



On Bayne and Chalmers' Phenomenal Unity Thesis (or: Much Ado about Nothing)

Guus Duindam¹ 

Received: 14 November 2018 / Revised: 11 April 2019 / Accepted: 28 July 2019 /

Published online: 19 August 2019

© Springer Nature B.V. 2019

Abstract

According to the Phenomenal Unity Thesis (“PUT”) – most prominently defended by Tim Bayne and David Chalmers – necessarily, any set of phenomenal states of a subject at a time is phenomenally unified. The standard formulation of this thesis is unacceptably vague because it does not specify what it is to be a subject. In this paper, I first consider possible meanings for ‘subject’ as used in PUT and argue that every plausible candidate definition renders the thesis trivially true. I consider and reject Tim Bayne’s proposal that ‘subject’ means ‘human being’. Then I argue, contra Bayne and Chalmers, that PUT is not incompatible with any major theory of consciousness, and contra Michael Tye, that split-brain patients do not provide evidence against PUT. I close by considering some nontrivial alternatives to PUT.

Keywords Phenomenal Unity thesis · Phenomenal consciousness · David Chalmers · Tim Bayne · Split-brain patients · Phenomenal unity · Subject

1 Introduction

Multiple phenomenal experiences characterize what it is like for me to write this introduction. For instance, I have a visual experience of the screens of my computer; an auditory experience of the music playing through my speakers (Beecham’s recording of *La Bohème*), a tactile experience of my cat sitting on my lap, etc. It seems to me that these experiences are *unified*, for there is something it is like to be *me* as I have all of them. According to an influential formulation of the **phenomenal unity thesis** (“PUT”), such unity is a necessary feature of all phenomenal consciousness. Thus, according to PUT, “*necessarily*, any set of phenomenal states of a subject at a time is phenomenally unified (Bayne and Chalmers 2010).” Tim Bayne and David Chalmers

✉ Guus Duindam
gduindam@umich.edu

¹ Department of Philosophy, University of Michigan, Ann Arbor, MI, USA

defend this thesis and claim that it is nontrivial, pointing out as evidence for this that philosophers exist who disagree with it.¹ It is true that several philosophers have denied that unity is a necessary feature of phenomenal consciousness.² For instance, Michael Tye claims that phenomenal unity breaks down in split-brain patients.³ These disagreements notwithstanding, I argue here that PUT is trivially true.

The phenomenal unity thesis in its standard formulation is ambiguous, for it is never made clear what it is to be a ‘subject’. Hence it is not clear what it would take for PUT to be true. I begin by considering possible clarifications of PUT. In sections 2 and 3 I argue that all plausible definitions of ‘subject’ as used in PUT render that thesis analytically true. Then, in section 4, I discuss Tim Bayne’s recent proposal that ‘subject’ is equivalent to ‘human being’. In section 5, I show *contra* Bayne and Chalmers that PUT is not incompatible with any major theory of consciousness, and *contra* Tye that split-brain patients do not show that PUT is false. Given PUT’s analyticity, these results ought to be unsurprising. Section 6 concludes by suggesting some alternative theses about the unity of consciousness that, unlike PUT, would be both substantive and interesting. I argue that we should shift our focus accordingly.

2 The Phenomenal Unity Thesis

According to Bayne and Chalmers, phenomenal unity is a necessary feature of phenomenal consciousness. They defend:

Phenomenal Unity Thesis Necessarily, any set of phenomenal states of a subject at a time is phenomenally unified.⁴

Some mental state is *phenomenally conscious* just if “there is something it is like to be in that state.”⁵ Some *set* of mental states is *phenomenally unified* just if there is something it’s like to be in the *set*, that is, just if there’s something it’s like to be in the conjunctive state which contains all the phenomenal states in the set.⁶ At the risk of stating the obvious, note that it is incompatible with PUT that there be *more than one* thing-it’s-like to be in the conjunction of all of one’s phenomenally conscious states. For suppose that there are exactly two distinct things-it’s-like to be in some set *s* of phenomenal states: p_1 and p_2 . Then either there is *one* thing-it’s-like to experience p_1 and p_2 at once – that is, to experience $(p_1 \& p_2)$ – or there is not. If there *is* one thing it’s like to experience $(p_1 \& p_2)$ then there was after all one unique thing-it’s-like to experience *s*, namely, $(p_1 \& p_2)$. If, on the contrary, there’s nothing it’s like to experience $(p_1 \& p_2)$ then *s* is not unified. If some set of states is phenomenally unified there can be only *one* thing it’s like to experience that set at some time, i.e., the thing-it’s-like to experience the conjunction of all the phenomenal states in the set.⁷ Thus PUT is equivalent to:

¹ Ibid., 510. “[I]t is nontrivial that there will be *something it is like* to be in the conjunctive state. This can be seen from the fact that some philosophers deny the total unity thesis, or at least entertain its denial.”

² See, e.g., Dennett 1992, Tye 2003.

³ Tye 2003, 126–29.

⁴ Bayne and Chalmers 2010, 510–11.

⁵ Ibid., 509.

⁶ Ibid.

⁷ Bayne and Chalmers themselves confirm this at Ibid., 503.

PUT (2) Necessarily, for any subject at a time, for any set of states-there’s-something-it’s-like-to-have, there is one thing it’s like to be in the conjunction of those states.

We have now reformulated the phenomenal unity thesis to specify both what it means to be a phenomenal state, and what it means for such states to be unified. But its meaning remains vague, since it is not clear what it is to be a subject. Remarkably, Bayne and Chalmers fail to define ‘subject’ even when they discuss the possibility of “subject unity.”⁸ They are not the only philosophers of mind who have left the meaning of ‘subject’ ambiguous while discussing subject unity. Indeed, the term is unfortunately frequently used without explanation.⁹ Until it is explained, however, it is not at all clear what it would mean for PUT to be true.

Since Bayne and Chalmers do not provide a definition of ‘subject’, it is worth seeing whether any common-language meaning of the term fits. The *Oxford English Dictionary* lists 17 distinct uses of the noun, but only one could possibly apply:

A being (or power) that thinks, knows, or perceives (more fully conscious subject, thinking subject); the conscious mind, esp. as opposed to any objects external to it. In later use also more broadly: the person or self considered as a conscious agent.¹⁰

‘Subject’ as used in PUT cannot plausibly be read to mean ‘person’, for (1) it seems that only humans are (uncontroversial) persons, while PUT allegedly asserts a truth about *all* conscious beings, not merely humans (Bayne and Chalmers 2010), (2) personhood does not seem to be a necessary condition for consciousness (or at least, it is not so by definition) and (3) most philosophers in the literature use ‘subject’ as explicitly distinct from ‘person’, for instance, when they discuss the possibility of a split-brain ‘*subject*’ being two distinct *persons* (Bayne 2004). Using ‘subject’ to refer to ‘a being that thinks’ or ‘a being that knows or perceives’ likewise definitionally limits the class of potentially conscious beings; it seems at least conceivable that some being could be conscious without thinking, much less knowing or perceiving. If PUT is to be a substantive thesis about the nature of consciousness in general, rather than a contingent one about a particular class of conscious things, then its formulation must be understood as applicable to all phenomenally conscious things. Adopting a limiting definition of ‘subject’ is just incompatible with the aim of PUT’s defenders. It seems most plausible, then, to read ‘subject’ to mean the neutral ‘conscious mind’ or ‘being’.

Just what does it mean to be a conscious mind? Since we are discussing a thesis about the nature of phenomenal consciousness, presumably we mean *phenomenally* conscious mind.¹¹ A mind is phenomenally conscious only when there is something it’s like to be that mind. Again, we can specify that there can be only *one* thing it is like to be that mind at some time *t*. After all, if at any time there were two things-it’s-like to be

⁸ *Ibid.*, 510.

⁹ See, e.g., Rosenthal 2000, 2003, Bayne and Chalmers 2010, 503; Tye 2003, 12–13.

¹⁰ *Oxford English Dictionary*, s.v. “subject” Other definitions are obviously inapplicable since they refer to ‘subject’ in its legal (e.g. citizen), linguistic (e.g. the thing of which a property is predicated), or topical (the matter of thought or inquiry) senses.

¹¹ Indeed, Bayne later claims that “phenomenal consciousness” is pleonastic because *all* consciousness is phenomenal. See Bayne 2010, 7.

some ‘mind’, and nothing it’s-like to be the conjunction of those two things, then we are dealing with *two* distinct minds. Unless we legislate our favorite solution to the mind-body problem by building it into the *definition* of ‘mind’ (e.g. by defining ‘mind’ as ‘brain’) we have *no* nonarbitrary means of individuating minds beyond asking whether there is *one and only one thing-it’s like* to be that mind at a time. Likewise for ‘conscious being’ or any other plausible definition of ‘subject’. As Elizabeth Schechter notes in her criticism of Bayne’s pluralist attitude to the individuation of streams of consciousness:

There is a concept of the subject of experience...that is *so closely tied* to the concept of a stream of consciousness, that any account that says how many streams of consciousness there are, appears to equally constitute an account of how many subjects of experience there are (Schechter 2013).

We individuate subjects by individuating unified streams of consciousness, or unified sets of phenomenal states. And this makes clear that the phenomenal unity thesis in the formulation under consideration expresses a merely analytic truth. The phenomenal unity thesis means nothing over and above:

PUT (3) Necessarily, for any unified stream of consciousness at time t , for any set of states-there’s-something-it’s-like-to-have, there is one thing it’s like to be in the conjunction of those states.

which is equivalent to the more obviously trivial:

PUT (4) Necessarily, any unified stream of consciousness at time t is unified.

Something x is a distinct subject (conscious mind, being, stream of consciousness) at time t just when there is *one and only one* thing it’s like to be x at t . If, at t , x has multiple phenomenal experiences, and x is a subject, then the conclusion that there is one thing it’s like to have the *set* of x ’s phenomenal experiences is strictly entailed; denying it forces us to deny that x is a (one distinct) subject. If, at t , x has only one phenomenal experience, and x is a conscious mind, then trivially there is one thing-it’s-like to be in the conjunction of all of x ’s experiences. Similarly if x experiences nothing at t . PUT is not a substantive claim about consciousness; it is trivially true.

At the end of their paper, Bayne and Chalmers briefly discuss a conception of ‘subject’ like the one now under discussion. Considering the possibility that PUT is “a deep conceptual truth,” they note that “our notion of a subject of experience” could “somehow [be] premised on unity.”¹² To this proposal they object that “it cannot explain why states of consciousness come packaged into unified phenomenal fields in the first place.”¹³ That, it seems to me, is true. It shows that this is not a helpful way to explain the deep truth Bayne and Chalmers believe to be expressed by PUT. But the reason for this is not that ‘unified stream of consciousness’ is an unacceptable definition for ‘subject’. If – as I have suggested – PUT does not express *any* ‘deep truth’, then *a fortiori* none could be discoverable by conceptual analysis. This is not to say there are no interesting and meaningful questions in the neighborhood. To the contrary, as

¹² Bayne and Chalmers 2010, 537.

¹³ *Ibid.*, 538.

Bayne and Chalmers suggest, we might well wonder why there *are* unified phenomenal fields (or, to put the same in other words: why there are *subjects*). But, again, since we are here merely attempting to clarify what is *meant* by PUT (that is, what it would *take* for it to be true), it is no argument against such a clarification to say that it does not answer these questions.

3 Intransitive Unification

Using common-language definitions of ‘subject’ renders PUT analytically true, then, and no technical definition of the term as-applied is provided by Bayne and Chalmers. There is another way, however, in which we might go about attempting to reconstruct PUT as a substantive thesis about consciousness. This is to imagine a set of connected phenomenal states for which PUT would be false and try to reverse-engineer a concept of ‘subject’ such that it would be compatible with this scenario. This does not seem to me a particularly promising way to construct a definition of ‘subject’ that is attractive independently of its ability to rescue PUT. Setting this aside, however, I argue that it can anyway not help us settle on a definition of PUT that would render the claim substantive.

We can imagine a network of phenomenal states that are only intransitively unified.¹⁴ Suppose, for instance, that network *N* contains phenomenal states *P1*, *P2*, and *P3*. It is conceivable that the phenomenal states in *N* are only intransitively phenomenally unified. In other words, we can imagine that there is one-thing it is like for *N* to experience *P1&P2*, and one-thing it is like to experience *P2&P3*, but nothing it is like to experience *P1&P3*. On this view, there would be no single thing-it-is-like to be *N*. Hence, if *N* is a subject, PUT would be false for *N*.¹⁵ But, of course, PUT could not be false for *N* and trivially true at the same time. So, perhaps defining ‘subject’ as ‘network of phenomenal states’ will provide a satisfactory account of PUT.

As it stands, however, this definition is not acceptable. It leaves entirely unclear *in virtue of what* *P1–3* are parts of the ‘network’ *N*. And as we have already noted, a definition of ‘subject’ which cannot provide any nonarbitrary way of *individuating* subjects is not a definition at all. Hence, proponents of this view must clarify what constitutes the network of phenomenal states. I am not convinced this can be persuasively done in a way that also retains PUT’s nontriviality. Prominent accounts of subjects in terms of networks of phenomenal states – such as Shoemakers functionalist account – still take subjects to be partially *identifiable* by their mental unification.¹⁶ Lack of phenomenal unity between mental states is, on such accounts, still a reason to deny that they are states of the same subject.¹⁷ Such views therefore leave PUT’s defenders with the same problem as definitions of subjects in terms of unified streams of consciousness did.

¹⁴ Thanks to the anonymous reviewer for this Journal for pressing this objection.

¹⁵ I should note here that I am happy to grant *arguendo* the possibility of such disunity *in human beings*. Certainly, I do not see any compelling reason to rule out its possibility for humans a priori. But whether this kind of partial unity is possible or conceivable in human beings is not relevant for the discussion at hand, which is about phenomenal unity in *subjects*. As I argue below (section IV), a phenomenal unity thesis about human beings is substantive but very different from PUT. Thanks to the anonymous reviewer for this Journal for pressing this objection. 1996.

¹⁶ See, e.g., Sydney Shoemaker

¹⁷ *Ibid.*

This is not to say that we cannot provide a more minimalist account of subjects in terms of phenomenal networks. We could just hold that ‘subject’ is a network of pairwise unified phenomenal states which may or may not be subsumed under total phenomenal unification. States form such a network, we can stipulate, if every phenomenal state in the network is unified with at least one other phenomenal state. I grant that this definition would finally succeed in rendering PUT substantive, for it is compatible with intransitive unity. But, I think, it should be rejected. Interestingly, this definition is only plausible – that is, it can only provide individuation – if we rule out that any single phenomenal experience could be had by more than one subject. In other words, this account would rule out *by definition* that a token phenomenal state could be shared between subjects. Though it is possible that after substantive argument we may conclude that such sharing does not or cannot occur, plausible accounts of subjects, it seems to me, ought not rule this possibility out by definition.

First, the possibility of token phenomenal states shared between subjects – i.e. the possibility of overlapping subjects of experience – is a live option in the literature (Roelofs 2016, 2018). Luke Roelofs has convincingly demonstrated that it is not strongly inconceivable, i.e., it is not logically impossible.¹⁸ Indeed, he has provided an account, couched in terms of co-presentation, that aims to explain the phenomenology of shared phenomenal states.¹⁹ It seems to me this account successfully establishes that overlapping subjects are a conceptual possibility. What is more, it seems to me to provide an attractive way forward for proponents of the partial-unity model of split-brain patients, which explains the split-brain patient data by appealing to intransitive unification of their phenomenal states. If overlapping subjects of experience are a possibility – as I think Roelofs shows they are – then the split-brain phenomenon could be analyzed in terms of two unified streams of consciousness which share some token mental states between them.

It is not necessary for the purpose of this paper to provide a full defense of the possibility of overlapping subjects of experience, however. The minimalist definition of subject which would rescue PUT from analyticity stipulates that possibility out of existence. Even if it turns out that overlapping subjects of experience are in fact impossible, they are not *by definition* impossible, that is, they are not strongly inconceivable. Hence, though a fully defended account of subjects could conclude that subjects do not overlap, this is a substantive question which ought not be resolved by definitional fiat. Therefore, this definition of ‘subject’ must be rejected as candidate for ‘subject’ in PUT. It cannot save PUT from analyticity.

4 Human Unity Theses

The reader might now object that we could use ‘subject’ to refer to the bearer of a conscious state. On this view we would use ‘subject’ to refer to a (meta)physical rather than to a merely mental entity.²⁰ For instance, we might use ‘subject’ to refer to ‘human

¹⁸ Roelofs 2016, within Subjects and between Subjects.”

¹⁹ Ibid., 3214–18; See also Roelofs 2018

²⁰ Using ‘subject’ as ‘bearer of a conscious state’ without adding (meta)physical content to the concept is no solution, since it will be just as impossible to individuate ‘bearers’ of conscious states in a principled way as it is to do this for minds.

being’ or ‘animal’: the kinds of creatures generally accepted to be the bearers of conscious states. Interestingly, this is the approach Bayne adopts in his later book on the unity of consciousness, where he retains the term ‘subject’ in the unity thesis but assumes “that we can count subjects of experience by counting human beings (Bayne 2010).”

It should be clear that ‘human being’ cannot be the definition of ‘subject’. After all, we surely cannot want to rule out *by definition* that non-human creatures could be subjects of experience. It is therefore rather misleading to defend a thesis allegedly about “any conscious subject of experience”²¹ by assuming “in the interests of keeping the discussion manageable”²² that subjects *just are* human beings. This is just to give up on defending a thesis about the nature of consciousness and defend instead a thesis about the empirical manifestations of human consciousness. Indeed, although Bayne references his earlier formulation of the unity thesis,²³ it is clear that he is now defending a much weaker claim. Thus he explicitly notes that his “only claim is that we have no good reason to think that any such division [of phenomenal consciousness] has actually occurred in the members of our own species.”²⁴ Not only is this a clear move away from the metaphysical necessity at issue in the earlier paper, this claim seems no longer to contain any appeal to necessity at all. Instead it becomes simply a contingent fact about human psychology.

Puzzlingly Bayne holds on to (equivalents of) his old formulation of the unity thesis despite this significant shift. But he cannot have his cake and eat it too. Either his goal is to defend a thesis about “any conscious subject,”²⁵ or it is to defend a contingent empirical claim about the psychological necessity of unified consciousness in human beings. If it is the first – as it clearly is in the article in which he and Chalmers propose PUT – then ‘subject’ in the unity thesis cannot be read to mean ‘human being’ but must rather mean something generally applicable to all possible experiencers. Such mental uses of ‘subject’ render the thesis analytically true, for the only nonarbitrary way to count mental subjects is to count unified streams of consciousness. If it is the second – as it is in his later book – then his formulations of the unity thesis are misleading and should be revised. For instance, he could defend any of the following:

1. **Human Unity Thesis (HUT):** no human being has ever experienced phenomenal disunity.
2. **Strong HUT:** no human being *could ever* experience phenomenal disunity.
3. **Human Single-Subject Thesis:** no human being could ever be two subjects of experience.

It seems that the Bayne of 2010 is only after defending (1): he explicitly denies accepting the stronger (2) or (3).²⁶ Notice that (2) and (3) are logical equivalents, for to argue that

²¹ Bayne, *The Unity of Consciousness*, 16.

²² *Ibid.*

²³ *Ibid.*

²⁴ *Ibid.*, 17.

²⁵ *Ibid.*, 16.

²⁶ *Ibid.*, 17: “I do not claim that it is a conceptual or metaphysical truth that our conscious states are always unified; indeed, I do not even claim that the unity of consciousness is grounded in the laws of nature. Perhaps there are surgical innovations or evolutionary developments that could bring about a division in the stream of consciousness; perhaps there are other species in which the unity of consciousness can be lost.”

some human organism contained two disunified streams of consciousness *just is* to argue that some human organism contained two subjects of experience. Now these claims are all substantive, for it is at least conceivable that human physiology could support two subjects of experience – indeed, some claim that this is what occurs in split-brain cases (Schechter 2010, 2012). But the very conceivability of one human being supporting two subjects of experience shows that we cannot equate ‘subject’ and ‘human being’. Bayne’s claims about human consciousness are substantive and interesting, but they are nontrivial precisely because they are *not* equivalents of PUT as he had previously defended it.

5 PUT and Theories of Consciousness

Taking stock: I have argued that PUT is analytically true, and that Bayne’s most recent defense of it avoids triviality only by making significant changes to the thesis. If I am right, no conceivable claim about consciousness could provide an argument against PUT. But this may seem problematic, because (1) according to Bayne and Chalmers, PUT is incompatible with several major theories of phenomenal consciousness, viz., higher-order thought (“HOT”) theories and representationalism, and (2) according to Tye, split-brain patients provide evidence against PUT.

It would be rather surprising if a trivial claim about conscious subjects turned out to be incompatible with any live theory of consciousness, for that theory would then be analytically false. But Bayne and Chalmers do not show that PUT is incompatible with either Tye’s representationalism or Rosenthal’s higher-order thought theory. According to higher-order thought theory, some mental state is conscious if and only if a subject has a higher-order thought about it.²⁷ Bayne and Chalmers take this to be incompatible with PUT because “[i]t is surely possible for a subject to think “I am in *A*” and “I am in *B*” without connecting these into a thought ‘I am in *A* and *B*.’”²⁸ But the HOT-theorist, it seems to me, can easily bite the bullet and reject this, since higher-order thoughts are not themselves conscious unless we also have a (third-order) thought about *them*. If HOT-theory is true, then our consciously thinking “I am in *A*” and “I am in *B*” is possible only if we have the unconscious higher-order thought “I am in *A* and *B*.” Since the higher-order thought is unconscious, we will not be aware of this connecting thought. The HOT-theorist can thus accept Bayne and Chalmers’ point that we can consciously think “I am in *A*” and “I am in *B*” without *consciously* thinking “I am in *A* and *B*” while holding onto the view that an unconscious HOT with that content is necessary. This is sufficient to render HOT-theories consistent with PUT.²⁹ Rosenthal’s view that HOTs can function “in a wholesale way” with an “undifferentiated bunch” of phenomenal states as their object lends further plausibility to this response, since on this view the conjunctive higher-order thought need not be particularly detailed even when it covers the conjunction of all of one’s phenomenal states.³⁰

²⁷ See, e.g., Rosenthal 2003

²⁸ Bayne and Chalmers 2010, 532.

²⁹ Of course, this solution remains vulnerable to the standard objections to HOT-theory, in particular that on this account consciousness requires implausibly extensive mental resources. But whether or not this objection is successful, it does not show that HOT-theory is incompatible with PUT.

³⁰ Rosenthal 2003, 328–29.

In any case, even if HOT-theories were to fail entirely to account for the apparent unity of *human* consciousness, this would not render them incompatible with PUT. If HOT-theory entailed that human consciousness is sometimes disunified, it would not thereby entail the falsity of PUT, but only the claim that human beings sometimes contain multiple subjects of experience. That may be implausible, but it is not analytically false.

Turning to Representationalism: it is true that according to Tye, *experiences* are necessarily *disunified*. That may seem directly incompatible with PUT, but it has this appearance only due to Tye's atypical usage of the term 'experience.' In *Consciousness and Persons* Tye argues that the totality of our phenomenal field at a time generally contains only one experience.³¹ PUT just specifies that for any subject, there must be one thing it's-like to experience the conjunction of all of its phenomenal states. This is clearly true even when the subject only has one such state.

Following Marks (1980), Tye does argue that split-brain patients have a disunified consciousness in experimental settings, though they do not become two persons.³² But this, too, is compatible with PUT: it merely entails that split-brained persons sometimes contain two subjects of experience, viz., two separate streams of consciousness.³³ (As we have seen, this analysis can also be worked out in terms of overlapping subjects). Perhaps that is an implausible view – indeed Bayne's switch-model on which split-brain persons remain single subjects provides an attractive alternative.³⁴ But Tye's account of split-brain cases does not contain an argument against PUT.

6 Conclusion

Several philosophers have argued at some length in favor of, or against, the phenomenal unity thesis, on the supposition that it contained an important, substantive claim about the nature of consciousness. I have argued in this paper that it instead expresses a merely trivial truth. In the process, I hope to have illustrated that claims like PUT simply do not *mean* anything unless we know what such claims are meant to be *about*. Hence, the frequent use of 'subject' in the literature without explanation is highly problematic. Certainly, not all claims meant to be about 'all subjects' will turn out to be analytic. But until their defenders clarify what such claims are meant to be about, discussion about them is unlikely to be productive.

The implications of this criticism are not just negative. To the contrary, in recognizing the problem with PUT we make room for more fruitful alternatives. While PUT is trivial, there are interesting claims in the neighborhood that are not. As we have seen, Tim Bayne later in effect defends the *Human Unity Thesis*, a substantive but contingent claim about human consciousness. And we need not limit ourselves to contingent claims to capture the kind of hypothesis that PUT was meant to express. As Derek Parfit's seminal discussion in *Reasons and Persons* illustrates, one important reason to be interested in the unity of consciousness is because of its implications for questions of

³¹ Tye 2003, 25–35.

³² Tye 2003, 36.

³³ For this response to Tye's account of split-brain patients as an objection to PUT see also Alter 2010.

³⁴ See Bayne 2008, 2010.

personal identity Parfit (1984). Accordingly, we could consider a *Personal Unity Thesis*. Such a thesis would not be analytic, for the standard conceptions of personhood do not by definition entail total phenomenal unity.³⁵ And it would be interesting, for our answer to it has important consequences. Does (single) *personhood* persist in split-brain cases, for instance? Or in patients with severe schizophrenia? Discussions of a *Personal Unity Thesis* informed by the medical literature would help adjudicate such questions, while discussions of PUT cannot.

A closely related possibility is to consider unification in *agents*, where agents are entities capable of moral responsibility and holders of obligations. An *Agential Unity Thesis* would hold, in effect, that a fully unified phenomenal field is a necessary condition for agency. This claim too is clearly substantive, for there is nothing about the concept of a holder of moral obligations that rules out phenomenal disunification. And it is interesting. Discussion of *Agential Unity* would add a new and enriching angle to the existing literature on the nature and requirements of agency. In addition, if empirical evidence continues to suggest the possibility of phenomenal disunification in some humans (e.g., in cases of severe illness or surgical intervention), the truth or falsity of the *Agential Unity Thesis* would have far-reaching consequences for the moral responsibilities of such patients.

In conclusion: though PUT is trivial, there are several substantive and interesting claims in the neighborhood that are not. Philosophers of mind interested in the unity of consciousness would do well to shift their focus accordingly.

Acknowledgements I am grateful to Eric Lormand for many helpful discussions on the topic of this paper, and to Jacob Berger and the anonymous reviewer for this Journal for their thorough and helpful comments on an earlier version of this paper.

References

- Alter, T. (2010). A defense of the necessary Unity of phenomenal consciousness. *Pacific Philosophical Quarterly*, 91(1), 19–37. <https://doi.org/10.1111/j.1468-0114.2009.01356.x>.
- Bayne, T. (2004). Self-consciousness and the Unity of consciousness. *The Monist*, 87(2), 219–236.
- Bayne, T. (2008). The Unity of consciousness and the Split brain syndrome. *Journal of Philosophy*, 105(6), 277–300.
- Bayne, T. (2010). *The Unity of consciousness*. New York: Oxford University Press.
- Bayne, Tim and David J Chalmers. (2010) “What is the Unity of consciousness?” In The character of consciousness, 497–539.
- Dennett, D. (1992). The self as a Center of Narrative Gravity. In F. Kessel, P. Cole, & D. Johnson (Eds.), *Self and consciousness: Multiple perspectives*. Hillsdale: Erlbaum.
- Marks, C. E. (1980). *Commissurotomy, Consciousness, and Unity of Mind*. Montgomery: Bradford Books.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Clarendon Press.
- Roelofs, L. (2016). The Unity of consciousness, within subjects and between subjects. *Philosophical Studies*, 173(12), 3199–3221. <https://doi.org/10.1007/s11098-016-0658-7>.
- Roelofs, L. (2018). Seeing the invisible: How to perceive, imagine, and infer the minds of others. *Erkenntnis*, 83(2), 205–229. <https://doi.org/10.1007/s10670-017-9886-2>.
- Rosenthal, D. M. (2000). Metacognition and higher-order thoughts. *Consciousness and Cognition*, 9(2), 231–242. <https://doi.org/10.1006/ccog.2000.0441>.

³⁵ For a concise overview see *The Oxford Handbook of Philosophy*, ed. Ted Honderich, s.v. “person.”

- Rosenthal, D. M. (2003). Unity of consciousness and the self. *Proceedings of the Aristotelian Society*, 103, 325–352 <http://linkinghub.elsevier.com/retrieve/pii/S1053810000904410>. Accessed 10 Apr 2019.
- Schechter, E. (2010). Individuating mental tokens: The Split-brain case. *Philosophia*, 38(1), 195–216. <https://doi.org/10.1007/s11406-009-9187-3>.
- Schechter, E. (2012). The switch model of Split-brain consciousness. *Philosophical Psychology*, 25(2), 203–226. <https://doi.org/10.1080/09515089.2011.579417>.
- Schechter, E. (2013). The Unity of consciousness: Subjects and objectivity. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 165(2), 671–692.
- Shoemaker, S. (1996). *The first person perspective and other essays*. Cambridge: Cambridge University Press.
- Tye, M. (2003). *Consciousness and persons: Unity and Identity*. Cambridge: MIT Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.