

# **Divided Brains & Unified Phenomenology: An Essay on Michael Tye's "Consciousness and Persons"**

Tim Bayne  
University of Oxford and St. Catherine's College  
Manor Road  
Oxford OX1 3UJ  
United Kingdom  
tim.bayne@gmail.com

This paper was published in (2005) *Philosophical Psychology*, 18/4: 495-512. Please consult the published version for purposes of quotation.

---

## **1. From the one-experience view to the received model**

Tye begins *Consciousness and Persons* by reminding us that there are a number of unity relations in consciousness. There is the unity of belonging to one and the same subject of experience (subject unity), the unity of being directed at a single object (object unity), and the unity of seeming to occupy a single space (spatial unity). Standing over and against all these unity relations is the experienced togetherness of the contents of phenomenal consciousness. As I write this, I hear the low murmur of cars as they travel along the road beneath my window, I taste the pleasant acidity of my morning coffee, I am aware of a dull ache in my calves due to yesterday's jog, I see the leaves fall from the tree outside my window, and I have a vague sense of anxiety about an impending trip to the dentist. These experiential contents stand in various relations of object unity, spatial unity, and subject unity. They are also related by being experienced together within a single field (or stream) of consciousness. Tye calls this relation of phenomenal togetherness "phenomenal unity."

If there is a central motif to *Consciousness and Persons* it is that the unity of consciousness—and phenomenal unity in particular—has been misconceived. Most philosophers have approached the unity of consciousness in terms of the following question: "How are various experiences (auditory experiences; visual experiences, bodily experiences, and so on) unified?" Tye calls this approach the *received view*, and argues that it is mistaken [1]. Sense-specific experiences, he says, "...are the figments of philosophers' and psychologists' imaginations (p. 28). Instead of thinking of the unity of consciousness in

terms of unity relations between experiences, Tye urges us to think of it in terms of unity relations between experiential *contents*.

The core intuition, lost in the usual way of stating the problem of unity, is that, in normal cases, simultaneously experienced perceptual qualities—the loudness of a sound, the smoothness of a surface, the sweetness of a taste, the pungency of a smell—are experienced together and thus are phenomenologically unified ... phenomenological unity is a matter of simultaneously experienced perceptual qualities entering into the same phenomenal content. (p. 36)

Tye calls his approach to the unity of consciousness the “one experience model”. As I read it, the one experience model consists of the conjunction of the following three claims:

- (1) The only experiences that (normal?) human beings have are entire streams of consciousness, where a stream of consciousness is “a period of consciousness between one state of unconsciousness and the next” (p. 97);
- (2) Experiences cannot have experiences as proper parts (p. 40);
- (3) Phenomenal unity should be analysed in terms of the closure of phenomenal content under conjunction (p. 36).

Although it is clear that Tye is committed to all three claims, it is not clear how he sees them as related. I suspect that he thinks that there are important connections between them, but he does not spell these out in any way. (1) and (2) seem to be closely connected: if an entire stream of consciousness is an experience in its own right, and if experiences cannot have experiences as proper parts, then no proper part of a stream of consciousness could be an experience. And the conjunction of (1) and (2) *might* be thought to provide indirect support for (3): if phenomenal unity is a relation that holds between the parts of a stream of consciousness (my current bodily sensations and my visual sensations, for instance), and if the parts of a stream of consciousness are not experiences in their own right, then the traditional analysis of phenomenal unity, according to which it is a relation between experiences, cannot be right. This does not, of course, imply that Tye’s conjunctive analysis phenomenal unity is right, but it does suggest that we might want to take a serious look at it.

Of course, one might also reject (1) and (3) but endorse (2). One might argue that the stream of consciousness is not itself an experience, but rather a medium in which experiences are ‘contained’ as unified. It is natural to assume that if the unity of consciousness involves relations between experiences then those unity relations must be mereological (that is, part-whole in nature), but this assumption is not unquestionable.

I will return to the question of how the three components of Tye's model might be related, but for now I want to examine each of these three claims separately, beginning with the claim that experiences are *entire* temporally extended streams of consciousness.

### 1.1. Individuating experiences

As Tye is well aware, the claim that the only experiences we have are entire streams of consciousness is at odds with common sense: we certainly seem to have more than one experience within uninterrupted periods of consciousness! Indeed, it is plausible to suppose that there are times during which we have more than one experience at once. Proponents of the received view of the unity of consciousness have rarely been explicit about how they are conceiving of experiences, but I think they have tended to assume that any account of experience ought to begin with these common sense commitments. One conception of experience that would comport with such commitments is a view on which experiences are individuated in terms of instances of phenomenal properties (give or take a bit). Roughly, a subject has distinct experiences when she instantiates distinct phenomenal properties. For example, an experience of tasting a Merlot is distinct from the experience of an itch in one's left leg because these experiences involve distinct phenomenal properties. The qualifiers are needed because there is reason to think that experiences do not stand in a one-to-one correspondence with instances of phenomenal properties. Consider a visual experience of (say) scarlet at a certain location in the visual field: corresponding to what is intuitively a single experience are instances of multiple phenomenal properties (e.g., phenomenal red, phenomenal scarlet, phenomenal scarlet of such-and-such shade, and so on). But despite this (and other worries), individuating experiences in terms of phenomenal property instances has a certain appeal; more to the point, I suspect that it is the approach that proponents of the received view have by and large assumed.

Tye's account of the unity of consciousness clearly puts him at odds with the phenomenal property approach: if experiences just are instances of phenomenal properties (give or take a bit), then a normal stream of consciousness clearly contains multiple experiences at any one time. How then does Tye individuate experiences? His approach is functional. Experiences, says Tye, are representations with poised, abstract, nonconceptual, intentional content (PANIC). How does a PANIC approach to consciousness individuate experiences? What makes it the case that we have an experience with content  $\langle A \& B \rangle$  as opposed to two experiences with unconjoined contents  $\langle A \rangle$  and  $\langle B \rangle$  respectively? Tye's answer, I think, is to ask whether the conjunctive content  $\langle A \& B \rangle$  is poised for control of the reasoning system, or whether only the contents  $\langle A \rangle$  and  $\langle B \rangle$  are poised for control of the reasoning system. So, Tye's claim that normal human consciousness involves but one

experience seems to reduce to this: the contents of normal human consciousness are conjointly poised for direct input into the reasoning system.

We might have achieved some clarity over how experiences might be individuated, but in so doing we have problematized the contrast between Tye's one-experience view and the received view. Given that they appear to be operating with different conceptions of experience the dispute between them threatens to be merely verbal. Perhaps the central disagreement here concerns the question of how experiences ought to be individuated: should they be individuated in terms of instances of phenomenal properties (give or take a bit), or in functional terms? But is this really a substantive issue? If so, what hangs on it? There is, I think, a case to be made for the view that how one ought to individuate experiences depends on one's theoretical needs, and if that is right, then it might be argued that both the one-experience model and the received view have some legitimacy.

## 1.2. Experiences and the parts thereof

Suppose that we were to follow Tye and individuate experiences in functional terms. Would it follow that an experience cannot have experiences as parts? I think not. Take an experience ( $e_1$ ). According to Tye,  $e_1$  is a PANIC representation. As such, it seems possible for  $e_1$  to have as one of its proper parts another PANIC representation,  $e_2$ . But since  $e_2$  is a PANIC representation, and experiences are PANIC representations, it seems to follow that  $e_2$  will also be an experience, and hence that experiences can have experiences as parts.

Tye responds to this objection by denying that every PANIC representation is an experience (p. 40). In order to be an experience, Tye says, a PANIC representation must be *maximal*, where by "maximal" he seems to mean that the representation is not itself a proper part of another PANIC representation.

This stipulation blocks the representationalist objection to the one-experience model, but the victory seems hollow. Tye has merely *defined* "experience" so as to ensure that his account holds. Why should a representationalist—or anyone else for that matter—accept this definition? Why must representations be maximal in order to be experiences? I am unsure of Tye's answer to this question, but I *think* it has something to do with the thought that the only natural boundaries around PANIC representations are those around maximal PANIC representations. In a couple of passages, Tye compares experiences to clouds and statues of clay (pp. 30, 99). Clouds have as proper parts collections of water molecules, but collections of water molecules are not themselves clouds. Similarly, statues of clay contain smaller chunks of clay as proper parts, but such chunks do not (typically) constitute statues in their own right (although they can). Tye grants that experiences can

have stages (or parts), but asserts that experience stages are no more experiences than undetached cloud stages are clouds or undetached statue stages are statues (p. 99).

But there are important differences between cloud and statue stages on the one hand, and experience stages on the other. It seems to me that there is little point in making distinctions between the parts of a cloud, whereas there is much to be said in favour of distinguishing parts of a phenomenal stream from each other. We know that there are functional and phenomenal differences between the parts of a phenomenal stream. There is the distinction between attentional figure (or focus) and attentional background. There are also differences in the cognitive roles of phenomenal contents; for instance, it is possible that one is conscious of a brown dog and the taste of coffee, but only the first content is currently impacting on one's belief-forming system. These points motivate the thought that experience stages—that is, parts of phenomenal streams—have some claim to be thought of as experiences in their own right, and that the phenomenal divisions within an experience are not arbitrary.

Another reason to think that experiences can have experiences as parts derives from what we might call “phenomenal articulation”. I use this term to refer to the fact that phenomenal contents can occur in different phenomenal contexts. For example, one can hear a dog barking while: experiencing an itch in one's right leg; experiencing an itch in one's left leg; or having no bodily sensations at all. This sort of combinatorial structure might not apply to all phenomenal contents—perhaps there are experiential contents that can occur only in each other's presence—but it seems to be fairly widespread (see Bayne & Chalmers 2003; Dainton 2000 for more discussion of this issue).

How might phenomenal articulation provide an argument for thinking that experiences have parts? The intuitive explanation of why one can experience a dog barking without experiencing an itch and vice-versa is that the experience of the dog barking is distinct from the experience of the itch. It seems natural to explain articulation in the contents of consciousness by appeal to articulation in the vehicles of those contents. I do not claim that phenomenal articulation provides a knockdown objection to the one-experience view, but I think it does provide *prima facie* motivation for it.

Finally, one might have questions about the *modal* status of Tye's position. Although he does not address this issue directly, it is natural, I think, to read him as claiming that it is a necessary truth that experience cannot have other experiences as parts. I say this because Tye's arguments for the model (and against the received view) tend to be of an *a priori* and armchair nature. But this seems to get the modal status of the one-experience claim wrong. Even if it turns out that experiences do not have parts, they seem to be the sorts of

things that could have had parts. Consider Tye's analogy between experiences on the one hand and clouds and statues on the other: surely clouds and statues are the kinds of things that could have had objects of their own kind as proper parts. (Indeed, no doubt some statues actually have statues as proper parts!)

### 1.3. Conjunctive closure

I turn now to Tye's conjunctive analysis of phenomenal unity. Tye claims that what it is for experiential contents  $\langle A \rangle$  and  $\langle B \rangle$  to be unified is for the subject of experience in question to have an experience with content  $\langle A \& B \rangle$  (p. 36f.).

I am well disposed to conjunctive analyses of phenomenal unity; indeed, David Chalmers and I have defended such an analysis of phenomenal unity (Bayne & Chalmers 2003). We argued that phenomenal unity can be understood in terms of the closure of phenomenal *character* under conjunctive closure. What it is for phenomenal states  $e_1$  and  $e_2$  to be phenomenal unified is for the subject to have a state whose phenomenal character entails the phenomenal characters of  $e_1$  and  $e_2$  (see Bayne and Chalmers 2003 for details). But this analysis is importantly different from Tye's. Because Tye identifies phenomenal character with a certain type of representational content, his analysis is, at bottom, an analysis in terms of the closure of representational content under conjunction. I do not think that this will work. Although there is something to be said in favour of the idea that phenomenal unity often brings with it conjunctive closure of phenomenally conscious representational content, I am not persuaded that phenomenal unity can be *reduced* to such a relation. Indeed, I think it is doubtful that phenomenal unity is even *co-extensive* with the conjunctive closure of representational content.

First, phenomenal unity might be possible in the absence of conjunctive closure of representational content. Consider a person who has two conscious beliefs at the same time. It seems conceivable that these beliefs could be phenomenally unified without the person in question also having a belief whose content is the conjunction of the contents of the beliefs in question. Admittedly, such cases might be rare—indeed, they might even be psychologically impossible for creatures like us—but they do not seem to be conceptually incoherent, and the mere conceptual possibility of such a situation suffices to cast doubt on Tye's analysis of phenomenal unity (which is, I take it, a *conceptual* analysis).

I doubt that Tye would be moved by this objection, for he denies that propositional attitudes can be phenomenally conscious. Tye claims that the only phenomenology involved in conscious thought derives from whatever imagery accompanies it (p. 79) [2]. I am not persuaded by Tye's position, but I need not rest my case against his analysis of phenomenal unity on any controversial assumptions about the phenomenality of

propositional attitudes, for the same objection can be run in terms of perceptual content [3]. It seems possible for two perceptual states to be phenomenally unified—in the sense that there is something it is like to instantiate them both—without their contents being unified. In fact, this description might apply to visual agnosia, a condition in which the patient appears to have phenomenally unified experiences of the parts of visual object, but lacks any coherent (contentfully unified) experience of the object as a whole (Farah 1990) [4].

Not only might we have phenomenal unity without conjunctive closure, we might also have conjunctive closure without phenomenal unity. To see this, suppose that it is possible for a subject of experience to have two streams of consciousness at once. (I do not myself think that this is possible, but Tye does, as we will see.) Now suppose that one stream of consciousness includes an experience with content  $\langle A \rangle$  and the other stream of consciousness includes experiences with contents  $\langle B \rangle$  and  $\langle A \& B \rangle$  respectively. Here we have the closure of phenomenal content under conjunction, but we do not have full phenomenal unity, for the experience with content  $\langle A \rangle$  is not unified with the experiences with contents  $\langle B \rangle$  and  $\langle A \& B \rangle$ .

Of course, Tye might respond by arguing that this objection simply begs the question against him: if a subject has states with contents  $\langle B \rangle$ ,  $\langle A \rangle$  and  $\langle A \& B \rangle$  at the same time, then these states *must* be mutually phenomenally unified. My objection is sound only if I have given some reason for thinking that these states need not be phenomenally unified, and—he might continue—I have not done this.

It is certainly true that I have provided no reason for thinking that a subject could have states  $\langle B \rangle$  and  $\langle A \rangle$  in one stream of consciousness and  $\langle A \& B \rangle$  in another, but I do not think that this renders the objection impotent. Given that Tye allows for the possibility of a subject with two streams of consciousness at once, it seems to me that the burden is on him to explain what would rule the above scenario out. Why should the contents of one stream of consciousness be (conceptually) dependant on the contents of the other stream of consciousness? If a subject could think  $\langle A \rangle$  in one stream and  $\langle B \rangle$  in their other stream, why could they not think  $\langle A \rangle$  and  $\langle B \rangle$  (only) in one stream and  $\langle A \& B \rangle$  (only) in another stream?

A further problem with Tye's conjunctive analysis of phenomenal unity is that it seems to be incompatible with the claim that experiences are entire streams of consciousness. It seems plausible to suppose that the contents of an experience are mutually phenomenally unified: that, after all, is what gives an experience its unity. Can Tye accept this view? Only if the contents of an entire stream of consciousness—that is, the period of

consciousness between one state of unconsciousness and the next—are closed under conjunction. And that seems very unlikely. I had an experience of tasting coffee this morning, and this evening I am currently experiencing a fine Merlot: these contents are had by a single experience, since they are part of an unbroken period of consciousness. But it is *clearly* not the case that this experience includes the content <having an experience of coffee and tasting a good Merlot>. This content was not, and is not, poised for direct control of my reasoning system.

## 2. Against the received view

Rather than arguing in favor of his one-experience account of the unity of consciousness, Tye focuses on objections to the received view. He paints with a wide brush here: ignoring differences of detail between various positions, Tye attempts to show that no variant of the received view can succeed. He offers three main arguments for this thesis [5].

### 2.1. The explanatory regress objection

Tye's first argument is as follows:

If what it is like to undergo the overall or maximal experience is different from what it is like to undergo the component sense-specific experiences, E1-E5, then there must be a unifying relation between the latter experiences that is itself experienced. The experience of the unifying relation is not itself a sense-specific experience. But it is an experience nonetheless; for if there were no experience of the unifying relation, then there would be nothing it is like to have the sense specific experiences unified. There are, then, it seems, six experiences: the five sense-specific experiences, and the experience of unity. However, the maximal experience isn't just a conjunction of experiences. It is a genuinely new unified experience with its own phenomenology. So, what unites the six experiences together? It seems there must be a further unifying relation that binds the experiences. This relation, however, must itself be experienced. For the unity is phenomenal. And now a regress has begun to which there is no end (p. 22).

I take it that the supposed regress is *explanatory*. The worry is that there can be no phenomenal unity relation between experiences that explains how the experiences in question are unified, for any such relation would have to be experienced, thus giving rise to the question of how it—the relation of co-consciousness—is unified with those experiences that it unifies.

Versions of this objection have been circulating in the literature for some time, and a number of responses to it have emerged (see Bayne 2001; Dainton 2004; Hurley 1998). One

option that has been suggested is that although co-consciousness is an experience it is a peculiar sort of experience: unlike other experiences, self-consciousness is a self-binding experience. Perhaps co-consciousness can bind experiences together to form unified phenomenal wholes without itself needing to be so bound.

A second and more attractive response is to deny that co-consciousness is itself an experience. Tye assumes that because co-consciousness is a phenomenal relation it must have its own phenomenology: "if there were no experience of the unifying relation, then there would be nothing it is like to have the sense specific experiences unified" (p. 22). This inference is tempting but perhaps it should be resisted. Perhaps there can be something it is like to have a set of unified experiences without that "what it's like" subsuming or involving an experience of the unity relation that binds the experiences in question together. I think the lesson to draw from the explanatory regress objection is that co-consciousness is a phenomenal relation in the sense that it makes a phenomenal difference, but not in the sense that it is itself an object of experience.

## 2.2. The phenomenal bloat objection

Tye has a second regress argument against the received view:

There is also a real question as to whether there is a maximal, unifying experience in the first place. For consider three simultaneous unified experiences,  $e_1$ ,  $e_2$  and  $e_3$ . If the unity or experienced togetherness of any two experiences requires that there be a unifying experience [that subsumes them? – TB], then the unity of  $e_1$  and  $e_2$  requires that there be a further experience  $E$  that includes them. (p. 22).

Whereas Tye's first argument concerns an explanatory regress, I take this argument to concern a regress of content. The idea seems to be that if one thinks of the unity of consciousness as involving relations between distinct experiences, and if one allows that experiences can have as parts other experiences, then one faces the question of what makes it the case that an experience is unified with its parts. So take an experience ( $e_1$ ) and one of its parts ( $e_2$ ). The unity of these two experiences seems to demand a further experience ( $e_3$ ) that is distinct from both  $e_1$  and  $e_2$ . But however "big" an experience is, we can always ask what makes it the case that it is unified with its parts. And now we appear to have embarked on another regress – a regress of content. For every experience we have, we appear to need a further experience which might explain how the first is unified with its parts. We might call this the *problem of phenomenal bloat*.

In responding to this objection the proponent of the standard view might consider rejecting the view that experiences are unified only when they are co-subsumed.

Although that is certainly a strategy worth considering, I myself am not much tempted by

it. I think the phenomenal bloat objection can be deflated by noting that the unifying experience need not be distinct from the unified experiences. Take a complex experience ( $e_3$ ) that subsumes two other experiences ( $e_1$  and  $e_2$ ). Experiences  $e_1$  and  $e_3$  are unified in virtue of being subsumed by a single experience, but that single experience need be none other than  $e_3$  itself. There is nothing incoherent in the idea of a maximal unifying experience.

### **2.3. The objection from introspection**

Whereas Tye's first two arguments against the standard model attack its very coherence, his third argument attempts to undermine its epistemic basis [6]. The argument appeals to a certain conception of introspection. Introspection, says Tye, gives one access only to the contents of one's experiences. Experiences are transparent—we see through them to the world as represented by them. And, says Tye,

If we are not aware of our experiences via introspection, we are not aware of them as unified. The unity relation is not given to us introspectively as a relation connecting experiences. Why, then, suppose that there is such a relation at all? (p. 25)

One might respond to this objection by challenging Tye's claim that introspection provides one with access only to the contents of experience. Arguably introspection provides access to properties of an experience that do not fall within its content, such as whether the experience is a mental image or a perceptual state. And if introspection has access to the mode of an experience perhaps it can also provide us with access to its unity relations.

Of course, the issues just raised lie at the heart of the representationalist account of phenomenal consciousness that Tye assumes here, and I lack the space to pursue them in any detail. Instead, I want to suggest that we have reason to posit phenomenal unity relations between experiences even if introspection affords us access only to the contents of consciousness. Suppose that you are introspectively aware of hearing a dog bark and seeing it run around. It is natural to assume that since this experience has a complex phenomenal character, it involves an experience of hearing a dog bark and an experience of seeing it run around, and that these experiences are unified. How could one be having an experience with auditory and visual character if one's experiences of hearing it were not unified with one's experience of seeing it run around?

No doubt Tye would challenge the assumption that if one is having an experience of hearing a dog bark and seeing it run around then one must be having an experience of seeing it bark and an experience of seeing it run around. Given his conception of experiences Tye would be within his rights to level such a challenge, but, as I point out

above, the proponent of the received view is under no obligation to endorse Tye's conception of experiences. If the proponent of the received view were to individuate experiences in terms of phenomenal properties, then she would have a strong case for thinking that any experience with complex phenomenology will involve simpler experiences that are bound together by relations of co-consciousness. In short, the proponent of the received view can respond to the argument from introspection even if she is prepared to grant that introspection provides one with direct access only to the representational contents of consciousness.

### **3. The Unity of Consciousness, Personhood, and the Split-Brain**

I turn now to Tye's discussion of the split-brain syndrome and personhood.

#### **3.1. Mapping the terrain**

Tye begins his account of the split-brain with the following taxonomy (p. 112f.):

- (a) Split-brain subjects are two persons with two separate minds;
- (b) Split brain subjects have a single unified stream of consciousness; the responses produced by the right hemisphere are those of an automaton;
- (c) It is indeterminate how many persons split-brainers are;
- (d) Split-brain subjects have a unified phenomenal consciousness but disunified access consciousness;
- (e) Split-brain subjects are single persons who undergo two separate streams of consciousness.
- (f) Split-brain subjects are single persons whose phenomenal consciousness is briefly split into two under experimental conditions but whose consciousness at others times is unified.

Tye argues that there are serious problems with models (a) through (e), and that (f) is left as the most plausible account of the split-brain.

I have reservations about using this taxonomy to structure discussion of the split-brain, for it conflates questions about personhood with questions about the structure of consciousness. Of course, some theorists hold that the structure of consciousness (partially) determines personal identity – Nagel assumed this position in his famous discussion of the split-brain, and I myself have some sympathy for it. But this view is contentious. Many theorists—including Tye himself—deny that the structure of consciousness has any direct bearing on the individuation of persons. Given this, it seems

best to work with a taxonomy of split-brain models that clearly distinguishes questions concerning the structure of consciousness in the split-brain from questions concerning the number of persons split-brain patients are (or constitute, or whatever).

A second problem with Tye's taxonomy is that it is incomplete. Missing from it is the partial unity model. The partial unity model was inchoate in some of the early neuroscientific work on the split-brain (see Trevarthen 1974a; Trevarthen 1974b), but as far as I know was first explicitly developed and defended by Michael Lockwood in his *Mind, Brain and the Quantum* (see also Lockwood 1994). Lockwood suggested that one could account for the behavior of split-brain patients under experimental conditions by supposing that (synchronic) co-consciousness is not transitive. On such a view, split-brain subjects have three simultaneous experiences where two of the experiences are co-conscious with the third experience but not with each other. In a partially unified consciousness co-consciousness has a branching structure at a single time. In and of itself Tye's omission of partial unity models would not be a cause for concern. What makes it problematic is that Tye argues for his account of the split-brain on the grounds that none of the alternative models is acceptable. Successful prosecution of this kind of argument demands that one canvass all of the plausible alternatives to one's preferred view [7].

Tye does discuss and endorse the notion that phenomenal unity is nontransitive (pp. 129-32), but it turns out that what he means by the nontransitivity of phenomenal unity is something quite different from what is more commonly meant by the term. Tye discusses a scenario in which a subject has an experience of greenness that is unified with an experience of a pinprick, an experience of redness that is unified with an experience of a pinprick, but no experience of redness that is unified with an experience of greenness. His account of this scenario is as follows:

S has two multimodal experiences,  $E_1$  and  $E_2$ .... The pricking is phenomenally unified with redness by entering into the same phenomenal content – the phenomenal content of  $E_1$ . The pricking is phenomenally unified with greenness in like manner, but this time the common content is the phenomenal content of  $E_2$ . Since S has no experience whose phenomenal content has entering into it both greenness and redness, the two colours are not phenomenally unified (p. 130f.).

Tye's conception of what it would be for non-transitivity to fail is clearly very different from Lockwood's. Lockwood introduced the notion of non-transitivity in order to describe experiences that are neither fully unified in one consciousness nor contained in completely separate consciousness. He described such experiences as (merely) partially unified: such experiences occur within a single stream of consciousness, but that stream is

only partially unified, for it contains experiences that are not co-conscious with each other. Tye, however, is interested in relations between the contents of experience rather than relations between experiences themselves. Because of this, it is possible both for transitivity, as Tye thinks of it, to fail in a fully unified consciousness, and it is possible for it to hold in a partially unified consciousness. (At least, a proponent of the received view will think that such scenarios are possible – Tye himself might not.)

### 3.2. Tye's model of the split-brain

Although most of Tye's attention is devoted to undermining other models of the split-brain, he does offer one positive line of argument for his view. The argument is an inference to the best explanation: the behavior of split-brain subjects is best explained by holding that their consciousness is divided under (certain) experimental conditions but unified in everyday contexts [8].

This model of the split-brain is attractive, but it is not without its difficulties. One issue on which Tye is noticeably silent concerns the mechanisms by means of which the patient's phenomenology is divided and then reunified. One might well wonder how the transition between everyday and experimental contexts could (dis)unify consciousness. If phenomenal unity supervenes only on the functional integration of those neural regions involved in generating phenomenology, how could the presence and absence of experimental conditions affect neural integration?

One response to this question is to go externalist about the vehicles of phenomenal consciousness and deny that phenomenology supervenes only on neural activity; indeed, I suspect that this is the *only* plausible line of response to this question. Perhaps Tye should follow Hurley in identifying the unity of consciousness with a "dynamic singularity in the field of causal flows that is centred on but not bounded by a biological organism" (1998, p. 207). If something like this were the case, then one might be able to divide and reunite a person's consciousness by interfering with their ability to interact with their environment in an integrated way. I have no idea whether Tye would wish to sign up to vehicle externalism for phenomenal consciousness, but I do think he is pushed in that direction. Whether or not that is a problem for his view depends, of course, on the plausibility of vehicle externalism.

A second problem for Tye's account is that it is not clear that his inference to the best explanation really is an inference to the *best* explanation. The central alternative to Tye's account of the split-brain is what we might call the two-streams view, according to which split-brainers always have two-streams of consciousness. Two-streamers explain behavioural integration by reference to the *duplication of content*. They hold that there is

every reason to expect behavioural integration outside of the context of tachistoscopic experiments, for each of the person's two streams of consciousness will have almost identical contents, and identical experiential contents will generate behavioural integration (see e.g., Puccetti 1981; Moor 1982).

Can Tye reject this explanation of behavioural integration in the split-brain? I'm not sure that he can. His own discussion of the non-transitivity of phenomenal unity not only assumes that split-brain subjects can have duplicate experiences (one in each stream of consciousness), it also seems to demand that the duplication of content across streams would not impair the subject's behavioural integration. Of course, it is not obvious that near duplication of content across streams would lead to behavioural integration: brains are chaotic, non-linear systems, and even slight differences in phenomenal content might manifest themselves in behavior. But whether or not this is right, it is at best an open question whether *Tye* has reason to endorse his own model of the split-brain rather than the two-streams alternative.

Let me conclude this section by pointing out an apparent tension between Tye's analysis of the split-brain and his account of the unity of consciousness. Tye holds that split-brain patients have a divided consciousness in certain experimental conditions. Now, consider such a patient who has a unified consciousness since waking up at 8am, a divided consciousness between 9am and 10am (when they are in experimental conditions), and a unified consciousness between 10am and 11am, when they fall asleep. How many streams of consciousness has this person had? Tye's answer is "two" (p. 103). These two streams of consciousness overlap between from 8-9am, diverge between 9-10am, and overlap again between 10-11am. Now, why should we say that there are two streams of consciousness (experiences) here? After all, we have but one uninterrupted period of consciousness. Well, the intuitive answer is that there are experiential contents that occur between 9am and 10am (while the patient is in a split-brain experiment) that are not phenomenally unified. Take two such contents, <A> and <B>: what makes it the case that these two contents are not phenomenally unified? Tye's conjunctive closure account will suggest that <A> and <B> are not phenomenally unified because they are not accompanied by an experience with content <A&B>. Now, take an experiential content that occurs just after 8am (<C>) and another than occurs just before 11am (<D>). What makes it the case that these two contents occur within a single stream of consciousness? It seems highly implausible to suppose that these two contents will be phenomenally unified: they occur almost three hours apart from each other! Of course, they are linked by chains of phenomenal unity, but so too are <A> and <B>. Tye says that his model of the split-brain supports his account of the unity of consciousness (p. 113). I do not see the connection; in

fact, I think there is a serious tension between his account of the split-brain and his account of the unity of consciousness.

### **3.3 Persons, subjects and the unity thesis**

One of the reasons that the split-brain syndrome has generated so much interest among philosophers is that it appears to threaten an attractive conception of how the unity of consciousness is related to the notion of a subject of experience. According to what we might call the unity thesis, the simultaneous experiences of a subject of experience must be mutually phenomenally unified (Bayne & Chalmers, 2003).

Tye rejects the unity thesis. As we have seen, Tye holds that split-brain patients are single subjects of experience who have two streams of consciousness, at least in certain experimental conditions. More generally, Tye holds that there is no close connection between subjects of experience and the unity of consciousness. He argues that subjects of experience are intentional systems—what he calls “psychological frameworks” (chapter 6). Such a conception of the subject of experience does not sit well with the unity thesis, since it is plausible to suppose that a psychological framework can support more than one stream of consciousness at a time. To this end, Tye owes us an account of why the unity thesis (or something very much like it) has the force that it does. Tye needs to explain away its appeal.

As Tye observes, a popular route to the unity thesis proceeds via the notion of projectibility. One might argue that disunified subjects of experience are impossible on the grounds that we cannot project ourselves into the phenomenal perspective of such a subject. Tye rejects this argument. He thinks that we *can* project ourselves into the perspective of a disunified subject:

in one sense, we can project ourselves into the mental lives of split-brain patients: we know just what it is like to experience each experience a split-brain patient undergoes. What we cannot do is imagine ourselves undergoing both the visual experiences a split-brain subject experiences at any given moment while the experimental controls are in place and simultaneously being aware, all in one introspective act, of the content of each. In this respect we are not different from the split-brain patient; for he cannot imagine that either (p. 20f.).

Just what Tye means when he suggests that we project ourselves into each of the experiences of the split-brain patient is somewhat unclear. Remember that Tye holds that the only experiences human beings have are entire streams of consciousness. Is he really

suggesting that we can project ourselves into an entire stream of consciousness? I suspect not. I think that he is using “experience” loosely here, and is only suggesting that we can project ourselves into one of the split-brain patient’s experiential *stages* or *phases*.

But the real worry with this passage is that it seems to presuppose an overly simplistic conception of what is required for phenomenal projection. Tye seems to think that in order to project oneself into *S*’s perspective one need only project oneself into each of *S*’s experiences. That does not seem right. Arguably, in order to project oneself into *S*’s phenomenal perspective one needs not only to project oneself into each of *S*’s experiences, one also needs to capture the fact that the experiences in question are *S*’s experience (and, perhaps, that they occur simultaneously). Of course, it is not entirely how one captures this fact and much more could be said here, but there seems to be something to the thought that Tye has failed to meet the conditions needed for successful phenomenal projection [9].

Although I do not endorse Tye’s reasons for rejecting the argument from projectibility, I too think the argument fails. I think it goes wrong by assuming that projective inaccessibility is a reliable guide to phenomenal impossibility. It is not. Our inability to project ourselves into the phenomenal perspective of a bat would not give us any reason to think that there isn’t anything it is like to be a bat. Our projective abilities are constrained by the structure of our own phenomenal perspective, and there is little reason to think that this perspective limns the space of phenomenal possibility. Proponents of the unity thesis need to do better than rest their case on the argument from projectibility.

Luckily, they can. Rather than appealing to our powers of imaginative projectability, proponents of the unity thesis can appeal to our powers of conceivability. Arguably we can see that there is something incoherent in the idea of a subject of experience with a disunified consciousness. Call this the *unity intuition*. Tye appears not to share the unity intuition; or rather, if he shares it, he appears willing to let it be defeated. There is, he says, something it is like for the split-brain subject to have each of their experiences, but there is no one thing it is like for them to have all of their experiences [10].

What are we to make of this clash of intuitions? Perhaps Tye and I simply have different concepts of what it is to be a subject of experience. One reason for thinking that this is so is that, as best I can tell, Tye identifies the concept of a person with the concept of a subject of experience (see p. 134), whereas I am quite willing to allow that there is a notion of personhood that is distinct from the notion of a subject of experience. I think there is quite a lot to be said for the view that *persons* are intentional systems—what Tye calls “psychological frameworks” (chapter 6)—and it seems entirely possible for a single

intentional system to have (support) two streams of consciousness, at least for short periods of time. But is it possible that a *subject of experience* could have two streams of consciousness at once? I am inclined to think that it is not. My conception of what it is to be a subject of experience is constitutively bound up with the unity thesis.

So we might want to rethink the assumption that persons are subjects of experience. We need not deny that in the normal course of things persons and subjects of experience are intimately related, but intimate relations are one thing, identity (and indeed constitution) is another. In the same way that many of us have to come to distinguish human beings from persons, so too we might want to distinguish persons from subjects of experience.

## Acknowledgement

I am very grateful to Barry Dainton, Tim Crane, David Chalmers, Uriah Kriegel and Cory Wright for their comments on earlier drafts of this paper.

## Notes

[1] Tye does not say who he regards as endorsing the received view, but I suspect that the following would make it on to his list: Bayne & Chalmers (2003), Brook (2001), Dainton (2000), Lockwood (1989), Hurley (1998), Rosenthal (2003), Shoemaker (1996), Shoemaker (2003).

[2] There is an interesting and important question here of how the phenomenal imagery that accompanies propositional attitudes is unified with the propositional attitudes themselves. Tye says that the phenomenology of conscious thought derives from the phenomenology of an associated auditory experience (p. 81), but what is it for auditory experience to be associated with a propositional attitude? It surely does not reduce to the fact that the imagery is simultaneous with the propositional attitude, nor would it seem to involve a causal relationship between the two states. It is, I think, quite unclear how “association” is to be understood here.

[3] Even if Tye is right in claiming that *our* propositional attitudes lack phenomenal character, it would seem to be possible for a creature to have propositional attitudes with phenomenal character. At any rate, I can't see any contradiction or incoherence in the idea that a propositional attitude might have phenomenal character. If non-conceptual representations can have a phenomenology what is to prevent conceptual representations from having a phenomenology? And the mere *possibility* of propositional attitudes having a phenomenology seems sufficient to create a problem for Tye's analysis of phenomenal unity.

[4] One might also ask whether the contents of normal phenomenology are conjunctively unified. Consider just one component of phenomenology: the awareness of one's body. It seems to me that one can have a simultaneous experience of each of one's limbs without having an experience

whose contents represent the totality of one's body parts. I do not claim that this is obviously possible but it does seem to me that it is plausible.

[5] In chapter five of *Consciousness and Persons* Tye gives what might be read as a fourth argument for his account of the unity of consciousness. He suggests that his account provides a better explanation of the fact that phenomenal unity is a non-transitive relation than do standard accounts of the unity of consciousness. According to Tye, the proponent of the standard view will have to say that it is simply a brute fact that co-consciousness is a non-transitive relation. But, says Tye, brute facts such as this should be avoided wherever possible (p. 131). I think two responses are appropriate. First, not everyone will agree that (synchronic) phenomenal unity is non-transitive; I myself have doubts about this, and I am not alone (see Dainton 2000; Lockwood 1994). Secondly, Tye's own account seems to be committed to a certain number of brute facts.

[6] Tye's argument from objection from introspection bears a family resemblance to Hurley's introspection-based objection to "what it's like" accounts of the unity of consciousness (Hurley 1998; Hurley 2003. For a brief reply to Hurley's argument see Bayne and Chalmers 2003.

[7] Presumably Tye thinks that he can ignore the question of whether co-consciousness is transitive on the grounds that he has already argued that there neither is nor could be any such phenomenal unity relation between experiences.

[8] This model of the split-brain is not novel. To the best of my knowledge Marks (1980) is the first person to have defended it in print. Moor (1982) also defends it.

[9] I thank Cory Wright for prompting me to say more here. For discussion of whether we can project ourselves into the phenomenal perspective of a disunified subject see Lockwood (1994, p. 95), Peacocke (1994, p. xx) and Dainton (2000, p. 98).

[10] Those who are sceptical about the reliability of modal intuitions might accept that we have the unity intuition but deny that it provides us with good reasons to endorse the unity thesis. Obviously addressing the reliability of modal intuitions in general and the unity intuition in particular goes beyond the current scope of this paper.

## References

- BAYNE, T.J. (2001). Review of B. Dainton, "Stream of Consciousness." *Journal of Consciousness Studies*, 8, 79-92.
- BAYNE, T.J. & CHALMERS, D. (2003). What is the unity of consciousness? In A. CLEEREMANS (Ed) *The unity of consciousness: Binding, integration, and dissociation*. Oxford: OUP.
- BROOK, A. (2003). The unity of consciousness. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/consciousness-unity/>
- DAINTON, B. (2000). *The stream of consciousness*. London: Routledge.

- DAINTON, B. (2004). Reply to Gilmore: Unity and introspectibility. *Psyche*, 10 (1).
- FARAH, M. (1990). *Visual agnosia*. Cambridge, MA: MIT Press.
- HURLEY, S.L. (1998). *Consciousness in action*. Cambridge, MA: Harvard.
- LOCKWOOD, M. (1989). *Mind, brain and the quantum*. Oxford: Basil Blackwell.
- LOCKWOOD, M. (1994). Issues of unity and objectivity. In C. PEACOCKE (ed.) *Objectivity, simulation and the unity of consciousness* (pp. 89-95). Proceedings of the British Academy, 83. Oxford: OUP.
- MARKS, C. (1980). *Commissurotomy, consciousness, and unity of mind*. Cambridge MA: MIT Press.
- MOOR, J. (1982). Split-brains and atomic persons. *Philosophy of Science*, 49: 91-106.
- PEACOCKE, C. (1994). The issues and their further development. In C. PEACOCKE (ed.) *Objectivity, simulation and the unity of consciousness* (pp. xi-xxvi). Oxford: OUP.
- PUCETTI, R. (1981). The case for mental duality: evidence from split-brain data and other considerations. *Behavioral and Brain Sciences*, 4, 93-123.
- ROSENTHAL, D. (2003). Unity of consciousness and the self. *Proceedings of the Aristotelian Society*, 103, 325-52.
- SHOEMAKER, S. (1996). Unity of consciousness and consciousness of unity. In S. SHOEMAKER *The first-person perspective and other essays* (176-97). Cambridge: CUP.
- SHOEMAKER, S. (2003). Consciousness and co-consciousness. In A. CLEEREMANS (Ed) *The Unity of Consciousness: Binding, Integration, Dissociation*. Oxford: OUP.
- TREVARTHEN, C. (1974a). Analysis of cerebral activities that generate and regulate consciousness in commissurotomy patients. In S.J. DIAMOND & J.G. BEAUMONT (eds.) *Hemisphere function in the human brain* (pp. 235-61). London: Elek Science.
- TREVARTHEN, C. (1974b). Functional relations of disconnected hemispheres with the brain stem, and with each other: Monkey and man. In M. KINSBOURNE & W. L. SMITH (eds.) *Hemispheric disconnection and cerebral function* (pp. 187-207). Springfield: Ill: Charles C. Thomas.
- TYE, M. (2003). *Consciousness and persons*. Cambridge, MA: MIT Press.