

*Can Phenomenal Concepts Explain The Explanatory Gap?*¹

Abstract: The inference from conceivability to possibility has been challenged in numerous ways. One of these ways is the so-called *phenomenal concept strategy*, which has become one of the main strategies against the conceivability argument against physicalism. However, David Chalmers has recently presented a dilemma for the phenomenal concept strategy, and he has argued that no version of the strategy can succeed. In this paper, I examine the dilemma, and I argue that there is a way out of it. I conclude that Chalmers has not posed any serious problem for the phenomenal concept strategy to succeed in blocking the conceivability argument. In doing so, my aim is not only to show that Chalmers' arguments have not refuted the phenomenal concept strategy, but also to clarify what any version of the strategy should achieve in order to be successful.

One of the most important arguments against physicalism is the so-called *conceivability argument*. Intuitively, this argument claims that since certain statements concerning the separation of the physical and the phenomenal are conceivable, they are possible. This inference from conceivability to possibility has been challenged in numerous ways. One of these ways is the so-called *phenomenal concept strategy*, which has become one of the main strategies against the conceivability argument. David Chalmers says it “is perhaps the most attractive option for a physicalist to take in responding to the problem of consciousness”.² Certainly, in the recent years, a multitude of proposals of that sort have been proposed and developed.³ However, Chalmers (2006) has recently argued that no version of the phenomenal concept strategy can succeed. In what follows, I will examine his argument for that conclusion, and I will argue that it is not sound. I will conclude that he has not posed any serious problem for the phenomenal concept strategy to succeed

¹ [Acknowledgments]

² Chalmers 2006: 2.

³ See, for instance, Loar 1999 and 2003, Hill 1997 and Hill & McLaughlin 1999, Papineau 2002, Sturgeon 2000.

in blocking the conceivability argument. In doing so, my aim is not only to show that Chalmers' arguments have not refuted the phenomenal concept strategy, but also to clarify what any version of the strategy should achieve in order to be successful.

1. Conceivability arguments and the phenomenal concept strategy

The intuitive idea behind the conceivability argument is that we can conceive of a physically identical world where there is no consciousness, and therefore, the argument goes, this implies that consciousness is something over and above the physical facts.

A bit more precisely, the conceivability argument can be formulated in this way: Let P be a complete physical description of the actual world. Let Q be any phenomenal truth, for instance the truth that I am seeing something red, or that I am phenomenally conscious. It seems that we can imagine a scenario where P holds but Q does not. For example, we can conceive of a possible world physically identical to our world that lacks consciousness. This strongly suggests that P (the physical description of the world) does not a priori entail phenomenal truths about the world (Q). That is, $P \supset Q$ is not a priori true. In other words, the sentence $P \& \sim Q$ is *conceivable* (at least in the minimal sense that it is *not a priori false*).

Physicalism is committed to the claim that the conditional $P \supset Q$ is necessary. But as we have seen, it is not a priori true, and therefore its negation, namely $P \& \sim Q$, is conceivable. From this, it is inferred that $P \& \sim Q$ is possible, and therefore the conditional is not necessary. If so, physicalism would be false.

The most controversial step in the conceivability argument is the inference from the conceivability of $P \& \sim Q$ to its possibility. It is controversial because the fact

that a sentence or proposition is conceivable does not always imply that such a sentence or proposition is true at some possible world. These are some typical examples that are used to illustrate this problem:

- (i) Water = H₂O
- (ii) Heat = Molecular motion

These sentences are conceivably false, but there is no possible world where they are false. So we seem to have clear counterexamples to the claim that the conceivability of a sentence entails its possibility.

Chalmers (1996) and other advocates of the conceivability argument reply as follows: In cases such as (i) and (ii), when we conceive that those sentences are false, we can observe that there is always a corresponding scenario where what we conceive, or something suitably related, is true. That is, when we take the negations of those sentences, not only do they seem conceivable but they also seem to be *possible*, at least in some sense, which they characterise as follows. Let's take the negation of (i): it seems that there is at least a possible world, namely the Twin Earth, such that if we took it to be our actual world, it would follow that the negation of (i) would be true. This is so because if we were living at the Twin Earth, 'water' would refer to XYZ, not to H₂O, (since the former would be the stuff with the superficial properties of water, that is, XYZ would be the "watery" stuff in that world), and therefore (i) would be false at that world.

Likewise, when we conceive of (ii) being false, we can think of a possible world where the phenomenon that causes sensations of heat is other than molecular motion. If that world was actual, then it would follow that (ii) would be false, because

'heat' would not refer to molecular motion. In other words, (ii) is false at that possible world considered as actual.

The advocates of the conceivability argument, such as Chalmers (1996), think that these examples generalize: they claim that all standard cases of a posteriori necessities satisfy a certain model, namely, it is always the case that when we can conceive of a statement being true, there is a possible world such that, if it was *actual*, the statement would be true in that world.

The next step is to apply that model to the conditional $P \supset Q$. If this conditional is conceivably false, that implies, according to the conceivability argument, that there is some possible world considered as actual where $P \supset Q$ is false. And this would entail that the conditional is not necessary, and therefore physicalism is false.

The phenomenal concept strategy attacks the conceivability argument in the following way. As we have seen, there is an epistemic gap between physical truths and phenomenal truths: P does not entail Q a priori. The phenomenal concept strategy claims that the epistemic gap is not due to any ontological gap, as advocates of the conceivability argument suggest, but rather to the special features of phenomenal concepts.

The phenomenal concept strategy claims that it is because phenomenal concepts have certain key features that physical concepts and phenomenal concepts are not *a priori* connected. This lack of a priori connection means that we can conceive of a given phenomenal concept applying to something without any physical concept applying to it, and *vice versa*. That is, thoughts that consist in the joint application of a phenomenal and physical concept to an individual are always *a posteriori*. And this is exactly what the strategy aims to explain.

For instance, Hill and McLaughlin (1999) argue that the deployment of phenomenal concepts involves a different psychological faculty from that involved in the deployment of physical concepts, and because of this, phenomenal and physical concepts are not a priori connected. According to them, these two different psychological faculties have to do with the different ways in which the reference of physical and phenomenal concepts is fixed. As they say, “if a subject fixes the reference of two concepts in different ways, then it will generally be *a priori* possible to analyze this difference purely in terms of a difference between the psychological states that are involved in the two reference-fixings.” (1999: 453). We can then appeal to those different psychological states in order to explain why physical and phenomenal concepts are not a priori connected, so that we obtain a *psychological* explanation of the lack of a priori connection between physical and phenomenal concepts. The two different reference-fixing mechanisms that Hill & McLaughlin cite are the following:

It is plausible, we maintain, that the reference of the concept of pain is fixed by the fact that subjects have a commitment (or a disposition) to apply the concept to internal states that are experienced directly as having a certain qualitative feel. Further, it is plausible that the reference of (say) the concept of C fiber stimulation is fixed by a stipulation involving a description of the form “the neural process that has such-and-such a structure and that is responsible for such-and-such experimental effects in the actual world.” Under the assumption that the reference of the two concepts in question is fixed in these very different ways, we can account for the fact that it is impossible to see *a priori* that the concepts have the same reference in purely psychological terms. (1999: 453)

Here we have a purely psychological explanation of the fact that we cannot see a priori that a certain phenomenal concept and a certain physical concept are co-

referential. In this way, we can construct general explanations of why physical and phenomenal concepts are not a priori connected.

The general upshot is that phenomenal concepts are not a priori connected with any physical or theoretical concepts, that is, phenomenal concepts are *conceptually isolated* from physical-theoretical concepts.⁴ Likewise, other versions of the phenomenal concept strategy appeal to other features of phenomenal concepts to explain the conceptual isolation of phenomenal concepts.

The task of the phenomenal concept strategy is to provide an alternative explanation of why the conditional $P \supset Q$ is not a priori true (or why $P \& \sim Q$ is conceivable). According to the strategy, the conceptual isolation of phenomenal concepts explains why sentences involving phenomenal concepts (Q) cannot be a priori deduced from sentences involving only physical concepts (such as P). Therefore, $P \supset Q$ cannot be a priori true. If the conditional was true, it would be a posteriori true. The crucial fact about this alternative explanation of the aposteriority of the conditional is that it is independent of the conditional's modal status: the explanation appeals to some features of the phenomenal (and physical) concepts that constitute the conditional, not to whether the conditional itself is necessary or not. Therefore, if that explanation of the a posteriori character of $P \supset Q$ is a plausible alternative, then the conditional's being a posteriori does not entail that it is not necessary. For being a posteriori is compatible with being necessary.

So this is the general structure of the phenomenal concept strategy. In the next section, we turn to Chalmers' argument against the strategy.

2. *A Dilemma for the Phenomenal Concept Strategy*

⁴ The claim is that phenomenal concepts are conceptually isolated from physical concepts, but this is compatible with there being conceptual ties among phenomenal concepts.

According to Chalmers, any successful version of the phenomenal concept strategy should have the following structure:

Proponents put forward a thesis C attributing certain psychological features –call these the key features- to human beings. They argue (1) that C is true: humans actually have the key features; (2) that C explains our epistemic situation with regard to consciousness: C explains why we are confronted with the relevant distinctive epistemic gaps; and (3) that C itself can be explained in physical terms: one can (at least in principle) give a materialistically acceptable explanation of how it is that humans have the key features. (2006: 8)

Chalmers argues that while the three elements (1) to (3) are essential for the strategy to work, no account of phenomenal concepts can satisfy both (2) and (3). That is, he argues that no account of phenomenal concepts (call it C) that an advocate of the phenomenal concept strategy might propose can simultaneously satisfy the three requirements that are essential for the strategy to succeed. In particular, no account of phenomenal concepts can both be explicable in physical terms, and strong enough to explain our epistemic situation with respect to consciousness. Therefore, Chalmers concludes that any version of the phenomenal concept strategy is irredeemably condemned to fail in resisting the inference from the epistemic gap to an ontological gap.

So the main problem with the phenomenal concept strategy, according to Chalmers, is that it cannot simultaneously satisfy requirements (2) and (3). The argument for this conclusion has the form of a dilemma:

- I. If $P \& \sim C$ is conceivable, then C is not physically explicable.
- II. If $P \& \sim C$ is not conceivable, then C cannot explain our epistemic situation.

III. Either C is not physically explicable, or C cannot explain our epistemic situation.⁵

The dilemma starts in the following way: for any account of phenomenal concepts C, either it is entailed a priori by a complete physical description of the world (P) or it is not. That is, either the conditional $P \supset C$ is a priori true or not. If the conditional is a priori true, then $P \& \sim C$ is not conceivable. If the conditional is not a priori true, $P \& \sim C$ will be conceivable.

How can we decide whether $P \& \sim C$ is conceivable or not? In order to establish that, we can try to imagine some physical duplicate of the world, that is, a world that satisfies P, which does not satisfy the strategy's account of phenomenal concepts, that is, C. If we can imagine such a world, then P does not entail C a priori, and therefore $P \& \sim C$ would be conceivable. In particular, we can focus on zombie-worlds: physically identical duplicates of this world where no-one is conscious. According to the proponent of the phenomenal concept strategy, zombie-worlds are conceivable (even if they are not possible). Zombie-worlds would clearly satisfy P. Then, we can ask ourselves whether zombie-worlds would also satisfy C. If they do not, then $P \& \sim C$ is conceivable. If, on the other hand, we cannot imagine a physically identical world where C does not hold, not even a zombie-world, then $P \& \sim C$ is not conceivable. These two options represent the two horns of Chalmers' dilemma, that is, the two premises of his argument. I will examine these two premises in turn.

The *first premise* explores what happens if $P \& \sim C$ is conceivable. According to Chalmers, the problem for the phenomenal concept strategy in this case is that P does not entail C a priori, and this has the consequence that we cannot *explain* that account of phenomenal concepts in physical terms, in Chalmers' sense of explanation.

⁵ Chalmers 2006: 11.

According to Chalmers, a reductive explanation of X in physical terms requires that P entails X a priori. Therefore, P has to entail C a priori, if we want a reductive explanation of our account of phenomenal concepts in physical terms.

Chalmers has argued that in order to explain some macro-phenomenon in microphysical terms, we have to show how microphysical facts give rise to that macro-phenomenon, and this requires that microphysical facts transparently show how the macro-phenomenon obtains. If microphysical facts did not entail a priori the instantiation of the macro-phenomenon, then we would require some additional explanatory materials in order to explain such an instantiation, and this would show that the original explanation was incomplete.⁶

Someone might wonder why it is necessary for the phenomenal concept strategy to offer a reductive explanation of their proposed account of phenomenal concepts in physical terms. According to Chalmers, in order for the phenomenal concept strategy to work, it has to show that physicalism is compatible with the epistemic gaps, and therefore, its proponents have to show how physical facts alone give rise to the key features of phenomenal concepts which in turn give rise to the relevant epistemic gaps. If they do not provide a physically acceptable explanation of the key features of phenomenal concepts, then they will not have motivated the view that their explanation of the epistemic gap is indeed compatible with physicalism.

The *second premise* of the argument examines the possibility that P&~C is *not* conceivable, that is, that a complete physical description of the world (P) a priori entails the corresponding account of phenomenal concepts C. In this case, even zombie-worlds satisfy C: every conceivable physical duplicate of the actual world

⁶ For the purposes of this paper, and for the sake of discussion, I will assume that reductive explanation of X in physical terms requires that P entails X a priori. This assumption is controversial, and for the record, I do not really accept it, but I will grant it here. For some doubts concerning this account of reductive explanation, see Block & Stalnaker 1999.

satisfies C. This horn of the dilemma represents the possibility that we are really capable of offering a physically acceptable explanation of the key features of phenomenal concepts (according to Chalmers' standards). So we would not have the problem of the previous horn.

Unfortunately, Chalmers argues, this second horn of the dilemma faces another problem: if our account of phenomenal concepts is tame enough to be physically explicable (in Chalmers' sense), then it could not explain our epistemic gap. Let me elaborate why this is so, according to Chalmers, in some detail.

If P entails C a priori, then even zombies satisfy C. But Chalmers argues (and this is the crucial point) that zombies do *not share our epistemic situation* with regards to consciousness. If so, then C does not entail a priori that someone is in our epistemic situation with regard to consciousness, because we can conceive of some beings (zombies) that satisfy C but do not share our epistemic situation. And therefore, C cannot explain our epistemic situation (again, assuming Chalmers' account of explanation as a priori entailment).

Why does Chalmers think that zombies would not share our epistemic situation? Well, zombies seem to have a very different epistemic position with regards to consciousness from ours. We are conscious, but they are not. We have true beliefs about our own phenomenal states, but they do not. We have knowledge of phenomenal facts, but they do not. Our epistemic position is very different from zombies' epistemic position. Chalmers characterizes someone's epistemic situation as including "the truth values of their beliefs and the epistemic status of their beliefs (as justified or unjustified, and as cognitively significant or insignificant)" (Chalmers 2006: 14). He assumes we can draw a correspondence between our beliefs and zombies' beliefs. Then, he adds: "A zombie will share the epistemic situation of a

conscious being if the zombie and the conscious being have corresponding beliefs, all of which have corresponding truth-value and epistemic status” (14). So in order to evaluate whether zombies share our epistemic situation we have to compare the truth-values and epistemic status of our beliefs and their corresponding beliefs. If there is any difference, our epistemic positions differ.

Chalmers considers some of zombies’ beliefs such as ‘I am phenomenally conscious’. He suggests that this belief does not seem to be true of zombies. This could seem a very trivial point, since zombies are not conscious by definition. But actually, this issue turns out to be more complicated, since the truth-values of such zombies’ corresponding beliefs depend on what their contents are, and it is not clear what their contents could be. Since zombies lack consciousness, how could their corresponding beliefs be about consciousness?

Chalmers realizes that judgements about the truth-value of zombies’ corresponding “phenomenal” beliefs are controversial, and therefore he does not want to rely on them. So he focuses on the issue of the justification of zombies’ corresponding “phenomenal” beliefs. In order to do this, he uses Jackson’s example of Mary (the colour-scientist who had spent all her life in a black-and-white room, until one day she was released and was able to see a red object for the first time).⁷ He also introduces a new character, *Zombie-Mary*, that is, a physical duplicate of Mary who does not have any conscious mental state. Chalmers asks us to consider the following two utterances, by Mary and *Zombie-Mary* respectively:

Mary (after being released): ‘I am having a red sensation.’

Zombie-Mary (after being released): ‘I am having a red sensation.’

⁷ See Jackson 1982.

Chalmers argues that the beliefs expressed by these utterances are not equally justified. Mary's judgement is justified by her own phenomenal states, and arguably she gains new substantive phenomenal knowledge by having that sensation for the first time. Zombie-Mary's judgement does not have the same kind of justification, since she is not enjoying any kind of phenomenal state. Maybe she is gaining some new sort of knowledge (indexical, for instance), but she is not gaining new *phenomenal* knowledge. So, according to Chalmers, zombies' beliefs have a different justification status, and therefore, they do not share our epistemic situation, even if they satisfy C. Hence, C cannot explain our epistemic situation.

Chalmers concludes that for any account of phenomenal concepts, either it cannot be physically explicable (first horn of the dilemma) or it cannot explain our epistemic situation with regards to consciousness (second horn). This, then, is Chalmers' new response to the phenomenal concept strategy. In the remainder of the paper, I will argue that it fails.

3. The second horn of the dilemma: explaining the epistemic gap

3.1. Epistemic situation vs. epistemic gap

In my response, I will focus on the second horn of the dilemma. I think that, even if we agree with Chalmers' argument that zombies' "phenomenal" beliefs have a different justificatory status from ours, Chalmers' second horn of the dilemma does not pose a serious challenge for the phenomenal concept strategy. The problem is that Chalmers has misidentified the explanatory goal of the phenomenal concept strategy. On my view, what an account of phenomenal concepts has to explain is simply the

epistemic gap between physical and phenomenal truths, not our *entire* epistemic situation with regards to consciousness. And as I will argue, Chalmers has presented no obstacle to the idea that an account of phenomenal concepts might explain the epistemic gap in that sense.⁸

I think it is clear, from the presentation of the conceivability argument and the phenomenal concept strategy in section 1, that what C should explain is the *epistemic gap* between the physical and the phenomenal, because it is precisely this epistemic gap which is supposed to imply an ontological gap. The aim of the phenomenal concept strategy is to show that there is an alternative explanation of the epistemic gap in terms of certain features of phenomenal concepts, an explanation that does not entail that there is an ontological gap. The strategy is not supposed to explain *our full epistemic access* to consciousness.

The question at issue, then, is the following: how should we characterize the epistemic gap between physical and phenomenal truths that the phenomenal concept strategy aims to explain? In section 3.2, I will introduce the characterization of the epistemic gap that I favour, and I will argue that zombies do not pose a problem to the task of explaining the epistemic gap in that sense. Secondly (in 3.3) I will explain Chalmers' doubts about characterizing the epistemic gap in that way, and I will present his alternative understanding of the epistemic gap to be explained by the strategy. I will then argue that an advocate of phenomenal concept strategy has no good reason to accept Chalmers' characterization.

⁸ One might also challenge Chalmers' assumption that zombies "phenomenal" beliefs have a different justificatory status from ours. For instance, Carruthers and Veillet (forthcoming) defend the view that zombies do share our epistemic situation. They argue that zombies' corresponding "phenomenal" beliefs can be both true and justified, they just have a different content: they are about zombies' quasi-phenomenal states. However, I am not going to base my defence of the phenomenal concept strategy on the view that zombies' "phenomenal" beliefs are true and justified. Instead, I will focus on the view that what the strategy has to explain is the epistemic gap between physical and phenomenal beliefs, which can be characterised independently of those phenomenal beliefs' truth and justification. Nevertheless, if it did turn out to be the case that zombies' corresponding beliefs are true and justified, that would be compatible with my own arguments against Chalmers' dilemma.

3.2. *The epistemic gap, properly understood*

My suggestion is that what C has to explain is the fact that P (the complete physical description of the world) does not entail Q (any phenomenal truth) a priori. I think that the phenomenal concept strategy does not have to explain our whole epistemic situation but just why there is an *inferential disconnection* between physical truths and phenomenal truths, that is, why the latter are not a priori inferable from the former.

So the relevant question, then, if we want to know whether an account of phenomenal concepts C entails a priori our epistemic gap, is the following: can we imagine beings that satisfy C but fail to satisfy the inferential disconnection between P and Q? But, what is it for a being to *fail* to satisfy that inferential disconnection? Well, since in order to satisfy the inferential disconnection, a subject has to be such that she is *not* able to infer Q a priori from P, it seems that she will fail to instantiate such a disconnection when, in effect, she is *able* to infer Q a priori from P.

Chalmers argues that zombies do not share our epistemic gap. If we understand the epistemic gap as an inferential disconnection between P and Q, then the question at stake would be this: can we imagine zombies that satisfy C *and* are able to infer Q a priori from P?

We can easily see that there cannot be such zombies. The reason is that zombies cannot even entertain Q, since they do not possess the concepts required to understand Q, namely, phenomenal concepts. They do not have phenomenal states, which are essential in order to possess (most of) our phenomenal concepts, so they cannot have phenomenal beliefs (with the same content as ours). Chalmers seems to agree with this point: “It is plausible that a nonconscious being such as a zombie cannot have beliefs with exactly the same content as our beliefs about consciousness”

(2006: 14-15). They can have corresponding beliefs but with different content: in the case of our phenomenal beliefs, zombies can have corresponding *quasi-phenomenal* beliefs (about their corresponding quasi-phenomenal states). Then, it follows that zombies cannot entertain Q, and therefore, zombies are *not able* to infer Q from P a priori. For in order to infer Q a priori from P, first you need to entertain Q. So zombies cannot infer Q a priori from P. That is to say, they *do satisfy* the inferential disconnection between P and Q. Hence, they do satisfy the epistemic gap in the sense I am advocating. Therefore, there is no reason to think that C does not entail a priori the epistemic gap.

Someone could reply that what is at issue here is rather whether zombies would be able to infer a priori their *corresponding* “phenomenal” belief Q* from their corresponding belief P. So let’s examine this question: can zombies a priori infer Q* from P? I think that they cannot. The reason is the following. Recall that we are considering the second horn of Chalmers’ dilemma, according to which P&~C is not conceivable, and therefore even zombies satisfy C. Then, zombies’ corresponding phenomenal concepts satisfy the account of phenomenal concepts C, so those phenomenal concepts are not a priori connected to physical concepts. So P will not entail Q* a priori.

We can see this point a bit more clearly by focusing on a particular account of phenomenal concepts. For instance, according to Hill & McLaughlin, phenomenal concepts and physical concepts play very different psychological roles, and this is what explains the lack of a priori connection. We could characterise these psychological roles in purely *functional* terms, and therefore a zombie (that is, a functional duplicate of us) would also have concepts that played those different roles. So we could talk about zombies’ corresponding quasi-phenomenal concepts (those

concepts that are functionally equivalent to our phenomenal concepts) and zombies' corresponding physical concepts. Since these zombie-concepts also play different roles, they will not be a priori connected, and therefore, sentences involving quasi-phenomenal concepts cannot be a priori inferred from sentences involving only physical concepts.⁹

Therefore, we can conclude that, if we understand the epistemic gap as an inferential disconnection between physical and phenomenal beliefs, then there is no evidence that C might hold without the epistemic gap holding. In particular, Chalmers' alleged counterexample, namely, zombies, is not a real case of beings that satisfy C but not the epistemic gap, since there is no *relevant* epistemic gap that they fail to satisfy, even if they do not instantiate all the aspects of our epistemic situation. Hence, zombies do not represent a problem for the claim that C can explain the epistemic gap.

3.3. Chalmers' reply: the knowledge-involving epistemic gap

Chalmers seems to accept that there is some kind of inferential disconnection that is also present in the zombies' case, but he claims that to explain the inferential disconnection in that sense is not enough for explaining our epistemic gap. He agrees that zombies are not able to entertain Q, which entails that they are not able to infer Q from P a priori, and he also agrees that they are not able to a priori infer Q* from P either: "it is plausible that a zombie's physical and quasi-phenomenal beliefs are no more inferentially connected than a conscious being's beliefs" (2006: 24). Still, he

⁹ Different versions of the phenomenal concept strategy explain this lack of a priori connection in different ways, but they all agree that phenomenal concepts are such that the relevant sentences involving them are not a priori inferable from sentences involving physical concepts only. See, for instance, Sturgeon 2000 or Loar 1999, 2003.

thinks that zombies fail to instantiate the *relevant* aspects of *our* epistemic gap. Let's see why:

Whereas the inferential disconnection strategy might physically explain an inferential disconnection between physical and phenomenal *beliefs*, the anti-physicalist's crucial epistemic gap involves a disconnection between physical and phenomenal *knowledge*. (2006: 24)

In the anti-physicalist's arguments, the relevant epistemic gap (from which an ontological gap is inferred) is characterised in a way that truth and knowledge are essential. [...] It is crucial to the conceivability argument that one can conceive beings that lack phenomenal states that one actually has. And it is crucial to the explanatory gap that one has cognitively significant knowledge of the states that we cannot explain. (2006: 23)

What Chalmers is claiming at this point is that the inferential disconnection that the phenomenal concept strategy should explain is a disconnection between physical truths and phenomenal beliefs that are *true and justified*. In other words, C has to explain how we can conceive of beings that are physically identical to us but lack the phenomenal states of which *we have* substantive knowledge. For Chalmers, I take it, a subject has substantive knowledge of phenomenal states if and only if she has phenomenal states. Therefore, Chalmers is claiming here that an account of phenomenal concepts has to explain the epistemic gap in the following sense: a subject will instantiate the epistemic gap when *she has phenomenal states* (so that she can have justified true beliefs about them) and in addition she cannot infer a priori those beliefs about her own phenomenal states from a physical description of the world.

However, I want to argue that C merely has to entail a weaker version of the epistemic gap, where the crucial bit is that there is an inferential disconnection

between physical and phenomenal (or quasi-phenomenal) beliefs. It is not necessary that subjects actually have phenomenal states (nor believe that they have them, nor know that they have them) in order for them to instantiate the epistemic gap.

I think that the demand for an explanation of the epistemic gap in the stronger sense is not motivated. In section 3.3.2, I will explore the argument that Chalmers presents for the claim that the phenomenal concept strategy should explain the epistemic gap in the strong sense, and I will argue that it does not provide sufficient motivation. Before that, in section 3.3.1, I will argue that there is a more straightforward problem with Chalmers' claim that the phenomenal concept strategy should explain the epistemic gap in the strong sense: this is clearly too strong, because if we accept Chalmers' characterization of the epistemic gap, then it will follow that in order for C to explain the epistemic gap, the epistemic gap must be closed. Firstly I will explain why Chalmers' characterization has that consequence and secondly I will explain why it is a problematic consequence.

3.3.1. Closing the epistemic gap?

Recall that Q can be any phenomenal truth, so let Q be a proposition of the form 'I am phenomenally conscious'. And recall that Chalmers' truth-and-knowledge-involving epistemic gap requires that any subject experiencing this epistemic gap is herself conscious, that is, her corresponding Q has to be true. Then, we can characterise this epistemic gap as follows:

Chalmers' Epistemic Gap (E): Q & It is not a priori that $(P \supset Q)$.

I take it that this characterization is a priori true. For the epistemic gap has been *defined* as crucially involving phenomenal knowledge, and, plausibly, the having of phenomenal knowledge a priori entails the having of phenomenal states. That is, E entails Q a priori. And remember that, on this horn of the dilemma, C is a priori entailed by P. So, if we assume that C a priori entails E, this will follow:

P a priori entails C

C a priori entails E

E a priori entails Q

Therefore, *P a priori entails Q*

This would mean that there is no epistemic gap between P and Q. That is, if we characterize the epistemic gap as E, then, on this horn of the dilemma, in order for an account of phenomenal concepts to explain the epistemic gap, there cannot be an epistemic gap!

This is problematic, for the following reason. Chalmers himself claims, elsewhere in the paper, that we should characterize both C (the account of the key psychological features of phenomenal concepts) and the *epistemic situation* to be explained by C in “*topic-neutral* terms: terms that do not explicitly attribute phenomenal states or concepts that refer to them. [...] This allows the possibility that even if consciousness cannot be physically explained, we might be able to physically explain the key psychological features and our epistemic situation” (2006: 12).

Well, I think that it is clear that Chalmers’ characterization of the epistemic gap does not allow the possibility that even if consciousness cannot be physically

explained, the epistemic gap could still be explained. As we have seen above, if we characterize the epistemic gap as E, then for C to entail E a priori, it has to be the case that P entails Q a priori. But if P does not entail Q a priori (that is, if consciousness cannot be physically explained) then C cannot entail E a priori either, and therefore C cannot explain E (in Chalmers' sense of explanation). So Chalmers' characterization of the epistemic gap has the consequence that, if consciousness is not physically explainable, the epistemic gap will not be physically explainable either. In other words: Chalmers' characterization of E is not really *topic-neutral*, because it involves phenomenal states, as we can see above. Therefore, it is a bad characterization.¹⁰

My conclusion is that we should reject Chalmers' characterization of the epistemic gap as E, and we should endorse a characterization that really is topic-neutral. My characterization of the epistemic gap as an inferential disconnection between physical and phenomenal beliefs does satisfy that constraint, and therefore there are good reasons to prefer it. And as I have shown earlier, versions of the phenomenal concept strategy that take the second horn of the dilemma, that is, that

¹⁰ In discussion, Chalmers has suggested that, on his view, it is not true *by definition* that satisfying the epistemic gap requires having phenomenal states: as a matter of fact, the epistemic gap requires having phenomenal states, but this is not part of the definition of epistemic gap. That is, he denies that E entails Q a priori. However, he says in his paper that "In the anti-physicalist's arguments, the relevant epistemic gap (from which an ontological gap is inferred) is characterised in a way that truth and knowledge are essential. [...] It is crucial to the conceivability argument that one can conceive beings that lack phenomenal states that one actually has." (2006: 23) Here, the epistemic gap does seem to require having *phenomenal states* (by definition), and therefore, this is not a topic-neutral characterization of the epistemic gap. More importantly, though, if we drop the requirement that E entails Q a priori, then it is not clear what Chalmers' response to my argument in section 3.2. would be. If the epistemic gap just involves the inferential disconnection between P and Q, and not the having of phenomenal states *per se*, then zombies can satisfy the epistemic gap too. Chalmers' argument requires a more substantial characterization of the epistemic gap, such that zombies cannot instantiate the epistemic gap in that sense, but it is topic-neutral. It is not clear what this notion could be. Chalmers' suggestion seems to be to characterise the epistemic gap as the inferential disconnection between physical and phenomenal (or pseudo-phenomenal) beliefs, where the latter are *true and justified*, but where this notion of justification does not involve having phenomenal states by definition, although non-conscious beings cannot, as a matter of fact, have that kind of justification. In response, I think that, first, it is not clear to me whether such a notion is coherent, and secondly, even if it was, it is not clear what the motivations for preferring it over my own characterization of the epistemic gap are. For more on this second point, see section 3.3.2.

offer accounts of phenomenal concepts that are physically explicable, do not have problems in explaining the epistemic gap in the preferred sense.

3.3.2. *The conceivability argument and the epistemic gap*

Given these difficulties, why does Chalmers hold that we have to explain the epistemic gap in the strongest sense, that is, one that involves the presence of phenomenal states themselves? Chalmers says that the epistemic gap in that sense (the truth-and-knowledge-involving epistemic gap) is the one which is supposed to entail an ontological gap, according to the conceivability argument. Therefore, he argues, it is the epistemic gap in that sense that the phenomenal concept strategy has to explain. He adds: “If one characterised these gaps in a way that were neutral on the truth of phenomenal beliefs, the arguments would not get off the ground” (2006: 14). Why should we believe that the truth of our phenomenal beliefs is crucial for getting the conceivability argument off the ground? I think that this has to do with the fact that the conceivability argument is a valid argument against physicalism only if at least some phenomenal beliefs are true. Let’s review the argument, in order to see why:

- (a) It is conceivable that $P \& \sim Q$.
- (b) If it is conceivable that $P \& \sim Q$, then it is possible that $P \& \sim Q$.
- (c) If it is possible that $P \& \sim Q$, then physicalism is false.
- (d) So physicalism is false.¹¹

As Chalmers (forthcoming) argues, premise (c) is plausible only if there is some phenomenal belief Q that is true of the actual world. That is, physicalism is committed

¹¹ This formulation of the argument draws on Chalmers (forthcoming).

to the claim that any physical duplicate of the world is such that all *truths* about the world are true there as well. If there were no phenomenal truths, then the possibility of $P \& \sim Q$ would not be a problem for physicalism at all.

So we can agree that in order to get the conceivability argument *against physicalism* off the ground, Q has to be true. However, this does not force us to maintain that the *epistemic gap* (premise a) requires the truth of some phenomenal belief Q. Recall that the phenomenal concept strategy's aim is to block the inference from the conceivability of $P \& \sim Q$ to its possibility, that is, they want to refute (b). And as I have explained, the premise that requires the truth of some phenomenal belief Q is (c). That is, the epistemic gap can be formulated independently of whether Q is true or not. Of course, the *argument* itself can get off the ground only if Q is true. But this is irrelevant for the phenomenal concept strategy, whose main aim is to establish that there is an alternative explanation of the epistemic gap that does not have the consequences that (b) asserts. This task is completely independent of whether (c) is true or not. Therefore, there is no good reason here to believe that we should characterise the epistemic gap in a sense that involves the truth of Q.

4. Conclusion

Chalmers has presented a dilemma for the phenomenal concept strategy. He argues that no matter what horn is taken, there will be serious problems. I have replied that, on the contrary, an advocate of the phenomenal concept strategy could take the second horn of the dilemma, and still have a successful response to conceivability arguments. Of course, if she accepts the existence of an explanatory gap between the physical and the phenomenal, then she cannot provide a physical explanation (in

Chalmers' sense) of the epistemic situation that *conscious beings* are in. But I have argued that this is not relevant at all for the success of the strategy. For it is enough, I believe, that the strategy presents an account of phenomenal concepts which explains why physical truths do not entail phenomenal truths a priori; the account does not have to explain how we can *know* phenomenal truths. Therefore, I think that there is at least one way out of Chalmers' dilemma. So the phenomenal concept strategy is very much alive and well.

Word-count: 7,111

References:

Block, N. and Stalnaker, R. 1999. "Conceptual analysis, dualism, and the explanatory gap". *Philosophical Review* 108: 1-46.

Carruthers, P. and Veillet, B. (forthcoming). "The phenomenal concept strategy". Available at <http://www.philosophy.umd.edu/Faculty/pcarruthers/Phenomenal%20Concept%20Strategy.htm>

Chalmers, D. 1996. *The Conscious Mind*, Oxford University Press.

Chalmers, D. 2006. "Phenomenal concepts and the explanatory gap". Forthcoming in T. Alter & S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge: Essays on Consciousness and Physicalism*, Oxford University Press, December 2006. Now available at:

<http://consc.net/papers/pceg.html>

Chalmers, D. (forthcoming). "The two-dimensional argument against materialism". In D. Chalmers (ed.) *The Character of Consciousness*, Oxford University Press. Now available at:

<http://consc.net/papers/2dargument.html>

Hill, C. 1997. "Imaginability, conceivability, possibility, and the mind-body problem". *Philosophical Studies* 87: 61-85.

Hill, C. & McLaughlin, B. 1999. "There are fewer things in reality that are dreamt of in Chalmers' philosophy". *Philosophy and Phenomenological Research* 59: 445-54.

Jackson, F. 1982. Epiphenomenal qualia. *Philosophical Quarterly* 32: 127-36.

Loar, B. 1999. On David Chalmers' *The Conscious Mind*. *Philosophy and Phenomenological Research* 59: 465-72.

Loar, B. 2003. Qualia, properties, modality. *Philosophical Issues* 13 (*Philosophy of Mind*): 113-29.

Papineau, D. 2002. *Thinking about Consciousness*, Oxford University Press.

Papineau, D. 2006. Phenomenal and perceptual concepts. Forthcoming in T. Alter & S. Walter (eds.) *Phenomenal Concepts and Phenomenal Knowledge: Essays on Consciousness and Physicalism*, Oxford University Press, December 2006. Now available at: <http://www.kcl.ac.uk/ip/davidpapineau/Staff/Papineau/OnlinePapers/Phen%20and%20Perceptual%20Concepts.htm>

Sturgeon, S. 2000. *Matters of Mind*, Routledge.