

On Understanding a Theory on Conscious Experiences

Abstract. McGinn claims, among other things, that we cannot understand the theory that explains how echolocationary experiences arise from the bat's brain. One of McGinn's arguments for this claim appeals to the fact that we cannot know in principle what it is like to have echolocationary experiences. According to Kirk, McGinn's argument fails because it rests on an illegitimate assumption concerning what explanatory theories are supposed to accomplish. However, I will argue that Kirk's objection misfires because he misapprehends McGinn's argument. Further, I will articulate and briefly assess some ways in which McGinn's argument can be blocked.

Keywords. The mind-body problem, Concepts of consciousness, What it is like to be a bat, Colin McGinn, Robert Kirk

McGinn (1989) claims, among other things, that we cannot understand the theory that explains how echolocationary (or batty) experiences arise from the bat's brain. One of the influential arguments McGinn develops for this claim appeals to the fact that we cannot know in principle what it is like to have batty experiences. According to Kirk (1991), McGinn's argument fails because it rests on an illegitimate assumption concerning what explanatory theories are supposed to accomplish. However, my main aim in this note is to show that Kirk's objection misfires because he misapprehends McGinn's argument. The objection that I will consider is somewhat old but I hope this does not by itself detract from its significance. The sort of misunderstanding of McGinn's argument that is encapsulated in Kirk's objection has not been sufficiently recognized in the literature, which might explain at least in part the general tendency many philosophers seem to have of rejecting McGinn's overall account out of hand. After answering Kirk's objection, I will articulate and briefly assess some ways in which McGinn's argument can be blocked.

1. It is widely assumed that we human beings cannot know in principle what it is like to have batty experiences. There is a clear sense in which the characteristic qualitative aspect, or the phenomenal character, of the experiences bats have when they navigate the environment by using their echolocationary techniques appears to be irredeemably beyond our cognitive grasp.

Batty experiences will never be intelligible to us, it seems, in the way our experiences like smelling a skunk or tasting coffee are.¹

What follows from our lack of access to batty experiences? In particular, does our failure to access batty experiences provide any support for the following thesis?

(T) We cannot understand the theory that explains how batty experiences arise from the bat's brain.

McGinn argues that (T) is supported by the fact that we cannot access the phenomenal character of batty experiences:

Call this type of experience [batty experience] *B*, and call the explanatory property that links *B* to the bat's brain *P_I*. By grasping *P_I* it would be perfectly intelligible to us how the bat's brain generates *B*-experiences, we would have an explanatory theory of the causal nexus in question...But then it seems to follow that grasp of the theory that explains *B*-experiences would confer a grasp of the nature of those experiences: for how could we understand that theory without understanding the concept *B* that occurs in it? How could we grasp the *nature* of *B*-experiences without grasping the *character* of those experiences?...Our concepts of consciousness just *are* inherently constrained by our own form of consciousness, so that any theory the understanding of which requires us to transcend these constraints would ipso facto be inaccessible to us. (1989, pp. 355-6)²

McGinn's argument here is, roughly, this: (fully) understanding a theory that explains how *B*-experiences arise from the bat's brain requires us to grasp the concept *B* that (ineliminably) occurs in that theory,³ which in turn requires us to grasp the character of those experiences; however, we just cannot grasp the character of those experiences; therefore, we cannot understand the theory that explains how *B*-experiences arise from the bat's brain.

¹ Nagel's seminal paper (1974) played the main role in forcefully bringing to the attention of philosophers the significance of our epistemic position with respect to batty experiences for the mind-body problem.

² All McGinn references are to this work.

³ McGinn does not explicitly state but presumably takes for granted that the qualifications in the parentheses ('fully' and 'ineliminably') are necessary for the argument to get off the ground. If understanding were taken as *partial* understanding, or if *B* were a concept that *eliminably* occurs in the theory, the first premise of McGinn's argument would be obviously false. Having noted that, I will suppress these qualifications in the remainder of the paper.

In a more explicit form, the argument runs as follows:

(1) Understanding a particular theory requires grasping all the concepts that occur in that theory.

(2) The concept *B* occurs in the theory that explains how *B*-experiences arise from the bat's brain.

From (1) and (2), it follows that

(3) Understanding the theory that explains how *B*-experiences arise from the bat's brain requires grasping the concept *B*.

We also independently have the following:

(4) We can grasp the concept *B* only if we can grasp the character of *B*-experiences.

(5) We can grasp the character of *B*-experiences only if we can have *B*-experiences.⁴

(6) We cannot have *B*-experiences.⁵

From (3)-(6), it follows that

(T) We cannot understand the theory that explains how *B*-experiences arise from the bat's brain.

Let's call (1) *the grasping requirement* (GR); and, by taking special note of (1), let's call this argument *the argument from grasping*.⁶

How is the argument from grasping related to the argument McGinn develops for the sort of "mysterianism" he is best known for? According to McGinn's mysterianism, we cannot understand (or are "cognitively closed" with respect to) the theory that explains how *our* experiences arise from our brains. Clearly, a straightforward adjustment of the argument from

⁴ It is clear that the argument could have been stated by simply having 'We can grasp the concept *B* only if we can have *B*-experiences' as a premise instead of having both (4) and (5). However, I here stick with the way McGinn seems to prefer to state it.

⁵ Premises (5) and (6) are not explicitly stated in the passage quoted from McGinn, but the context surrounding the passage leaves no doubt that McGinn holds them.

⁶ I claim that the argument from grasping is one plausible and textually supported interpretation of McGinn's argument in the relevant passage, while I do not wish to claim that it is *the* correct one. Thanks to an anonymous reviewer for pressing on this issue.

grasping cannot support McGinn's mysterianism simply because we possess concepts of our experiences and thereby know what our experiences are like. McGinn's argument for his mysterianism, which I call *the closure argument by elimination*, runs roughly as follows: introspection and perception are the "two possible avenues open to us in our aspiration to identify *P* [the brain property that is responsible for our consciousness]" (p. 397), and neither can help us identify *P* – therefore, we cannot identify *P*, in which case we cannot solve the mind-body problem in the case of humans. The argument from grasping is presented by McGinn as "a further point" (p. 355) in (and hence it is a digression from) his main discussion of whether introspection can enable us to get to *P*; and as such it stands on its own and is independent of the closure argument by elimination. Given the implications of its conclusion, the closure argument by elimination is bound to be more controversial than the argument from grasping⁷; however, it seems to me that the latter also brings out some interesting issues and deserves a separate and focused investigation.

What, one might reasonably wonder, if the argument from grasping succeeds in establishing (T)? What is the significance of the purported truth that we cannot understand the theory that explains how *B*-experiences arise from the bat's brain? As McGinn sees it, nothing less than the possibility of *our* achieving "a general solution to the mind-body problem" (p. 356) is at stake. If (T) is true, then, according to McGinn, "even if we could solve [the mind-body problem] for our own case, we could not solve it for bats or Martians" (p. 356). Of course, such a result would be especially worrisome for a variety of reductionist views on consciousness: if *B*-experiences are nothing but some physical features of the bat's brain, as physicalism claims them to be, or if they are nothing but some causal-role properties of the bat's brain, as functionalism claims them to be, then what can possibly obstruct our path to

⁷ For a critical discussion of McGinn's mysterianism, see Sacks (1994).

solving the mind-body problem for bats? The question, I believe, is forceful: it certainly seems that the truth of (T) would be a mystery if some form of reductionism were true.

I believe McGinn's account has not received the due attention it deserves, and I am largely in agreement with Kriegel's following observation: "The literature on mysterianism has so far been somewhat dogmatically dismissive. Critical discussions of the merits and demerits of the view are few and far between. In particular, McGinn's argument is rarely if ever engaged" (2009, p. 455). The current paper may be read as a modest attempt to remedy this unfortunate situation by focusing on a particular but significant strand in McGinn's position. In what follows, I will argue against the objection Kirk develops against McGinn's argument. Once a misunderstanding like Kirk's is eliminated, the ground is cleared for drawing out the full implications of McGinn's argument. I will conclude by suggesting a trilemma, one that captures the options available for resisting the argument and thereby functions as an invitation for the potential dissidents to clarify their stand.

2. Kirk argues that McGinn's argument in the passage quoted above assumes that "a satisfactory theory of explaining the nature of subjective experience must be capable of actually conferring concepts of experience on those who start off without them" (1991, p. 20).⁸ Let us call this assumption *the conferring requirement* (CR). So, we have the following:

(CR) A theory of conscious experience is explanatorily satisfactory for us (or for cognitive beings in general) only if it is capable of conferring a grasp of the concepts of conscious experience involved in that theory to (those of) us (or cognitive beings in general) who start off without (or who do not have an independent grasp of) them.

By Kirk's lights, CR "is illegitimate" (p. 19) because it requires the explanatory theory to achieve something no theory can possibly achieve and can therefore be reasonably expected to achieve. Grasping concepts of experience requires "an actual sort of experience of the right sort," which is "something no theory could supply" (p. 19). Holding that the explanatory theory

⁸ All Kirk citations are to this work.

should confer a grasp of such concepts as *B* is, Kirk argues, setting up “an insuperable hurdle” (p. 21) for that theory; and, if we fail to get over it, then it is not our cognitive powers but that very hurdle itself that should deserve the blame.⁹

According to Kirk, if CR were legitimate, i.e. if it were required for a satisfactory theory of conscious experience that it confer concepts of experience on those who start off without them, then the proper conclusion to be drawn would be that no theory can meet that requirement and hence there cannot *be* such a satisfactory theory, one that is a possible object of our understanding. Therefore, Kirk finds it “puzzling” (p. 19) that McGinn assumes the legitimacy of that requirement while holding that there *is* a satisfactory theory of conscious experience.

Let us assume, for the moment, that McGinn endorses CR. How, then, would the argument for (T) proceed? It is evident that in the passage above, McGinn’s main intention is to develop an argument for (T), but it is not at all clear how the argument can possibly be intended to move from CR. CR states that the explanatory theory in question should confer a grasp of the concept *B*, and (T) states that we cannot understand that theory. However, if the explanatory theory *confers* a grasp of *B*, as CR says it must, then what reason can we possibly have to think that (T) is true? If anything, just the opposite appears to be correct: holding CR provides a good reason for thinking that (T) is *not* true. That the explanatory theory confers a grasp of *B* supports the conclusion that we *can* understand that theory.

⁹ I would like to note in passing that Patricia Churchland attributes to McGinn some other requirement that is relevantly similar to CR. In a rather belligerent response to McGinn’s (2014) review of her (2013), Churchland writes that McGinn’s account suffers from “a whopping flaw” and that “no causal explanation for a phenomenon...should be *expected* to actually produce that phenomenon” (2014, emphasis original). On a natural interpretation, Churchland attributes to McGinn the requirement that a satisfactory explanation of a particular subjective experience must actually produce that experience in those who understand that explanation (*the producing requirement*, PR). PR follows from CR on the reasonable assumption that an explanation of a particular subjective experience can confer concepts of experience on those who start off without them only if that explanation can produce that subjective experience in them. I will not discuss Churchland’s attack on McGinn’s position separately but I am confident that what I have to say below about Kirk’s critique applies *mutatis mutandis* to it.

The moral is that if McGinn were really to endorse CR, then what is, from his point of view, a formidable obstacle to our understanding the theory explaining batty experiences (namely, our apparent failure to grasp the concept *B* that occurs in that theory) would be overcome by what that theory confers. So, assuming that McGinn endorses CR faces the difficulty of giving a reasonable (and charitable) account of how McGinn's argument for (T) can possibly proceed.

The argument intended by McGinn for (T), I claim, is the argument from grasping, and as such, it has nothing much to do with (more specifically, neither assumes nor needs to assume) CR. So, regardless of whether Kirk shows the illegitimacy of CR, he fails to properly address McGinn's argument from grasping. Furthermore, if McGinn neither assumes nor needs to assume CR in his argument for (T), then there need not be anything "puzzling" in McGinn's commitment to the thesis that there *is* a satisfactory theory of conscious experience.¹⁰

3. Despite what I have argued for above, there is a particular statement in the passage above that might well give the impression that McGinn endorses CR. To quote again, McGinn writes: "it seems to follow that grasp of the theory that explains *B*-experiences would *confer* a grasp of the nature of those experiences." Kirk (p. 19) places a special emphasis on this statement and takes it as evidence for the claim that McGinn endorses CR.

¹⁰ It might be objected that McGinn is committed to the thesis that CR holds (if not for us) at least for *some* beings (e.g., those beings that are "cognitively open" to the relevant theory) and therefore that Kirk's objection to McGinn's account stands untouched by my point above. However, there is at least one good reason to think that McGinn is not committed to the idea that CR holds for some beings. It is important here to note the distinction McGinn makes between *absolute* and *relative* cognitive closure: "A problem is absolutely cognitively closed if no possible mind could resolve it; a problem is relatively closed if minds of some sorts can in principle solve it while minds of other sorts cannot" (p. 360). McGinn also writes: "It certain seems to be at least an open question whether the problem is absolutely insoluble; I would not be surprised if it were" (p. 361). Now, if McGinn grants the possibility of *absolute* cognitive closure, as he clearly does, then there might well be no minds that are cognitively open to the theory that explains *B*-experiences. However, given that McGinn holds that there is such a theory, then from his point of view, there being such a theory has nothing to do with its potential to endow *some* beings with a grasp of *B* because there might well be no such beings.

However, McGinn's statement at hand can be interpreted in a different way, and in a way that gives further support to interpreting McGinn's argument along the lines of the argument from grasping. In the relevant passage, McGinn can be plausibly taken as raising the following question: "Assuming that we cannot have *B*-experiences, how can we possibly understand the theory that involves the concept *B*?" Having *B*-experiences would enable us to form the concept *B* that occurs in that theory. Barring that, McGinn argues, the relevant explanatory theory can be understandable *by us* only if it confers the concept *B on us*. According to McGinn, we can understand a theory only if we can grasp the concepts it involves (GR); and, if we do not have an independent grasp of some of those concepts (i.e. if we do not have a grasp of those concepts prior to our exposure to the theory), then we can understand the theory only if it confers on us such a grasp. That theories must confer on us a grasp of some concepts that we do not have an independent grasp of if they are to be understandable by us follows from GR (more on this below). Therefore, GR is the claim that Kirk needs to attack if he wishes to undermine McGinn's argument for (T).¹¹ What McGinn has to say about conferring ultimately depends upon the GR he takes for granted.

Under this interpretation, the statement that Kirk focuses on does not commit McGinn to CR. CR is a requirement that theories that include some concepts that we don't have an independent grasp of must satisfy in order for them to be *explanatorily satisfactory*. According to Kirk, McGinn commits "the error of assuming that a *satisfactory* theory must actually endow us with concepts for characterizing experience" (p. 22, emphasis mine). However, McGinn's statement at hand is about a requirement that those theories must satisfy in order for them to be *understandable by us* (*the understandability requirement, UR*). So, we have the following:

¹¹ Interestingly, Kirk explicitly accepts at one point GR. He writes: "Having concepts such as *B* is *necessary* in order to understand the theory that explains the character of *B*-experiences" (p. 21). Given this, one might find it "puzzling" that Kirk accepts GR but still attacks what McGinn has to say about conferring. Of course, the puzzle dissolves once we realize that what McGinn has to say about conferring has nothing much to do with the CR Kirk attributes to him. This provides further support to my claim above that Kirk has misapprehended McGinn's statement on conferring.

(UR) A given theory is understandable by us (or cognitive beings in general) only if it is capable of conferring a grasp of those concepts involved in that theory to (those of) us (or cognitive beings in general) who start off without (or who do not have an independent grasp of) them.

From UR, we can derive the following as one of its specific instances:

(URE) A theory of conscious experience is understandable by us (or cognitive beings in general) only if it is capable of conferring a grasp of those concepts of conscious experience involved in that theory to (those of) us (or cognitive beings in general) who start off without (or who do not have an independent grasp of) them.

For McGinn, if a theory that includes some concepts that we don't have an independent grasp of does not confer a grasp of those concepts on us, then that theory is not understandable by us. However, the fact that a theory is not understandable by us does not mean that it does not satisfactorily explain what it is intended to explain. So, a theory that does not confer on us a grasp of the concepts that we do not have any independent grasp of, according to McGinn, is *not* understandable by us but might still *be* explanatorily satisfactory.¹²

4. As stated above, Kirk's principal reason against CR is that there are no theories on conscious experiences that can possibly fulfill and can therefore be reasonably expected to fulfill it. Since McGinn does not embrace CR, I will not spend time on Kirk's argument against it. However, one might develop an analogous argument against URE and argue that because there are no theories on conscious experiences that can possibly fulfill and can therefore be reasonably expected to fulfill URE, it is *illegitimate*. If Kirk's reason against CR is forceful, and if it works equally effectively against URE, as one might reasonably argue, then pointing at the fact that Kirk misidentifies the requirement McGinn embraces does not fully circumvent the problem that Kirk thinks afflicts McGinn's account. If so, it might be maintained that the real issue with McGinn's position has not yet been brought into relief.

¹² One might wonder what the argument from grasping would look like once McGinn's UR (or URE) is taken on board. The revision required is minimal: replace (5) above by (5') We can grasp the character of *B*-experiences only if either we can have *B*-experiences or the theory can confer on us a grasp of the character of *B*-experiences, and add as a new premise (7) The theory *cannot* confer on us a grasp of the character of *B*-experiences.

By Kirk's lights, CR is "illegitimate" because it requires an explanatory theory of conscious experience to achieve something no theory of conscious experience can possibly achieve and can therefore reasonably expected to achieve: given that *no* explanatory theory can confer on us concepts of conscious experience we do not have an independent grasp of, then CR sets an "insuperable hurdle" for such an explanatory theory to be satisfactory, which renders it illegitimate. The question I am concerned with now is how compelling an objection to URE, which is similar in spirit and form to Kirk's objection to CR, would be. The objection is this: given that *no* theory of conscious experience that involves some concepts of conscious experience that we do not have an independent grasp of can confer a grasp of those concepts on us, URE sets up an "insuperable hurdle" for such a theory to be understandable by us. Therefore, the objection goes, URE cannot be a requirement that we can reasonably expect theories of conscious experience should meet in order for them to be understandable by us.¹³

In response to this, I would like to concede two points. First, it is reasonable to derive the illegitimacy of a particular requirement from the fact that it cannot possibly be met by those on which it is placed as a requirement. This is, roughly, for the familiar reasons why many philosophers hold that *ought* entails *can*. Second, there are no theories on conscious experiences that can possibly satisfy URE. It is also worth noting that McGinn does not make any claim contradicting the former point and, more importantly, explicitly grants the latter point (p. 356) (just as would be expected given his intention to argue for (T): if there were theories that confer on us such a grasp, then that would be a reason for thinking that (T) is *not* true).

However, I would like to argue that it does not follow from these two concessions that URE is illegitimate. The crucial point here is that URE is derived from UR and UR is intended to be a *general* requirement for all theories but not a local requirement solely for theories on

¹³ I would like to thank an anonymous reviewer for the comment that led to a clearer statement of this objection.

conscious experiences. And, as such, the central rationale for UR derives from an eminently plausible general requirement about the conditions under which a theory is understandable by us, viz. GR: if we can understand a theory only if we can grasp all the concepts that occur in it (GR), then we can understand a theory that contains some concepts we do not have an independent grasp of only if it confers such a grasp on us (UR). Hence, pointing at the putative fact that there are no theories *on conscious experiences* that can possibly satisfy a particular instance of UR (viz. URE) cannot by itself show that UR as a general principle or URE as a specific instance of UR is illegitimate. The argument for the illegitimacy of URE needs to target the general principle from which it is derived (viz. UR) and needs to demonstrate that there are no theories on *any areas of inquiry* that can possibly satisfy it. And, the putative fact that UR cannot be satisfied by theories on *a particular area* cannot show that UR is illegitimate just as the fact that a particular person is incapable of fulfilling a particular moral requirement cannot show that the requirement is illegitimate. If Dexter has an irresistible urge to commit murder and is incapable of acting in accordance with a commandment like ‘Thou shall not kill!’, then the proper conclusion to be derived is not that the commandment is illegitimate. Similarly, if theories on conscious experiences cannot satisfy the relevant specific version of UR (viz. URE), then the proper conclusion to be derived is not that that specific version is illegitimate.

Further, there *are* certainly theories that satisfy UR. Take for instance Galileo’s theory of motion, where, on a standard history of science, the concept *acceleration* is clearly introduced for the first time in physics by distinguishing it from *velocity*. Assuming that we do not have an independent grasp of the concept *acceleration*, we can understand Galileo’s theory of motion only if it confers on us a grasp of that concept that occurs in it. And the theory actually confers

on us such a grasp by defining it in terms of some concepts that we already have an independent grasp of such as *velocity* and *time*.¹⁴ Plainly, examples can easily be multiplied indefinitely.

The moral I draw is that UR is a general requirement that is supported by GR, which is itself another general requirement, and therefore, from the fact that there are no theories *on conscious experiences* that can satisfy one of its specific instances, it does not follow that that specific instance is illegitimate.

5. Let me highlight the central points that have emerged in our discussion. Firstly, the CR that Kirk thinks McGinn endorses is not an assumption of the argument intended by McGinn, namely the argument from grasping. The argument from grasping stands untouched even if Kirk is right that CR is illegitimate. Secondly, there is good reason to think that McGinn does not endorse CR. The truth of CR would count against the truth of (T). Thirdly, McGinn holds UR, i.e., that a theory that involves some concepts we do not have an independent grasp of is understandable by us only if understanding it confers on us a grasp of those concepts. The constraint placed here by McGinn on theories concerns their understandability by us and can be plausibly taken as following from GR; and as such, it does not entail CR. Fourthly, McGinn holds that there are not any theories of conscious experience that can satisfy UR.

The upshot is that McGinn does not endorse the CR Kirk attributes to him but endorses UR. Further, McGinn does not think the latter requirement is satisfied by theories on conscious experiences. All in all, Kirk's attack leaves McGinn's argument for (T) unscratched.

6. I would now like to close the paper with articulating and assessing the ways I find most plausible to block McGinn's argument from grasping. Despite its sketchiness, I believe that this

¹⁴ In this particular case, conferring takes the form of defining: the theory confers a grasp of a particular concept on the subject through defining it in terms of some of its other concepts. An interesting question is whether there are other forms the conferring relation in question might take, i.e., whether a theory can fulfill UR without defining the problematic concepts in terms of some other concepts. I myself cannot conceive any other way. That being said, however, I can fortunately sidestep the question for the purposes of this paper.

will be of value for at least one good reason. The discussion below will further attest to the force of the argument from grasping and invite the potential dissidents to clarify their stand.

It is clear that the argument is valid. Further, I hold that premises (1) and (6) are virtually unassailable – any attempt to block the argument by denying one of these premises is *too* desperate:

- (1) Understanding a particular theory requires grasping all the concepts that occur in that theory.
- (6) We cannot have *B*-experiences.

This leaves us with premises (2), (4) and (5) as possible targets:

- (2) The concept *B* occurs in the theory that explains how *B*-experiences arise from the bat's brain.
- (4) We can grasp the concept *B* only if we can grasp the character of *B*-experiences.
- (5) We can grasp the character of *B*-experiences only if we can have *B*-experiences.

There are basically two different ways to attack these premises. First, one might deny (2). The most plausible way of attacking (2) is, it seems to me, to argue that it is based on the mistaken assumption that there *are* “concepts of consciousness” (McGinn) such as *B*. Concepts of consciousness, as McGinn conceives them, are formed through the introspective attention of the experiencing subject to the qualities of her experiences. One might reasonably endorse eliminativism about concepts of consciousness and argue that there are no concepts satisfying the conditions McGinn articulates¹⁵ (perhaps because, one might say, the whole idea of

¹⁵ An alternative way of denying premise (2) is to maintain that there are concepts of consciousness but they are not part of the relevant theory. As I see it, the idea here is not substantially different from eliminativism that I mention above. More specifically, this attack on (2) is eliminativism *so far as the theory that explains how B-experiences arise from the bat's brain is concerned*. Therefore, at least for the purposes of this paper, the difference between the claim that there are phenomenal concepts but they are no part of the relevant theory and the claim that there are no concepts of consciousness *tout court* is not big enough to justify a separate treatment of the former.

concepts of consciousness stems either from a mistaken picture of introspection as “turning one’s gaze inward” or from a faulty view about the qualities of experiences).¹⁶

Another way to attack the argument from grasping is to deny what is entailed by (4) and (5), *viz.* that we can grasp the concept *B* only if we can have *B*-experiences. According to this objection, there *are* concepts of consciousness, as McGinn takes them to be, but grasping them does *not* require having the experiences that they refer to. The most viable version this objection might take, I believe, endorses reductionism about concepts of consciousness and maintain that those concepts can be analyzed in terms of some other concepts the grasping of which does not require one to have any specific experiences. Take, for instance, those varieties of functionalism according to which there are causally-based synonyms of concepts of consciousness. On such views, grasping a concept of consciousness is nothing more mysterious or demanding than grasping a concept describing a causal role. And, since grasping causal-role concepts does not require one to have any specific experiences, as one might reasonably argue, phenomenal concepts *qua* causal-role concepts do not also require one to have any specific experiences (including those experiences that they refer to). That is, if functionalism (or some form of reductionism about concepts of consciousness in general) is true, then either (4) or (5) is false.

So far as I can see, there are no other objections to the argument from grasping that are even remotely plausible. What I suggest, then, is a trilemma: *either* eliminativism *or* reductionism about concepts of consciousness, *or* (T). In slightly different words, if there are

¹⁶ Note what McGinn says about how concepts of consciousness are related to introspection: “Our acquaintance with consciousness could hardly be more direct; phenomenological description thus comes (relatively) easily. ‘Introspection’ is the name of the faculty through which we catch consciousness in all its vivid nakedness. By virtue of possessing this cognitive faculty we ascribe concepts of consciousness to ourselves; we thus have ‘immediate access’ to the properties of consciousness” (p. 354). McGinn’s description of “introspectively ascribed concepts” (p. 354) is anything except uncontroversial. See, for instance, Dennett (1988).

concepts of consciousness such as *B* that cannot be analyzed in terms of some other concepts that we possess, then (T) is inevitable – or so I have argued.¹⁷

References

- Churchland, P. (2013). *Touching a nerve*. New York: Norton.
- Churchland, P. (2014). In response to: storm over the brain. *The New York Review of Books*.
- Dennett, D. (1988). Quining Qualia. In A. Marcel & E. Bisiach (Eds.) *Consciousness in contemporary science*. Oxford: Oxford University Press.
- Kirk, R. (1991). Why shouldn't we be able to solve the mind-body problem? *Analysis*, 51 (1), pp. 17-23.
- Kriegel, U. (2009). Mysterianism. In T. Bayne, A. Cleermans, & P. Wilken (Eds.), *The Oxford companion to consciousness*. Oxford: Oxford University Press.
- Loar, B. (2004). Phenomenal states. Originally published 1990/7. In P. Ludlow, Y. Nagasawa, & D. Stoljar (Eds.), *There is something about Mary*. Cambridge: MIT Press.
- McGinn, C. (1989). Can we solve the mind-body problem? *Mind*, 98 (391), 349-366.
- McGinn, C. (2014). Storm over the brain. *The New York Review of Books*.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83 (4), 435-450.
- Papineau, D. (2004). *Thinking about consciousness*. Oxford: Clarendon Press.
- Sacks, M. (1994). Cognitive closure and the limits of understanding. *Ratio*, 7 (1), pp. 26-42.
- Stoljar, D. (2005). Physicalism and phenomenal concepts. *Mind and Language*, 20, 469-494.

¹⁷ It is worth noting that the trilemma in question raises a serious challenge for a popular physicalist strategy to block various arguments for dualism, often called “the Phenomenal Concept Strategy” (PCS), a definitive thrust of which is a rejection of both conceptual eliminativism and conceptual reductionism (see for instance Loar (2004), Papineau (2004), Stoljar (2005), and REDACTED). The challenge for PCS is to show either that the trilemma I have just stated is a false one (i.e., it does not exhaust all the (plausible) options that one might have concerning the argument from grasping) or that, despite appearances, (T) can be adequately accommodated by a sort of physicalism, a doctrine that has enough bite to deserve that title. It appears that in either way, PCS faces a tall order.