

- no a priori route to possibility. Unlike the model we favor, Kripke keeps the a priori tie between imaginability/conceivability and possibility. However, he rejects the idea that we can know a priori (i.e., in an empirically indefeasible way) what we are imagining or conceiving. Thus, he denies that we are really imagining that there is water but no H₂O when we seem to ourselves to be imagining that there is water but no H₂O. Our model is compatible with first-person authority in such case.
13. It should be noted that one can reject dualist intuitions while embracing other modal intuitions as trustworthy. Not all modal intuitions stand or fall together. A selective skepticism is a legitimate option. For further discussion, see Hill (1997) and Hill and McLaughlin (1999).
 14. A more appropriate term would have been *new wave materialism/function-alism*. Hill and I, however, favor a strictly materialist version: we hold that sensory states have neural correlates and that they are identical with their neural correlates. See Hill and McLaughlin (1999).
 15. The difference between *phenomenal property* and *sensory property* and between *phenomenal concepts* and *sensory concepts* is merely verbal.
 16. Horgan and Tienson (p. 315) point out that the argument is "similar in spirit" to Stephen White's (1986) property dualism argument. (His argument was inspired by "objection five" in Smart [1962].) I respond to White's property dualism argument in McLaughlin (1997).
 17. In general, from the fact that a representation R of something, X, does not represent X as F, it does not follow that R represents X as not-F.
 18. Compare the objection that Colin McGinn (Chapter 14, this volume) raises to a certain way of trying to solve the mind-body problem, a way that looks on the surface (see Note 19) like the new wave materialist solution to the mind-body problem for phenomenal properties.
 19. From the perspective of new wave materialism, McGinn (Chapter 14, this volume), or those who try to solve the mind-body problem in the way he criticizes, makes the mistake of assuming that phenomenal concepts phenomenologically present the properties to which they refer. (New wave materialists do not try to solve the mind-body problem in the way McGinn criticizes.)
 20. Keep in mind here the point made in Note 7.
 21. New wave materialists wish to be materialists across the board, so to speak, and so of course require an account of how concepts refer that is compatible with materialism.

In Physicism and Its Discontents, ed. Carl
Gillet & Barry Loewer. Cambridge UP,
2001.

Physicalism Unfalsified: Chalmers's Inconclusive Conceivability Argument

ANDREW MELNYK

Let a *conceivability argument* be any argument that aims to refute physicalism by showing that some claim that physicalism must treat as necessary is not in fact necessary because its negation is conceivable and hence genuinely possible. According to conceivability arguments in the recent tradition, it is conceivable that so-and-so iff the thought that so-and-so can be entertained without explicit contradiction or the sense of conceptual blockage one has (albeit to an unusually intense degree) in attempting to entertain the negation of the claim that two is a number. According to such arguments, the claim that it is assumed physicalism must treat as necessary is a certain more or less immediate consequence of any physicalist type-identity claim whereby some nonphysical property – some property expressible in nonphysical vocabulary – is identified with either a physical property or a physically realized functional property. So suppose the (nonphysical) property of being N is identified with the (physical or functional) property of being P. Then, on the assumption that the identity statement is formed using rigid designators and that such identity statements are necessary, it follows that, necessarily, an object is N iff it is P. But if, as the argument claims, it is conceivable, and hence possible, that something should be N but not P (or P but not N), then it is false that, necessarily, an object is N iff it is P, and hence false that being N just is being P.

But traditional conceivability arguments of this sort fail. Even though it is surely conceivable, in the sense specified, that something should be N but not P, the real possibility that something should be N but not P just does not follow. Someone competent with the concept of table salt could gain competence with the concept of NaCl in a chemistry class, which omitted to mention that NaCl is the same stuff as table salt. It would then be conceivable for this person, in the sense specified, that his food be sprinkled with salt but not with NaCl. He could entertain the thought that this is so without explicit contradiction or any sense of

conceptual blockage. But it is not really possible that his food be sprinkled with salt but not with NaCl, because salt and NaCl are the very same stuff, and it is not possible that his food be both sprinkled and not sprinkled with the very same stuff. So the conceivability, in the sense specified, of a proposition does not in general entail its possibility.

The explanation for this failure of entailment is that the reference of our concepts is not an a priori matter. One does not know simply on the basis of one's competence with a concept what property that concept picks out. Accordingly, one might be competent with two concepts without realizing that in fact they pick out the very same property.¹ Someone might therefore entertain the formally consistent thought, "My food is sprinkled with salt but not with NaCl," and in doing so feel no sense of conceptual blockage, even though, given that 'salt' and 'NaCl' corefer, and that no food is both sprinkled and not sprinkled with the same stuff, this thought expresses no possible state of affairs. Similarly, someone might entertain the formally consistent thought, 'Something is N but not P,' and in doing so feel no sense of conceptual blockage, even if, because being N is the same property as being P, and it is impossible that something should both have and not have the very same property, it is impossible that something should be N but not P. So it is because the reference of our concepts is an a posteriori matter that conceivability (in the sense specified) does not entail possibility.² Someone can infer that her consistent and blockage-free thought that something is N but not P expresses a genuine possibility only if she makes an a posteriori assumption that conceptual competence alone does not guarantee: that 'being N' and 'being P' do not pick out the very same property. But this assumption begs the question against physicalism, because it is tantamount to assuming that being N is not the same property as being P.

Traditional conceivability arguments fail, therefore, because they rest on a false assumption in the philosophy of mind and language: that one knows, simply on the basis of one's competence with a concept, what property the concept picks out. Physicalism is safe from any traditional conceivability argument so long as whatever turns out to be the correct view of the determination of the reference of our concepts implies the falsity of this assumption.³

Now in his recent book, *The Conscious Mind*, David Chalmers offers a conceivability argument for the property-dualist conclusion that the property of having a red sensation, like all other phenomenal properties, is in no sense physical: it is neither type-identical with any physical property nor is it a functional property that is, as it happens, always realized physically.⁴ But his argument is not of the traditional sort I have been

discussing. It employs a different and proprietary notion of conceivability according to which the inference from conceivability to possibility is unproblematic; it is not targeted against the physicalist claim that each phenomenal property is necessarily coextensive with some physical property, and, most importantly, it is quite consistent with the very claim that undermines traditional conceivability arguments, namely that one can be competent with a concept and still not know what property it refers to.

My question is whether Chalmers's conceivability argument refutes physicalism. But because he introduces an elaborate technical machinery with which to express it, I cannot evaluate it either fairly or effectively while ignoring this machinery. In this chapter, then, I aim to reconstruct the argument as sympathetically as possible, taking his technical machinery as seriously as he does, but to argue that even so it is inconclusive, because it rests on a nonobvious psychosemantic assumption that we have been given no reason to believe. Having explained Chalmers's technical machinery in my first section, I use it in my second to reconstruct his antiphysicalist argument. In the next two sections, I identify the psychosemantic assumption on which the argument rests, and argue that it is in fact unsupported. Finally, I venture some conclusions about conceivability arguments in general.⁵

I

Chalmers claims (pp. 52–62) that an important part of the meaning of any concept is, or can be represented by, a pair of functions, which he calls the primary intension and the secondary intension of the concept. Each function is a mapping from possible worlds to referents in worlds, but not the same mapping. A concept's *primary* intension (henceforth, *I-intension*) is a mapping from each possible world, considered as *actual*, to the concept's referent in that world. It tells us not only what the concept actually refers to, given the way the actual world is, but also what it would have referred to, if the actual world had turned out to be different in various ways. "The [primary] intension," he says, "*specifies* how reference depends on the way the external world turns out" (p. 57). The *I-intension* of 'water,' for example, which for expository purposes Chalmers suggests that we treat roughly as 'watery stuff,' maps each world onto whatever stuff in that world (if any) happens to be watery, regardless of whether the stuff that happens to be watery in that world is H₂O, or XYZ, or whatever. A concept's *secondary* intension (henceforth, *2-intension*), by contrast, is a mapping from each possible world, considered as *counterfactual*, to the concept's referent in that world. It tells us what the concept would refer to in each such world. If

a concept is a rigid designator, its 2-intension maps each possible world onto the set of objects, in that world, which constitute the referent of the concept in the actual world (as determined by the 1-intension plus the way the actual world happens to be). So, given that 'water' is rigid, the 2-intension of 'water' maps each world onto H_2O in that world, because, given the way the actual world is, the 1-intension of 'water' maps the actual world onto H_2O .

Moreover, there is a crucial difference in epistemological status between 1- and 2-intensions. Every concept, Chalmers insists, is such that its 1-intension is accessible a priori to anyone who is competent with it (e.g., pp. 57–9). But the 2-intension of some concepts – those that are rigid designators – can only be known a posteriori, because to know them one needs to know not only their 1-intensions but also what the actual world happens to be like. Notice that, since knowing the 1-intension of a concept only tells you what the concept refers to in the actual world, if you also know – a posteriori, of course – what the actual world is like, Chalmers's view implies that one does not know, simply on the basis of one's competence with a concept, what property the concept picks out. He therefore rejects the assumption on which I diagnosed traditional conceivability arguments as relying.

A statement is necessary, according to Chalmers (pp. 62–6), iff it is true in all possible worlds, and possible iff true at some possible world. But because the truth of a statement at a world can be evaluated in light of either the 1- or the 2-intensions of its constituent concepts, we can distinguish between the necessity (i.e., truth in all possible worlds) of a statement according to 1-intensions (henceforth, *1-necessity*) and its necessity according to 2-intensions (henceforth, *2-necessity*); also between the possibility of a statement (i.e., truth at some possible world) according to 1-intensions (henceforth, *1-possibility*) and its possibility according to 2-intensions (henceforth, *2-possibility*). So, for example, to say that the conditional 'If P, then Q' is 1-necessary is to say that, when the truth value of P and Q is evaluated at each world by reference to the 1-intensions of the concepts deployed in P and Q, all the worlds at which P is true are worlds at which Q is true; and to say that 'P and not-Q' is 2-possible is to say that there is some possible world at which 'P and not-Q' is true when evaluated according to 2-intensions. And, according to Chalmers, 1- and 2-necessity and 1- and 2-possibility are the only kinds of necessity and possibility we need to postulate, even in order to handle Kripke-Putnam a posteriori necessities (see his critique of 'strong' meta-physical necessity on pp. 136–8).

The distinction between 1- and 2-intensions also generates a distinction between two kinds of conceivability. *1-conceivability* and *2-*

conceivability: "The conceivability of a statement," Chalmers says (p. 67), "involves two things: first, the conceivability of a relevant world, and second, the truth of the statement in that world." If the statement comes out as true at the conceived world when evaluated according to 1-intensions, then the statement is said to be 1-conceivable; but if the statement comes out as true at the conceived world when evaluated according to 2-intensions, then it is said to be 2-conceivable.

How, on Chalmers's account, is conceivability related to possibility? If a statement is 1-conceivable, then there is some conceivable world correctly describable as one in which the statement is true according to 1-intensions, from which it follows trivially that the statement is 1-possible (p. 67). And because 1-intensions are accessible a priori, the 1-conceivability of a statement is in principle an a priori affair, and so at least in principle the 1-possibility of any 1-possible statement can be determined a priori. In practice, however, judgments of conceivability can err, because although 1-intensions are accessible a priori, they may not actually have been accessed, because of insufficient reflection, and so a conceived world may be misdescribed (pp. 67 and 99). Similarly, the 2-conceivability of a statement entails its 2-possibility, though because 2-intensions are not in general accessible a priori, 2-conceivability is not even in principle an a priori affair, and so the 2-possibility of a 2-possible statement cannot in general be determined a priori. With this account, Chalmers can give a neat diagnosis of where traditional conceivability arguments fail. If being N just is being P, then what follows is that it is 2-necessary that something is N iff it is P. Now it is certainly 1-conceivable that something be N but not P. But because even the ideal 1-conceivability of a statement does not entail its 2-possibility, this does not entail that it is 2-possible that something should be N but not P, and so the 2-necessity that something is N iff it is P, and hence the claim that being N just is being P, is safe. According to Chalmers, however, the very same technical machinery that permits this diagnosis of the error in traditional conceivability arguments can still be used to refute physicalism.

II

But how, exactly? Surprisingly, in view of the poor track record of conceivability arguments, Chalmers presents nowhere in his book a formalized version of his argument sufficiently detailed to answer this question. But, for textual reasons I have no space to give, I offer the following reconstruction, where we let 'P' express every fact, past, present

or future, including every nomic fact, expressible in the proprietary vocabulary of physics and let 'Q' be the (true) statement that sometimes people have red sensations:

- (1) It is 1-conceivable that [P and not-Q].
- ∴ (2) It is 1-possible that [P and not-Q].
- (3) All concepts in the proprietary vocabulary of physics are such that their 1-intensions map each possible world onto the same referents as do their 2-intensions.
- ∴ (4) There is a possible world in which 'P and not-Q' is true when 'P' is evaluated in light of 2-intensions and 'not-Q' is evaluated in light of 1-intensions.
- (5) If physicalism is true, then there is no possible world in which 'P and not-Q' is true when 'P' is evaluated in light of 2-intensions and 'not-Q' is evaluated in light of 1-intensions.
- ∴ (6) Physicalism is false.

This argument can be seen at once to differ from a traditional conceivability argument. First, it employs a notion of conceivability that definitionally guarantees the validity of the step from (1) to (2), the inference from conceivability to (one sort of) genuine possibility. Secondly, what is claimed, by its first premise, to be 1-conceivable is not (a) that someone should be in physical or functional state B without actually having a red sensation (where B is whatever the physicalist identifies with having a red sensation), but (b) that the *totality* of physical facts, including but not limited to the tokening of B, should be exactly as they actually are and yet that no one should have a red sensation (p. 131). In fact, Chalmers holds that both (a) and (b) are 1-conceivable, but it is only the conceivability of (b) that generates an argument against physicalism. Contrast his position as regards water, where he holds that while the analog of (a) is 1-conceivable, the analog of (b) is not: it is 1-conceivable that my cup should contain H₂O but no water, but it is not 1-conceivable that the totality of physical facts, including but not limited to my cup's containing H₂O, should be exactly as they actually are and yet that there should be no water (see ch. 2, sect. 6). That is why Chalmers is not committed to an argument for the nonphysicality of water that parallels the one for the nonphysicality of phenomenal states.

The step from (2) and (3) to (4) is valid, given the definition of '1-possibility'. Premise (3) is defended on the grounds that denying it is tantamount to claiming that physical entities and properties have hidden essences of which we know nothing (pp. 134–6).⁶ Premise (5) is obviously crucial, for it claims that physicalism is committed to the impossibility of

precisely what conceivability has (allegedly) shown to be possible. It is equivalent to the claim that, given physicalism, 'If P, then Q' is true in all possible worlds when 'P' is evaluated in light of 2-intensions and 'Q' is evaluated in light of 1-intensions, in other words, that the physical facts in a certain sense necessitate a certain class of nonphysical facts. Now most physicalists will acknowledge that, given physicalism, the truth that P necessitates the truth that Q, but insist that this necessitation is entirely a posteriori, perhaps explainable (as I would explain it) by appeal to a posteriori necessary identities holding between each nonphysical property and some functional property whose associated role is specific in physical or topic-neutral terms. They will reject (5), only conceding that, given physicalism, 'If P, then Q' is 2-necessary. But physicalists should concede (5), according to Chalmers. If physicalism is true, then the way we have discovered the world to be physically should fix or determine all the (positive) nonphysical facts there are. And if 'Q' is true at all in the actual world, which by hypothesis it is, then it is true in the actual world when evaluated according to 1-intensions, because the 1-intension and 2-intension of any concept must map the actual world onto the same referent. So 'Q,' when evaluated in light of 1-intensions, expresses a fact, and hence one of those facts that should be necessitated by the physical facts, if physicalism is true (see pp. 132–3). From (5) and (4), the antiphysicalist conclusion obviously follows.

III

But why should we suppose that (1) is true, in other words, that it is 1-conceivable that P and not-Q? Let us first clarify the meaning of this premise. For Chalmers (p. 67), "The conceivability of a statement involves two things: first, the conceivability of a relevant world, and second, the truth of the statement in that world." So premise (1) is claiming that some possible world is conceivable at which 'P and not-Q' is true when evaluated according to 1-intensions. But what is it for a possible world to be conceivable? Chalmers surely does not intend the implausible and unmaterialist (hence question-beggingly) idea that conceiving a world is a matter of directing some alleged faculty of modal perception upon some segment of extramental modal reality. All that he can plausibly mean is that conceiving a possible world is somehow *mentally representing* a possible world, in other words, hosting mental representations of a possible world⁷ – conscious or nonconscious, propositional or iconic – and doing so in a way functionally appropriate to conceiving. So what (1) is claiming is that some possible world is conceivable, in the sense of somehow mentally representable, at which 'P and not-Q' is true

when evaluated according to 1-intensions.⁸ And I shall be arguing that Chalmers has provided no good reason for supposing that this is true, which leaves the physicalist free to resist his argument by refusing to accept its first premise.

Chalmers's one and only reason for claiming that it is 1-conceivable that P and not-Q seems to be that it is conceivable in the *ordinary* sense that P and not-Q. He infers the 1-conceivability of 'P and not-Q' from the fact that, as revealed by thought-experiment, we can entertain the thought that P and not-Q without explicit contradiction or any sense of conceptual blockage.⁹ But Chalmers is not just equivocating here on senses of 'conceivable'; he must be making a theoretical identification, that is, taking it that what is *in fact* going on in us when in the ordinary sense we conceive that P and not-Q is, often enough, that we are 1-conceiving that 'P and not-Q', in other words, that we are somehow mentally representing a possible world that we correctly describe as one at which 'P and not-Q' is true when evaluated according to 1-intensions.

But if conceiving in the ordinary sense does involve correctly describing a mentally represented possible world as one in which some statement is true when evaluated according to 1-intensions, then, when we are conceiving in the ordinary sense, we must evidently have at least potential access to the 1-intensions of the constituent concepts of the statement evaluated. And, according to Chalmers, so we do, because the 1-intensions of concepts we are competent with are allegedly accessible a priori to us. Perhaps experience was needed to acquire those concepts in the first place; but once they have been acquired, no further experience is needed for their 1-intensions to be accessible. So anyone who understands the claim that P and not-Q, and hence is competent with its constituent concepts, has potential access to their 1-intensions and is therefore potentially in a position correctly to describe a mentally represented possible world as one in which it is true, according to 1-intensions, that P and not-Q. Of course, as noted, Chalmers allows that potential access is not always turned into actual access; but in principle, anyway, 1-intensions are always accessible a priori to those competent with the corresponding concepts.

Would it suffice for Chalmers's argument if the 1-intensions of concepts we are competent with were indeed accessible to us, but only a posteriori? Chalmers does insist that the accessibility is a priori, but yes, so long as we do in fact have potential access of some sort to the 1-intensions of the relevant concepts, even if it is a posteriori, Chalmers may still identify conceiving in the ordinary sense that P and not-Q with 1-conceiving that P and not-Q. But taking ordinary conceivers' accessibility to 1-intensions to be a posteriori presents problems. First of

all, Chalmers would need to supply some empirical evidence for thinking that we do in fact enjoy the alleged a posteriori accessibility. For if the accessibility is not understood as a priori, a consequence of our competence with the relevant concepts, then its presence cannot simply be inferred from our conceptual competence. And this empirical evidence is bound to be contested by physicalists, who will contend that, assuming the soundness of the rest of Chalmers's argument, the evidence for the physical character of phenomenal states is automatically evidence against thinking that we actually enjoy the a posteriori accessibility in question. Secondly, if the alleged accessibility to 1-intensions is a posteriori, then physicalists will rightly demand to know the evidence for so taking the 1-intensions of the constituent concepts of 'P and not-Q' that its 1-possibility is a consequence. And, as before, physicalists will insist that the evidence for physicalism is automatically evidence against so taking the 1-intensions of the constituent concepts of 'P and not-Q'. So an appeal to a posteriori accessibility opens more than one can of very messy worms. It is a priori accessibility to 1-intensions on which, in fact, Chalmers bases his conceivability argument, and it is the adequacy of this basis that I shall discuss.

The dialectical situation, then, is this: unless the 1-intensions of concepts we are competent with are accessible to us a priori, Chalmers cannot construe conceiving in the ordinary sense that P and not-Q as 1-conceiving that P and not-Q; and unless he can construe conceiving that P and not-Q in the ordinary sense as 1-conceiving that P and not-Q, he has provided no reason to think that it really is 1-conceivable that P and not-Q, in other words, no reason to think that the first premise of his argument against physicalism is true.

IV

So the key assumption on which the first premise of Chalmers's conceivability argument rests is that the 1-intensions of the concepts with which ordinary conceivers are competent are accessible to them a priori. Chalmers's idea, however, is not that, in applying concepts to hypothetical worlds considered as actual, we ordinary conceivers draw upon some previously existing database of explicitly represented knowledge of 1-intensions, for example, discursive definitions, on the pattern of 'watery stuff' for 'water'. Rather, the a priori accessibility to us of 1-intensions is embodied in a certain capacity he says we possess: the capacity to consider a possible way the world might have turned out to be, and then to say what the concepts we are competent with would have referred to, if the world had turned out to be that way (see pp. 57-60).

Using this capacity, we can say, for example, that 'water' would have referred to XYZ had the actual world turned out to be one in which the watery stuff was XYZ. Our knowledge of a concept's 1-intension is therefore merely implicit, implicit in our capacity to consider various ways the actual world might have turned out to be, and to say what the concept would have referred to, had the world turned out in those ways; and this knowledge becomes explicit only to the extent that we exercise this capacity and take note of its deliverances (see p. 59, on Kripke's apparent methodology). So we can now say with more precision what exactly Chalmers takes to be going on when we 1-conceive (and hence when we conceive in the ordinary sense) that P and not-Q: we are somehow mentally representing a possible world, and, by exercising the 1-intension-respecting capacity just described, we are able to say that the possible world we have mentally represented is one at which, according to 1-intensions, 'P and not-Q' is true.

The crucial question, therefore, becomes whether we ordinary conceivers do in fact possess the 1-intension-respecting capacity. I cannot prove that we do not, but it is remarkably difficult to come up with a good reason for thinking that we do.¹⁰ Notice, first, that our possession of such a capacity does not follow simply from the definition of the 1-intension of a concept. For the 1-intension of a concept is officially defined as what *in fact* the concept would have referred to, had the world turned out to be different in various ways. It is not defined as what competent concept-users *say* the concept would have referred to, had the world turned out to be different in various ways. Whether we can say what our concepts would have referred to had the actual world turned out to be different in various ways is what is at issue. Secondly, our possession of the 1-intension-respecting capacity does not follow from the undisputed fact that, when presented with a description of how the world might have turned out, we make confident pronouncements about what various concepts would then have referred to. For our confident pronouncements might nevertheless be wrong. Many people pronounce confidently but wrongly when asked in what direction a ball would move, if it were gently released from the back of a vehicle moving rapidly forward.¹¹ Similarly, one could pronounce confidently but wrongly about what a concept would refer to, if the world were discovered to have been different in a certain way.

The obvious suggestion at this point is that possessing a 1-intension-respecting capacity with respect to a concept is an inevitable consequence of being competent with the concept. But there is at least one influential way of understanding concepts that implies that although concepts definitely have 1-intensions, one can be competent with them

without possessing a 1-intension-respecting capacity to consider possible ways the world might have turned out to be, and then say what those concepts would have referred to, had the actual world turned out in those ways.

On the (Fodorian) view I have in mind, then, a (primitive) concept, 'F', perhaps to be understood as a word in the language of thought, refers to Fs (rather than Gs, say) in virtue of the fact that it stands in a certain complex causal or nomological relation – call it 'R' – to Fs (though not to Gs).¹² But then, because it is contingent that 'F' stands in this reference-constituting relation to Fs, there will be possible worlds in which it stands in R to nothing, or to Gs, or to Hs, and so on. Accordingly, 'F' will certainly have a (contingent) 1-intension: had the actual world turned out to be one in which 'F' stood in R to Gs, it would have referred to Gs; had the actual world turned out to be one in which 'F' stood in R to Hs, it would have referred to Hs, and so on. However, in order to be competent with this concept, the only capacity one needs is the capacity to host a symbol which *in fact* stands in R to Fs. In order to exercise this capacity, in other words to think an 'F'-thought, one needs only to host a token of a type that *in fact* stands in R to Fs. One does not need any awareness at all that 'F' refers to Fs in virtue of standing in R to Fs, or that any particular token of 'F' stands in R to Fs. Neither does one need to host the token, even in part, *because* one in any way represents to oneself that it is of a type that stands in R to Fs.¹³ Consequently, one could be presented with a description of a possible way the actual world might have turned out, a description that could even state explicitly that 'F' stands in R to Gs, and yet be quite unable to say what 'F' would have referred to, had that description turned out to be true of the actual world (unable to say, that is, that it would in that case have referred in fact to Gs). And, quite generally, if a concept is similar to 'F' one could be competent with it, and it could have a 1-intension, but one could still be unable to say what it would have referred to, had the actual world turned out differently in various ways.

Moreover, phenomenal concepts – certain concepts, picking out phenomenal states, that we deploy when, in higher-order introspective thoughts about phenomenal states, we ascribe phenomenal states to ourselves; concepts, that is, of the same sort as the one that is a constituent of 'P and not-Q' – are as good candidates as any to be primitive Fodorian concepts, in other words, to refer to phenomenal states in virtue of standing in some complex causal or nomological relation to those states, and to be such that their exercise requires no response to a representation. And if they are, then one could be given a complete description of the physical and functional state one is in while one has a

red sensation, and yet simply lack any capacity to say what 'is having a red sensation' would apply to were the situation described to turn out to be actual, even though, were you actually in that physical and functional state, you would apply to yourself the concept, 'is having a red sensation.' So if phenomenal concepts are primitive Fodorian concepts, competence with them does not guarantee that one possesses the relevant 1-intension-respecting capacity. But Chalmers offers nothing to show that they are not Fodorian. So, for all he says, we might not possess the 1-intension-respecting capacity he must suppose we possess in order to be able to construe the ordinary conceivability of 'P and not-Q' as evidence for its 1-conceivability.

But the difficulty in finding reason to think that we possess the relevant 1-intension-respecting capacity does not arise solely from the neglected possibility that phenomenal concepts might be primitive Fodorian concepts. Consider the following, more promising-looking view of concepts, one that forges a very close link between competence with a concept and its 1-intension: to be competent with a concept is to possess a certain *disposition to apply it* under various circumstances, actual and counterfactual; and the 1-intension of a concept – what it would *in fact* have referred to, had the world turned out to be different in various ways – is to be identified with what people possessing the competence-constituting disposition *would in fact have applied the concept to*, had the world turned out to be different in various ways.

But this alternative, dispositionalist view of concepts, even if true,¹⁴ does not help. Even if, to be competent with a given concept, I must have a disposition to apply it in certain ways, a disposition constitutive of its 1-intension, it just does not follow that I possess a 1-intension-respecting capacity with regard to it, in other words, that I can consider possible ways the world might have turned out to be and then say what the concept would have referred to had the world turned out in those ways. For I might have a disposition to apply a concept in a certain way, depending on circumstances, and yet simply not know how I would have applied it, had the world turned out in various ways. How I would in fact have applied the concept and how I say I would in fact have applied the concept clearly need not be the same. Nor is there any guarantee, or even presumption, that, in general, if I have a disposition, I know enough about it to enable me to say what I would have done had circumstances turned out to be different. Suppose I have a food allergy, that is, a disposition to react in various unpleasant ways to various foodstuffs, then (without consulting an allergist) I know only as much about it as I can infer from self-observation of actual reactions, and I might be quite wrong about whether I would have broken out in hives if the mystery dish I ate last night had

turned out to be sea slugs. Similarly, even on the alternative view of concepts, I might be quite wrong about what I would have called 'water' had the world turned out to be one in which XYZ is the colorless, odorless liquid that falls from the sky as rain, and so forth. One simply need have no special epistemic access to one's dispositions.

An advocate of the dispositionalist view of concepts might reply, however, that one's 1-intension-respecting capacity results from a sort of simulation, in particular, from allowing the dispositions that constitute one's conceptual competence to be triggered by *counterfactual* situations, 'off-line,' rather than by the *actual* situations that trigger them when, 'on-line,' one applies concepts in real life. So in the special case of dispositions that constitute conceptual competence, we have more to go on than inferences from self-observation of actual reactions. But if one could be enabled in this way to say what a given concept would have referred to had some counterfactual situation between actual, then one could not be doing it by literally allowing one's competence-constituting dispositions to be triggered by counterfactual situations, because counterfactual situations, being merely counterfactual, cannot literally trigger anything. One would have to be doing it by allowing one's dispositions to be triggered by *mental representations* of counterfactual situations. In that case, however, competence-constituting dispositions must have mental representations as their triggers. (Otherwise, why should exercising the disposition off-line in response to mere *representations* of counterfactual situations that you are not actually in cause any manifestation of the disposition at all?) So if one is to argue that we possess a 1-intension-respecting capacity in regard to phenomenal concepts by arguing from a dispositionalist account of conceptual competence, then one must argue that competence with phenomenal concepts is not merely a disposition to apply them, but a disposition to apply them that is triggered by mental representations. But nowhere does Chalmers attempt to do so. Nor, on the face of it, do phenomenal concepts seem very strong candidates to be concepts applied in reaction to mental representations. All the introspective considerations that have traditionally made knowledge of one's own phenomenal states seem peculiarly direct and unmediated seem to count against the idea. So, for all Chalmers says, competence with phenomenal concepts might be the possession of a disposition that is nonrepresentationally triggered, in which case, once again, competence with phenomenal concepts will not guarantee possession of the relevant 1-intension-respecting capacity.

Chalmers provides no reason at all to think that we actually possess the required 1-intension-respecting capacity, and it is not easy to think one up. But if we lack the capacity, as I suggest we might, then how can

we explain the fact that, when presented with a hypothetical situation nonphenomenally described, we make confident pronouncements about what *nonphenomenal* concepts would have referred to, had that situation turned out to be actual; whereas when presented with a hypothetical situation nonphenomenally described, we feel no inclination to make any pronouncements at all about what *phenomenal* concepts would have referred to, had that situation turned out to be actual? Why do we respond in these different ways to representations of hypothetical situations? An obvious possibility is that we are exercising off-line an *epistemological* capacity, a capacity to judge that a concept applies to the world – that the property it refers to is actually instantiated. However, whereas our application to the world of nonphenomenal concepts is always based on representations (possibly sensory) of the world, representations that constitute our evidence, our application to ourselves of phenomenal concepts is, as I just mentioned, notoriously direct and not based on any evidencing representations at all. We do not judge that a sensation feels *like this* on the basis of any evidence that we represent to ourselves, and we are at a loss if asked *how* we know we are in pain.¹⁵ Because our epistemological capacity to apply phenomenal concepts, when exercised in regard to the actual world, does not take representations as input,¹⁶ we should not expect it to produce any output when exercised in regard to representations of merely counterfactual situations. But if we nevertheless try so to exercise it, we will note its lack of output and find ourselves, by a natural movement of thought, with the sense that fixing the physical and functional facts leaves it wide open as to what the phenomenal facts are, at least when these are thought of via the deployment of phenomenal concepts.¹⁷

Similarly, if we do lack the I-intension-respecting capacity with regard to certain concepts, then how, if at all, can we determine what the I-intensions of those concepts are? By the usual methods of science, a posteriori. Consider for a moment the representational systems of nonhuman animals. The dances of honeybees, we think, represent the locations of pollen sources. Because these dances possess a sort of rudimentary compositional structure, the elements of these dances can be seen as strongly analogous to concepts (or terms). Now these elements presumably get to refer to features of the locations of pollen sources in virtue of meeting some naturalistically specifiable condition, a condition that would determine a I-intension. But because we are not honeybees, there is no question of our knowing a priori what that condition is; yet we can hope nevertheless to find out what it is – by a posteriori means. And what we can do for honeybees we can presumably do for ourselves.¹⁸

My conclusion, then, is that, for all that Chalmers has shown, it may be conceivable in the ordinary sense that P and not-Q, even though it is not I-conceivable that P and not-Q, the explanation for this being simply that conceiving in the ordinary sense lacks even the potential access to relevant I-intensions (via a I-intension-respecting capacity) that would be required for it to constitute I-conceiving and hence provide information about I-possibility. But in that case no reason to believe the first premise of Chalmers's conceivability argument – the I-conceivability of 'P and not-Q' – is provided by the (undisputed) conceivability in the ordinary sense of 'P and not-Q'. And because the conceivability in the ordinary sense of 'P and not-Q' is the only reason offered for believing this first premise, we are left with no reason at all to believe it. Chalmers's conceivability argument, as it currently stands, is therefore inconclusive.

Traditional conceivability arguments fail, it will be recalled, because we do not know a priori, simply on the basis of our competence with a concept, what that concept refers to. Now Chalmers explicitly allows that we do not know this a priori, but he still claims that we have a sort of implicit a priori knowledge – embodied in a I-intension-respecting capacity – of what the concepts we are competent with would have referred to had the actual world turned out to be different in various ways, of "how reference depends on the way the external world turns out" (p. 57). But Chalmers has not shown, and there is no obvious reason to think, that even this much of the semantics of our concepts is a priori. Nor, in consequence, has he shown that there is any sort of modal truth that is accessible a priori, even though, as he insists and I agree, all (nonnomological) modal truths are, in a broad but still worthwhile sense, conceptual.¹⁹

V

Any adequately presented conceivability argument must address these questions:

- (1) How does physicalism's commitment to some allegedly necessary truth arise?
- (2) How are we to understand the conceivability of something?
- (3) Why does the conceivability of something, understood in line with the answer to (2), entail, or even defeasibly support, its possibility?

Traditional conceivability arguments are clear on (1), clear enough, at least by implication, on (2), and deplorably silent on the all-important (3). It is a signal merit of Chalmers's conceivability argument that it is

tolerably clear on (1), and quite explicit on (2) and (3). But answering these questions opens up the possibility that adequate answers to one question may preclude adequate answers to others; and this possibility, or so I have claimed, is actual in the cases I have discussed. Traditional conceivability arguments respond to (2) by construing conceivability in such a way that their conceivability premises are very plausible; but then their implicit answer to (3) must rely on the implausible psychosemantic assumption that one knows, simply on the basis of one's competence with a concept, what property the concept picks out. Chalmers's argument addresses (3) by understanding conceivability in such a way that the conceivability of something evidently entails its possibility; but then it must answer (2) in such a way that the resulting conceivability premise is in fact groundless unless one makes the unsupported and nonobvious psychosemantic assumption that the I-intensions of concepts they are competent with are accessible a priori to ordinary, flesh-and-blood conceivers. A good question to ask about any conceivability argument is whether it can simultaneously return adequate answers to (2) and (3).

The psychosemantic assumptions on which traditional conceivability arguments and Chalmers's nontraditional argument rest certainly appear to be a posteriori. On the face of it, at least, it is a matter for empirical discovery what human conceptual competence is and hence what it requires. If it turns out that all conceivability arguments rely on a posteriori assumptions, then conceivability arguments can *at best* be as strong as the evidence supporting those assumptions; but then into the other pan of the scales must be thrown such evidence as supports physicalism, and even good conceivability arguments (if such exist) will lose the apodictic character I suspect they have often been thought to possess. Moreover, these assumptions are certainly no less controversial than physicalism. If that too is typical of the assumptions on which conceivability arguments rest, then conceivability arguments offer little promise of assisting any movement to rational consensus on the issue of physicalism, and those interested in physicalism would better spend their time discussing the bearing of actual scientific findings upon the prospects for physicalism. However, it has been altogether too common for the presenters of conceivability arguments – and not only their advocates – to proceed as if such arguments, especially when they move from conceivability to possibility, do not depend on any nontrivial assumptions at all. I hope that Chalmers's example encourages the view that launching a conceivability argument with any hope of success requires embedding it within ambitious positions in metaphysics and the philosophy of mind and language.²⁰

NOTES

1. Sophisticated thinkers obviously do know simply on the basis of their competence with the relevant concepts that 'salt' refers to salt, that 'NaCl' refers to NaCl, and so on; but these metalinguistic platitudes evidently provide no hint that 'salt' and 'NaCl' pick out the very same stuff.
2. Might conceivability nevertheless provide *prima facie* support for possibility? Only if, and to the extent that, it is reasonable to make the empirical presumption about our minds that distinct concepts generally pick out distinct properties; I cannot recall ever having seen any evidence offered for such a presumption.
3. Adopting a descriptivist theory of predicative concepts, incidentally, does not make knowledge of the coreferentiality of two concepts a priori in every case. Suppose competence with a concept requires associating it with some description uniquely satisfied by the referent. Then if two concepts have different, but consistent, associated definite descriptions, it is not a priori whether or not the two concepts corefer. For the very same referent might satisfy both unique descriptions. This point applies with especial force if a descriptivist, such as Frank Jackson in his "Reference And Description Revisited" (forthcoming), allows such descriptions as 'is the appropriate (actual) causal origin of uses of term T'. Who knows – a priori – what other terms' uses the object that meets this description might also be the appropriate (actual) causal origin of?
4. Chalmers (1996). All subsequent page references in the text are to this book.
5. I might add that I consider the representationalist theory of phenomenal states to provide an entirely satisfactory physicalist treatment of them; see, for instance, Dretske (1995) and Lycan (1996). I understand physicalism along the lines sketched in Melnyk (1994, pp. 221–41).
6. In fact, (3), which seems equivalent to denying that physical terms are rigid designators, strikes me as false: surely the I-intension of (say) 'electron' is something similar to 'particle causally responsible for that sort of track on the cloud chamber's screen,' whereas its 2-intension is as specified by physical theory, something similar to 'negatively-charged particle with so-and-so mass, spin, and so forth.' But Chalmers can easily reformulate his argument to avoid this objection. Given that we have in fact determined the essences of physical properties, which a physicalist opponent should probably concede, we can in fact conceive in accordance with the 2-intensions of physical terms. So Chalmers can replace (1) with the claim that 'P and not-Q' is conceivable when 'P' is evaluated in light of 2-intensions and 'not-Q' is evaluated in light of 1-intensions. This claim makes a large a posteriori assumption about the world, but it is not one that his physicalist opponent need treat as question begging. Because it immediately entails (4), premises (2) and (3) are not required in this reformulation.
7. Or, better, of part of a world, else no work is left for the second component in conceiving a statement, namely, correctly describing the conceived world in light of the relevant intensions.
8. It might be objected to this account of 1-conceivability that it now begs the question to claim that 1-conceivability entails 1-possibility; for this entailment will hold only if the originally conceived world is 1-possible, in other words, only if the representation of a possible world hosted by a conceiver

- before* she starts evaluating statements at that world in light of 1-intensions describes a genuinely 1-possible world. But Chalmers can avoid this difficulty in practice by letting the hosted representations be such as are thought to describe the actual world, so that their expressing a genuine 1-possibility will not be contested by any of his opponents.
9. Actually, Chalmers's appeal to conceivability in the ordinary sense (pp. 94–9) is not his only support. He also invokes other considerations (pp. 99–106). But they all depend on the basic consideration of conceivability in the ordinary sense. For example, Mary would know the phenomenal facts on the basis of just the physical facts if it were inconceivable that P and not-Q.
 10. In private correspondence, Chalmers concedes that in his book he does not say much about the warrant for assuming the a priori accessibility of 1-intensions, but he adds that "it seems to hold in just about every case we can think of." But this just assumes that we possess the 1-intension-respecting capacity; in other words, that, when we consider a possible way the world might have turned out to be, and then say what a concept would have referred to had the world turned out that way, we get it right, as defined by the concept's 1-intension. Why think a priori accessibility holds in even one case?
 11. It turns out that many people say that the ball would fall vertically, or even backward; in fact it would briefly travel forward. See McCloskey (1983, pp. 122–30).
 12. See Fodor (1990).
 13. Concepts of this sort are concepts that have what Joe Levine calls "non-ascriptive modes of presentation"; see his "Conceivability and the Metaphysics of Mind" (1998).
 14. On the face of it, the view seems to confuse semantics with epistemology, what it is to think an F-thought with what it is to judge that something is F, the latter reflecting not just conceptual competence but also substantive beliefs about the world.
 15. Phenomenal concepts may not be quite unique in this respect. Perhaps there is a singular term, in "mentalese," for referring to oneself, whose application to the world (i.e., to oneself) is also not based, epistemologically, on representations. If so, then its existence would predict that fixing all the physical facts would seem to leave undetermined whether some object is oneself.
 16. Except, of course, for those special representations that just are phenomenal states, on the representationalist view I endorse.
 17. In private correspondence, Chalmers asks how I can explain, on this account, our apparent ability to conceive *also* of a world physically and phenomenally just like ours. The answer is that representations of phenomenal states can be included by stipulation in the set of mental representations we host (in our conceiving boxes) when we conceive a possible situation. That is, they can be part of the mentally represented situation we allow our epistemological capacity to take as *input*.
 18. Jerry Fodor, in particular, has told me that he regards as a posteriori the task of discovering the reference-constituting relation for concepts. However, such views as Fodor's are inspired by those of Kripke, and Chalmers points out (p. 59) that Kripke, in arguing for his causal theory of reference, routinely appeals to what we say a term would refer to were a certain counterfactual situation actual, apparently assuming thereby the a priori accessi-

- bility of 1-intensions to competent concept-users. But if he does, then we have to say that Kripke's methodology is just wrong. In fact, it is not obvious why Kripke cannot be understood as construing the deliverances of intuition about counterfactual situations as feasible empirical evidence (which would, of course, open the cans of worms mentioned earlier in connection with the possibility of Chalmers's allowing 1-intensions to be accessible a posteriori only).
19. The only necessities and possibilities I am committed to are those that are logical or conceptual in Chalmers's sense – necessities and possibilities induced by 1- and 2-intensions. So I am not committed to what Chalmers calls "strong" metaphysical necessity. I am simply suggesting that, because not all 1-intensions are accessible a priori, certain 1-necessary claims, though conceptual, are not a priori. Latham (forthcoming), assumes that any a posteriori necessity would have to be strong, in this same sense, and hence also overlooks the possibility of a necessity that is conceptual but not a priori. Thanks to Terry Horgan, Michael Levin, Peter Markie, Jamie Phillips, Paul Weirich, and especially David Chalmers for comments on earlier drafts.
 - 20.