

CONCEIVING OF PAIN

I. INTRODUCTION

Kripke's modal argument against the mind-body identity thesis is by now familiar (Kripke 1980):

1. It is possible for there to be pain without C-fiber stimulation, and vice versa.¹
2. If pain = C-fiber stimulation, then it is not possible for there to be pain without C-fiber stimulation or vice versa.
- ∴ 3. Pain ≠ C-fiber stimulation.

(2) is supposed to follow from the fact that 'pain' and 'C-fiber stimulation' are rigid designators. If 'pain' picks out the same phenomenon in every possible world, and likewise for 'C-fiber stimulation', and these phenomena are identical, then there is no possible world in which pain occurs without C-fiber stimulation. (1) is supported by the apparent fact that we can conceive of pain without C-fiber stimulation and C-fiber stimulation without pain. Assuming that conceivability implies metaphysical possibility, this shows that pain without C-fiber stimulation and vice versa are metaphysically possible.

Broadly speaking, there are two ways that philosophers have responded to this argument:

- I. Keep conceivability as a reliable guide to possibility and attempt to provide a real possibility to explain away the apparent conceivability of the mental without the physical.

¹ 'C-fiber stimulation' is just a placeholder for whatever physical type corresponds to pain.

II. Deny that in the mental/physical case conceivability provides a reliable guide to possibility.

The first approach treats mental/physical identities on par with other scientific identities, on the model of water/H₂O or heat/mean kinetic energy. Standard attempts of this approach, however, have run into familiar problems. These problems throw into relief the real difficulty for the identity theorist who chooses strategy I. It is not simply that we can conceive of pain without C-fiber stimulation, but this fact coupled with the fact that in conceiving of pain, we do so in terms of one of its essential modes of presentation. In the light of this problem most of the action in this debate has shifted to the prospects for strategy II.²

But strategy I should not be written off just yet. In this paper we aim to see how far one can get in defending the identity thesis without challenging the inference from conceivability to possibility. Challenging that inference raises a host of complex issues about the metaphysics and epistemology of modality, issues that an identity theorist may well wish to avoid.

Our defense of the identity thesis consists of a dilemma for the modal argument, one horn of which involves a strategy I approach. The dilemma turns on whether the term 'pain' is rigid. If 'pain' is not rigid, then premise (2) of the modal argument should be rejected. On the other hand, if 'pain' is rigid then one must provide a semantic account of this term that accommodates its rigidity. The most plausible such account treats 'pain' as a natural kind term that refers to its causal or historical origin, namely C-fiber stimulation. If this is the correct semantics for 'pain', then pain itself is a natural kind, which, like any natural kind, has its microstructure essentially. It follows that any

² See, e.g. Loar (1990), Hill (1997) and the essays in Gendler and Hawthorne (2002).

phenomenon that is not C-fiber stimulation is not pain, even if it is qualitatively similar to pain. This means there could be phenomena that feel like pain but are not pain since they are not C-fiber stimulation. These possible phenomena can be used to explain away the apparent conceivability of pain without C-fiber stimulation, thus undermining the support for premise (1) of the modal argument.

We should emphasize that this is not a straightforward strategy I style response to the modal argument. Rather, it is a dilemma for the modal argument, one horn of which follows strategy I. This horn of the dilemma involves the claim that although feeling like pain is necessary for being pain, it is not sufficient. In order for something to be pain it must also be C-fiber stimulation. We realize that many philosophers will find this counterintuitive. Many find it compelling to think that if something feels like pain then it must be pain. But this is not a problem for an identity theorist who wishes to pursue our defense against the modal argument. We are not going to flatly argue that the feeling of pain is not sufficient for pain. Rather, we are going to argue that the modal argument faces a dilemma, one horn of which involves holding that the feeling like pain is not sufficient for being pain. Identity theorists who find this counterintuitive may simply grasp the other horn of the dilemma.

Our plan for this paper is as follows. First we will clarify what we will mean by 'conceivable'. This will lay some necessary groundwork for the strategy I response we'll develop later in the paper. Then we will present a version of strategy I given by John Perry in his recent book *Knowledge, Possibility and Consciousness* (Perry 2001a). Although we do not think Perry's response succeeds, we think it is interesting and it provides a useful foil for the presentation of our own view. After raising some problems

for Perry's response, we will develop the dilemma for the modal argument sketched above.

One additional preliminary matter must be addressed. In defending the identity thesis against Kripke's modal argument we are going to follow Perry in adopting a dialectical position he calls "antecedent physicalism". An antecedent physicalist is "someone who is committed to physicalism in the sense that she or he sees some compelling reasons for it and will not give it up without seeing some clear reason to do so," (Perry 2001a, 27). The antecedent physicalist's goal is *not* to convince a committed dualist of the truth of the identity thesis. To do that, one would have to rely solely on premises and principles that a dualist would accept. The antecedent physicalist's goal is rather to defend the identity thesis against attacks like Kripke's modal argument. In accomplishing this goal the antecedent physicalist is licensed in making use all the resources of physicalism.³ A dualist may not accept some of these resources, but that does not mean that they are unavailable in a defense of the identity thesis against challenges. In our defense of the identity thesis we will occasionally appeal to assumptions that a dualist probably would reject. This would be question begging if we were trying to convince the dualist of the truth of physicalism and the identity thesis. But that is not our aim in this paper.

³ We think Perry gets this exactly right. "In the dialectic of philosophy there are many different situations in which one can find oneself. Sometimes one is trying to persuade someone with quite different views and presuppositions that a certain thing is so. In such a case, one must set aside much of what one believes, all the controversial assumptions, and search for a common beginning point. The situation is quite different when one defends one's own view from a charge of inconsistency, incoherence, or inadequacy. In this case, it makes no sense to jettison your own view, the view whose consistency or adequacy you are defending. You want to rely on the distinctions and concepts that your view provides to counter the criticism, and you have every right to do so," (2001a, 27-28).

II. WHAT IS CONCEIVABILITY?

This is a vexed question. Stephen Yablo has distinguished six different notions of conceivability (Yablo 1993), David Chalmers has identified eight (Chalmers 2002), and Tamar Szabó Gendler and John Hawthorne have listed no fewer than fourteen (Gendler and Hawthorne 2002, 7-8). This proliferation of conceivability notions would be tolerable if there were a consensus about which notion or notions are relevant to the debates about the identity theory. This is of course not the case. Yablo, for example, holds that the relevant notion is *psychological*. For Yablo, conceiving that *p* in the relevant sense is “imagining that *p* by imagining a world of which *p* is held to be a true description,” (Yablo 1993, 29). On the other hand, Katalin Balog has urged that “the claim that zombies are *conceivable* does not have to do with our powers of imagination, or our *psychological* constitution in general, but rather with the nature of physical and phenomenal *concepts*,” (Balog 2000, 10). Balog’s preferred notion of conceivability is *semantic*: *p* is conceivable if not-*p* is not a conceptual (or analytic) truth. A third option is an *epistemic* notion of conceivability: *p* is conceivable if not-*p* is not *a priori* (see Chalmers 2002).

On either the semantic or epistemic construals of conceivability strategy I is a non-starter. A strategy I response denies that pain without C-fiber stimulation and vice versa are conceivable. If conceivability is semantic conceivability, this denial is the claim that it is a conceptual truth that pain and C-fiber stimulation co-occur. If conceivability is epistemic conceivability, the denial is the claim that we know *a priori* that pain and C-fiber stimulation co-occur. Neither of these claims is remotely plausible.

The only option for the strategy I theorist is to adopt a version of the psychological account of conceivability, according to which conceiving of p is imagining a scenario in which p is true. Our interest in this paper is to see whether there is a workable response to the modal argument that incorporates strategy I. Since the only viable form of conceivability for strategy I is psychological conceivability, we're going to employ this notion of conceivability in this paper.

There are two further reasons for understanding conceivability in this way, at least in the context of discussions of the identity theory and Kripke's modal argument. First, psychological conceivability must have been what Kripke had in mind when he gave the modal argument. Recall that Kripke *denies* that it is conceivable that heat is not molecular motion or that water is not H_2O .⁴ If Kripke had in mind semantic or epistemic conceivability, this denial would have been patently absurd. Second, the best case for the modal argument involves psychological conceivability. This is because the modal argument depends on a conceivability-possibility inference, and there are straightforward counterexamples to this inference on either the semantic or epistemic construals of conceivability. For example, it is semantically and epistemically conceivable but *not* metaphysically possible that water is not H_2O . On the other hand, as Kripke has argued, it is not psychological conceivable that water is not H_2O . We can imagine possible

⁴ "Can we then imagine a possible world in which heat is not molecular motion? ... It seems to me that any case which someone will think of, which he thinks at first is a case in which heat – contrary to what is actually the case – would have been something other than molecular motion, would actually be a case in which some creatures with different nerve endings from ours inhabit this planet (maybe even we, if it's a contingent fact about us that we have this particular neural structure), and in which these creatures were sensitive to that something else, say light, in such a way that they felt the same thing that we feel when we feel heat. But this is not a situation in which, say, light would have been heat, or even in which a stream of photons would have been heat ..." (Kripke 1980, 131-2).

worlds in which there are liquids with all the same surface features of water but which are not H₂O, but this is not to imagine a world in which water is not H₂O. The familiar counterexamples for the conceivability-possibility inference do not seem to arise on the psychological construal of conceivability. It seems, then, that fairness to the modal argument obliges us to adopt the psychological construal of conceivability.

To sum up, the only viable construal of conceivability for the strategy I defense of the identity theory is psychological conceivability. This is okay, though, since psychological conceivability must have been what Kripke had in mind, and this account of conceivability is the best case for the modal argument. Hence, for the purposes of this paper we will adopt the psychological construal of conceivability.

III. PERRY'S RESPONSE TO THE MODAL ARGUMENT

As we noted above, one strategy for rejecting Kripke's modal argument is to accept that pain without C-fiber stimulation and C-fiber stimulation without pain are conceivable but deny that conceivability is a reliable guide to genuine metaphysical possibility. This is not Perry's strategy. Perry defends the identity theory by denying the conceivability claim needed for premise (1) of the modal argument. Of course, a flat denial of conceivability is not very compelling in the absence of some plausible explanation for the appearance of conceivability. What is needed are some alternative possibilities that are plausibly mistaken for cases of pain without C-fiber stimulation and C-fiber stimulation without pain.

Perry's account of these alternative possibilities is couched in terms of his distinction between subject matter content and reflexive content. Perry introduces and explains this distinction using indexicals (Perry 2001b, ch.1). Compare:

(1) I am a computer scientist. (said by David Israel)

(2) You are a computer scientist. (said to David Israel)

(3) David Israel is a computer scientist.

According to Perry, (1)-(3) have the same subject matter content, the singular, Russellian proposition consisting of David Israel and the property of being a computer scientist. But Perry thinks that if we limit ourselves to subject matter content we cannot account for the difference in cognitive significance between utterances of these sentences. For example, we can learn something from an utterance of (1), namely that the speaker of that very utterance is a computer scientist, which we may not be able to learn from an utterance of (3). Similarly, upon hearing an utterance of (2) we gain information, that the addressee of that utterance is a computer scientist, which we may not gain from an utterance of (3). Perry takes this to show that there is another level of content, reflexive content, which accounts for these differences in cognitive significance. Where ι is the token of 'I' in (1), and τ the token of 'you' in (2), the reflexive contents of (1)-(3) are, respectively:

1_R. That the speaker of ι is a computer scientist.

2_R. That the addressee of τ is a computer scientist.

3_R. That the person named 'David Israel' to whom the use of this name in (3) refers is a computer scientist.

(1_R)-(3_R) are descriptive propositions; each contains a descriptive condition that determines David Israel. (1_R) contains the property of being the speaker of ι , (2_R) the property of being the addressee of τ , and (3_R) the property of being the person named

‘David Israel’ to whom the use of this name in (3) refers. These descriptive propositions are distinct from the singular proposition containing David Israel and the property of being a computer scientist.

The distinction between subject matter content and reflexive content corresponds to a distinction between possibilities. Compare:

4. Mark Twain is not Mark Twain.

5. Mark Twain is not Sam Clemens.

The subject matter contents of (4) and (5) are the same, necessarily false singular proposition. This subject matter content is correlated with a metaphysically impossible situation in which Mark Twain/Sam Clemens is not identical with himself. The reflexive contents of these sentences are different, and these reflexive contents are correlated with genuine metaphysically possibilities. The reflexive contents of (4) and (5) are, respectively (and roughly):

4_R. That the man named ‘Mark Twain’ to whom the first use of this name in (5) refers is not the same person as the man named ‘Mark Twain’ to whom the second use of this name in (5) refers.⁵

5_R. That the man named ‘Mark Twain’ to whom the use of this name in (5) refers is not the same person as the man named ‘Sam Clemens’ to whom the use of this name in (5) refers.

These reflexive contents are merely contingently false. Focusing on (5_R), there are possible worlds in which these uses of the names ‘Mark Twain’ and ‘Sam Clemens’ refer to different individuals. The proposition (5_R) is true with respect to these worlds. Perry calls the metaphysical possibility corresponding to these worlds the *reflexive* possibility

⁵ Cf. Perry’s solution to Kripke’s Paderewski puzzle, (Perry 2001b, 113-6).

corresponding to (5). While there is no subject matter possibility corresponding to (5), there is a genuine metaphysical possibility at the reflexive level.

On Perry's account, reflexive possibilities are the alternative possibilities that explain the apparent conceivability of the mental without the physical. Perry considers a subject, Elwood Fritchey, who believes that pain is not C-fiber stimulation (Perry 2001a, 181-2). Elwood's belief is based on the fact that he can imagine creatures who experience pain yet who lack C-fibers. According to Perry, what Elwood is really imagining is a situation in which the *concept* of pain refers to something other than what the concept of C-fiber stimulation refers to. That is, Elwood is imagining a world in which the concepts of pain and C-fiber stimulation do not co-refer. This possibility is the reflexive possibility corresponding to his belief that pain is not C-fiber stimulation. On the other hand, the subject matter possibility corresponding to his belief is the impossible situation in which pain/C-fiber stimulation is not identical with itself.

This involves an extension of Perry's subject matter content/reflexive content distinction to beliefs. The subject matter content of Elwood's belief that pain is not C-fiber stimulation is the necessarily false singular proposition that pain/C-fiber stimulation is not identical with itself. The reflexive content of Elwood's belief is the contingently false descriptive proposition that the concepts of pain and C-fiber stimulation do not co-refer. (In the case of beliefs, reflexive contents contain conditions about the references of concepts, whereas in the case of sentences reflexive contents contain conditions about the references of word tokens.) Elwood mistakenly believes that pain is not C-fiber stimulation because he can imagine possible worlds in which the reflexive content of his belief is true. In other words, he can conceive of the reflexive possibility corresponding

to this belief. What he cannot conceive is the metaphysically impossible subject matter possibility of this belief. It is easy for Elwood to mistake these two possibilities because they correspond to the two contents of his belief. According to Perry, Elwood's belief that pain is not C-fiber stimulation is necessarily false, since its subject matter content is necessarily false. Elwood mistakenly holds this belief because he can imagine situations in which its reflexive content is true and he easily mistakes reflexive content for subject matter content.

IV. CRITICISMS OF PERRY

1. Why reflexive content?

We think Perry's response to the modal argument is ingenious but carries some unnecessary baggage and is, ultimately, unpersuasive. The first thing to notice is that his distinction between subject matter content and reflexive content does no real work. The real work is being done by the possibility that the concepts of pain and C-fiber stimulation do not co-refer.

This is easy enough to see if we consider a parallel to the pain/C-fiber case. Imagine that you are teaching Hilary Putnam's "The Meaning of 'Meaning'". You've carefully presented his view, showing how the extension of 'water' depends on what liquid is in fact in the environment. You culminate by drawing the surprising conclusion that necessarily, water is H₂O. One of your vocal students immediately shoots up his hand and confidently declares, "That can't be true because I can imagine a world in which water is not H₂O". You ask him to explain. "Imagine," he says, "a world in which the inhabitants use the word 'water' to refer to gin. In that world, water is not H₂O." Has

the student produced a possible scenario in which water is not H₂O? No, he has not. He has produced a scenario in which the word 'water' is used to pick out something other than water, in which the word 'water' does not mean water. And sure enough, there is such a possible situation. The student has latched onto a real possibility, but has misdescribed that real possibility as one in which water is not H₂O.

Now, in making this response to the student, it is irrelevant whether we take this real possibility to correspond to part of the *content* of the student's claim that water is not H₂O. There are two choices here. One might say, with Perry, that the sentence 'Water is not H₂O' has two contents, one of which, the reflexive content, is merely contingently false. Or one might say that 'Water is not H₂O' has a single, necessarily false content, but this content is easily mistaken for the contingently false content of *another* sentence, 'Water does not refer to H₂O'. In terms of the student's belief that water is not H₂O, one might say with Perry that this belief has two contents, which the student is not careful to keep distinct, or that the student has two different beliefs, which he is not careful to keep distinct. As far as we can tell, nothing in the response to the student depends on speaking one way or the other.

This shows, we submit, that the force of Perry's response to the modal argument does not require his doctrine of reflexive content. Perry frames his response in terms of a single belief with multiple contents. But one could just as easily run the same response in terms of multiple beliefs. If so, Elwood has different beliefs that he is not careful to distinguish: the belief that pain is not C-fiber stimulation, and the belief that the concepts of pain and C-fiber stimulation do not co-refer. Perry would prefer to say that Elwood has a single belief with two different contents, and his mistake arises from not keeping

these contents separate. We prefer to say that Elwood mistakenly holds the first belief because he confuses it with the second belief. But as far as the response to the modal argument is concerned, this choice is irrelevant.

2. Confusing possibilities

What then ought we to make of the response, independently of its appeal to reflexive content? The response requires two components. First, there is a possible scenario in which the concept of pain picks out something other than what the concept of C-fiber stimulation picks out. Second, assuming that it is possible that the concept of pain does not co-refer with the concept of C-fiber stimulation, it is plausible to say, as Perry does, that this real possibility is what explains the apparent conceivability of pain without C-fiber stimulation.

In the next section, we shall examine the possibility that the concept of pain might not co-refer with the concept of C-fiber stimulation. For the moment, let's grant that possibility. Even so, is that possibility the one we are actually thinking of when we think that pain might not be C-fiber stimulation? Consider Peter, a philosopher. He imagines experiencing an intense pain, while simultaneously imagining that he has no brain whatsoever, and so, no C-fiber firings. Being a careful philosopher, Peter is well aware of the difference between pain and his concept of pain, and between C-fiber stimulation and his concept of C-fiber stimulation. Being careful with these distinctions, Peter is imagining (as he would put it) his experiencing pain, a token sensation, in the absence of any C-fiber firings.

But Perry must say that Peter, despite his protestations otherwise, is really imagining a situation in which his concept of pain does not co-refer with his concept of C-fiber stimulation and is mistakenly describing this possibility as a case in which there is a pain without C-fiber stimulation. But this strikes us as implausible, given Peter's facility with the requisite distinctions. It is one thing to offer this diagnosis for someone who is not careful to distinguish between concepts and things falling under concepts. It is quite another to offer it as a diagnosis for someone like Peter who is careful with this distinction.⁶

This shows that the possibilities that Perry cites are not plausibly taken to explain away philosophers' intuitions that there can be pain without C-fiber stimulation. A similar point can be made for non-philosophers. Elwood might imagine a world populated solely by primitive creatures who lack concepts altogether. In imagining these creatures feeling pain in the absence of C-fiber stimulation Elwood is not imagining himself or anyone else using the concept of pain differently than the way it is actually used. In other words, he is not imagining a world in which the concept of pain refers to something other than its actual world. Elwood is imagining some non-intelligent creatures *feeling* a certain way. Concepts don't come into at all.

This highlights an important constraint on strategy I style responses to Kripke's modal argument. To explain away the appearance of conceivability, it is not enough to find some possibility in the vicinity of the mental without the physical. One must find a

⁶ Our response here is related to Block's (2006) criticism of Perry. Block contends that by using 'topic neutral' concepts as the modes of presentation of phenomenal concepts, Perry evades the central difficulty. This is close to our objection here that even granting that Perry's diagnosis works for some cases, it still looks as if one can focus on the qualitative aspect of pain in the absence of C-fiber stimulation.

possibility that is plausibly mistaken, even by careful philosophers, for a case of pain without C-fiber stimulation. Such a possibility must include the presence of a certain painful feeling in the absence of C-fiber stimulation. This makes the prospects for a successful strategy I response look rather dim. After all, if something feels like pain, then isn't it pain? Later on we will argue that, at least on one plausible account of the semantics of phenomenal terms, the answer is not necessarily.

3. Concepts: denotation vs. reference

The first and most important component of Perry's response to the modal argument is the idea that it is possible for the concepts of pain and C-fiber stimulation to have different extensions. Perry is sensitive to the fact that many philosophers, Kripke included, are skeptical about this. After all, these concepts seem to determine their extensions via essential properties.

The crux of Perry's defense of this possibility is a distinction between what a concept *fits* or *denotes* and what it *refers* to (Perry 2001a, 185). A concept fits or denotes the entities that meet the conditions contained in the concept. For example, the concept of pain includes what Perry calls a Humean idea of pain. This idea resembles (in some extended sense) the way pain actually feels. In this way the concept of pain sets a condition that is met by pain itself. Hence, this concept fits or denotes the phenomenon of pain.

On the other hand, on Perry's account, reference is a matter of the source and applicandum of a concept. The source of a concept is its causal origin; its applicandum is the thing to which the concept is applied. Suppose, bizarrely, that in some possible world

the causal origin of the concept of pain is something other than pain, e.g. itchy feelings. Or, somewhat less bizarrely, suppose that in some possible world people apply the concept of pain to something other than the actual phenomenon of pain, e.g. D-fiber stimulation. According to Perry, these are worlds in which the concept of pain *refers* to something other than what it refers to in the actual world. We can be sure that we are still talking about our concept of pain because the identity conditions for a concept are given by the internal features that determine what the concept denotes, not what it refers to. We can hold these internal features fixed and consider different worlds in which the source or applicandum of the concept of pain varies.

The distinction between denotation and reference will be familiar to philosophers of language. A definite description *denotes* whichever object meets the conditions expressed by the description. A proper name *refers* to the object that is its causal or historical source. For Perry, a concept is a kind of amalgamation of a definite description and a proper name. On his view, concepts are like descriptions in that they set conditions that determine denotations, but in addition, concepts are like proper names in that they refer. The key point is that the denotation and the reference of a concept can come apart. With a nod to Kripke, Perry accepts that the concepts of pain and C-fiber stimulation *denote* the same phenomenon in every world. But he thinks that the *references* of these concepts can vary from world to world.

Note that this is *not* the claim that these concepts are non-rigid. Holding fixed the actual facts about the concepts of pain and C-fiber stimulation, these concepts refer to the same phenomenon with respect to every possible world. The claim that the references of these concepts can vary is the claim that that the facts about the causal or historical

origins of these concepts could have been different. If so, then there are possible worlds in which the concept of pain, or the concept of C-fiber stimulation, has a different referent than its actual referent. These are the possibilities that Perry needs for his response to the modal argument. According to Perry, in imagining a case of pain without C-fiber stimulation we are in fact imagining a possible world in which the references of the concepts of pain and C-fiber stimulation differ from their actual references. It is thus crucial for Perry's argument that there be such possible worlds.

The problem for Perry is that, as we will see, *nothing* both denotes and refers. If this is right then Perry must make a choice: the concept of pain either denotes or it refers. If it denotes, then there are no possible worlds in which it refers to something other than pain, since there are no possible worlds in which it refers. A better option for Perry is to hold that the concept of pain refers. But if so, then his account of the appearance of conceivability is rendered otiose, since a much more plausible explanation presents itself. Before we develop this point, however, we need to argue that nothing both denotes and refers.

Notice first that this point clearly holds of linguistic expressions. It is easy to see why this should be so. Whether or not a designator refers is a matter of the semantic mechanisms it employs for determining its extension.⁷ Some expressions determine their

⁷ In terms of propositions, a designator refers just in case it contributes its referent to propositions. A designator denotes just in case it contributes uniquely identifying descriptive conditions to propositions. Scott Soames has recently argued that some linguistic expressions, e.g. 'Professor Saul Kripke', contribute both their referents and descriptive conditions to propositions, in this case, the individual Saul Kripke along with the property of being a professor (Soames 2002). But these partially descriptive names are not counterexamples to our claim that no linguistic expression both denotes and refers. The descriptive conditions associated with partially descriptive names do not uniquely identify their referents, and hence these expressions do not denote.

extensions by virtue of mediating conditions; others determine them directly via causal or historical chains and the like. The distinction between these two kinds of semantic mechanisms for extension determination is *not* the rigid/non-rigid distinction. Even if a designator is rigid, if it determines its extension via mediating conditions then it denotes and does not refer. For example, rigid definite descriptions, e.g. ‘the sum of 5 and 7’, ‘the actual inventor of bifocals’, designate whichever objects satisfy the conditions they express. Even though they are rigid, they determine their extensions denotationally. No linguistic expression employs both the semantic mechanisms of denotation and reference, and for good reason.⁸ The fact that denotation and reference can come apart threatens to undermine a designator’s essential function, which is to uniquely identify something in the world.

Suppose the name ‘Gödel’ both denotes the man who proved the incompleteness of arithmetic and refers to Gödel. Suppose further that, contrary to popular belief, a man named ‘Schmidt’ proved the incompleteness of arithmetic (Kripke 1972, 84). If so, then the denotation of ‘Gödel’ is Schmidt, whereas the reference of ‘Gödel’ is Gödel. This is intolerable. Who would we be talking about when we use the name ‘Gödel’? What is someone saying when she says ‘Gödel lived in Princeton’? Suppose Schmidt lived in Berkeley. Is the sentence ‘Gödel lived in Princeton’ true or false? Is this sentence ambiguous? The point is *not* that we must answer these questions. It is precisely the opposite. Since ‘Gödel’ does not both denote and refer these questions do not arise. And the fact that they do not arise is crucial for the successful functioning of the name

⁸ Note that two-dimensional semantic models are not counterexamples to this claim, for on such models, the semantic mechanism is denotational, not referential. Indeed, the way to think of 2-D is as an attempt by advocates of denotation to accommodate the Kripkean insights about semantics without embracing referentialism.

‘Gödel’. The successful functioning of the name requires that it employs a single semantic mechanism for determining its referent.

Concepts share with linguistic designators the essential function of uniquely identifying something in the world. Like linguistic designators, a concept’s ability to uniquely determine its extension would be undermined if it both denoted and referred. Suppose the concept of a bachelor denotes adult unmarried males, but in addition, because of its historical origin, refers to young knights apprenticed to more experienced knights. What then is in the extension of the concept of a bachelor, adult unmarried males or knights in training? Furthermore, what is someone thinking when she thinks that Jones is a bachelor? Is it the thought that Jones is an adult unmarried male or that Jones is a knight apprenticed to another knight? It does not make sense to suppose that the thought is somehow ambiguous. Sentences can be ambiguous, but not thoughts. This highlights the fact that the problems for combining denotation and reference in concepts are more severe than they are for linguistic expressions. Combining denotation and reference in concepts threatens not only the possibility of uniquely determining an extension, but also the essential role of concepts in determining the identities of thoughts.

In his attempt to combine denotation and reference in the semantics of concepts, Perry undermines two essential functions of concepts: determining unique extensions and determining the identities of thoughts. As such, his semantics for concepts is straightforwardly unacceptable. We take this criticism to be the most damaging to his position. But the earlier criticisms apply as well. Perry’s appeal to reflexive content remains ultimately idle in the response. In addition, it is doubtful that the possibility Perry adduces can explain away all the cases in which someone (apparently) conceives of pain

without c-fiber stimulation. Perry's own attempt to respond to the modal argument ultimately fails.

V. A DILEMMA FOR THE MODAL ARGUMENT

The failure of Perry's response to the modal argument is instructive for two reasons. First, it highlights the fact that in order to give a plausible strategy I response to the modal argument, i.e. to plausibly explain away the apparent conceivability of pain without C-fiber stimulation, one has to find a possibility in which the *feeling* of pain, but not pain, occurs in the absence of C-fiber stimulation.

Second, consideration of Perry's view forces one to make a decision about the semantic mechanisms employed by terms like 'pain' and 'C-fiber stimulation'. (From this point on we'll focus on terms instead of concepts.) Whether or not these terms are rigid is obviously relevant to this decision. This leads directly to a dilemma for the modal argument.

The dilemma turns on whether terms like 'pain' and 'C-fiber stimulation' are rigid. In its most concise form, the dilemma is simply this: if the terms are not rigid, then premise (2) of the modal argument fails; if they are rigid, one can explain away the conceivability facts that provide the support for premise (1). Either way, the modal argument fails.

The first horn is easy enough to see. If the terms flanking the identity sign fail to be rigid, then there is no reason to think the resulting identity statement is necessarily true. If the designation of 'pain' or 'C-fiber stimulation' varies from world to world, then

even if the two designate the same thing in the actual world, there is no reason to think that they will continue to do so in other worlds where their designations are different.

The second horn of the dilemma requires some argument. Given that the term 'pain' is a rigid designator, two issues must be addressed. First, what semantic mechanism does 'pain' employ to determine its designation? Second, what does 'pain' designate? These questions could be taken in either order, but we think it is simpler to start with the first.

Given that 'pain' is rigid, what semantic mechanism does 'pain' utilize to fix its extension? As we have seen from the discussion of Perry, there are two choices here: either 'pain' denotes whatever phenomenon satisfies the descriptive conditions associated with this term or 'pain' refers to the phenomenon at its causal or historical origin.

The view that 'pain' refers to its causal or historical origin easily secures the rigidity of 'pain'. On a referential account of 'pain', the extension of this term is *not* fixed by way of satisfaction of associated descriptive conditions. Rather, the designation of 'pain' with respect to any world w is the causal or historical source of the word 'pain' in the actual world. In this way, 'pain' would designate the same phenomenon in every world.

On the other hand, at least on the most straightforward denotational semantics for 'pain', this term comes out non-rigid. Following Perry, let's assume that the term 'pain' is conventionally associated with a Humean idea of pain, so that any competent user of the term 'pain' will associate this Humean idea with this term. On a denotational semantics for 'pain', 'pain' would denote whatever state qualitatively resembles the Humean idea of pain. Since different states resemble this idea in different worlds, 'pain'

would come out non-rigid. On the assumption that ‘pain’ is rigid, this semantic account of ‘pain’ can therefore be ruled out. The same applies to any denotational account of ‘pain’ on which the descriptive conditions associated with ‘pain’ allow its extension to vary from world to world.

This leaves denotational accounts of ‘pain’ that secure *de facto* rigidity for ‘pain’ by employing descriptive conditions that end up determining the same extension in every world. For example, perhaps ‘pain’ denotes whatever phenomenon *actually* resembles the Humean idea of pain. This would be the analog for ‘pain’ of the rigidified descriptions view of proper names, i.e. the view that names are synonymous with rigidified descriptions containing the word ‘actual’. However, the rigidified descriptions view of proper names faces some serious problems. One problem has to do with the distinction between obstinate and persistent rigidity.⁹ Obstinate rigid designators designate the same thing in every possible world, including worlds in which that thing does not exist. Persistently rigid designators designate the same thing in every world in which that thing exists and nothing in worlds in which it does not exist. As Nathan Salmon (Salmon 1981) and David Kaplan (Kaplan 1989) have argued, there are good reasons for thinking that proper names are obstinately rigid whereas rigidified descriptions are merely persistently rigid. If so, the reference of a proper name cannot be fixed by the satisfaction of a rigidified description. Although the details of this go beyond the scope of this paper, these arguments are easily adapted to the case of terms like ‘pain’. Furthermore, as Scott Soames has argued (Soames 2002), if proper names are synonymous with rigidified descriptions then we cannot use names to make true

⁹ See (Salmon 1981) for the notions of obstinate and persistent rigidity.

propositional attitude ascriptions to occupants of counterfactual worlds. Once again, it is straightforward to apply Soames's argument to the case of the term 'pain', although this also takes us outside the bounds of the present paper. For these reasons, we are going to proceed on the assumption that a rigidified description account of 'pain' is unsatisfactory.¹⁰

Another possibility is to hold that the term 'pain' is synonymous with a description that is required to take wide-scope over modal operators. This is a way of simulating rigidity while maintaining a denotational semantics for 'pain'. However, as Soames has demonstrated (Soames 2002), this approach is even worse off than the rigidified descriptions view. The wide-scope approach makes wrong predictions about the validity of simple inferences involving sentences with modal operators, and it cannot deal with sentences containing both modal operators and propositional attitude verbs. It would take us too far afield to present these arguments in detail. Suffice it to say that we think they are sufficient for closing this off as a possible semantic account of the term 'pain'. We conclude, therefore, that the simplest and least problematic way of accommodating the rigidity of 'pain' is to hold that 'pain' refers to its causal or historical origin.

¹⁰ In any case, the present horn of the dilemma still succeeds even if 'pain' is synonymous with a rigidified description. If 'pain' rigidly denotes C-fiber stimulation then being C-fiber stimulation is an essential feature of pain, which has the consequence that feeling like pain is not sufficient for being pain. This allows for the possibility of states that feel like pain but that are not pain, which is the key to this horn of the dilemma. So although we are officially rejecting the rigidified descriptions account of 'pain', the success of the dilemma does not require that we do so.

This leads us to the issue of what ‘pain’ designates. Given that ‘pain’ is a rigid designator that refers to its causal or historical origin, there are only three remotely plausible candidates for its referent. They are:

- (i) The Functional Kind: Pain
- (ii) The Non-Physical Property of Being Pain
- (iii) C-fiber Stimulation

We will argue that (i) and (ii) ought to be rejected, leaving (iii) as the winning candidate. But as we will argue below, (iii) affords a straightforward explanation of the apparent conceivability of pain without C-fiber stimulation and thereby opens up room for a plausible denial of premise (1) of Kripke’s modal argument. We have a couple of tasks, then. First, argue that (i) and (ii) ought to be rejected. Second, show how an acceptance of (iii) allows for the denial of premise (1).

(i) Pain as a Functional Kind

One might grant that ‘pain’ refers, but instead of referring to its microstructural essence, maintain that it refers to a functional kind.¹¹ So, following Putnam (1967), one might argue that pain is the property of having some property or other that fills the pain role. Thus, even though a Martian and I share no physical property, we do share a functional property because we both have some property or other, C-fiber stimulation for me, inflated feet cavities for the Martian (say), that ‘realizes’ the functional role of pain.

We have two responses to this maneuver. First, we are dubious that there are any functional kinds of the sort defined. Go back to the Martian and me. We are told that we

¹¹ We would like to thank Derk Pereboom for raising this point.

share a property. It is not a first-order property, but a higher-order property, the property of having some property or other.... Why ought we to think in this case that there is such a higher order property at all? We are in agreement with Jaegwon Kim's argument that we have no more reason to think that we get higher-order properties from existentially quantifying over first-order properties than we get higher-order individuals from quantifying over 'first-order' individuals.¹²

Second, even granting that such a functional kind exists, were pain identical with it, pain would turn out to be, by definition, epiphenomenal. The problem is that any causal relation that pain might enter into looks as if it is fully explained by pain's realizer. So, for example, if C-fiber stimulation is the realizer of pain in human beings, then it looks as if C-fiber stimulation, and not any higher-order property, will be the property that figures in the causal etiology of pain behavior. The property that is active in pulling my hand back from a hot stove will be C-fiber stimulation. And even if a higher-order property is there, it looks as if it is preempted by its realizer from causally contributing to pain behavior. But this result, we contend, is incompatible with the strong intuition that pain matters causally to our behavior.¹³

(ii) Property Dualism

¹² See Kim (1997, 200-202; 1998, 103-106). For a response (which we judge to be ultimately ineffectual), see David (1997).

¹³ For a recent statement of this objection to functionalism, see McLaughlin (2006). An earlier version of the objection occurs in Block (1990). For an attempt to find a role for mental properties in explanation, if not causation, see Jackson and Pettit (1990). See Kim (1998, 72-77) for a response to Jackson and Pettit.

One might argue that there is a unique non-physical, purely mental phenomenon at the origin of the use of the term 'pain'. If so, then 'pain' would refer to this non-physical phenomenon rather than to C-fiber stimulation.

The problem, of course, is that the assumption that there are non-physical mental phenomena begs the question against the mind-body identity theory. Furthermore, the idea that 'pain' picks out a non-physical phenomenon goes against our ordinary practices of using this term. In talking about pain we take ourselves to be talking about a process that occurs inside the body. We take pain to have physical causes and effects, and therefore to be part of the causal nexus of physical events. If 'pain' designates a non-physical phenomenon then these practices and assumptions either will have to be rejected or the property dualist will have some serious explaining to do about how a non-physical property can matter to the physical world. The view that 'pain' designates a physical, bodily phenomenon best accommodates, we contend, our pre-theoretical judgments about pain's causal profile.

The property dualist might reply that by insisting that 'pain' designates a physical phenomenon we are the ones who are begging the question. But recall that we are arguing from the position of antecedent physicalism. We are not trying to convince the dualist of the truth of the identity thesis. Rather, we are defending the identity thesis against the problem raised by Kripke's modal argument. Our defense assumes that, regardless of whether the term 'pain' is rigid or not, there is no non-physical phenomenon for 'pain' to designate with respect to the actual world. A dualist will obviously reject this assumption. But that need not bother us, since it is not our aim to convince the dualist otherwise.

(iii) C-fiber Stimulation

We come then to (iii). If ‘pain’ refers, it refers to the physical phenomenon at its causal or historical origin, namely, C-fiber stimulation. This means, in effect, that ‘pain’ functions like other natural kind terms and hence that pain itself is a natural kind. If so, then like any natural kind, pain has its microstructure essentially. In other words, being C-fiber stimulation is an essential feature of pain. This doesn’t rule out the view that feeling like pain is also an essential feature of pain. Indeed, on any plausible account of pain, feeling like pain is necessary for being pain. But it does rule out the view that feeling like pain is *sufficient* for being pain. If being C-fiber stimulation is essential to being pain, then even if something feels like pain, if it is not C-fiber stimulation then it is not pain. This means that there could be phenomena that feel like pain but are not pain since they are not C-fiber stimulation. It would be quite easy to mistake conceiving of these phenomena for conceiving of pain without C-fiber stimulation. Hence, the view that ‘pain’ is rigid opens up space for exactly the kinds of possibilities needed for denying that pain without C-fiber stimulation is conceivable.

To sum up, the dilemma for the modal argument is as follows. The term ‘pain’ is rigid or it is not. If it is rigid, then the best account of its semantics treats it as a natural kind term that refers to its causal or historical origin, C-fiber stimulation, and hence pain is essentially C-fiber stimulation. This provides possibilities that are plausibly mistaken for cases of pain in the absence of C-fiber stimulation, thus undermining premise (1) of the modal argument. On the other hand, if ‘pain’ is not rigid then premise (2) of the modal argument is false. On either horn, therefore, the identity theorist can safely avoid

the modal argument, and can do so without rejecting the inference from conceivability to possibility.

VI. OBJECTIONS AND REPLIES

1. The Argument Recast

One response to our argument would be to accept that the dilemma defeats Kripke's original modal argument, but hold that a simple revision avoids the dilemma — just substitute 'the sensation of pain' for 'pain':

1'. It is possible for there to be the sensation of pain without C-fiber stimulation, and vice versa.

2'. If the sensation of pain = C-fiber stimulation, then it is not possible for there to be the sensation of pain without C-fiber stimulation or vice versa.

∴ 3'. The sensation of pain \neq C-fiber stimulation.

The conclusion of this argument should be just as troubling to the identity theorist as the conclusion of the original modal argument.

It should be clear, however, that essentially the same dilemma arises for this revised argument. The description 'the sensation of pain' is either rigid or not rigid. If it is not rigid, for example, if it is a semantically complex description that denotes, with respect to a world w , whichever sensation in w qualitatively resembles pain, then premise (2') is false. On the other hand if 'the sensation of pain' is rigid, e.g. if it semantically similar to descriptions like 'the city of Paris' or 'the color red', then premise (1') can be rejected. If 'the sensation of pain' rigidly designates pain, that is, C-fiber stimulation, then it is easy to explain away the apparent conceivability of the sensation of pain in the absence of C-fiber stimulation. In conceiving of the sensation of pain in the absence of

C-fiber stimulation one is in fact conceiving of something that qualitatively resembles the sensation of pain but which is not the sensation of pain, since it is not C-fiber stimulation.

2. Property Dualism (Again)

Another objection is that on our argument, property dualism sneaks back in. Consider the horn of the dilemma on which ‘pain’ is a rigid natural kind term that refers to C-fiber stimulation. We have argued that if this is the case then there are possible phenomena that qualitatively resemble pain but which are not pain. Consider one such phenomenon, call it p_1 , and its qualitative character, q_1 . By hypothesis, p_1 and pain are not identical. The question is whether q_1 is identical to the qualitative character of pain. If it is, then it looks as if this qualitative character is itself a non-physical property shared by p_1 and pain. But that is just property dualism.

The identity theorist who opts for this horn of the dilemma is going to have to deny that p_1 and pain have exactly the same qualitative character. But it is enough for this horn of the dilemma that p_1 and pain are qualitatively *similar*. They need not be exactly the same down to the finest phenomenological detail. So long as there is the slightest phenomenological difference between p_1 and pain, the problem about property dualism sneaking back in is avoided.

It could be replied, however, that this defeats the purpose of this horn of the dilemma. The relevance of phenomena like p_1 to the modal argument is that they can be used to explain away the apparent conceivability of pain in the absence of C-fiber stimulation. The claim is that in conceiving of pain without C-fiber stimulation, one is in fact conceiving of a phenomenon like p_1 that is qualitatively like pain but which is not

identical to pain. But if p_1 and pain are not qualitatively identical, it becomes less plausible to think that this explains away the conceivability of pain in the absence of C-fiber stimulation. Can't we conceive of something that feels *exactly* like pain but which is not C-fiber stimulation? If we can, then the version of strategy I employed on this horn of the dilemma is useless.

Recall, however, that in the context of this debate, conceivability means *psychological* conceivability, which is a kind of imaginability. To psychologically conceive of something is to imagine a world in which it is the case. Given this account of conceivability, it is open to the identity theorist to deny that we can conceive of something that feels exactly like pain down to the finest phenomenological detail. Our imaginative powers are too coarsely grained, indeterminate and fuzzy. We can imagine creatures who have sensations that resemble pains, but our imaginative capacities are not finely grained enough to fill in all of the phenomenological details of these sensations. The imagined sensations are qualitatively indeterminate in a way that genuine pains are not. At best, therefore, we can conceive of creatures having sensations that are qualitatively similar to pain, but due to the indeterminacy in our imaginative capacities, the similarity is not exact.

3. Conceiving of C-fiber Stimulation

So far we have focused on conceiving of pain in the absence of C-fiber stimulation. But what about the conceivability of C-fiber stimulation in the absence of pain?

What is one conceiving when one, allegedly, conceives of C-fiber stimulation, but in so conceiving, does not conceive of pain? If pain is C-fiber stimulation, indeed, if necessarily pain is C-fiber stimulation, to conceive of the one is to conceive of the other. It is tempting to say that one can conceive of C-fiber stimulation, but in failing to conceive of it under the phenomenal concept of pain, one fails to appreciate that one is conceiving of pain. Although attractive, this would be to abandon the strategy I response to the modal argument, in favor of a response that rejects the inference from conceivability to possibility. In adverting to a phenomenal concept of pain, one would be employing the phenomenal concept strategy,¹⁴ which, standardly understood, requires rejecting the move from conceivability to possibility.

A strategy I response requires finding some real possibility to explain away the apparent conceivability of C-fiber stimulation in the absence of pain. What real possibility could serve this purpose? One strategy would be to exploit the tenuous grasp on the concept of C-fiber stimulation. As we have defined it (and as it is usually understood), C-fiber stimulation is simply a dummy term for whatever physical type scientists discover is correlated with pain (in human beings). What are we thinking, then, when we think about C-fiber stimulation? Nothing all that determinate, we would suggest. We conceive of some neurons' firing, say, and imagine a human being with such neurons' firing not feeling any pain whatsoever. But of course, in conceiving of some neurons' firing, there is no reason to think that we have managed to think of the physical type with which pain is identical and so no reason to think that we are conceiving of that physical type without pain. We may be conceiving of some physical type without

¹⁴ For a statement of the strategy, see Stoljar (2005). For an argument against its viability, see Chalmers (2006).

conceiving of pain, but we see no reason to grant that it is the relevant physical type in that conception.

One might object that this response misses the point. True enough, the objector can grant, we do not know the physical state that is correlated with pain and we are using C-fiber stimulation simply as a stand in for whatever state our best physiological theory tells us is perfectly correlated with pain in humans. But these are just empirical details. For whatever that state turns out to be, the objector might urge, just so long as it is a broadly physical state, it looks as if one is going to be able to conceive of it without conceiving of pain. Else, the objector can end with a flourish, you will be relying on this physiological state to close the explanatory gap between the physical and the phenomenal.¹⁵ And that is surely to place high hopes on the future of physiological theory.

Maybe future physiological theory, or physical theory more broadly, will close the explanatory gap. We do not want to prejudge that issue. The progress of physical theory often surprises. Be that as it may, we do not want this response to the modal argument to require that in the future the explanatory gap be closed. Let's operate in the abstract then and imagine that the candidate physical state is N and we are imagining that N, whatever it turns out to be, is apparently conceivable without conceiving of pain.

Even in such a situation, there is room to maneuver. In conceiving of N, we conceive of it under some mode of presentation, call it μ . Taking a page from Perry, we can exploit the referential semantics for concepts and say that in (allegedly) conceiving of C-fiber stimulation, one is (actually) thinking of some state, using μ to conceive of it, but

¹⁵ For the classic statement of the problem of the explanatory gap, see Levine (1983).

in that situation, the given concept refers to some state besides C-fiber stimulation. Such a situation is *not* a situation in which one is conceiving of C-fiber stimulation without conceiving of pain. Rather, it is a situation in which one is thinking of some state in the way that one usually thinks of C-fiber stimulation without conceiving of pain. This latter possibility, which is metaphysically possible, is easily confused with the former, which is not.

4. Too Implausible

The final objection that we will consider is blunt: the view that the feeling of pain is not sufficient for pain is, quite simply, too implausible to be believed.

As we noted in our introduction, many philosophers find it compelling that feeling like pain just is pain. These philosophers, even if physicalists, need not reject our argument however. As we see our paper, we are laying out the options for an identity theorist's response to Kripke's modal argument, those options that do not question the inference from conceivability to possibility. If a philosopher finds the conclusion of the second horn of the dilemma too much to bear, she is free to exploit the resources of the first horn. Nothing we say commits a physicalist to pursuing the second horn, and so nothing we say commits a physicalist to holding that the feeling of pain is not sufficient for pain.

VII. CONCLUSION

The central idea in this paper is that Kripke's modal argument against the mind-body identity thesis faces a dilemma based on whether phenomenal terms like 'pain' are

rigid or not. Either 'pain' is rigid, in which case one can deny the conceivability of pain in the absence of C-fiber stimulation, or 'pain' is not rigid, in which case the identity of pain and C-fiber stimulation is not necessary. The interest of this dilemma is that it avoids having to deny the inference from conceivability to possibility. Severing the connection between conceivability and possibility raises many difficult and complex issues about the metaphysics and epistemology of modality. We think the dilemma we have posed should be attractive to identity theorists insofar as it avoids all of these issues.

In Kripke's original presentation of the modal argument he denies the conceivability of water that is not H₂O. Although we can conceive of a substance with all the surface features of water, this is not to conceive of water that is not H₂O. But, according to Kripke, the same cannot be said for pain. This is because he thought that to conceive of something that has the surface features of pain, that is, to conceive of something that feels painful, is to conceive of pain. Therefore, unlike the case of water or heat, one cannot explain away the conceivability of pain without C-fiber stimulation.

Our argument here has been to take issue with Kripke's selective appeal to natural kinds. If pain is a natural kind, then we should stand ready to treat it as we treat other natural kinds. In that case, its essence would not be exhausted by the way it feels, and hence conceiving of something that feels like pain would not be sufficient for conceiving of pain. On the other hand, if we want to say that 'pain' behaves differently, that it is not a natural kind, then rigidity is threatened, and with it, a key component of the modal argument. Either way, the identity theorist has ample resources to respond to Kripke's argument, even without wandering into the contentious territory of conceivability and possibility.

BIBLIOGRAPHY

- Balog, Katalin. 2000. "Conceivability, Possibility, and the Mind-Body Problem." *The Philosophical Review* 108: 497-528.
- Block, Ned. 1990. "Can the Mind Change the World?" In *Meaning and Method*, ed. George Boolos. Cambridge: Cambridge University Press.
- . 2006. "Max Black's Objection to the Mind-Body Identity." In *Phenomenal Concepts and Phenomenal Knowledge*, eds. Torin Alter and Sven Walter. Oxford: Oxford University Press.
- Chalmers, David. 2002. "Does Conceivability Entail Possibility?" In *Conceivability and Possibility*, eds. Tamar Gendler and John Hawthorne. Oxford: Oxford University Press, 145-200.
- . 2006. "Phenomenal Concepts and the Explanatory Gap." In *Phenomenal Concepts and Phenomenal Knowledge*, eds. Torin Alter and Sven Walter. Oxford: Oxford University Press.
- David, Marian. 1997. "Kim's Functionalism." *Philosophical Perspectives*, 11: 133-148.
- Gendler, Tamar and John Hawthorne, eds. 2002. *Conceivability and Possibility*. Oxford: Oxford University Press.
- Hill, Christopher. 1997. "Imaginability, Conceivability, Possibility, and the Mind-Body Problem." *Philosophical Studies* 87, 61-85.
- Jackson, Frank and Philip Pettit. 1990. "Program Explanations: A General Perspective." *Analysis* 50: 107-117.
- Kaplan, David. 1989. "Demonstratives." In *Themes from Kaplan*, eds. Joseph Almog, John Perry, and Howard Wettstein. Oxford: Oxford University Press, 481-563.

- Kim, Jaegwon. 1997. "The Mind-Body Problem: Taking Stock after 40 Years."
Philosophical Perspectives 11: 185-207.
- . 1998. *Mind in a Physical World*. Cambridge, MA: MIT University Press.
- Kripke, Saul. 1972. "Naming and Necessity." In *Semantics of Natural Language*, eds.
Donald Davidson and Gilbert Harman. Dordrecht, Holland: Reidel, 253-355.
- . 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Levine, Joseph. 1983. "Materialism and Qualia: the Explanatory Gap." *Pacific
Philosophical Quarterly* 64: 354-361.
- Loar, Brian. 1990. "Phenomenal States." *Philosophical Perspectives* 4, 81-108.
- McLaughlin, Brian. 2006. "Is Role-Functionalism Committed to Epiphenomenalism?"
Journal of Consciousness Studies, 13 (1-2): 39-66.
- Perry, John. 2001a. *Knowledge, Possibility and Consciousness*. Cambridge, MA: MIT
Press.
- Perry, John. 2001b. *Reference and Reflexivity*. Stanford: CSLI Publications.
- Putnam, Hilary. 1967. "The Nature of Mental States." Reprinted in *Mind, Language and
Reality: Philosophical Papers* Vol. 2. Cambridge: Cambridge University Press,
1975.
- Salmon, Nathan. 1981. *Reference and Essence*. Princeton: Princeton University Press.
- Soames, Scott. 2002. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and
Necessity*. New York: Oxford University Press.
- Stoljar, Daniel. 2005. "Physicalism and Phenomenal Concepts." *Mind and Language* 20:
469-494.

Yablo, Stephen. 2002. "Coulda, Woulda, Shoulda." In *Conceivability and Possibility*, eds. Tamar Gendler and John Hawthorne. Oxford: Oxford University Press, 441-492.