world.” But if a situation is a configuration of objects and *non-modal* properties within a single world, it would seem that one cannot imagine a situation in which, say, it is metaphysically necessary that all numbers exist: for

⁸⁸ Chalmers (2002a), p. 151.

to imagine this, one must either imagine every number in a single world having the modal property *necessarily existing*, or one must imagine the entire space of worlds and imagine that every number exists at every world. For these reasons, one might think that Chalmers's account of situations rules out the primary positive conceivability of statements of necessity and impossibility that are not *a priori*.

On reflection, however, it is clear that Chalmers must allow for the primary positive conceivability of statements of necessity and impossibility that are not *a priori*. Recall that, in giving his account of positive conceivability, Chalmers does not purport to be introducing a new kind of conceivability of his own discovery or invention. Instead, he takes himself to be describing a kind of conceivability with which philosophers are already familiar and of which they already have an intuitive understanding—namely, the kind involved in a typical philosophical thought experiment. Consequently, it is in principle possible for his account of positive conceivability to be incorrect, and it is legitimate for philosophers to use their preexisting understanding of the typical method of philosophical thought experiments to guide their interpretation of that account. Now, it is surely within the bounds of the typical method of philosophical thought experiments to imagine a situation in which all numbers necessarily exist, one in which God necessarily exists, or one in which it is necessary that every plurality has a sum. Thus, given that Chalmers's account of positive conceivability is supposed to describe the kind of conceivability involved in typical philosophical thought experiments, it is clear that one should not interpret his claim that a situation is a “configuration of objects and properties

within a world” in such a way that one turns out not to be imagining situations when one imagines these things. One should either understand the properties to which he refers as including modal properties, or one should deviate from the letter of his account of situations and allow a configuration of objects and properties *in one or more worlds* to count as a situation.

Thought experiments involving an imagined case in which a certain counterfactual conditional is true are particularly obvious examples of the use of coherently imaginable situations that include either modal properties or a plurality of worlds. Take a thought experiment in moral philosophy involving an imagined case in which, if Bill had not flipped the switch, Jimmy would not have been killed by the trolley. There are two ways to imagine this case. One can imagine Bill or Jimmy possessing certain modal properties (e.g., Jimmy possessing the property *being such that he would not have been killed by the trolley if Bill had not flipped the switch*). Or, if one accepts an account of counterfactual conditionals that relies on possible worlds (e.g., David Lewis’s⁸⁹), one can imagine a certain relationship holding between possible worlds (e.g., either that there are no possible worlds at which Bill does not flip the switch, or that the possible world closest to the actual world at which Bill does not flip the switch is one at which Jimmy is not killed by a trolley). Again, situations must either be allowed to include modal properties, or they must be allowed to be configurations of objects and properties in one or more worlds.

In light of all this, it is obvious that some statements of necessity or impossibility that are not *a priori* are primarily positively conceivable. It is

⁸⁹ Lewis (1973), p. 560.

coherent to imagine a situation which, when considered as actual, is correctly described as one in which there exists an object that exists at every metaphysically possible world and that is a sphere with a two-meter diameter at every metaphysically possible world. When considered as actual, this situation verifies the statement 'it is metaphysically necessary that a sphere exists'. After all, it is incoherent to suppose that the situation actually obtains without its being metaphysically necessary that a sphere exists. The statement 'it is metaphysically necessary that a sphere exists' is *a posteriori*, and it is a statement of necessity. Thus, at least some *a posteriori* statements of necessity or impossibility are primarily positively conceivable.

There is another way to arrive at this conclusion. For any statement *S*, let *S'* be the statement that results from applying an actuality operator to *S*, and let *S''* be the statement that results from applying a metaphysical necessity operator to *S'*. (Thus, if *S* is 'there is a cat on the mat', *S'* is 'actually, there is a cat on the mat', and *S''* is 'it is metaphysically necessary that actually, there is a cat on the mat'.) For any statement *S*, it is obviously *a priori* that *S* is true if and only if *S'* is true, and it is obviously *a priori* that *S'* is true if and only if *S''* is true. After all, it is obviously *a priori* that *S* is true just if it is true at the actual world, and that *S'* is true just if the exact same conditions are met—that is, just if *S*, the part of *S'* that is within the scope of the actuality operator, is true at the actual world. Furthermore, it is obviously *a priori* that *S''* is true just if *S'* is true at every metaphysically possible world, that *S'* is true if it is true at every metaphysically possible world (since the actual world is a metaphysically possible world), and that *S'* is true at every metaphysically possible world if it is true (since the

actuality operator refers rigidly to the actual world). It follows from this that, for any S , it is obviously *a priori* that S is true if and only if S'' is true. Thus, for any statement S , S'' is an obvious *a priori* consequence of S . As I mentioned earlier, if a coherently imaginable situation verifies a given statement, it also verifies any statements that are obvious *a priori* consequences of that statement. Thus, for any statement S , if a coherently imaginable situation verifies S , then it also verifies S'' . Plenty of statements S are verified by coherently imaginable situations: 'water is H_2O ', 'water is XYZ', and 'there exists a zombie', just to name a few. Thus, plenty of statements S'' are verified by coherently imaginable situations: 'it is metaphysically necessary that actually, water is H_2O ', 'it is metaphysically necessary that actually, water is XYZ', 'it is metaphysically necessary that actually, there exists a zombie', etc. All of these statements are statements of necessity or impossibility, and at least some of them are knowable only *a posteriori*, since the statements of which they are obvious *a priori* consequences are knowable only *a posteriori*. (For example, 'it is metaphysically necessary that actually, there is a cat on the mat' is *a posteriori* even though it is an obvious *a priori* consequence of 'there is a cat on the mat' because the latter statement is *a posteriori*.) Thus, some *a posteriori* statements of necessity or impossibility are primarily positively conceivable.⁹⁰

It turns out, then, that there is no general obstacle to the primary conceivability, positive or negative, of *a posteriori* statements of necessity or impossibility. The modal rationalist cannot reject all of the conceivability premises in one blow by arguing that every such statement is primarily

⁹⁰ I got this argument from John Hawthorne.

inconceivable. Nor does there appear to be any other feature shared by the statements mentioned in the conceivability premises that could plausibly show that all of these premises are false. I turn, therefore, to defending the conceivability premises one-by-one.

THE FIRST ARGUMENT'S CONCEIVABILITY PREMISES

Recall that M is a typical token of the English sentence type 'there exists a metaphysically necessary conscious being', and \sim M is a typical token of the English sentence type 'there does not exist a metaphysically necessary conscious being'. The first argument has four conceivability premises: (i) M is primarily negatively conceivable, (ii) M is primarily positively conceivable, (iii) \sim M is primarily negatively conceivable, and (iv) \sim M is primarily positively conceivable. Let us consider each of these premises in turn.

M is primarily negatively conceivable just in case \sim M is not *a priori*. It seems obvious, after extended reflection, that \sim M is not *a priori*, for such reflection reveals no way to reason *a priori* (from the concepts of existence, metaphysical necessity, and consciousness, or from anything else) to the conclusion that there does not exist a metaphysically necessary conscious being. By contrast, extended reflection reveals no plausible reasons to think that \sim M is *a priori*. Thus, we are entitled by (AP1)—or equivalently, (AP2)—to assume that \sim M is not *a priori* (and hence that M is primarily negatively conceivable) until a compelling reason to think otherwise is produced.

M is primarily positively conceivable just if one can coherently imagine a situation that verifies it when considered as actual. Upon reflection, it seems obviously possible to coherently imagine such a situation: just imagine a situation which, when considered as actual, is correctly described as one in which there is a conscious monad that exists at every metaphysically possible world. It seems impossible to find anything incoherent about this situation, no matter how hard one looks; and clearly, it is incoherent to suppose that this situation actually obtains without M's being true. By (AP3), we are entitled to assume that M is primarily positively conceivable until a compelling reason to think otherwise is produced.

If (CP-) were *a priori*, a compelling reason to think that M is not primarily negatively conceivable could arguably be produced, provided that we assume that \sim M is primarily negatively conceivable. As I pointed out earlier, if M is true, then it is metaphysically necessary. Moreover, one of M and \sim M is metaphysically necessary and the other is metaphysically impossible. Thus, if M is true, then \sim M is metaphysically impossible and (because it is semantically neutral) primarily impossible. Thus, M implies that \sim M is primarily impossible. Now, if (CP-) is true, then given that \sim M is primarily negatively conceivable, \sim M is primarily possible. Since M implies that \sim M is not primarily possible, given that \sim M is primarily negatively conceivable, M implies that (CP-) is false, and hence that \sim (CP-) is true. Thus, if M is primarily negatively conceivable (it could be argued), then \sim (CP-) is primarily negatively conceivable. However, the argument would go, \sim (CP-) is not primarily negatively conceivable, because (CP-) is *a priori*. Thus, M is not primarily negatively conceivable.

If (CP+) were *a priori*, then provided that we assume that $\sim M$ is primarily positively conceivable, a parallel argument could arguably show that M is not primarily positively conceivable. For the same reasons given above, M implies that $\sim M$ is primarily impossible. Now, if (CP+) is true, then given that $\sim M$ is primarily positively conceivable, $\sim M$ is primarily possible. Since M implies that $\sim M$ is not primarily possible, given that $\sim M$ is primarily positively conceivable, it implies that (CP+) is false, and hence that $\sim(\text{CP+})$ is true. Thus, if M is primarily positively conceivable (it could be argued), then $\sim(\text{CP+})$ is primarily positively conceivable. However, it would be argued, $\sim(\text{CP+})$ is not primarily positively conceivable. For (CP+) is *a priori* (it would be said), and this implies both that $\sim(\text{CP+})$ is not primarily negatively conceivable and (via the posneg principle) that $\sim(\text{CP+})$ is not primarily positively conceivable. Thus, M is not primarily positively conceivable.⁹¹

If they are sound, then the two parallel arguments just presented give us compelling reasons to abandon the premises that M is primarily negatively conceivable and that M is primarily positively conceivable. But it seems obvious, after extended reflection, that they are not sound, and we are entitled to assume that they are not sound until compelling reasons to believe otherwise are produced. For it seems obvious, after such reflection, that neither (CP+) nor (CP-) is *a priori*, and we are entitled to assume this until we are given compelling

⁹¹ In responding to Yablo's abbreviated version of my first argument (in which he relies on the conceivability of the existence of a metaphysically necessary god), Chalmers argues that because one cannot conceive of the existence of such a god without thereby conceiving of the falsity of modal rationalism, and because modal rationalism is *a priori* and therefore not conceivably false, the existence of a metaphysically necessary god is inconceivable. (See Chalmers (1999), pp. 484-85.) Thus, Chalmers would presumably approve of the two arguments I just presented.

reasons to believe otherwise. It is difficult to see how one could undermine the two particular conceivability premises under consideration, if not by illegitimately assuming the *a priori* of (CP+) and (CP-).

~M is primarily negatively conceivable just in case M is not *a priori*.

Certain philosophical views commit one to the *a priori* of M (e.g., the view that God is a metaphysically necessary being and that it is *a priori* that God exists), but such views are implausible precisely because it is plausible that M is not *a priori*. Even after extended reflection, it seems perfectly coherent to suppose that no metaphysically necessary conscious being exists, so it seems obvious after extended reflection that M is not *a priori*. Such reflection reveals no plausible reasons to suppose that M is *a priori*. So by (AP1) and (AP2), we are entitled to assume that ~M is primarily negatively conceivable until compelling reasons to believe otherwise are produced.

~M is primarily positively conceivable just in case one can coherently imagine a situation that verifies it when considered as actual. Even after extended reflection, it seems obviously possible to coherently imagine such a situation: just imagine a situation which, when considered as actual, is correctly described as one in which every conscious being is extended and no extended being exists at every metaphysically possible world. The primary positive conceivability of ~M seems as evident as that of any statement, and extended reflection does not reveal any plausible reasons to suppose that it is not primarily positively conceivable. By (AP3), we are entitled to assume that ~M is primarily positively conceivable until compelling reasons to believe otherwise are produced.

THE SECOND ARGUMENT'S CONCEIVABILITY PREMISES

The second argument relies on two conceivability premises: (i) $\sim(\text{CP-})$ is primarily negatively conceivable, and (ii) $\sim(\text{CP+})$ is primarily positively conceivable.

If (CP+) and (CP-) were *a priori*, then it would follow from the definition of primary negative conceivability that $\sim(\text{CP+})$ and $\sim(\text{CP-})$ are not primarily negatively conceivable, and it would follow from this and the posneg principle that $\sim(\text{CP+})$ is not primarily positively conceivable. But we are entitled to assume that (CP+) and (CP-) are not *a priori*. Thus, we are entitled to assume that $\sim(\text{CP-})$ is primarily negatively conceivable, and that the primary positive inconceivability of $\sim(\text{CP+})$ cannot be established in the manner just described.

Even after extended reflection, it seems obvious that $\sim(\text{CP+})$ is primarily positively conceivable, and there seem not to be any plausible reasons to think that it is not primarily positively conceivable. It seems obviously possible to coherently imagine a situation that verifies it when considered as actual: just imagine a situation which is correctly described, when considered as actual, as one in which a particular typical token T of the English sentence type 'not every conscious being has parts' is primarily positively conceivable (as it actually is) but is such that its primary intension is false at every metaphysically possible world. This situation verifies $\sim(\text{CP+})$ when considered as actual, for it is incoherent to suppose that it actually obtains without it being the case that

$\sim(\text{CP}+)$ is true; and there does not seem to be anything incoherent about this situation, no matter how hard one looks. On the other hand, extended reflection produces no plausible reasons to think that $\sim(\text{CP}+)$ is not primarily positively conceivable. By (AP3), we are entitled to the assumption that $\sim(\text{CP}+)$ is primarily positively conceivable until compelling reasons to think otherwise are produced.

VIII. Conclusion

If we take for granted the claim that every expression token has both a primary and a secondary intension, and if we assume that not every statement is primarily possible, then Chalmers's modal rationalism—the conjunction of (CP+) and (CP-)—is an intelligible theory that makes interesting claims about the relationship between metaphysical possibility and the *a priori* domain. Chalmers's positive case for modal rationalism is entirely unpersuasive, however. Indeed, there are no good reasons to think that either (CP+) or (CP-) is true.

It seems obvious, both initially and after extended reflection, that neither (CP+) nor (CP-) is *a priori*, and hence that modal rationalism is not *a priori*. By contrast, extended reflection produces no plausible reasons to suppose that either (CP+) or (CP-) is *a priori*. Thus, by (AP1), we are entitled to assume that neither (CP+) nor (CP-) is *a priori*, and to use this assumption in further reasoning, until compelling reasons to believe otherwise are produced. Each of my two arguments shows that (CP+) is false if it is not *a priori* and that (CP-) is false if it is not *a priori*, and each of them relies only on highly plausible premises. Unless someone produces a compelling reason to think that either (CP+) or (CP-) is *a priori*, there is an extremely strong case for thinking that (CP+) and (CP-) are both false, and thus that Chalmers's modal rationalism is false.

Works Cited

- Bealer, George (2002), 'Modal Epistemology and the Rationalist Renaissance', in Tamar Szabo Gendler and John Hawthorne (eds.), *Conceivability and Possibility* (Oxford: Clarendon Press), 71-125.
- Chalmers, David (1996), *The Conscious Mind* (New York: Oxford University Press).
- _____ (1999), 'Materialism and the Metaphysics of Modality', *Philosophy and Phenomenological Research*, 59:473-96.
- _____ (2002a), 'Does Conceivability Entail Possibility?' in Tamar Szabo Gendler and John Hawthorne (eds.), *Conceivability and Possibility* (Oxford: Clarendon Press).
- _____ (2002b), 'On Sense and Intension' *Philosophical Perspectives*, 16:135-82.
- _____ (2005), 'The Two-Dimensional Argument Against Materialism', <<http://consc.net/papers/2dargument.html>>
- _____ (2006), 'The Foundations of Two-Dimensional Semantics' in Manuel García-Carpintero and Josep Macià (eds.), *Two-Dimensional Semantics* (Oxford: Clarendon Press).
- García-Carpintero, Manuel and Josep Macià (eds.) (2006), *Two-Dimensional Semantics* (Oxford: Clarendon Press).
- Jackson, Frank (1998), *From Metaphysics to Ethics* (Oxford: Clarendon Press).
- Kaplan, David (1989), 'Demonstratives' in Joseph Almog, John Perry, and Howard Wettstein (eds.), *Themes from Kaplan* (New York: Oxford University Press).
- Lewis, David (1973), 'Causation', *The Journal of Philosophy*, 70:556-567.
- Plantinga, Alvin (1974), *The Nature of Necessity* (Oxford: Clarendon Press).
- Rosen, Gideon (forthcoming), 'The Limits of Contingency', in F. MacBride and C. Wright (eds.), *Being Committed* (Oxford: Oxford University Press).
- Salmon, Nathan (1986), *Frege's Puzzle* (Cambridge, MA: The MIT Press).
- _____ (2005), 'The Logic of What Might Have Been' in *Metaphysics, Mathematics, and Meaning: Philosophical Papers* (Oxford: Oxford University Press), 129-49.
- Soames, Scott (2005), *Reference and Description: The Case against Two-Dimensionalism* (Princeton: Princeton University Press).
- van Inwagen, Peter (2001), 'Modal Epistemology', in *Ontology, Identity, and*

Modality (Cambridge: Cambridge University Press), 243-58.

Yablo, Stephen (1999), 'Concepts and Consciousness', *Philosophy and Phenomenological Research*, 59:455-63.