CHAPTER 1

THE SELF AND THE PHENOMENAL

Barry Dainton

1. Selves and streams

The foundations of classical physics were shaken when the sixteen year old Einstein imagined being able move as fast as he wished, and setting off in pursuit of a passing ray of light. But although the thought-experiment is justly renowned, its import is by no means self-evident. Why is it, exactly, that a passing ray of light is harder to overtake than a speeding bullet? My topic is the metaphysics of the self, not relativistic physics, and I will not consider these matters further, but a variant of Einstein's thought experiment can shed useful light on this topic.

Suppose, with Einstein, that you can move in any way that you wish, and are constrained only by what is clearly imaginable, rather than by what is nomologically possible. Spend a few moments imagining yourself zooming off to visit distant stars and galaxies, past and future centuries, parallel universes – anywhere that takes your fancy. Then ask yourself this: even so empowered, could you displace yourself in such a way that you leave your *self* behind? Obviously not. The notion is plainly nonsensical; since you and your self are one, where you go it goes, and vice versa. Now consider this: could you move so fast that you leave your *stream of consciousness* behind? The suggestion seems as absurd as its predecessor. We can outlast our current stream of consciousness simply by continuing to exist when unconscious, but we cannot move from it while remaining conscious, no matter how hard we try; any envisaged attempt merely involves an extension of our current stream, not a successful escape from it. Finally, consider a third variant: could your stream of consciousness continue on but fail to take you with it? This is a little trickier, but if we stipulate that as your stream continues on it has much the same subjective character as it usually has, and does not branch, then the answer once again seems perfectly clear. The suggestion that you could cease to exist in such circumstances is as absurd as the suggestion that you ceased to exist at some point during the past

few seconds of your conscious life. Whatever else we might be, we are not the kind of being that can cease to exist whilst our consciousness continues on – or so it is natural to suppose.

As is very clear, the connection between ourselves and our streams of consciousness is strong indeed. It is stronger by far than the connections between ourselves and any of the other modes of continuity which characterize our lives, at least if we confine ourselves to the realm of the imaginable. It is easy to imagine ourselves moving from one physical body to another; we need simply imagine our stream of consciousness flowing smoothly on, its subjective character much as per usual – much as it has been for the past five minutes, say – as it is housed in (or sustained by) a succession of numerically different bodies. And as Locke and Kant noted, what goes for bodies and brains also goes for souls and any other self-sustaining thing, whether material or immaterial. It is equally easy to imagine ourselves continuing to exist whilst our psychological states are altered, removed or replaced. Once again, it suffices to imagine that our streams of consciousness flow on as the envisaged changes – in beliefs, memories, intentions and personality traits – take place. (If this does not strike you as obvious, consider how little of your total psychology is impinging on your current experience.)

Not everything that is imaginable is really possible, and no doubt some things that are not imaginable are possible, but the fact that we find it hard to imagine ourselves and our streams of consciousness going their separate ways, and that we find it easy to imagine ourselves surviving any amount of physical and psychological discontinuity, provided our streams of consciousness continue to flow on, is certainly suggestive. It suggests that an account of our existence and persistence conditions that is rooted in purely phenomenal unity and continuity – the sorts of unity and continuity found in our streams of consciousness – will be more compelling, more believable, than any account that is not.

Of course, the notion that we should adopt an experience-based or *phenomenalist* approach to the self is scarcely new, far from it, but the approach has endured a period of comparative neglect, and a number of avenues worth pursuing have yet to be fully explored.[1] Just as there are different ways in which physical

[1] 'Phenomenalist' because the approach is rooted in phenomenal consciousness, rather than Mill's reductive proposal for matter and material objects – though, as we shall see, the experience-based approach to the self *can* lead to a position that is analogous in certain respects to this view of Mill's.

and psychological approaches to the self can be developed, there are different ways of developing the phenomenalist approach. In what follows I will outline one such way.

Before doing so, however, I will spend some time on phenomenal unity. This is a sizeable topic in its own right, but my aims are limited. I will be primarily concerned to outline a case for the claim that the unity in our streams of consciousness, both at and over time, is sufficiently robust and autonomous to provide a solid foundation for a credible account of the self. I will also argue that one conception of the source of phenomenal unity is mistaken, and this in itself has implications for the kinds of things selves are, or could be.

2. Phenomenal unity

Some philosophers favour a narrow construal of the term *consciousness*. Starting with the seemingly innocent claim that what differentiates conscious from non-conscious states is that we are 'aware' of the former but not the latter as they occur, they then suggest (more contentiously) that we are aware only of those states that we think about (or attend to, or form beliefs about), and so conclude that our conscious states are restricted to those that we think about (or attend to, or form beliefs about).[2]

I think it is far more natural to construe 'consciousness' in a broader way. As I will be using the term, a state is conscious if there is something that it is like to be in it, irrespective of whether the state is the object of thought, belief or attention. This terminological preference is grounded in phenomenology. Generally speaking, at any one time, my own experience, and hence my own consciousness, is clearly not confined to whatever it is that I am paying attention to or thinking about. If I focus my attention on the fly I see crawling up the wall, my visual experience is not restricted to the fly, I continue to see the surrounding walls and floor – even though I do not explicitly notice them. I also continue to have a wide range of non-visual experiences – auditory experiences of the road works outside, along with bodily and emotional feelings – these too are present despite being largely unnoticed. (If this is not obvious, try to imagine what your experience

[2] E.g. Rosenthal (1986).

would be like if it *were* limited, at any given time, to whatever you happen to be paying attention to or thinking about!) I take all these background experiences to be constituent parts of my current total state of consciousness.

Let us turn to the main issue, and start by considering synchronic unity. From a phenomenological perspective, it is perfectly obvious that our typical streams of consciousness are unified at any given time. Our conscious thoughts and mental images do not occur separately from our bodily sensations, our bodily sensations do not occur separately from our visual experiences, our visual experiences do not occur separately from our auditory, tactile or olfactory experiences. On the contrary, all our experiences, of all types, typically occur together, in unified experiential ensembles. This unity or togetherness is itself something we experience. Focus your attention on one of your current bodily sensations, and then consider just how this sensation is related to part of your visual field, or something you can hear (or a mental image, or your current conscious thinking, etc.). Clearly, the bodily sensation and auditory experience are occurring together within your overall state of consciousness, and this togetherness is itself experienced: there is something distinctive that it is like to experience these two experiences together, rather than separately. This experienced togetherness – this *co-consciousness* as I will call it – does not, it should be noted, consist in a separate experiential ingredient, with its own distinctive phenomenal character; there is no additional experiential element occurring between the bodily sensation and the visual content. There is simply the bodily sensation and the sound, experienced together. Precisely the same relationship connects your bodily sensation with your visual experience, and your visual experience with your conscious thoughts. Quite generally, each and every part of a typical state of consciousness is co-conscious with each and every other part.

If we confine ourselves to the phenomenal level, and concern ourselves only with structures and features discernible in our consciousness, rather than their underlying mechanisms and causes, is there anything more to be said on the topic of synchronic unity? According to what I call the *Simple Conception* of consciousness, there is not. On this view, experiential unity is a primitive feature of consciousness. Unified states of consciousness simply consist of experiences related to one another by co-consciousness, and that is all there is to be said.

3. Pure awareness & bare loci

There is another a way of thinking about the synchronic unity of consciousness, one that has broader implications. The position I have in mind is a variant of the doctrine that consciousness requires, or is a form of, awareness. Not awareness in the form of belief-acquisition or conscious thought, but rather awareness in the form of *perception* (or *sensing*), where the latter is construed in the manner of the naïve realist, as an unmediated apprehension of an object or content. According to what I will call the *naïve perceptual* (NP) conception of consciousness, every episode of experiencing involves two components: an act or process of apprehension, and the various contents falling under, or apprehended by, this act. All the phenomenal qualities that feature in our consciousness reside at the level of content, and these qualities only become conscious when they are apprehended; the act of apprehension itself is devoid of qualities, and in this sense 'pure'. The 'contents' of consciousness thus include occurrent thoughts, mental images and emotional feelings, as well as the deliverances of our various senses. Since the contents falling under a subject's awareness can include peripheral or background elements in our experience, the NP-conception should not be equated with the 'narrow' conception of consciousness mentioned above.[3]

Given our present concerns, the NP-conception has a double interest. It yields an intuitively appealing account of the unity of consciousness: experiential contents are unified, or co-conscious, only if they fall under a single act of pure sensory awareness. But it also points towards a conception of what selves are. Perhaps conscious selves are nothing more than centres (or loci) of pure awareness. Following Mark Johnston (1987), I will call this the *bare locus* view of the self.

This view does have some appeal. I suspect many of us can make sense of the suggestion that we could, in extremis, be reduced to nothing more than a point of pure apprehension, gazing outward, all senses keenly alert but detecting nothing.

[3] So-called 'act-object' conceptions of consciousness were commonplace in the 19th and early 20th centuries, but versions of the doctrine continue to flourish among contemporary 'higher-order sense' theorists, see for example Armstrong (1997) and Lycan (1997); also Lockwood (1989: 162–3).

Nonetheless, when considered seriously, the doctrine quickly begins to look suspect.

Since a bare locus is nothing but a point-centre of featureless apprehension, there can clearly be nothing that it is like to *be* a bare locus in the absence of apprehended contents. The condition of being a bare locus that is not apprehending anything would be subjectively indistinguishable from not existing at all. Can we really make sense of the idea that an entity such as this could exist? Is the concept of a bare locus any more intelligible than the concept of a bare particular, the featureless *something* that allegedly remains when all the properties of an object are removed from it? I cannot see that it is. Nor does it help if we suppose that bare loci can only exist when apprehending some content or other. For again, given that bare loci have no qualitative contribution to make whatsoever, what reason do we have for supposing that anything about our conscious experience would be in any way different if they were absent? This invites the reply: 'Without the bare loci there would be no one to experience anything, and hence no experience!' But this does not take us any further, since what is in question is precisely how this could be possible![4]

Since bare loci are nothing more or less than centres of pure apprehension, these considerations also weigh against the NP-conception. However, the case against the latter would be stronger if the Simple Conception could be shown to have the resources to account for the intuitive appeal of the perceptual model. Why, if it is erroneous, can the latter seem so compelling?

Although the full story is complex, three of the more significant influences are not too difficult to discern. First and most obviously, much of our ordinary experience *is* perceptual, and much of our perceptual experience has a presentational character: the things we see and hear seem to be 'out there' in the world, as opposed to 'in here' with our thoughts, feelings and bodily experiences. The NP-conception misdescribes the reality of the situation – in our case at least, there is a definite something that it is like to *be* an observer of an external world – but it is no means a completely erroneous view of consciousness.

A second source of the perceptual model's appeal lies in the phenomenology of attention. We can attend to any part of our

[4] See Dainton (2002) for more on this theme.

experience, and when we do so it often seems like we are 'seeing' what is there for the first time. And in a sense we are: when some feature of what is there comes to our notice for the first time. But again, ordinary forms of attending have their own qualitative character, unlike the NP-theorist's pure apprehension, and are directed at only some of a subject's experiences, not all of them.

Thirdly, and perhaps most importantly, there are the so-called 'fringe feelings' that accompany our sensory experiences and conscious thoughts. Think of what it is like to see a face and know that you have seen it before, without being able to quite place it. Compare this with what it is like to know that a certain word just isn't right, or *is* right. Or think of what it is like to carry out a task with the reassuring feeling that everything is proceeding as it should – or with the feeling that it isn't. In each of these cases (and in many others of the same genre) there is a feeling or intuition with a quite specific character, a feeling whose significance is entirely unambiguous, but one that does not possess a qualitative character of a sensory kind – they have no definite spatial location, size or colour, for instance. Such feelings are easily overlooked, but when our attention is drawn to them it soon becomes apparent that they are commonplace; indeed, they are significant ingredients in our overall consciousness.[5]

As for their relevance to the topic in hand, when we ponder our natures as conscious beings most of us have no difficulty in accepting the claim that there is something more to us than the various contents we are currently aware of, and equally the claim that we are active apprehending *somethings* that are distinct from these contents. Having recognized the existence and significance of fringe feelings we can account for this without recourse to the additional (and introspectively invisible) level of consciousness posited by the NP-theorist. Such claims strike us as plausible because there is a sense in which they are true: there is indeed an additional something in consciousness; not a point-centre of pure awareness, but simply fringe feelings of various types, some more ubiquitous than others. Among the most significant in the present context are the feelings associated with conditions such as these: *being in a state of readiness, opening or straining one's senses, being pre-*

[5] William James' exploration of the fringe is to be found in the 'Principles of Tendency' section of the *Principles*. For more on the topic see Epstein (2000) and especially Mangan (2001).

pared for the unexpected, being determined or resolute, potency: having the sense that one is mentally or physically energized, that one is equipped for the task in hand, that one is about to act. It is, I suggest, the presence of these and similar feelings that, to a large extent, gives rise to the sense that we are active apprehending subjects. What else could possibly be needed? Imagine your consciousness being gradually emptied of all content until all that remains is a single solitary fringe feeling: *being ready for whatever happens next.* If we try to imagine ourselves existing in the form of a 'pure locus of awareness', don't we end up imagining something along these lines? If so, 'pure' does not really mean 'devoid of all content', it simply means 'devoid of all *non-fringe* content'.

There is a good deal more than could be said about this, but perhaps I have said enough to demonstrate that adopting the Simple Conception does not mean embracing a phenomenologically impoverished or obviously inadequate view of our consciousness. A recognizable picture of the latter can be constructed from experiential elements – of all kinds – related to one another by the primitive joining relation of co-consciousness.

4. Phenomenal continuity

Let us move on to streams of consciousness proper, and consider the ways in which their constituent parts are unified over time, from moment to moment.

The first point to note is the relevance of the distinction between narrow and broad conceptions of consciousness. From the narrow perspective it is implausible to suppose that consciousness is stream-like. Suppose it were the case that we are conscious only of what we are paying attention to; since our attention often flits quickly from one thing to another, and sometimes we are not attending to anything at all, our consciousness would exhibit a corresponding degree of discontinuity. But when matters are viewed from the broad perspective the situation is very different. Generally speaking, we are continually conscious throughout our waking (and dreaming) hours, in the sense that at each moment during such periods we are experiencing *something*. We may not be thinking about what we are experiencing, we may not be attending to what we are experiencing, we may not be conscious in a self-conscious way, but experience is there nonetheless. If this were not so, we would not have the impres-

sion of being continuously awake, nor would we have the impression that the world is continuously *there* before us in the way that it generally seems to be.

Our ordinary streams of consciousness may usually be composed of continuous (gap-free) stretches of experience, but their unity is more far-reaching. Yes, we experience continuously, but continuity is also something we experience, all the time. Think of what it like to view the passing countryside from the window of a moving train: trees, roads, buildings – you watch them all sliding by. This smooth continuous movement is something you actually *see*, not merely infer. It is much the same if you walk round a room, or turn your head; or listen to a rising violin tone, or your friend's conversation (or your own inner soliloquy) while enjoying the burning heat of a chilli pepper on your tongue; or do no more than acknowledge the relentless passing of time while gazing, unthinking, at an open expanse of sky. Quite generally, throughout our waking and dreaming hours, and in all regions of our consciousness, we are directly experiencing change or persistence. Our ordinary experience is not merely continuous, continuity is a ubiquitous presence within our experience.

Recognizing the two-sided character of phenomenal continuity is one thing, explaining precisely what it involves, even at the purely phenomenal level, is quite another, and this is not the place to attempt it. All I will do is outline a case for a strong constraint on any adequate account of phenomenal continuity. As is customary, I will use the (less than optimal) term 'specious present' to refer to those brief phases of our streams of consciousness during which we are directly aware of change and persistence.[6] Since a specious present has some (apparent) temporal depth, it has (seemingly) earlier and later parts; and since the transition between these parts is directly experienced, the parts are *phenomenally connected* – they are co-conscious, but diachronically rather than synchronically. If, as seems plausible, we regard a stream of consciousness as being composed of a succession of specious presents, a question arises. How are neighbouring specious presents in the same stream of consciousness related to one another? Two positions can be distinguished:

> Streams of consciousness are *partially connected*: there are diachronic phenomenal connections only within individual specious presents.

Streams of consciousness are *fully connected*: there are diachronic phenomenal connections within and between specious presents.

Although several prominent accounts of the diachronic structure of consciousness reject full connectedness in favour of partial connectedness, there are good reasons for doing the reverse.

Consider what it is like to hear a plane passing overhead: an ongoing roar that endures without interruption for a several minutes, before fading. Below is a pictorial representation of a few seconds of this experience (Figure 1):

This line does not reflect all aspects of the experience – e.g., the changes in volume and timbre – but it captures perfectly the feature which interests us: its sheer continuity. As the plane passes, we hear an uninterrupted flow of sound, with no seams or gaps of any kind.

Now suppose, as some have alleged, that our streams of consciousness are composed of non-overlapping pulses, each the duration of a single specious present, and that the impression we have that our consciousness is fully continuous is due to qualitative similarities between neighbouring pulses together with short-term memory-impressions.[7] It may seem that this theory can accommodate and account for the experience of phenomenal continuity, in the way depicted in Figure 2:

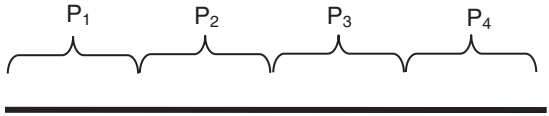

[6] How brief is 'brief'? Perhaps only around one second, or even a little less. The shortest interval within which typical human subjects can discern distinct temporally ordered stimuli (in all sense modalities) is around 30 milliseconds – Pöppel (1985) – this puts a lower bound on our specious presents, construed as temporal cross-sections of our entire streams of consciousness.

[7] Sprigge (1981, ch. 1) and Whitehead (1929) advocate this simple model, but the more complex accounts of time-consciousness developed by Broad (1938) and Husserl (1991) also reject full connectedness, and do so in an even more dramatic fashion. On these views, the specious present is a momentary experience with a content that *represents* a temporal interval, and does so in such a way that it appears to possess genuine temporal depth. In effect, accounts such as this merely *multiply* the number of discrete pulses, for it remains the case that there are no genuine experiential links between successive specious presents. See Dainton (2003) for a more detailed discussion, especially of Husserl's account.

Here each of P_1-P_4 represents adjoining specious presents, each containing a brief pulse of sound; the continuity of the sound is due to the fact that there are no gaps between the pulses. What more could possibly be required? Actually, a good deal. There are inter-experiential relations that the pulse theory ignores or misrepresents, a few of which are depicted in Figure 3.

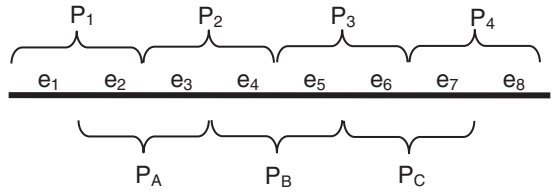

Here, e_1, e_2, e_3 . . . e_8 are the earlier and later phases of the specious presents P_1, P_2, P_3 and P_4. The pulse theory adequately captures some inter-experiential relations, e.g., e_1 is experienced as flowing into e_2, and e_3 is experienced as flowing into e_4. But what of e_2 and e_3? Or e_4 and e_5? Since by hypothesis these occur in distinct pulses, the transitions between them cannot be directly experienced. But this is manifestly at odds with the phenomenology. When listening to the passing plane what we hear is a seamless flow: e_1-flowing-into-e_2-flowing-into-e_3- . . . and so on. The diachronic phenomenal connections indicated by all the lower brackets have precisely the same phenomenal character as those indicated by the upper brackets – all seem equally real. And what goes for this auditory experience goes for the successive brief phases of our entire streams of consciousness. Since it denies this, the pulse theory is manifestly inadequate to the phenomena. The same will be true of any account of the diachronic structure of our streams of consciousness which fails to accept that our streams of consciousness are fully connected. There are direct diachronic phenomenal connections between *all* the adjoining brief phases of a single stream of consciousness. Unless we recognize this we cannot hope to do justice to the two-sided character of phenomenal continuity.

Or so I think it reasonable to believe. There is, however, a line of reasoning that casts doubt on the full-connectedness constraint, and before moving on I will indicate where I think it goes astray. Here is one way of formulating what I have in mind:

The Cut Argument Suppose that at precisely noon today, just as you are about to cross a busy street, you are vaporized, instantaneously, by a devastatingly powerful laser beam. Although your annihilation brings your stream of consciousness to an immediate halt, it does not affect in the slightest the character of your experience leading up to this event. How could it? Are we to suppose that your pre-noon experience is affected by events that have yet to occur, via some mysterious channel of backwards causation? Assuming not, it seems that your experience leading up to the moment of your annihilation is just as it would have been if the trigger on the laser had not been pressed, and you had continued about your business as per usual. The point generalizes. There is nothing special about you, and nothing special about that particular time. Any of us *could* be annihilated at any of our waking moments, and this could occur without impinging on the character of the immediately preceding experiences. And we can take a further step. We are all familiar with the sceptical claim that 'everything could be just as it is now even if the entire universe only came into existence five minutes ago'. If this is right, then surely the universe could have come into being five *seconds* ago – or even five milliseconds – without affecting the character of our experience now. This also applies to everyone, at all times. Putting these points together, and bearing in mind the fact that the 'cuts' could come anywhere in any stream, it is clear that the phenomenal character of any short stretch of experience *E* is logically independent of both later and earlier experiences in the same stream of consciousness: *E* could be just as it is, phenomenologically speaking, if the experiences immediately before and after it did not exist. It follows that a stream of consciousness consists of a succession of experiential phases, each of which is entirely self-contained and isolated, at least on the phenomenal level, from its immediate neighbours.

While the Cut Argument can seem persuasive, we have already seen that if our streams of consciousness really did consist of successions of isolated, phenomenally self-contained phases, then our experience could not exhibit the continuity that it typically does. So the Cut Argument must contain a flaw – but where?

The error is not difficult to discern. It is logically possible for a short-lived subject – a *two-second subject*, say – to have a stream of

consciousness which very closely resembles the two-seconds of consciousness that you (or anyone else) has just enjoyed. The mistake comes in supposing that this stream-phase is exactly the same as the corresponding stretch of your experience in *all* phenomenal respects. There are differences, and the differences are such that putting a succession of self-contained phases end to end, as it were, would *not* result in a stream of consciousness of the kind we typically enjoy.

To illustrate, suppose you listen to someone repeatedly playing a C-major scale on a piano: *C-D-E-F-G-A-B-C*. . . . Let us also suppose that on some distant (but logically possible) world a short-lived duplicate of you comes into being, in an exactly similar room, containing a similar piano, and hears just the *D* segment of a similarly pitched C-scale, before ceasing to exist. Would your ephemeral duplicate's *D*-experience exactly resemble your own? In intrinsic phenomenal respects, perhaps. The auditory qualities of the *D*-tone you each experience might be indistinguishable; you might also both hear *D* while experiencing a short-term acoustic memory-image of a preceding *C*, and an anticipation of hearing an *E*. But so far as diachronic phenomenal *relations* are concerned, your duplicate's experience cannot exactly resemble yours. You hear *C*-running into-*D*, and *D*-running into-*E*, but since neither *C* nor *E* are experienced by your duplicate, neither are these phenomenal relations. Hence in one key respect, your duplicate's experience is necessarily different from your own.

Once diachronic phenomenal connections are brought into the picture it is obvious where the Cut Argument goes astray. An experience's non-relational phenomenal characteristics may be logically independent of what comes before or after, but it is otherwise with regard to its relational features. Once this is recognized, it is evident that a *stream* of consciousness does not consist of a succession of self-contained chunks or pulses of experience, laid end to end like a row of bricks. Or at least, if we do think of a stream in this way, we must not forget the cement which holds the bricks together. In addition to phenomenal connections *within* stream-phases, we must also recognize the phenomenal connections *between* them. The plausibility of the Cut Argument derives from a conception of the diachronic features of consciousness that is as incomplete as it is impoverished. By way of a final illustration of this point, consider Figure 4 below, where four brief stream-segments are shown.

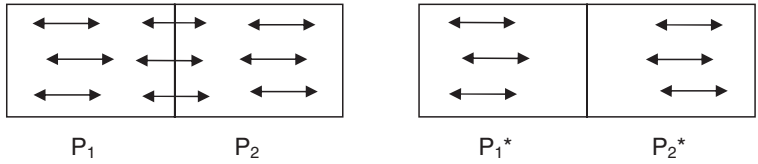

Thanks to the presence of the diachronic phenomenal connections linking P_1 and P_2, this pair of stream-phases form a phenomenally continuous stretch of experience. In this instance, the depicted spatial proximity between these phases accurately reflects experiential reality. It is otherwise with respect to P_1* and P_2*. Given the absence of phenomenal connections between these phases, they might as well belong to different subjects, or exist in different universes. In short, the Cut Argument presupposes that our ordinary experience is as depicted on the right, when in reality, it is as shown on the left.

There is, of course, more to be said on this topic. Just how long is a specious present? What sort of experiential connection gives rise to diachronic phenomenal relations? But we need not address these matters here.[8] For present purposes it is sufficient that any adequate account of phenomenal continuity must recognize that own typical streams of consciousness are fully, rather than merely partially, connected.

5. Consciousness and self: a parting of ways

The claim that experiences within a typical stream of consciousness are bound together by purely phenomenal relations has solid phenomenological support. We do not need to look beyond the phenomenal to determine which experiences belong to which streams: simultaneous experiences are 'co-streamal', part of the same stream, only if they are related by synchronic co-consciousness; non-simultaneous experiences are co-streamal only if they are diachronically co-conscious, either *directly*, i.e., they occur within a single phenomenal present, or *indirectly*, i.e., they are part of a chain of overlapping specious presents.

[8] The simplest way of accommodating full connectedness is to permit neighbouring specious presents to overlap by sharing common parts, in the manner advocated Russell (1984) and Foster (1979). See Dainton (2000, 2003) for more detailed treatments of these topics.

What use is all this in regard to the self? Given the intimacy of the connection between selves and streams of consciousness we noted at the outset, and using 'consubjective' to mean 'belongs to the same self or subject', the following seems plausible:

The S-thesis: Co-streamal experiences are consubjective.

Bearing in mind what has just been noted concerning stream membership, we also have:

The C-thesis: Experiences that are co-conscious, directly or indirectly, are consubjective.

This is fine as far as it goes, but it doesn't take us very far. Any attempt to construct an account of the self appealing only to phenomenal relations faces an obvious hurdle. While we are undeniably conscious some of the time, we are not conscious all of the time. Or at least, if we are conscious all of the time, this is by no means obvious. It seems to most of us that we lose consciousness completely at least once every twenty four hours, when we slide off into dreamless sleep. If this is right, then a typical person can expect to have several thousand distinct streams of consciousness during the course of their life.

There are two fundamentally different ways of approaching the 'gap' problem available to the phenomenalist, each rooted in a different conception of the kind of thing a self is. Each of these conceptions is a way of developing the idea that a self is essentially a thing that has experiences. Here is the first:

The Essentially Conscious Self [ECS]: a self is a thing whose nature it is to *be* conscious; a self is experiencing at every moment at which it exists. A self cannot lose consciousness and continue to exist.

According to this conception, selves cannot exist when unconscious, and so they cannot persist through periods of unconsciousness. Given the seemingly intermittent nature of our conscious lives, advocates of ECS have two options. They can adhere to the common sense view of things, accept that we each have a succession of distinct streams of consciousness, and provide some account of what renders such streams consubjective. Alternatively, they can take the more radical option of denying the appearances, and hold that a self has but a single uninterrupted stream of consciousness.

Neither of these options is very palatable. The only obvious way of developing the 'many streams' approach is to hold that distinct streams of consciousness are consubjective in virtue of some feature that is internal to the streams. But what feature? Do all your streams of consciousness have some unique phenomenal feature that could serve to mark them out as yours and yours alone? Perhaps on occasion, but all of the time? Every day? As an infant? It seems unlikely.[9] The 'single stream' approach fares little better. Advocates of this doctrine face a dilemma. They can embrace the hard-to-believe option of holding that typical human selves exist for no more than a few hours,[10] or they can follow Descartes in denying that we ever truly lose consciousness during the normal course of our lives – what seems like dreamless sleep is really low-level consciousness that we are later unable to recall. While this suggestion may not be so absurd as to merit the scorn habitually heaped upon it, as things stand it is at best speculative – and of course Descartes was not acquainted with modern day general anaesthetics.[11]

Phenomenalists are not obliged to opt for any of these proposals. The notion that selves are essentially experiencers can be construed in a different way:

> **The Potentially Conscious Self [PCS]:** a self is a thing whose nature it is to be *capable* of being conscious; a self has the *capacity* for consciousness at every moment at which it exists, and it possesses this capacity essentially. A self can lose consciousness provided it retains the potential to be consciousness.

[9] Not everyone would agree with this verdict: see Nathan (1997) for a recent defence of the similarity approach. Nathan is sympathetic to the view that our overall states of consciousness are always characterized by a *tone* or *flavour*, of a subtle and easily overlooked kind, and that since we each of us have a unique phenomenal essence, as it were, one that is unlike anyone else's, this feature can serve to distinguish our own experiences, whenever and wherever they might occur. This contention is not susceptible to direct empirical refutation, but equally, it lacks direct empirical support, and so far as I can see, the considerations which might incline one to accept the speculation are outweighed by those point in the other direction.

[10] Strawson favours the more radical position that selves only endure for a second or so, on the grounds that 'strongly unified' or 'hiatus free' episodes of experience only last this long. I agree that episodes of what Strawson calls 'Self-experience . . .' – which he describes as '. . . the experience that people have of themselves as being, specifically, a mental presence' (1999: 104) – are usually short-lived. But this is compatible with consciousness in the broad sense being continuous, and strongly unified from moment to moment, for much longer periods. See also Dainton (2000: §5.2).

[11] There are further radical options: to hold that interruptions in a subject's consciousness exist in objective but not subjective time, the Berkeleian view also defended by Evans (1970) – or to deny that selves exist in time at all – see Robinson (forthcoming).

This conception retains the tight link between selves and consciousness, but has the merit of conforming to the common sense view that consciousness is something that we can lose and regain. Also, since on this view a self continues to exist when unconscious, we can simply say that what renders distinct streams consubjective is the fact that they belong to the same self – thus avoiding the untempting stratagems forced upon those who subscribe to the doctrine of the essentially conscious self.

But while it brings definite advantages, the PCS doctrine faces difficulties of its own. What manner of existence do non-conscious selves enjoy? If conscious and non-conscious selves are instances of the same kind of being, *what* kind of being?

6. Phenomenal substances

One way forward is to start with the general notion of 'a thing which can have or produce experiences', and specify the identity conditions of such things relying solely on the intimate relationship between selves and their experiences encapsulated in the S- and C-theses. The following concept provides a point of departure:

> **Experience Producer** [EP]: any object or system that can directly produce experience in response to external or internal changes.

This characterization is neutral with regard to structure and constitution. Human brains are EPs, but so are feline brains (probably), as are immaterial souls, if there are or could be such things. What makes something an experience producer is the capacity, grounded in natural law, to produce unified experience – it is this capacity that matters, nothing else. As we know from our own case, such a capacity can be activated by external influences (as in perception), or as a consequence of internal processes (mental images, memories, internal dialogue, etc.).

The concept of an EP is a useful starting point, but it does not provide us with everything we need. Most significantly, it does not provide us with a conception of an entity whose persistence conditions are framed in phenomenal terms. Being an EP is a non-essential temporary property of many of the objects that possess the relevant capacities. Most human brains are EPs, but their identity conditions are biological rather than experiential, and a given

brain can gain and lose the capacity to produce experience. Since selves of the potentially conscious variety cannot do likewise, identifying selves with EPs of this type is not an option. And of course, we want to leave room, conceptually at least, for a single self to be sustained by a succession of such supporting systems.

There is a further difficulty of a different kind. Suppose it were the case that your auditory and visual experiences were the products of two distinct neural systems, each of which is capable of functioning independently of the other. Given that a typical field of auditory experience is a synchronically unified ensemble, and a typical field of visual experience likewise, by the lights of the definition supplied above, each of these neural systems would count as an EP in its own right. But clearly, EPs like this are not selves, they are merely parts of selves.

Noting these deficiencies points the way forward. It would be a mistake simply to equate selves with EPs, but it may nonetheless be possible to state the identity conditions of the former in terms of relationships among the latter, at and over time. We can start thus:

> **E-linkage:** at any time *t*, one or more EPs are *E-linked* if and only if (i) they are active, and the experiences they are producing are mutually synchronically co-conscious, or (ii) they are not active, but if they were, the experiences that would result would be mutually synchronically co-conscious.

E-linked EPs are akin to groups of 'co-targeted' spotlights: the members of such a group need not be turned on, but if they *were*, the resulting light-rays would converge on the same spatial location. In similar fashion, the members of a group of E-linked EPs need not all be active and producing experience, but if they are, all the resulting experiences are co-conscious. (I allow a single EP to be E-linked so as to accommodate the case of a self composed of just one EP at a given time.)

It is certainly very plausible to suppose that E-linked EPs are consubjective, and it is hard to see how normally functioning EPs at a given time could be consubjective unless they were E-linked. Since we are interested in *all* the EPs that belong to any one self at a given time, it will prove convenient to have a label for such collections:

> **E-system:** any maximal collection of E-linked EPs at a given time *t* compose an E-system.

A collection *C* of E-linked EPs is less than maximal if at the time in question there are EPs that are not in *C* but that are nonetheless E-linked to the members of C.

We now have the beginnings of a plausible account of how EPs at a given time must be related if they are to belong to, or compose, a single self. We still need an account, in phenomenal terms, of how EPs at different times must be related in order to be consubjective. Here we can rely on the diachronic aspect of the C-thesis, on the fact that experiences can be co-conscious over time as well as at times. As we saw earlier, the temporal depth of diachronic co-consciousness is not great – only experiences within a single specious present are so related – but the relationship is just as real as its synchronic counterpart. Putting this relationship to work, it is plausible to hold that E-systems existing at different times that can contribute to the same specious present belong to the same self. So too, over the longer term, E-systems that can contribute to the same extended stream of consciousness belong to the same self. More formally:

(1) E-systems S_1 at t_1 and S_2 at t_2 are *directly streamally linked* if, and only if, (i) both systems are active and producing experiences that are directly diachronically co-conscious, or (ii) if both systems *were* active, the experiences they would produce would be directly diachronically co-conscious.

The notion of 'direct streamal linkage' is simply the diachronic counterpart of E-linkage. As for the longer term:

(2) E-systems S_1 at t_1 and S_2 at t_2 are *indirectly streamally linked* if, and only if, they are at either end of a chain of E-systems, the neighbouring members of which are directly streamally linked.

Together, (1) and (2) allow us to state the conditions under which a single E-system persists through time:

(3) E-systems S_1 at t_1 and S_2 at t_2 are the same enduring E-system if, and only if, S_1 and S_2 are streamally linked, directly or indirectly.

Alternatively, for those who prefer to think of persistence in a four-dimensional framework:

(3)′ E-systems S_1 at t_1 and S_2 at t_2 are temporal parts of the same perduring E-system if, and only if, S_1 and S_2 are streamally linked, directly or indirectly.

Bringing these formulations somewhat closer to home:

(4) Members of a collection of experiences are consubjective if, and only if, they are produced by EPs belonging to the same E-system.

And finally:

(5) E-systems can be regarded as a distinctive kind of entity: *phenomenal substances*.

(6) E-systems and selves are one and the same.[12]

A simple picture underlies these formulations: the tracing of selves through time via the tracing of capacities for phenomenally continuous experience. Our brains are (probably) E-systems for most if not all of their careers, but the identity conditions of the latter are sufficiently general so as to permit a wide range of very different entities also to count as E-systems. There is no need for the relevant systems to be organic, physical or spatially continuous: a succession of immaterial souls would constitute an E-system if they were capable of sustaining a single stream of consciousness. Since E-systems have distinctive and well-defined identity conditions, there is no reason why they should not be regarded as perfectly respectable objects or substances in their own right. Since these identity conditions are framed in experiential terms, to call them *phenomenal* substances seems appropriate.

We could stop here, but a case can be made for taking a further step. E-systems are composed of EPs, and although EPs all have at least one thing in common – the capacity to produce experience – in other respects they may be very different. It would surely be unwise to rule out even the possibility that conscious beings may come in very different forms. We can cut through these variations, and isolate what is common to all conscious beings, at least of the PCS variety, by defining their identity conditions in terms of what is common to all EPs. Rather than taking experience-producing *objects* as our basic ingredient, we take certain *properties* of these objects. As for the sort of property, there is only one candidate: the property common to all EPs: the power to generate experiences.

It is as easy to formulate a set of identity conditions in these new terms as it was in the old. In fact, we can keep the formula-

[12] Endurantists need not regard themselves as being numerically identical with any particular collection of EPs – they can regard E-systems as *composed* or *constituted* from particular collections of EPs, a view which permits an E-system to be made up of different EPs than it actually is at any one time – the definition of an E-system leaves this option open. Most perdurantists will be happy to accept the identity.

tions and simply interpret them in a new way. To start with, we allow the abbreviation 'EP' to refer to *experiential power* rather than *experience producer*. An EP in the new sense is simply a nomological capacity, no doubt usually possessed by some object or system, to produce experience of one or more kinds in response to appropriate triggering circumstances. Next, we re-interpret the definitions of 'E-linkage' and 'E-system' in the same way. An E-system is now a maximal collection of experiential powers, unified in their ability to produce experiences that are co-conscious. To avoid confusion I will henceforth use the expression *E*-system* to mark the change. The persistence of these systems can be defined as in (1)–(3), in terms of the capacity to produce extended streams of experience. We can then draw revised versions of our earlier conclusions:

(5)* E*-systems can be regarded as a distinctive kind of entity: *phenomenal substances*.

(6)* E*-systems and selves are one and the same.

Although the E*-proposal has some counterintuitive features – thinking of ourselves as nothing more than a collection of properties may take a bit of getting used to – it also has a good deal to recommend it. The account is general enough to apply to most of the things we might ever be tempted to call 'selves', it tells us what all these entities have in common, and it is permissive enough to permit a single self to survive the most radical transformations, provided its capacity for experience is not eliminated.

The account combines generality with a degree of specificity. The most obvious way of defining the identity conditions of an arbitrary experiential power *P* is in terms of three components: (i) the kind of experience *P* produces when activated, (ii) the type of circumstance that will trigger *P*, (iii) the ground of *P*. The 'ground' of a particular token capacity is what possesses that capacity. Some capacities are grounded in objects, or in certain (usually non-dispositional) properties of an object, others may be grounded in nothing more than regions of space – electrical fields can be construed thus. Our own EPs appear to be grounded in various regions of our brains, but there may be selves whose EPs are differently grounded. By virtue of the fact that EPs *have* grounds, and that differently grounded EPs will usually have different triggering conditions, there is plenty of scope for the

detailed descriptions of different selves to vary. But of course, if systems of very different kinds can possess or constitute EPs, variations of this sort are only to be expected. And these differences in detail are quite compatible with all selves having the same high-level existence and persistence conditions.

Some may object: 'The E*-theory may have certain merits, but it signally fails to do justice to the fact that, whatever else we may be, we are definitely *things*, and whatever else they may be, clusters of properties are not things.' This objection overlooks the fact that on one standard metaphysical view non-abstract objects of all kinds are nothing more than clusters of properties! But setting aside this point, in one respect the objection is warranted. The claim that E*-systems are substances, albeit of the phenomenal variety, sits uneasily with the doctrine that the distinguishing mark of a genuine substance is that it does not depend for its existence on anything else. But while it may well be that no E*-systems have this degree of independence, it is worth noting that there is another traditional criterion of genuine object-hood: unity. The grounds of an E*-system need not exhibit material or spatial cohesion, but it remains the case that unity is the defining characteristic of E*-systems. To qualify as members of such a system experiential powers must be capable generating co-conscious experiences. And as we saw earlier, the unity to be found in our consciousness is very real indeed.

7. Selves and subjects

So much for the E*-account in bare outline. There are many details to be filled in, and further issues to explore. EPs are dispositional properties; there are several competing accounts of such properties, all with implications for the kind of entity E*-systems might be. We need an account of the different ways in which capacities can be rendered inoperative – how does deep sleep differ from death? How do we draw the boundaries? Can an E*-system undergo fission? If so, what happens to a self which divides? Even if our essential cores are nothing more than clusters of experiential capacities, there is more to a typical self than its essential core; we also have bodies and psychologies, and we need some account of how E*-systems are related to these. These are all interesting topics, but I cannot hope to do justice to any of them here, so will conclude with a few general observations.

There is a sense in which things have come full circle. I began by advocating a phenomenalist *approach* to the self; I have concluded by advocating a phenomenalist conception of what a self *is*. Whereas Mill maintained that material things were nothing more (or less) than enduring potentials for experience, I am suggesting that this is what we ourselves are. Phenomenalism of this variety is less obviously objectionable than phenomenalism about material things.

I have not suggested any limit on the kind of experiential potential that is needed to sustain a self in being. This omission was deliberate. As noted at the outset, it is hard to envisage one's stream of consciousness flowing on but failing to take oneself with it, and this holds irrespective of how broad or deep one's stream is. If I imagine the capabilities of my own E*-system being progressively diminished, with the result that my consciousness becomes progressively simpler in character, it seems obvious that I continue to exist, even if I am incapable of experiencing anything other than a few primitive bodily sensations. In the absence of actual experience, the continued existence of the capacities for such experience is enough to keep me in existence – or so I am inclined to believe.

Not everyone will find this minimalist approach to survival congenial, but those who do must address this question: When your experiential capabilities are reduced to a very low level, do you remain a *self*? The idea that you might remain in existence but fail to remain a self may seem very odd. If by 'self' we simply mean 'subject of consciousness' or 'thing that can have experience', then the idea is nonsensical. But there are those who equate selves with beings who are capable of being aware of themselves *as* selves.

Self-awareness comes in different forms, but there is no need to attempt a inventory here: anyone reduced to the minimal level of sentience will be incapable of self-awareness in any interesting sense. So anyone attracted by the minimalist approach to survival *and* a high-level reflexive notion of selfhood should be prepared to acknowledge that they are not essentially selves, only subjects, and that as such they share their identity conditions with more primitive beings (though of course more advanced sentient beings, if such exist, are subjects of experience too).

Irrespective of whether we think of ourselves subjects or selves, we should try to think of ourselves in the right way. We are things

that are capable of *having* experiences. What exactly does this involve? If the various doctrines I have been propounding here are along the right lines, subjects are not things which passively apprehend their experiences, nor are they things which experiences impinge upon**.** They are things which *make* or *generate* experiences. This too can seem somewhat odd. But then, given the character of our experiences, of the experiences we make, it is not surprising that this seems odd. As has often been observed, selves are rather elusive.[13]

References

Armstrong, D. (1997). 'What Is Consciousness?' in Block *et al.*

Block, N., Flanagan, O. and Güzeldere, G. (eds.) (1997). *The Nature of Consciousness: Philosophical and Scientific Debates*. Cambridge: MIT Press.

Broad, C. D. (1938). *An Examination of McTaggart's Philosophy*. Cambridge: Cambridge University Press.

Dainton, B. (2000). *Stream of Consciousness*. London: Routledge.

——. (2002). 'The Gaze of Consciousness'. *Journal of Consciousness Studies*, 9, No. 2, pp. 31–48.

——. (2003). 'Time in Experience: Reply to Gallagher'. *Psyche* 2003. http://psyche.cs.monash.edu.au/symposia/dainton/index.html

Epstein, R. (2000). 'The Neural-Cognitive Basis of the Jamesian Stream of Thought', *Consciousness and Cognition*, 9, 550–575.

Evans, C. O. (1970). *The Subject of Consciousness*. London: George Allen & Unwin Ltd. Online at: http://www.mentalstates.net/SOC.html

Foster, J. (1979). 'In *self*-defence' in G.F. Macdonald (ed.) *Perception and Identity*, London: Macmillan.

Husserl, E. (1991). *On the Phenomenology of the Consciousness of Internal Time (1893–1917)*, edited and translated by J. B. Brough. Dordrecht: Kluwer.

James, W. (1952). *The Principles of Psychology*. Chicago: Encyclopedia Britannica Inc.

Johnston, Mark. (1987). 'Human Beings,' *Journal of Philosophy* 84: 59–83.

Locke, J. (1959). *An Essay Concerning Human Nature*, Vol. 1. New York: Dover

Lockwood, M. (1989). *Mind, Brain and the Quantum*. Oxford: Blackwell.

Lycan, W. (1997). 'Consciousness as Internal Monitoring', in Block *et al.*

Mangan, B. (2001). 'Sensation's Ghost: The Non-Sensory "Fringe" of Consciousness' Psyche, 7(18), October 2001.

Nathan, N. M. L. (1997). 'Self and Will'. *International Journal of Philosophical Studies*, Vol. 5 (1), pp. 81–94.

Pöppel, E. (1985). *Mindworks: Time and Conscious Experience*, New York: Harcourt Brace Jovanovich

Robinson, H. (forthcoming). 'Personal identity, self and time', in A. Batthany (ed.) *Mind: its Place in the World. Non-reductionist Approaches to the Ontology of Consciousness*, Vienna.

Rosenthal, D. (1986). 'Two Conceptions of Consciousness', *Philosophical Studies*, 94: 3, pp. 329–359.

Russell, B. (1984). 'On the experience of time', in *The Collected Papers of Bertrand Russell Volume 7*, London: George Allen and Unwin.

[13] My thanks to Tim Bayne, Graham Nerlich and Galen Strawson for helpful comments and discussion.

Sprigge, T. (1983) *The Vindication of Absolute Idealism*. Edinburgh: Edinburgh University Press.
Strawson, G. (1999). 'The Self and the SESMET', *Journal of Consciousness Studies*, 6, No. 4, pp. 99–135.
Whitehead, A. N. (1929). *Process and Reality*. Cambridge: Cambridge University Press.