

Consciousness as a Guide to Personal Persistence

Barry Dainton
Department of Philosophy
University of Liverpool
7 Abercromby Square
Liverpool, L69 7WY
bdainton@liv.ac.uk

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

This paper appeared in (2005) *The Australasian Journal of Philosophy*, 83/4: 549-71. Please consult the published version for purposes of quotation.

That with which the consciousness of this present thinking thing can join itself, makes the same person, and is one self with it, and with nothing else; and so attributes to itself and owns all the actions of that thing as its own, as far as that consciousness reaches, and no farther; as every one who reflects will perceive.

Locke

1. Introduction

In this paper we put the case for an approach to personal identity that has been largely ignored in recent years. The approach in question is Lockean (or neo-Lockean) to the extent that it construes personal identity in terms of relationships between mental states and capacities, but it parts company with mainstream Lockean accounts of personal identity in focusing on *phenomenal* rather than *psychological* continuity. Whereas the latter

is a causal relationship, the former is purely experiential. Roughly speaking, experiences are related by phenomenal continuity when they belong to unified streams of consciousness of the sort we generally enjoy.

We believe the shift to an experiential framework is not just another option for Lockeans, it is an option that standard Lockean methodology renders obligatory. In motivating their account Lockeans rely on imaginary cases. When presented with a scenario in which bodily continuity goes one way and mental continuity goes another, many of us find it intuitively obvious that the person goes with the latter rather than the former. Lockeans conclude that we should define our persistence conditions in terms of mental continuity. While not disagreeing with this, we think this line of reasoning can be taken a step further: when phenomenal and psychological continuities are put to the same sort of test, psychological continuity fares as badly as physical continuity.

Of course, there are those who question the use of imaginary cases in evaluating accounts of personal identity. Of particular importance here is a famous puzzle due to Bernard Williams, in which our intuitions are pulled in entirely different directions by different presentations of the same putative mind-transfer scenario. We argue that Williams' scenarios are able to generate conflicting responses only because no mention is made of phenomenal continuity, and that when the latter is clearly distinguished from psychological continuity all ambiguities concerning the outcomes of the procedures are eliminated. Recognizing that phenomenal continuity is a *sui generis* form of mental continuity thus serves a double function: it offers the prospect of a more compelling variant of a Lockean account of our persistence conditions, and it neutralizes a serious threat to Lockean methodology.

Establishing this much takes up the first half of the paper. In the second half we explore the prospects of developing a viable phenomenal account of personal identity. Of particular importance is the question of how to deal with interruptions in phenomenal continuity. After considering and finding fault with proposals in the (meagre) literature, we propose a solution of our own. But rather than defend any one account of personal identity in any detail, our main aim is to establish that the phenomenal approach merits the attention lavished on the more orthodox alternatives.

2. A Method and a Menace

By separating in thought features that cannot be parted in practice, we can reasonably hope to distinguish what is important from what is not, or what is essential from what is merely contingent. Hence the widespread use of imaginary cases in the literature on personal identity. Various different continuities – material, mental, and biological – run

through a typical human life. These continuities may be inseparable in practice, but they are eminently separable in thought, and so we are invited to reflect on imaginary scenarios in which a person is subjected to some unusual treatment, e.g. brainwashing, teletransportation, a brain or body transplant, the effect of which is to remove one of the normal concomitants of personal persistence. If it seems evident that the person in the scenario survives the treatment, then it can be argued that the element removed is not a necessary condition of personal persistence. The hope is that the repeated use of this method, with a wide variety of scenarios, will finally yield conditions for personal persistence which all will agree are not only necessary but sufficient.

Although using imaginary cases in this way – we will call it ‘the Method’ – may seem to have a clear rationale, we believe it would be a mistake to expect too much from it. There are several reasons for this, but one consideration looms larger than the rest. We certainly want an account of our persistence conditions that is consonant with our deepest instincts about what is survivable and what isn’t – and theories that are in harmony with our intuitive responses to imaginary scenarios will obviously score more highly in this regard than those that are not – but we also want an account that meshes with our wider philosophical commitments. A theory that is attractive on the intuitive level will not be acceptable if we cannot bring it into line with our other beliefs about the nature of things. People who bring different metaphysical commitments to the table when assessing accounts of personal identity may well reach different verdicts on their overall merits, even if they agree on the degree to which each account is intuitively appealing. Since opinion on many central philosophical issues remains divided, and is likely to remain so, it would be unwise to expect a convergence of views on personal identity any time soon.

Nonetheless, while the Method may not be able to give us everything, it can give us something. By providing us with a general indication of the kind of account which will prove most appealing at the intuitive level, the Method can tell us something about the most appropriate *basis* for further philosophical investigation and development; it can supply us with a point of departure and an indication of the direction in which to move.

However, there are those who doubt the Method can serve even the modest function we propose. In his influential paper ‘The self and the future’ Bernard Williams describes a series of imaginary cases in which mental and bodily continuities diverge, in orthodox Lockean fashion. What is novel – and, indeed, threatening – about Williams’ cases is that our intuitions are pulled in *opposing* directions at once.

The following pair of scenarios is in certain respects simpler than those employed by Williams, but they suffice to illustrate the point.

(S1) After a long and successful career as a subversive, you have finally been apprehended by the authorities, who are eager to interrogate you about your accomplices. Unfortunately for you, the authorities in question prefer to use traditional methods: brutal but effective physical torture. You are informed that in order to avoid leaving incriminating marks on your body, you will be relocated in a different body; the torture will then be carried out; you will be returned to your original (and unblemished) body once a satisfactory confession has been extracted. Thanks to recent neuro-technical advances, the body-transfer no longer requires a brain-transplant: a brain-state transfer device will do the job instead. This machine is able to copy the psychological states (memories, beliefs, intentions, personality traits, and so on) from one brain to another brain. A helmet is placed on your head, and the switches are thrown. You wake up. Although a little nauseous, and clearly in a different body, you feel very much like your usual self. The torture, when it comes, is as bad as you feared.

(S2) Your long and successful career as a subversive is about to end: you realize that your arrest is imminent. You also know what to expect when apprehended: brutal torture. Your collaborators tell you not to worry. They have got their hands on a brain-state transfer device. They tell you that thanks to this device, when the torture commences your brain will no longer house your memories, beliefs or personality traits. Your psychology will be put into storage, and your brain will be imprinted with a psychology copied from someone wholly ignorant of your doings. You are not greatly consoled by this prospect. Having a different set of beliefs and memories will surely not prevent *you* feeling the pain inflicted on your body. How could it? At best, if your own memories and beliefs are restored, you will not be able to remember the pain, but this will do nothing to alleviate it when it is inflicted. If you follow the advice of your well-meaning friends, it seems you will face a double trauma: torture compounded with drastic psychological manipulation – a complete brainwashing.

Given an appropriate narrative context, the stipulation that a brain-state transfer device shifts a person from one body to another can seem as natural and plausible as the

stipulation that it merely effects a drastic form of brainwashing on a subject who remains in their original body.¹ What should we make of this *embarrass de richesses*?

Rovane advocates a partial abandonment of the Method: 'we cannot suppose that the issue of personal identity is resolvable by intuition alone, as all Lockean thought experiments try to do' [1998: 44]. Gendler suggests that Williams has shown that our ability to find unusual scenarios *intelligible* can outrun our ability to make *judgments* about them: 'We may well feel that a scenario is perfectly coherent, without knowing what we would do or say were we to encounter it. In such circumstances, our evaluation of the case is likely to depend on how the case is presented.' [1998: 604].

Williams himself tentatively favours the brainwashing interpretation, arguing that our ability to make sense of this interpretation suggests that deep down we recognize that our identity is fundamentally bodily. This is why our fear of future pain can extend through any amount of psychological disruption – why we can envisage ourselves undergoing the most radical alteration of beliefs, personality and memories yet suffering thereafter. While there is clearly something to be said for Williams' line, it does have the drawback of downplaying the fact that our intuitions are also pulled in the other direction.

There is an alternative response, one which gives equal weight to each the different directions in which our intuitions can be pulled. As Mark Johnston has noted [1987], if we assume *each* of Williams' scenarios reveals a *part* of our understanding of our persistence conditions, we are led to the conclusion that neither bodily nor psychological continuity is either necessary or sufficient for our survival. According to Johnston, what Williams has stumbled across is the fact that our 'pure concept' of a person, the concept revealed by our reactions to the full range of imaginary cases, is the concept of a 'bare locus' of mental life, the persistence of which is unconstrained by both physical and psychological factors. It is *because* there are no determinate conditions on the tracing of any bare locus through any imaginable psycho-physical vicissitude that we can make sense of the different interpretations of Williams' scenario. But Johnston further maintains that we should reject the bare locus conception despite its intuitive attractiveness. Given the lack of constraints on the movements of bare loci, how can we be justified in supposing our

¹ The identity of the participants is stipulated in this (deliberately) question-begging way in Williams' second scenario, but not the first; we have opted for symmetry. So far as the Method is concerned, it does not matter whether the identity or non-identity of the participants in an imaginary case is stipulated in the course of the presentation, or left open. If identity claims are made, readers will make their own judgments as to their plausibility. We have opted for a symmetrical treatment to minimize needless complication.

ordinary evidence for personal persistence, which takes the form of mental and physical continuities, is in the least bit reliable? Johnston concludes that we should abandon the Method and recommends starting with 'scientifically validated common sense', the position that we are animals of a particular species.

3. A Third Continuity

We believe that none of the responses just outlined is obligatory. To appreciate the real import of the Williams scenarios it is necessary to bring into consideration a form of mental continuity that has been largely overlooked in the personal identity literature: phenomenal (or experiential) continuity.

Psychological continuity, in the form favoured by contemporary neo-Lockeans, needs no introduction. The base ingredients are psychological states, such as beliefs, desires, intentions and memories, and the diachronic unifying agent is causation. Roughly speaking, a later person is psychologically continuous with an earlier person only if their psychological states are causally dependent upon the earlier person's in the right sorts of ways. Speaking slightly less roughly, an earlier person is 'psychologically continuous' with a later person if and only if they are linked by overlapping chains of direct 'psychological connections'. The latter consist of earlier and later instances of psychological states that possess the same psychological character, and are appropriately causally related. Note the lack of reference to *material* continuity in the definition of psychological continuity. It is the essentially causal character of the latter that secures its conceptual (or at least imaginative) independence from the former. We have already encountered an illustration of this. Williams' brain-state transfer device does not physically remove neural circuits from one brain and implant them in another; rather, it forms a representation of the psychological condition of one person's brain and, on this basis, causes another brain to enter into a similar psychological condition. Psychological connectedness is thus preserved in the absence of material (neural) continuity.

Phenomenal continuity is entirely different. The base ingredients are experiences (phenomenal states), rather than dispositional states (such as beliefs, intentions, and memories), and the unifying agent is *phenomenal connectedness*, rather than causal dependency.

In speaking of phenomenal connections and connectedness we are simply referring to the ways in which conscious states occur together within streams of consciousness. This mode of connectedness is perhaps most obvious in the synchronic case. Think of your current conscious thoughts, bodily sensations (such as pains or itches) and your experiences (such as the sound of the radio, and the visual experience of these words).

These experiences do not simply occur on their own as isolated states; rather, they occur together within your overall consciousness as parts or aspects or components of an overall stream of consciousness. This experienced togetherness, this unity-within-consciousness, is how phenomenal connectedness is manifest at any one time. In saying that an auditory and a visual experience are 'experienced together' we do not mean to imply that their unity depends on the occurrence of some additional higher-order experience. Higher-order states of consciousness do exist, but our experiences at any one time are unified even when we are not thinking about them. Phenomenal connectedness is simply that relationship of experienced togetherness that holds between all the diverse contents of a typical state of consciousness at a given time, higher-order thoughts included.

Phenomenal connectedness also obtains *over* time. A typical stream of consciousness is not a succession of discrete experiential atoms, far from it. Each brief phase of a stream of consciousness is *experienced* as flowing into the next. Think of what it is like to suffer a prolonged toothache, or to hear an extended tone played on a flute, or to watch a balloon float slowly across the sky. Each phase of your experience merges seamlessly with the next, and the next – indeed, so seamless is the flow that the division of experiences of this kind into distinct phases is often entirely arbitrary. Diachronic phenomenal connections of this sort are a pervasive feature of our streams of experience.

Of course, diachronic phenomenal connectedness is a short-lived phenomenon – probably not more than a second or so, the duration of the so-called 'specious present' – whereas a typical stream of consciousness lasts for hours. But, echoing the distinction between psychological continuity and connectedness, we can say that experiences at different times that are not phenomenally connected are nonetheless phenomenally *continuous* with one another provided they are linked by an overlapping chain of direct phenomenal connections. A *stream* of consciousness is simply any collection of experiences whose simultaneous members are related by synchronic phenomenal connectedness, and whose non-simultaneous members are related by phenomenal continuity.

Note that phenomenal unity and continuity cannot be reductively analyzed (at least in any remotely obvious way) in terms of beliefs or memories. Since, as we have just noted, our experiences remain unified irrespective of whether we are thinking about them or not, it is implausible to suppose that synchronic phenomenal unity is a product of conscious beliefs. It is even more implausible to suppose that it is a product of unconscious beliefs. So far as the diachronic unity of consciousness is concerned, some philosophers – e.g., Horwich [1987: 35], Mellor [1998: 122] – have maintained that this is entirely a product of short-term memory, but from a phenomenological perspective, such accounts are highly unrealistic. Although we can certainly remember experiencing

change and persistence, we can also *experience* change and persistence – and we do so all through our waking hours.

This brief sketch of synchronic and diachronic phenomenal unity leaves a good many questions unanswered – for differing approaches to synchronic unity see Dainton [2000] and Bayne and Chalmers [2003] – but for present purposes the answers to these questions are largely irrelevant. What matters is that phenomenal connectedness be recognized as a distinctive form of unity that is conceptually distinct from both physical and psychological modes of unity. Whereas psychological continuity is a purely *causal* relationship, its phenomenal counterpart is a purely *experiential* relationship.²

Having arrived at a reasonably clear understanding of the differences between these distinct forms of mental continuity a new spectrum of imaginary cases becomes available. As we have already seen, it is easy to imagine psychological and physical continuities coming apart: Williams' brain-state transfer device accomplishes precisely this. It is equally easy to imagine psychological and phenomenal continuities coming apart. Perhaps the easiest way of doing so is to imagine a brain-state transfer device being used to implant a new psychology into your current brain whilst you remain fully conscious throughout, and your stream of consciousness continues on without any rupture in experienced continuity. At the start of the procedure you have your current memories, beliefs, intentions and personality traits; by the time the procedure is complete your consciousness is being informed by a numerically distinct and causally unrelated set of memories, beliefs, intentions and personality traits. Since the replacement psychology is not causally dependent upon your original psychology, psychological continuity is ruptured, but phenomenal continuity is not.

What would it be like to undergo such a procedure? There are various possibilities, depending on how quickly the replacement occurs, and how different your new psychology is from your original psychology. In the limiting case, the replacement psychology could exactly resemble your own, in which case your stream of consciousness

²Or at least, this is how psychological continuity is widely construed in the personal identity literature. There is, of course, a perfectly respectable sense in which experiential relationships are themselves 'psychological', in a more general sense of the term. It might be objected that the distinction between psychological and phenomenal continuity collapses in the context of purely causal analyses of phenomenal states and properties. Since we find such reductive analyses implausible, we do not find this objection worrisome. But even those inclined to accept this form of reductionism must still recognize continuities in *experience* and continuities in *beliefs, memories and intentions* as being distinct.

might be entirely unremarkable: you would continue to have a normal range of perceptual experiences, you would not experience any emotional upheaval, your thinking would be just as coherent as it usually is. Provided that at each moment you have a functioning psychological and/or cognitive system, it is hard to see why the absence of the normal relationships of diachronic causal dependency between the elements of this system should necessarily impinge on the character of your experience in any way. (It would be otherwise if these causal relationships themselves had an introspectibly discernible character, but as Hume observed, they do not.) If, by contrast, your new psychology were markedly different from the original, some disruption in cognitive functioning would be inevitable. You might find yourself unable to think clearly for a while; you might find yourself thinking first in one language, then another. You might find you entirely lose the capacity for conscious thought. But this cognitive fragmentation could occur within a stream of consciousness that in other respects is phenomenally continuous: we can, for example, suppose that your auditory, visual and tactile experiences exhibit their usual continuity.

We can take things a stage further by imagining phenomenal and physical continuities coming apart. Having recognized the significance of phenomenal continuity we can help ourselves to a second fictitious device: a *streamal diverter*. This device can be used to divert the flow of consciousness from one brain to another, in an instant. Although it preserves phenomenal continuity it has no impact on psychological continuity; whether a subject's psychology accompanies their stream of consciousness into a new body and brain depends on whether a brain-state transfer device is used along with a streamal diverter. Again, quite what it would be like to have one's stream of consciousness shifted in this fashion will depend on the character of the new body, and whether or not one's original psychology accompanies one's consciousness. The transfer could well be an extremely disturbing and disorienting affair. But this is not a bar to phenomenal continuity being fully preserved: as most of us know, it is perfectly possible to *experience* extreme confusion and disorientation.

The fact that it is possible to imagine scenarios in which physical, psychological and phenomenal continuities come apart does not entail that such separations really are possible, nomologically or logically. Different positions in the philosophy of mind have different implications with regard to this issue, and we are remaining neutral on these matters. Such scenarios can nonetheless be of use, or so Lockeans traditionally hold, in isolating the kind of continuity which is most closely associated with our continued personal existence. That phenomenal continuity trumps both physical and psychological

continuity in this key respect strikes us as obvious – indeed, we tacitly assumed as much over the last couple of paragraphs – but the issue requires more careful treatment.

4. The Inseparability Thesis

Consider the following scenario:

(S3) Brain state transfer devices and streamal diverters are used in tandem: Smith's psychology and stream of consciousness are transferred to your brain, your psychology and stream of consciousness are transferred to Smith's brain.

The re-locating of both forms of mental continuity constitutes a *full mental transfer*; it has the same mental consequences as a brain transplant, but with none of the mess. A transfer of this type leaves no room for doubt: you move into Smith's body, Smith moves into yours, and you both remain fully conscious during the transfers. If prior to the procedure you are told that a painful treatment will be inflicted on Smith's body shortly after the transfer is completed, you have every reason to be afraid: the resulting pain sensations lie directly upstream of your current experiences. Smith, by contrast, has nothing to fear at all; he knows that when the torture begins *his* consciousness will be elsewhere. Now consider:

(S4) Brain state transfer devices are used to transfer Smith's psychology to your brain, and your psychology to Smith's brain. Streamal diverters are not used; phenomenal continuity is unaffected; your stream of consciousness continues to be sustained by your original brain, and Smith's by his.

The outcome of this *partial mental transfer* is, we suggest, also entirely unambiguous: each person remains with their original body. The envisaged procedure amounts to no more than a form of brainwashing, a rapid psychological change that is carried out while both subjects remain fully conscious. We have already outlined the case for supposing that phenomenal continuity is compatible with psychological discontinuity, and we have seen that it is intuitively plausible to suppose that psychological discontinuity can be survived. The second Williams scenario suggests as much, and any lingering doubts on this score are wholly dispelled by the stipulation that phenomenal continuity remains uninterrupted all through the psychological manipulations. We can suppose that throughout the procedure both you and Smith are watching a news channel on television (on different sets, in different rooms); the transfer induces feelings of puzzlement and unease, but no loss or disruption of phenomenal continuity. It is, we suggest, highly implausible to

suppose that a subject's stream of consciousness could continue without taking that subject with it.³ This point is sufficiently important to merit a label of its own, thus:

The Inseparability Thesis: self and phenomenal continuity cannot come apart: all the experiences in a single (non-branching) stream of consciousness are co-personal.

If you are still not entirely convinced, put yourself in Smith's position, in advance of the transfer being carried out. You know that your body will soon be tortured, and that agonizing pain will ensue. You also know that these pain sensations will be phenomenally continuous with your current experiences; they will be connected to the latter by an unbroken chain of directly experienced transitions: your current experience is going to *flow right into* all this pain. Knowing this, can you begin to take seriously the notion that these experiences will be experienced by anyone else *but* you? Is your conviction on this score affected in the least by the knowledge that when the torture begins, someone else who *believes* themselves to be Smith will not be feeling the ensuing pain? Surely not. How can someone else's delusions lessen your agonies?⁴

The contention that personal persistence and the capacity for conscious experience are closely linked can be further supported by reflecting on *zombification*. Suppose that we have a process that preserves a person's psychology – or at least the bulk of it – but removes their potential for phenomenal consciousness. The person's psychology will continue on, but henceforth none of their mental states will be phenomenally conscious. Now, suppose that you are faced with imminent zombification. How will you react? We suspect that many will respond to this fate in much the way that they respond to the prospect of imminent death: one aspect of your mind might continue, but you won't! (Of course, some argue that zombification is not possible. They might be right, but all we need here is the claim that zombification is conceivable – and it certainly seems to be.)

Let us move on. S3 suggests that the combination of phenomenal continuity and psychological continuity is person preserving, and that physical continuity by itself is not. S4 suggests that the combination of phenomenal continuity and physical continuity is

³ What about a situation in which a stream of consciousness fissions? We consider this possibility below.

⁴ Of course, delusions about phenomenal continuity are possible. So far as the Inseparability Thesis is concerned, what matters is the presence (or absence) of an uninterrupted flow of consciousness, of real phenomenal relations, not beliefs about such relations.

person preserving, and that psychological continuity by itself is not. Should we conclude that phenomenal continuity plus *either* psychological *or* physical continuity is person preserving? While this is certainly an option, so far as we can see, it is not one that has any intuitive appeal. Consider:

(S5) Streamal diverters are used to transfer Smith's stream of consciousness to your brain, and your stream of consciousness to Smith's brain. Brain state transfer devices are not used; psychological continuity is unaffected; your original psychology continues to be sustained by your original brain, and Smith's by his.

Having isolated phenomenal continuity we can easily make sense of the idea that a person can leave both their *body* and their *psychology* behind. When considering the psychological disruption wrought by the brain state transfer device in S4 it seemed clear that *you* would remain with your stream of consciousness, rather than moving over, with your psychology, into Smith's body and brain. In reaching this verdict, the fact that your stream of consciousness continued to be sustained by your original brain was surely of negligible significance, for just imagine: if in fact your body had been entirely annihilated during the transfer process, but your stream of consciousness flowed on just as before, maintained by the machinery of the streamal diverter, would *your* existence have been threatened? How could it have been? As is clear, the crucial influence in play here is the intuition enshrined in the Inseparability Thesis: selves and streams of consciousness cannot come apart. If the Inseparability Thesis overpowers all else in the context of S4, why not S5 also? At a purely phenomenological level, the streams of consciousness in S5 might well be exactly similar to those in S4: in both cases, your stream comes to be informed by Smith's psychology and his stream comes to be informed by your psychology. The upshot of *both* cases is identical: *you* stay with your stream of consciousness, and Smith stays with his.

5. Taking Stock

We are now in a better position to evaluate the conundrum posed by Williams. Our ability to find the original scenarios, S1 and S2, intelligible indicates that we are committed to the following pair of theses:

T1 Mental continuity is person-preserving, even across changes in brain and body.

T2 Loss of psychological continuity is survivable. Or as Williams put it: 'one's fears can extend to future pain whatever psychological changes precede it'.

Although both theses are intuitively compelling, they can also be brought into conflict, as Williams shows. If we equate mental continuity with psychological continuity, then a mental transfer which can seem to result in two persons exchanging bodies can *also* seem to amount to nothing more than two people remaining with their original bodies while undergoing a brainwashing. However, once the distinction between phenomenal and psychological continuities is recognized, the conflict disappears: T2 is certainly true, but so too is T1, provided that by 'mental continuity' we mean *phenomenal* continuity, either by itself or combined with psychological continuity. If phenomenal consciousness and self-identity are intimately linked, as we have argued, Williams' scenarios *should* be baffling; we are unsure what to make of them because we are left in the dark about what really matters from the point of view of one's continued existence. The reader fills in the details about phenomenal continuity for themselves, and the narrative structure of S1 and S2 leads them to fill in the relevant details in different ways in the two scenarios. But as the Smith cases illustrate, clarifying the fate of a subject's stream of consciousness also removes any doubt about the fate of the subject themselves: they invariably follow the flow of experience.

Once this point is taken on board, Rovane's Williams-inspired doubts as to the viability of the Method look to be unfounded. Likewise Gendler's suggestion that our ability to make sense of exceptional situations exceeds our ability to make reliable judgments about them. Perhaps there are such situations, but psychological transfer situations of the type Williams was concerned with are not amongst them, at least when fully and properly described: once *all* the relevant facts about mental continuity are supplied, we not only find such cases intelligible, we can also make confident judgments about their outcomes.

The situation with regard to Johnston's diagnosis is more nuanced. We are suggesting that our intuitive responses to the augmented (S3-S5) Williams cases point consistently and unambiguously in the direction of the Inseparability Thesis. The latter is plainly inconsistent with the bare locus view, at least if the latter is construed as the doctrine that there are no physical or mental constraints on our persistence. However, the Inseparability Thesis is perfectly compatible with a weaker version of the bare locus view, the doctrine that there are no physical or *psychological* constraints on our persistence. The fact that Johnston generally equates mental continuity with psychological continuity might lead one to conclude that he had the weaker doctrine in mind. But equally, this could be a mistake. In most of his formulations of the bare locus view he insists only on an absence of *necessary* conditions on the persistence of such things. Moreover, when arguing that the bare locus view, if true, would render our ordinary judgments about our persistence unreliable, he focuses the problem of what happens to us when we *lose*

consciousness. Although he fails explicitly to mention phenomenal continuity, the adoption of this tactic suggests that Johnston is fully aware of the force of the force of the Inseparability Thesis. His failure to draw attention to this is regrettable, but understandable: had he done so, his Williams-inspired argument against the viability of the Method would have been largely undermined.⁵

6. Moving Forward: Issues and Bridges

We have argued that rather than undermining the Method, Williams' thought-experiment motivates a phenomenological, experience-based, conception of personal identity. But perhaps this conception of ourselves is mistaken, or brings with it indefensible philosophical baggage. In the remainder of the paper we shall explore some of the issues one faces in attempting to develop and defend a 'phenomenalist' approach to personal persistence. The issues here are complex and wide-ranging, and we can do little more than identify some of the more central questions and promising avenues for further investigation. Note that the phenomenalist approach is not alone in facing tough questions: approaches based on psychological and physical continuities also face well-known objections. Those who find the phenomenal approach appealing should be granted as much of an opportunity to develop their position as proponents of other approaches have had.

Three questions strike us as particularly pressing. The first is this: How simple can a stream of consciousness get before it leaves its owner behind? How primitive can the contents of a stream become before it ceases to support one of us? Call this the 'depth problem'.

The first thing to note about the depth problem is that standard psychological accounts of personal identity are confronted by a precisely parallel question. Psychological theorists need to say how primitive one's mind can become before one ceases to exist (and at what point in mental development a person comes into existence). If the depth problem is a problem for the phenomenalist then it – or its analogue – is also a problem for psychological accounts.

As for solutions, no doubt opinions will differ. Unger [1990] is unable to conceive of himself surviving if his cognitive abilities were reduced to canine levels, even if his capacities for sensory consciousness remained largely intact. We incline to the view that

⁵ There are, of course, objections to the use of imaginary cases in the context of personal identity that are not inspired by Williams, e.g., Wilkes [1988] – see Snowdon [1991] for a reply. We do not deal with these here.

no cognitive sophistication is necessary for our survival, and that we could survive with a consciousness of the simplest of forms, e.g. a few basic bodily feelings (it is arguable that we all enjoyed a consciousness of this form prior to birth). The Inseparability Thesis is relevant here: provided your stream of consciousness flows on, wouldn't *you* go with it, irrespective of how simple its character becomes? But we are not going to try to defend this line here. We simply note that there is room for debate.

A second issue that must be addressed is an objection of principle. In taking phenomenal continuity as our guide to self-persistence, we are assuming that streams of consciousness can be identified and individuated, at and over time, without prior reference to their owners. Some find this problematic. According to van Inwagen, 'to imagine whether a certain situation contains a continuous consciousness we have to find out first whether a certain situation contains a continuously existent thinker. We can't do things the other way around. We can't find out whether the situation contains a continuously existent thinker by first finding out whether it contains a continuous consciousness' [1990: 206]. Van Inwagen seems to think that the only sense to be made of the notion of a continuous consciousness is in terms of the notion of a subject of consciousness being continuously conscious. And his reason for this is that experiences are not independent entities, but modifications of substances – conscious subjects.⁶

We noted at the outset that accounts of personal identity need not only be intuitively plausible but also metaphysically viable in order to be acceptable. But we also noted that views on fundamental matters of metaphysics differ. This is one of the occasions on which differences in metaphysical approaches makes a difference. In contrast to van Inwagen, we are not at all convinced that one can identify a continuous consciousness (or a continuous capacity for consciousness) only by first identifying a continuously conscious subject of experience. Indeed, our discussion of phenomenal continuity had as one of its aims the goal of motivating the thought that phenomenal continuity can be understood independently of the notion of a subject of experience. In saying this we are not suggesting that phenomenologists must deny that subjects of experience exist. They are, however, committed to the claim that the persistence conditions of subjects experience can be characterized in phenomenal terms. And although we are as yet unsure as to precisely what form this characterization should take – see below – we have yet to see a compelling case for supposing an account along these general lines to be in principle impossible.

⁶ The complaint is by no means uncommon; see, for example, Strawson [1967: 169] and Lowe [1996: 25].

A third challenge for the phenomenalist, one we shall be concerned with for the remainder of this paper, is to overcome the problem posed by interruptions in consciousness. How must streams of consciousness separated by a temporal interval and not linked by phenomenal continuity – streams such as the one you had yesterday and the one you have now – be related if they belong to the same person? Can such streams belong to the same person, if we take consciousness as our guide to personal persistence? We call this the *bridge problem*.

Solutions to the bridge problem can be divided into two camps: those that try to construct a bridge (constructivists), and those that do not (non-constructivists).

Non-constructivists claim that we exist only so long as we are conscious. Clearly, if individual persons are confined to uninterrupted streams of consciousness, bridges linking streams are not required. How long does a typical human self endure on this view? Common sense suggests that our streams of consciousness are usually around sixteen hours duration, but non-constructivists have generally found common sense wanting in this regard. Some say we are conscious only for very brief periods (less than a minute), others tend to the other extreme, and subscribe to the view that we can remain continually conscious, in some form or other, for many years (perhaps infinitely many).⁷

We are not much attracted to non-constructivism. Very short-term versions of the position give rise to what is obviously a counter-intuitive view of our life-spans: few of us are inclined to view the onset of dreamless sleep with the trepidation with which we view death! Longer-term non-constructivists leave us with a more appealing view of ourselves, but their view is hostage to claims about the character of our consciousness that are at best highly speculative. In light of these difficulties we think that phenomenalist would do well to consider bridge building.

7. Conventions and Similarities

Perhaps the most minimalist approach to bridge building appeals to nothing more than our practices. Perhaps there is nothing more to two streams of consciousness being co-personal other than the fact that we take them to be. Waking up one morning, I identify a

⁷ Strawson [1997] argues that we exist for at most a few seconds. Descartes and Leibniz are perhaps the best-known philosophers who have claimed that we remain (dimly) conscious during periods of so-called unconsciousness. Perhaps surprisingly, Husserl was also of this opinion – indeed, he believed we are all dimly conscious for an infinite period of time preceding our (bodily) birth, and will remain in the same condition for an infinite period of time after our (bodily) deaths – see Smith [2003: 94-4, 200-10].

certain previous stream of experience as mine. My confidence in this identification is bolstered by the fact that other people invariably concur with my choices. But there is nothing underlying or justifying this practice, it is simply the way we do things. Our inter-streamal survival is grounded in nothing more than convention.

This view has had few takers, and this is no surprise.⁸ While there are many arbitrary conventions, it is hard to believe that our conventions for assigning experiences to persons fall into this category. It is natural to suppose, rather, that personal identity is founded on connections that are independent of our conceptual resources and linguistic conventions. The boundaries which mark the beginning and end of our existence are more akin to the boundaries of a continent than they are to the boundaries of a nation – or so, at least, we are inclined to think.

A second approach to bridge building can also be regarded as minimalist, despite being non-conventionalist. The approach in question, one which has found favour in some quarters, appeals not to any form of external connection between earlier and later streams, but to a purely *internal* relation, namely similarity in qualitative phenomenal character. James wrote: 'even when there is a time-gap the consciousness after it feels as it belonged together with the consciousness before it, as another part of the same self' [1952: 154]. He suggested that each of our current states of consciousness possesses a highly distinctive personal quality, a qualitative feeling we are all intimately familiar with, and which serves to bind streams into co-personal ensembles:

whatever past feelings appear with those qualities must be admitted to receive the greeting of the present mental state, to be owned by it, and accepted as belonging together with it in a common self. This community of self is what the time-gap cannot break in twain, and is why a present thought, although not ignorant of the time-gap, can still regard itself as continuous with certain portions of the past. [James 1952: 155]

In more recent years the similarity account has been defended by Sprigge [1988] and Nathan [1997].

If qualitative identity makes for numerical identity it is logically impossible for there to be an exact duplicate of you who *isn't* you. We find this consequence objectionable. Setting this point aside, one might worry whether our conscious lives possess a distinctive

⁸ Among the few: Braddon-Mitchell and Miller [2004]; space precludes a proper assessment of their arguments here.

subjective *flavour* of this kind? Perhaps in some cases, some of the time. But in all cases, all of the time? We have our doubts. And what evidence could one possibly have that one's own stream of consciousness possessed a unique phenomenal character?⁹

8. Phenomenal Joinability?

Another approach worth considering begins with the notion that earlier and later streams are co-personal only when the two streams *could* have formed a single uninterrupted stream. John Foster has proposed solutions to the bridge problem along these lines.¹⁰ One version runs thus:

Given two total streams A and B, where A is earlier than B, and where there is a temporal interval between them, let us say that A and B are directly joinable if and only if there is something which ensures (whether logically or nomologically) that, with B held constant, a hypothetical continuation of A to the time when (or just after) B begins would join up with B; and let us say that two streams are indirectly joinable if and only if they are connected by a series of total streams whose successive members are directly joinable. [1991: 251]

We can then say that streams are co-personal if and only if they are directly or indirectly joinable.

The approach is elegant, but does it work? We have our doubts. The central difficulty lies with the notion of potential joinability: is this potential logical or nomological? In his earliest attempt at bridge-building Foster took the continuity of the subject to rest on laws of nature [1979: 179]. But do the natural laws governing *our* minds conform to this requirement? It seems unlikely. Consider Jim who woke at eight o'clock this morning feeling unusually fresh after seven hours of sound sleep. According to Foster's theory in its nomological guise, it seems that if on waking Jim is the person he remembers being the night before, there must be natural laws which would have ensured that the experiences he actually had on waking up would have been just as they were, even if he had spent the night without sleep. But in fact the reverse is the case: had Jim spent the night drinking coffee natural laws would have *prevented* his morning experiences being as they were.

The problem here is a quite general one: for given any pair of putative co-personal streams of (human) consciousness, it is unlikely that the later member of the pair would

⁹ We thank a referee for this point.

¹⁰ See Foster [1979; 1982; 1985; 1991]. We cannot here supply a full account of the differences between Foster's various proposals, which occur in quite different metaphysical contexts.

exist unchanged if the earlier member were extended to meet it. Indeed, in the majority of actual cases, it may well be nomologically impossible for one's earlier stream of consciousness to be extended while one's later stream remains in place. In which case, since two such streams are not potentially joinable (or co-conscious), they would not be co-personal, and the majority of the streams we would ordinarily take to belong to the same person would be nothing of the sort!

Let us turn to consider the second option. What purely logical factor might ensure that stream A would meet stream B, if the latter remained the same and the former were extended? As Foster himself notes, there is one compelling answer to this question, namely that A and B would be so related if they belonged to one and the same subject. Given certain assumptions, this would undoubtedly be the case. But the account is now manifestly circular: to say that streams of consciousness are potentially joinable only in virtue of being co-personal is to employ the very relationship we are seeking to elucidate in the elucidation itself. If we want an informative account of streamal co-personality in our world we must look beyond logical joinability.

9. Phenomenal Substances

Despite their differences, the similarity and joinability solutions to the bridge problem share a common feature: each attempts to solve the problem by finding a relationship that holds between temporally separated streams of consciousness, a relationship which suffices to render such streams co-personal. Solutions of this type do not construe experiential subjects as *persisting through* periods of unconsciousness; in effect, subjects who lose and gain consciousness enjoy an intermittent existence. This is clearly a counterintuitive result – most of us do not equate losing consciousness with ceasing to exist, even if only temporarily. A more attractive solution would construe persons as a distinctive kind of being, a kind whose members are able to exist and persist in both conscious and non-conscious modes. But can the notion that persons are things that can persist through periods of unconsciousness be reconciled with the doctrine that consciousness is our guide to personal persistence?

The reconciliation is possible. The key lies in a change of perspective. Rather than regarding persons as primarily things that *are* conscious, we regard them as things that *are capable* of being conscious, as beings that possess *capacities* for experience. A typical human person possesses a vast range of experiential capacities, only a few of which are active at any one time. When a person becomes unconscious, none of their experiential capacities are active, but the capacities nonetheless remain in existence: the irretrievable loss of the capacity for consciousness is what differentiates being merely unconscious

from being dead. Or at least, this is a plausible view for anyone who inclines to an experience-based account of what persons are.

From this new perspective, defining personal persistence conditions in experiential terms looks to be comparatively straightforward: we can trace the persistence of a person by reference to the existence and persistence of capacities for a unified consciousness, capacities that are sometimes exercised, sometimes not. This core idea can be developed in different ways; all we can provide here is an outline of just one of these.

We can start with the very general notion of an *experience producer* [EP], i.e., any object or system which is capable of generating experience, of one or more specified kinds, when appropriately stimulated, by virtue of the laws of nature. Your brain currently qualifies as an EP, but so too (probably) are lesser parts of it, e.g., the right and left halves. Since to qualify as an EP an object must actually be able to produce experience, your brain is not essentially an EP: it seems unlikely that a typical human brain is capable of producing experience during the first few weeks of its (foetal) existence, and adult brains can survive damage which completely eliminate their experiential capacities. Of course there may well be EPs very different from your brain or any of its parts.

Now, EPs that are active and generating phenomenally connected experiences should obviously be regarded as co-personal, but what of dormant EPs? While we cannot appeal to experiences they *are* producing, we can appeal to the experiences they *would* produce if they were active: EPs that would produce phenomenally connected experiences if they were active should also be regarded as co-personal. Since phenomenal connectedness, actual and potential, obtains diachronically as well as synchronically, we have all the ingredients we need to solve the bridge problem.

Spelling things out in a little more detail, let us say that for any time t , a collection C of EPs are *synchronically E-linked* if and only if the experiences produced by the members of C if they were active and experience-producing at t would be phenomenally unified, or co-conscious. Call any synchronically E-linked collection of EPs an *E-system*. Let us further say that an E-system S_1 at t_1 and an E-system S_2 at t_2 are *directly streamally linked* iff the experiences each system would produce if some or all of their component EPs were active would be diachronically phenomenally connected (i.e., occur within a single specious present). Since diachronic phenomenal connectedness is a short-term affair, so too is direct streamal linkage. But this is of no consequence: let us say that E-systems at distinct times are *indirectly streamally linked* iff they are connected to one another by overlapping chains of direct streamal linkage. Finally, call EPs that are synchronically E-linked, or that belong to E-systems that are streamally linked, directly or indirectly, *E-related*.

Experiences and experience producers can now be assigned to persons in a natural and straightforward way: (i) EPs are co-personal only if they are E-related, and (ii) experiences are co-personal only if they are the product of E-related EPs. We have something more. E-related EPs can be regarded as constituting *entities* of a distinctive kind: objects whose existence and persistence conditions are specified in phenomenal terms. We shall call objects of this kind *phenomenal substances*.

The notion that persons are phenomenal substances has a good deal to be said for it.¹¹ Since their existence and persistence conditions are well-defined, phenomenal substances look to be genuine *things*, and they are also things which can persist through periods of unconsciousness – just like us. Moreover, since their persistence conditions are defined in phenomenal terms, it is at least conceptually possible for phenomenal substances to survive any amount of physical and psychological disruption. Since we can easily envisage ourselves surviving both types of disruption, as Williams' scenarios illustrate, it is hard to see how we could be anything *but* phenomenal substances!

This said, there are further issues concerning the precise relationship between persons and EPs, and EPs themselves. Consider the case of a person *P* with a moderately modular brain; at any one time *P*'s mind consists of a number of distinct but E-related EPs, and these EPs are replaced at regular intervals. Call the complete collection of EPs which belong to *P* during the course of the latter's life '*C*'. How is *P* related to *C*? The simplest option is to hold that *P* and *C* are numerically identical. However, there is an obstacle to this identification. If we understand 'collection' in the usual way, it makes no sense to suppose that *C* could have had a different membership. Is it plausible to suppose that *P* is similarly constrained? Not in the least. *P* could have died earlier, and so possessed fewer EPs; equally, *P* could have undergone more frequent module replacements and so possessed a larger collection of EPs. *P* and *C* thus have different modal properties. The difficulty is a familiar one – it arises in connection with all compound continuants, e.g. statues and their clay parts – and there are two main strategies for dealing with it. One option is to accept that objects can exactly coincide by sharing parts, but nonetheless remain distinct. On this view, it is true to say that a statue *is* the collection of clay pieces from which it is made, but only if we take 'is' to mean 'is composed of' or 'is constituted from', rather than 'is identical with'. The second option is to hold that objects which coincide in this fashion are in fact numerically identical, albeit only contingently so. We do not take a stance on this issue, we note simply that the relationship between *P* and *C*

¹¹ Though we stress again that here we are only providing the outlines of an account – many issues (e.g., the synchronic transitivity or otherwise of E-linkage) require further treatment.

can be construed in either of these ways: either *P* is merely composed of *C*, or *P* is numerically identical with *C*, albeit only contingently.¹²

There is a second issue on which neo-Lockeans are divided. Where do the synchronic boundaries of a person lie? In less abstract terms, in the case of embodied mental subjects such as the typical human being, is a person at any one time identical with (or composed of) their entire body, or just that part of it which actually possesses experiential (or mental) capacities? Some may incline to view that, strictly speaking, persons are their essential cores, others may incline to the view that persons are what they appear to be: embodied thinking beings. We tacitly adopted the 'core' view earlier, when we suggested that our own EPs are our brains, or parts of them. Another option is to hold that human EPs typically consist of human brains *and* bodies.

For those inclined to identify persons with their essential cores, there is a further option worth noting. We have suggested that a person persists through a period of unconsciousness by virtue of retaining the capacity for phenomenal consciousness; provided a person's capacities for consciousness endure, so too does the person. Now of course, generally speaking, capacities are possessed by *objects*, and it is this assumption that leads to the view that persons are phenomenal substances in the sense just specified: the latter being nothing other than objects which possess the capacity for generating phenomenally continuous consciousness.

Bearing all this in mind, a question can now be posed: So far as one's own personal survival is concerned, what matter most, the object that possesses the capacities, or the capacities all by themselves? The answer seems clear: *only the capacities matter*. Provided your capacities for experience continue to exist, so too do you, irrespective of what possesses or grounds these capacities. In standard Lockean fashion we can envisage imaginary scenarios in which your capacities for consciousness are unchanged, at least so far as their actual and potential manifestations and triggering circumstances are concerned, despite swift and dramatic variations in their underlying categorical substrate.

Needless to say, this realization gives rise to a question. Can we provide an account of personal persistence in terms of experiential capacities, rather than experience producers? In fact, we have already done so, in all but name. Since experiential capacities are a species of causal power, we could as easily refer to them as *experiential powers*. The

¹² This issue often (but not invariably) divides along 'endurantist' and 'perdurantist' lines. See Shoemaker [1984], Johnston [1991] for defences of the constitution view, Lewis [1971] and Noonan [1992] for the alternative.

characterizations of E-linkage, E-systems and E-relatedness provided above all work just as well if 'EP' is allowed to stand for 'experiential power' rather than 'experience producer'.¹³

Construing selves as power-systems, as enduring fields of experiential potential, is unquestionably appealing from the point of view of ontological economy. It may also be the destination to which our intuitive responses to imaginary cases ultimately leads. But not everyone will be happy with the notion that they are nothing but fields of potential – are such things really *things*? Moreover, causal powers are controversial entities in themselves, and the precise relationship between dispositions and their bases is particularly contested. Quite what a fully developed field theory of the self would look like is unclear – we note merely that it is a further avenue well worth exploring.

10. Further Issues

Adopting the phenomenalist approach may give rise to further options, but it also gives rise to further questions and potential difficulties. We can only address a few of these here, and then only briefly.

Suppose, for example, that the brain of Julius Caesar was removed from his body shortly before his death, wiped clean of all psychological traces and maintained in minimally functioning condition for the next two thousand years. Throughout this period Caesar's brain has the capacity to produce simple forms of sensory experience, but this capacity is never exercised. At the end of this period this brain's sensory systems are reactivated, and the resulting stream of consciousness is diverted into a new body and brain, where it is informed by the psychology of a typical 20th century person. Is the resulting person Caesar, or somebody else? Given the radical material and psychological discontinuities you may incline to the latter view. We incline to the former view. As we have already seen, there is reason to think the most drastic physical and psychological discontinuities *are* survivable, provided the capacity for phenomenal consciousness is preserved – as it is in the envisaged case. Are longer periods of dormancy harder to survive than shorter periods, other things being equal? It seems implausible. And of course, analogous problem cases for the physical and psychological continuity theories are easily devised.

A different kind of hard case is more problematic, for it threatens to undermine the claim that our persistence conditions are to be defined in terms of continuously existing potentialities for consciousness. Suppose there were a race of creatures who outwardly resemble human beings, but whose nervous systems undergo a restorative deep

¹³ See Dainton [1995] for a more detailed sketch of this type of account.

restructuring once per year, and that during this time these creatures no longer possess any experiential capacities. According to the definitions supplied earlier, the EPs either side of such an interruption are not E-related, and hence do not belong to the same person. Is this not highly implausible?

It is indeed, and the problem is of more than merely abstract interest, for we cannot be sure that our own EPs do not occasionally lapse in the way envisaged. One way of enabling survival in such cases would be to allow material or psychological continuities (if such remain) to take over when E-relatedness fails. Although this is an option, it has the counterintuitive consequence that persons are entities with disjunctive persistence conditions. A better option would be to find a way of securing identity in such cases without appealing to anything other than phenomenal relations and capacities. There is at least one way in which this might be done: we can appeal to *second-order* experiential capacities. Roughly speaking, a thing has a second-order experiential capacity if it doesn't have a first-order capacity, but will acquire one if left to its own devices, under normal background conditions. We can then say that a first-order capacity C_1 at t_1 and a second-order capacity C_2 at t_2 are streamally linked iff the experience C_1 would produce if it were active would be co-conscious with the experiences C_2 would produce if it were (contrary to actual fact) a first-order rather than a second-order capacity, and active. This modification is well within the spirit of the original proposal, which itself depends heavily upon counterfactuals.¹⁴

The hardest of hard cases is perhaps that of fission. Earlier we suggested that streams of consciousness and their subjects are inseparable. The possibility of personal fission puts this Inseparability Thesis under pressure. Suppose it were possible for a stream of consciousness to smoothly divide, without loss of phenomenal continuity. Let us further suppose that each of the resulting streams belongs to a separate *person* – as might well be the case, for example, after the successful transplanting of each of your cerebral hemispheres into a new body. If the post-fission people are distinct from one another – as they appear to be – they cannot both be identical with the pre-fission person. Given this, isn't it obvious that phenomenal continuity can fail to be person-preserving?

Not necessarily. There are three general interpretations of fission the phenomenalist can choose between, and only one of these undermines the Inseparability Thesis, and then only in a limited way. Some neo-Lockeans, e.g. Parfit [1971] and Shoemaker [1970], adopt

¹⁴ Accepting this modification has implications for the beginnings of personal existence: human foetuses (probably) only come to possess first-order capacities when they are several months old, they (arguably) possess second-order capacities from the moment they come into existence.

the view that post-fission persons would be numerically distinct from one another and the person who divided into them. They advocate a *non-branching psychological continuity* theory, according to which psychological continuity is always person-preserving provided unless it takes a branching form. A similar position is open to the phenomenalist, who can say that E-relatedness is person-preserving unless it takes a branching form. The Inseparability Thesis is thus conserved, except in the special case of fission. But on other views fission is *not* fatal. On Lewis' view [1976], fission merely marks the separation of two pre-existing people whose lives partially overlap. This interpretation of fission puts no pressure on the Inseparability Thesis, but it does undermine the claim that a single stream of consciousness at a given time belongs to a single subject. This admittedly counterintuitive result can be avoided by holding that fission involves a single person dividing into two people who are numerically identical with one another and the person who divided into them – see Dainton [1992]. Each of these views has its merits, but we leave their evaluation for another occasion.

12. Conclusion

Rather than defending any particular version of the phenomenalist account of personal identity, we have merely tried to identify some promising lines for further inquiry. Our main aim has been to establish that the phenomenalist approach ought to be taken seriously, alongside psychological and physical approaches. Whether it can overcome the various obstacles to which we have drawn attention in the last few sections (and no doubt there are others) is an open question. But what does seem clear is that the phenomenalist approach can resolve the puzzle generated by Williams's cases – something that neither physical nor standard psychological accounts of personal identity can do.¹⁵

References

- Bayne, T. and Chalmers, D. 2003. What is the Unity of Consciousness? in *The Unity of Consciousness: Binding, Integration, Dissociation*, ed. A. Cleeremans, Oxford: Oxford University Press: 23-58.
- Braddon-Mitchell D. and Miller, K. 2004. How to be a Conventional Person, *The Monist*: 89 (4).
- Dainton, B. 1992. Time and Division, *Ratio*, NS 5: 102-28.

¹⁵ We are very grateful to the editor and several referees of the AJP for their helpful comments.

- Dainton, B. 1995. Survival and Experience, *Proceedings of the Aristotelian Society*, 1995-6: 17-36.
- Dainton, B. 2000. *Stream of Consciousness: Unity and Continuity in Conscious Experience*, London: Routledge.
- Foster, J. 1979. In *Self-Defence*, in *Perception and Identity: Essays Presented to A. J. Ayer with His Replies*, ed. G.F. Macdonald, London: Macmillan Press.
- Foster, J. 1982. *The Case for Idealism*, London: Routledge and Kegan Paul.
- Foster, J. 1985. *Ayer*, London: Routledge.
- Foster, J. 1991. *The Immaterial Self*, London: Routledge.
- Gendler, T. S. 1998. Exceptional Persons: On the Limits of Imaginary Cases, *Journal of Consciousness Studies*, 5: 592-610.
- Horwich, P. 1987. *Asymmetries in Time*, Cambridge MA: MIT Press.
- James, W. 1952. *The Principles of Psychology*, Chicago: William Benton.
- Johnston, M. 1987. Human Beings, *Journal of Philosophy*, 84: 59-83.
- Johnston, M. 1992. Constitution is not Identity, *Mind*, 101: 89-105.
- Lewis, D. 1971. Counterparts of Persons and their Bodies, *Journal of Philosophy*, 68: 203-11.
- Lewis, D. 1976. Survival and Identity, *Philosophical Papers*, vol. 1, Oxford: Oxford University Press: 55-77.
- Lowe, E.J. 1996. *Subjects of Experience*, Cambridge: Cambridge University Press.
- Mellor, H. 1998. *Real Time II*, London: Routledge.
- Nathan, N.M.L. 1997. Self and Will, *International Journal of Philosophical Studies*, 5 (1): 81-94.
- Noonan, H. 1991. Indeterminate Identity, Contingent Identity and Abelardian Predicates, *Philosophical Quarterly*, 41: 183-193.
- Parfit, D. 1971. Personal Identity, *Philosophical Review* 80: 3-27.
- Rovane, C. 1998. *The Bounds of Agency*, Princeton: Princeton University Press.
- Shoemaker, S. 1970. Persons and their Pasts, *American Philosophical Quarterly*, 7: 269-85.
- Shoemaker, S. 1984. Personal Identity: A Materialist Account in S. Shoemaker and R. Swinburne, *Personal Identity*, Oxford: Blackwell.
- Smith, A.D. 2003. *Husserl and the Cartesian Meditations*, London: Routledge.

- Snowdon, P.F. 1991. Personal Identity and Brain Transplants, in *Human Beings*, ed. D. Cockburn, Cambridge: Cambridge University Press: 109-26.
- Sprigge, T.L.S. 1988. Personal and Impersonal Identity, *Mind*, 97 (385): 29-49.
- Strawson, G. 1997. The Self, *Journal of Consciousness Studies*, 4 (5-6): 405-28.
- Strawson, P.F. 1967. *The Bounds of Sense*, London: Routledge and Kegan Paul.
- Unger, P. 1990. *Identity, Consciousness and Value*, Oxford: Oxford University Press.
- Van Inwagen, P. 1990. *Material Beings*, Ithaca: Cornell University Press.
- Williams, B. 1970. The Self and the Future, *Philosophical Review*, 79: 161-80.
- Wilkes, K. V. 1988. *Real People: Personal Identity without Thought Experiments*, Oxford: Clarendon Press.