From Phenomenal Selves to Hyperselves

BARRY DAINTON

Abstract
The claim that we are *subjects of experience*, i.e. beings whose nature is intimately bound up with consciousness, is in many ways a plausible one. There is, however, more than one way of developing a metaphysical account of the nature of subjects. The view that subjects are essentially conscious has the unfortunate consequence that subjects cannot survive periods of unconsciousness. A more appealing alternative is to hold that subjects are beings with the *capacity* to be conscious, a capacity which need not always be exercised. But this view can itself be developed in more than one way. The option I defend here is that subjects are nothing more than capacities for consciousness, a view I call the 'C-theory'. Although the C-theory supplies us with a potentially appealing account of the nature of subjects (and hence ourselves), there are challenges to be overcome. Olson has argued that identifying ourselves with what are, in effect, *parts* of human organisms leads to a variety of intolerable problems. I suggest that these problems are by no means insuperable. Bayne and Johnston have argued that identifying subjects with experience-producing systems is confronted with a different difficulty. What if these systems can produce multiple streams of consciousness at once. Whatever else they may be, aren't subjects the kind of thing that can have *just one* stream of consciousness at a time? In response I argue that this is true in one sense, but not in another. Once this is appreciated, the notion that a subject could have several streams of consciousness at once no longer seems absurd, or impossible.

1.

The metaphysical thesis that we are *subjects of experience* is in many ways an appealing one. Or at least, I have generally found it so, and when it comes to assessing philosophical accounts of our own nature, this is of no little importance. When I engage in a philosophical inquiry into the nature of the self, what I am seeking is an answer to the question 'What sort of thing am I?' For an answer to this question to be maximally credible it needs to be metaphysically coherent – it needs to make sense philosophically – but it also needs to be as intuitively plausible as possible. So in assessing a putative answer to the question if I ask myself 'Could I really be a thing of *that* sort?' I want to be able to respond 'Yes, of course, very easily

doi:10.1017/S1358246115000065 ©The Royal Institute of Philosophy and the contributors 2015

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

indeed'. The claim that what I am, fundamentally, is a subject of experience meets this requirement fully.

By a 'subject of experience' I mean simply this: a thing whose sole essential property is the ability to enjoy unified states and streams of consciousness. If I know anything about myself right now – or whenever I am awake and capable of thinking anything at all – it is that I am having some experiences, and these experiences form parts of a continuous, flowing stream of consciousness. My experiences may be misleading me as to the real character of other aspects of reality, but I can't seriously doubt that I'm really having the experiences I currently seem to be having. By virtue of this fact, if subjects of experience exist then I obviously qualify as one simply by virtue of being one of the things that is currently having experiences. Moreover, since the only requirement on subjects is their ability to have experiences, should it turn out that appearances are *very* misleading it need not and should not affect how I conceive of myself. Perhaps I am in fact a non-human entity on some distant planet who (for some reason) is suffering a series of vivid hallucinations. Since this alien is having experiences it too qualifies as a subject of experience.

Taking myself to be a subject has further advantages. In my search for an account of my own nature that I can readily believe, one of the things I am looking for is an account which makes it possible for me to survive undergoing changes or transformations which it seems obvious (to me) that I would in fact survive, if they were to occur. A key advantage of taking myself to be a subject is that I would then be able to survive readily conceivable transformations which would *not* be survivable if I were an entity of a different kind.

Vampires are non-human creatures of the night, who have the ability to turn ordinary human beings into vampires by feeding them their blood (according to some tales anyway). Suppose I have just imbibed some vampire blood, and am in the midst of the – quite agonizing – transformation. Do I survive? If the stream of consciousness that I currently enjoying continues on without any interruption in the felt flow of thoughts, perceptions and feelings – and we can stipulate that this is the case – it seems clear to me that I do survive this transformation. The idea that the stream of consciousness that is now mine could continue on, without interruption, but fail to take *me* with it, is not one I can take seriously. But this has a significant consequence. Since by the end of the transformation I am no longer a human being (even though I outwardly resemble one) it seems that I am not essentially a human being. For no human being can

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

survive being transformed into something that is very definitely not a human being.

Radical physical transformations needn't involve change of species. Why is it so easy for me to imagine myself existing in the form of an envatted brain? For two reasons. First, because like many others I assume that all my mental capacities and faculties – my capacities for consciousness included – reside solely in my brain. Provided it is kept in a healthy fully-functioning condition, a human brain doesn't need the rest of a human body in order to sustain a conscious human mind. (It might, of course, require its sensory nerves to be stimulated in appropriate ways if it is to enjoy anything resembling normal perceptual experience.) Second, because it seems obvious to me that provided these mental capacities continue to exist then *I* will continue to exist, even if I end up losing the rest of my body. I'm glad that in actual fact I have the rest of my body – I'd be sorry to lose it – but lose it I could. In being reduced to the condition of an envatted brain my existence as a subject would not be threatened in the slightest, for if my brain retains the ability to sustain a fully conscious mind it also retains the ability to sustain a subject in existence.

Or at least, it wouldn't be threatened provided the surgery does not damage the brain. Certain forms of damage would be inconsequential so far as my continued existence is concerned. Damage to neural tissues which are not relevant to the production of experience would not threaten my existence as a subject of experience. But if, say, some of my brain's visual systems were damaged, leaving me entirely incapable of visual experience, then I would be touched – diminished – at my very core, since a form of experience I was previously able to enjoy is no longer available to me. Even so, since I retain the capacity for other forms of experience, I continue to survive, my identity entirely undiluted and unimpaired. However, should my brain suffer damage which permanently obliterates its ability to produce experience in any shape or form, then I too would have ceased to exist. Once again, taking myself to be a subject of experience allows me to make sense of how a range of possible changes could have precisely the effects on *me* that I believe they would in fact have.

Consider a different case, one where the envisaged transformation is psychological rather than biological. Let's suppose – let's not enquire why – that I am about to be subjected to a radical brainwashing, a procedure which over a period of an hour or so will remove all my memories, my values, my personality traits, my likes and dislikes, my hopes and fears, and replaces them with quite psychological traits that are quite different. Provided my current stream of consciousness

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

continues on without interruption – provided I continue to have a unified stream of sensory and emotional experiences throughout, even though my thought processes become fragmented and inchoate – it seems absurd to me to suppose that this process would be fatal to me. This fact alone suggests that accounts of the self which require a strong degree of psychological continuity if we are to persist through time – and here the contemporary neo-Lockean view comes to mind – cannot be right. In contrast, if we are subjects of experience, there is no psychological constraint on our persistence at all. Provided we remain capable of experiencing – experiencing anything we remain in existence.

2.

There is more to be said about all the scenarios briefly outlined above, but we have seen enough for this at least to be clear: the notion that we are each of us, fundamentally, subjects of experience has considerable appeal, at the intuitive level. But we mustn't overlook the other half of the equation: can we make philosophical sense of this concept? From a metaphysical perspective, what *are* subjects of experience? Out of what are they constructed? What does it take for a single subject to persist through time?

There are different ways of answering these questions Although the doctrine that subjects are immaterial substances whose essential nature is to be conscious has a long and venerable history, it holds little appeal for those who believe that their conscious states are entirely the product of physical activity in our brains. The Cartesian view also has the unfortunate – or at least implausible – consequence that subjects are by their nature conscious throughout their lives, and so cannot survive periods of unconsciousness. Since few of us contemplate periods of dreamless sleep with the sort of dread which would accompany a visit to the dentist's chair, let alone the executioner's chamber, it seems most of us do not believe that periods of unconsciousness are invariably fatal – or that we are the kind of thing which cannot emerge from such periods unscathed.

In fact, the difficulties confronting the Cartesian view are of little import, for in making metaphysical sense of subjects of experience we do not need to appeal to immaterial substances of any sort. There are other possibilities, and of these some look to have considerable promise.

One option is to hold that subjects are things which essentially possess the *capacity* to be conscious. The notion that subjects and

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

capacities for experience are intimately linked is itself venerable. Locke certainly appreciated the connection, criticizing Descartes' position thus:

> … the perception of ideas being (as I conceive) to the soul, what motion is to the body; not its essence, but one of its operations. And therefore, though thinking be supposed never so much the proper action of the soul, yet it is not necessary to suppose that it should be always thinking, always in action. … We know certainly, by experience, that we sometimes think; and thence draw this infallible consequence, that there is something in us that has a power to think. But whether that substance perpetually thinks or no, we can be no further assured than experience informs us.[1]

As Locke clearly realized, holding that subjects are essentially possessors of capacities for consciousness has a very significant advantage. Capacities can be active during some periods and entirely dormant during others. A car has the capacity for self-powered motion, but most cars only exercise this capacity from time to time. Nevertheless, a well-maintained car continues to possess its capacity for self-powered motion even when it is not exercising it. Consequently, if we say that a particular subject continues to exist for as long as its capacities for consciousness exist, the notion that subjects can survive periods of complete unconsciousness is no longer problematic. A subject persists through such periods provided their capacities for consciousness continue to exist, completely dormant though they be.

Recognising that subjects and capacities for experience are intimately related is an important step, but there remains the task of establishing precisely what kind of thing subjects are. One option is to hold that subjects are *physical objects*, albeit of a distinctive kind: they are material things which possess capacities for consciousness. A position along these general lines has been elaborated and defended at considerable length by Peter Unger in *Identity, Consciousness and Value* (1990), and more recently – and succinctly – by Christopher Peacocke in *The Mirror of the World* (2014).[2]

[1] John Locke, *An Essay Concerning Human Understanding*, ed. P. Nidditch (Oxford, Oxford University Press, 1690/1975), Book II, chapter 1, §10.

[2] Peacocke holds that a subject of experience x is identical with a subject y if and only if x and y have the same S-generator – or in his terms 'same material integrating apparatus' (*The Mirror of the World*, 66). But he does not claim that a subject is numerically identical with an integrating

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

Our own mental capacities reside in our brains, but straightforwardly equating subjects with *human* brains is not an appealing step to take. After all, we know that there are (very) many non-human brains with the capacity to produce unified streams of consciousness – the brains of monkeys, dogs and dolphins, for example – and these non-human brains all sustain subjects of experience. Since it might be possible for subjects to be sustained by non-biological physical systems – some advanced form of computer, perhaps – it would also be a mistake to equate subjects with purely biological organs of any kind. What we need is a more general notion. Accordingly, let's use the term *S-generator* to refer any physical object or system that in its current form is capable of generating a single unified streams of consciousness. With this more refined concept in hand, should we take ourselves to *be* S-generators?

This position has its advantages. It can encompass subjects of many different shapes, sizes and constitutions. Also, a system with the capacity to produce two streams of consciousness – as may well be the case with human split-brain patients – will support two S-generators, and hence two subjects, which seems the right result. But this position also suffers from some serious drawbacks, at least in the form defended by Unger and Peacocke.

The latter both hold that the persistence conditions of S-generators, and hence ourselves, are broadly similar to those of other complex physical objects, such as cars, computers and caravans. Just as a car (or biological brain) can survive a complete changes in its constituent parts provided the changes don't occur too quickly – over a period of several days or weeks, say – so too can S-generators. But equally, just as a particular car (or biological brain) cannot survive if its parts are replaced very rapidly – in a matter of minutes or seconds, say – the same applies in the case of S-generators.

To see why this last claim is problematic, it's useful to consider a spectrum of cases. At one end of the spectrum we have a brain all of whose parts get changed in a gradual fashion – atom by atom, neuron by neuron – over a period of five years. At the other end of

---

apparatus, since the latter possess properties (such as being spatially localized in the brain) that the subject does not – he takes subjects to be metaphysically basic entities. But for present purposes what matters is that he does hold that 'the identity of a subject over time consists in the identity of the apparatus that integrates states and events in such a way that a single subject has, or may have, perceptions, sensations, thoughts, action-awareness, and the rest, both at a time and over time.' (Ibid., 65)

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

the spectrum, we have a brain all of whose parts are all replaced very rapidly indeed: near-instantaneously.[3] In between these extremes there is a vast range of intermediate cases, where the part-replacement happens more or less quickly. Importantly, in *each* of these cases in our spectrum, or so we stipulate, the mental life that is produced by the physical system(s) in question is subjectively indistinguishable: during their waking hours the subject(s) involved enjoy streams of consciousness that are as continuous as the stream you are currently enjoying.

Now, because they want subjects to have the same general persistence conditions as any other compound material object, proponents of the S-generator view are committed to the claim that somewhere along this spectrum of cases, the part-replacement is fatal to the original subject. If the replacement is carried out over a couple of days or weeks, for example, there is no problem: the subject survives; but if the replacement is carried out over a couple of seconds, the subject most definitely does *not* survive.

It is this last claim which I find incredible. If my own brain were suffer the envisaged part-substitutions, given that in *all* the cases along the spectrum, the flow of my consciousness – or at least, the consciousness being produced by the physical systems in question – would be entirely unperturbed, I find it impossible to believe that in *some* of these cases I would survive, but in others would cease to

[3] Achieving an instantaneous change of parts would be impractical using anything resembling standard surgical techniques, but this does not mean it is physically impossible, just rather improbable. Cosmologists have recently started worrying that the most numerous conscious subjects in the universe might be *Boltzmann brains*, i.e. brain-like physical systems that are produced by random quantum fluctuations; these systems only enjoy a brief period of experience: shortly after their creation they vanish again, as their particles disperse. The spontaneous creation of experience-sustaining brains is immensely improbable, but if the universe is infinite – as currently seems likely – there will nonetheless be vast numbers of these entities overall. If Boltzmann brains *are* physically possible, then there is no obvious reason why these same processes couldn't cause a brain at a certain time and location to suddenly vanish, and then – a moment or so later – be replaced by an exact duplicate (this newly created brain does *not* vanish straight away, but endures). Needless to say, this sequence of events is far less probable than the simple creation of a single briefly-existing Boltzmann brain, but improbable is not the same as impossible. If the universe is infinite, it probably happens from time to time. For more on Boltzmann brains, see Sean Carroll, 'Boltzmann's Anthropic Brain', *Discover*, August 1, 2006.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

exist. Given that the physical alterations to brain have no impact whatsoever at the experiential level, they do not threaten my existence in the slightest.

A variant of this scenario reveals a further potential difficulty for the S-generator approach. Consider a case in a region of the cell-replacement spectrum where the entire process takes place in around half an hour. We have been assuming thus far (albeit tacitly) that whenever a cluster of neurons gets ever-so-carefully removed from my brain it is simply discarded. Let's suppose that this is not the case. Instead, when each cluster of neurons is removed, it is (ever-so-carefully) frozen. Let's further suppose that a couple of days after the replacement procedure has been completed the frozen neurons are carefully un-frozen and re-assembled, to yield a fully functioning brain – a brain that is composed of the same particles – the same flesh and blood – as my original brain possessed, before the cell-replacement treatment commenced. Let's call the latter the 'original matter' or *OM-brain*, and the brain which is composed of new matter but which remains continuously functional throughout the *CF-brain*.

This scenario is of course a neural variant of the classic 'Ship of Theseus' puzzle. The latter puzzle concerns the relationship between material objects and their constituent parts. So far as the 'What happens to the original ship?' question goes, although there are different views as to what the correct answer is, it is generally accepted that our intuitions are pulled in both directions. It seems plausible to think that the CF-ship (as we can call it) is the original, but it can also seem plausible to hold that the OM-ship is the original. Now, if the fundamental premise of the S-generator view is correct, and (a) we are subjects of experience, and (b) subjects are themselves material objects, then we should experience a similar divergence of intuitions in the neural variant. It should seem perfectly intelligible to suppose that the original subject lives on sustained by the CF-brain, but the claim that the original subject is in fact the OM-brain should also seem entirely plausible. The latter, after all, is composed of the same matter in the same configuration as the pre-replacement-process brain. Peacocke for one agrees with this. Discussing such cases he writes 'we are liable to have conflicting intuitions about identity in any case that is structurally analogous to that of Theseus' ship'.[4]

However, as far as I can see, this is not plausible in the least. Given that the CF-brain exists, and continuously sustains a capacity for

[4] Peacocke, *The Mirror of the World*, 69.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

conscious – a capacity which could conceivably be exercised without interruption for the entire period – the claim that this subject is one and the same as the original subject is completely compelling. In contrast, the claim that the original subject – or something with a strong claim to be such – pops back into existence when the newly re-constituted OM-brain starts to sustain consciousness does not seem remotely plausible. The tiny brain-fragments that are removed from the CF-brain and then deep-frozen do not posses any capacity for consciousness, singly or collectively, and so manifestly do not constitute a subject. When the fragments are thawed and re-assembled into a working brain, the latter does sustain a subject. But this subject's claim to be numerically identical with the original subject is negligible compared with the strength of the subject sustained by the CF's brain's claim to this status. There is no ambiguity or lack of certainty here – it seems *just obvious* that these verdicts as to what happens to whom are true. In this respect at least, subjects do not seem to be the same kind of thing as compound material objects such as ships, trees or houses.

3.

Evidently, the S-generator approach is seriously problematic, at least as an intuitively appealing account of what subjects of experience are. It is perverse, to put it mildly, to suppose that an uninterrupted (and non-branching) stream of consciousness could fail to belong to a single unique subject, but as we have just seen, the S-generator approach delivers precisely this result.

Is it possible to do better? I believe so. In *The Phenomenal Self* I elaborated an account of subjects which is as naturalistically oriented as the S-generator view, but which does not identify subjects with material objects of any kind – and so does not suffer from the deficiencies we just encountered.[5] On this alternative view – which I called the 'C-theory' – subjects are what I call *C-systems*, and C-systems are *collections of experiential capacities* that possess some highly distinctive attributes. The C-system of a typical human subject (or so it is natural to think) is composed of a sizeable number of different sorts of experiential capacities. Although these capacities are for very different forms of experience – visual, auditory, bodily sensations, conscious thoughts, and so forth – the capacities

[5] Barry Dainton, *The Phenomenal Self* (Oxford: Oxford University Press, 2008).

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

which belong to just one subject at a given time have a distinctive feature. If these capacities are active, all the experiences they are producing are *co-conscious*, i.e., they form parts of a single unified conscious episode. If at the time in question these capacities are not active – and typically, the bulk of our experiential capacities at any one time fall into this category – they nonetheless have a distinctive feature: they are such that if they *were* active, all the experiences they would be producing *would* be co-conscious with one another. In short, at any given time, it is the ability to contribute to a single unified state of consciousness which distinguishes experiential capacities which belong to a single subject from those which don't.

It is very plausible to think that this criterion for assigning experiential capacities to subjects at-a-time works well for human subjects. It is no less plausible to think that it works equally well for *all* subjects, irrespective of whether their C-systems are more complex than those of the typical human, or less complex.

Experiential capacities which belong to the same subject needn't exist at the same time. If this were the case, subjects would be condemned to a strictly momentary existence, which would be unfortunate! Happily, the C-theory does not curtail our anticipated lifespans. The solution is simple: C-systems which exist at different times belong to a single subject if their constituent capacities have the ability to contribute experiences to a single continuous *stream* of consciousness. Our streams of consciousness have a distinctive form of unity which exists both at and over time, a unity which is experiential in character: just think of the way each brief phase of a typical stream of consciousness flows – *is experienced* as flowing – into the next. Since experiences which belong to a single stream of consciousness obviously belong to a single subject, it seems equally obvious that capacities for experience which can together produce a single stream of consciousness, whether simultaneously or successively, should also belong to a single subject.

We can call C-systems which have this distinctive property *directly stream-related*. There may well be some subjects who are capable of remaining conscious throughout their entire life-spans, whether this be measured in hours, weeks, years or centuries. Since all the experiential capacities of subjects of this kind are capable of contributing to a single uninterrupted stream of consciousness, these capacities are all directly stream-related. But since our own experiential capabilities are more limited – we cannot stay continuously awake for more than a day or so – a typical human subject possesses experiential capacities which are not directly stream-related (e.g. those which are separated by a few week, or months, or years). This does

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

not affect the fundamentals. We need simply say that experiential capacities belong to the same C-system, and hence subject, if they are either directly or *indirectly* stream-related. Experiential capacities are indirectly stream-related if they are not directly stream-related to one another, but they are connected by an overlapping chain whose successive members consist of experiential capacities which are directly stream-related.

There are a number of further complications which I need not enter into here, but for present purposes all that matters is the basic picture, and this is an appealingly straightforward one. Subjects of experience are defined by their capacity to enjoy unified states and streams of experience – this capacity is the *only* essential property subjects possess. According to the C-theory, the distinguishing feature of experiential capacities which belong to a single subject is precisely their ability to contribute to unified states and streams of consciousness. Moreover, by taking subjects to *consist* of capacities for consciousness we strip away everything that is inessential to a subject's existence, and isolate precisely what all subjects have in common.

Last but not least, if we equate subjects with C-systems, the problems encountered by the S-generator view simply vanish.

Suppose some immensely powerful entity is tampering with your matter in your cranium. Over the past ten minutes it has been annihilating your brain every thirty seconds, and replacing it on each occasion with an exact duplicate; but as previously, this massive material disruption has had no experiential consequences: the succession of brains in your skull has been sustaining a fully continuous stream of consciousness throughout. Depending on how we opt to individuate experiential capacities, this process may well have resulted in the destruction of your original experiential capacities, and their replacement by qualitatively indistinguishable duplicates. So far as the C-theory is concerned, although you are a subject, and subjects are composed of experiential capacities, since the affected capacities have been producing experientially continuous streams of consciousness throughout, this sequence of events poses no threat to your continued existence. In contrast, as we have already seen, proponents of the S-generator account are committed to holding that you cease to exist, even though your stream of consciousness flows serenely on.[6]

[6] The identity of an experiential capacity depends on three elements: the kind(s) of experience it produces when triggered, the conditions which trigger it, and the particular *ground* or *base* it has. If our own experiential capacities are properties of our brains, then their bases will, I take it, be neural systems of one kind or another. Whether or not the very rapid

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

The C-theory delivers a similarly plausible verdict in the neural variant of the 'Ship of Theseus' case. If the gradual replacement of the neurons in the CF-brain does not interfere with its capacity to produce unified streams of consciousness – and this is what we are assuming – then the outcome of the procedure is intuitively obvious: you remain in existence throughout, and *your* experience is wholly the product of that brain. The C-theory concurs with this verdict: since the CF-brain possesses a single uninterrupted experiential potential throughout, it sustains a single subject (in this case, you). Prior to its thawing and re-construction, the OM-brain, consisting as it does of scattered frozen chunks of neural tissue, is wholly incapable of sustaining any consciousness whatsoever, and so it does *not* constitute or sustain a C-system or a subject.

One of the most appealing aspects of Locke's writings on personal identity and the self is his modesty, open-mindedness, and lack of dogmatism. Since he took the view that we had a great deal still to learn about the nature of matter and consciousness, Locke remained resolutely agnostic on the issue of whether conscious minds are material or immaterial in nature. What's more, his own account of personal identity is generally thought to be perfectly compatible with both hypotheses. Three centuries later, most would agree that we still have a great deal to learn about the real nature of the physical world, and how it is related to conscious experience. Accordingly, the fact that the C-theory is also entirely neutral on this issue is a very definite advantage. The claim that we are subjects, and subjects are composed of capacities for consciousness, can stand irrespective of whether these capacities are possessed by entities that are material or immaterial in nature, or whether experiences themselves are material or immaterial. Even if it should turn out that our own experiential

---

replacement of one's neurons with exact duplicates leads to a replacement of one's experiential capacities depends on the relationship between capacities (or dispositional properties) in general and their bases. If they are distinct, as some hold, then the original capacities may survive; if they are identical, as others hold, this won't be possible. But irrespective of which of these views is correct, the original subject's existence is not threatened provided there remains in existence an uninterrupted potential for fully continuous consciousness. (By way of analogy, think of the 'third' power-carrying rail of a transit system; there can be a continuous electrical potential between two locations x and y irrespective of whether the track in between consists of a single long piece of track, or several shorter laid end to end.) We will be looking at the relationship between capacities and their bases in more detail later on.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

capacities are housed in immaterial substances, it could easily remain the case that we are numerically distinct from these substances. The only way we subjects could be identical with immaterial substances is if the latter are composed of *nothing but* experiential capacities, and if the latter were the case, it is not obvious why the substances in question should be taken to be immaterial in nature at all.

4.

As I hope is now clear, the C-theory offers a promising account of what we fundamentally are. Or to be more accurate, it offers a promising *framework* for thinking about this issue, one which can be developed in different ways.

The C-theory rests on the claim that our streams of consciousness possess an inherent unity and continuity that is entirely experiential in nature. On other occasions I have defended the view that synchronic and diachronic phenomenal unity should both be understood in essentially the same way. On this view, streams of consciousness consist of (experiential) parts that are experienced together, or *co-conscious*, both simultaneously and successively. Co-consciousness in its synchronic and diachronic forms is a primitive inter-experiential relationship, or so I argue.[7] Although I believe this account has advantages over the alternatives, anyone who (for whatever reason) subscribes to a different account of the unity of consciousness – whether synchronic or diachronic – can still adopt the C-theory. All the latter requires is for our streams of consciousness to possess a genuine unity, of an experiential kind, both at and over time. The C-theory is not beholden to any one account of how this unity is best understood.

Similarly, whether or not consciousness essentially involves some form of *self-consciousness* continues to be debated. Locke famously held that a person is something which possesses the ability to *think about* itself as a persisting conscious being. To be capable of such thoughts a being requires a good deal of cognitive sophistication, and it seems very plausible that there might be subjects of experience – e.g. non-human animals, human infants – who lack this degree of

[7] See my *Stream of Consciousness* (London: Routledge, 2000/2006) for further elaboration. I respond to some recent challenges to my account of diachronic unity in 'Flow, Repetitions and Symmetries: Replies to Lee and Pelczar', in N. Oaklander (ed.) *Debates in the Metaphysics of Time* (London: Bloomsbury, 2014), 175–212.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

sophistication. Does the experience of even these quite primitive subjects involve some form of minimal self-consciousness? Some say yes, while others are more skeptical.[8] While the outcome of this debate is relevant to the forms of experience that are open to subjects, as far as I can see they do not impact at all on the ability of the C-theory to give us a believable account of the existence and persistence conditions of subjects.

Other relevant but as yet unresolved questions are metaphysical rather than phenomenological in character. There are, for example, very different accounts of the nature of dispositional properties, and the outcome of *these* debates will obviously influence how we think of our own natures, or at least they will if we are constituted of nothing but experiential capacities – which, after all, are a particular species of dispositional property. We encountered one example of this in the previous section, and we will be encountering another later on.

There are also issues of a more empirical or scientific kind which will influence the overall picture. One of the advantages of the C-theory, particularly over its Cartesian rival, is that it allows us to survive periods of complete unconsciousness. According to it, you will survive your next period of dreamless sleep provided your capacities for consciousness remain in existence during this period, in a dormant state. It certainly seems plausible to suppose that our brains retain their capacities for consciousness during periods of dreamless sleep – we can, after all, be woken in a second or two, and when it comes to their neural structures, waking brains and (dreamlessly) sleeping brain are really quite similar. Even so, we will have to wait for advances in neuroscience before we have a clear understanding of precisely what is involved in a sleeping – or anaesthetized – brain retaining its capacities for consciousness. Although advances are being made, the sleeping brain is largely as mysterious as its waking counterpart.

Although issues such as these are interesting and important, I will not be focusing on them here. Instead, I want to concentrate on some

[8] Following Brentano (and ultimately Aristotle) many classical phenomenologists have endorsed the doctrine that consciousness necessarily involves some form of self-consciousness, recent advocates include D. Zahavi, *Subjectivity and Selfhood* (Cambridge: Cambridge University Press, 2005); U. Kriegel *Subjective Consciousness: A Self-Representational Theory* (Oxford: Oxford University Press, 2009). For some skeptical doubts see Peacocke, *Mirror of Nature,* chapters 1–3, and Dainton, *The Phenomenal Self* , §2.4, §8.2.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

important challenges of the metaphysical variety, challenges which will have to be overcome if the C-theory is to prove viable. Unless these foundational challenges *can* be overcome, expending further work on developing the C-theory will simply not be worth the time or trouble.

5.

If we are subjects of experience, and subjects are C-systems, then it looks very much as though we are not human beings. A human being is a complex biological organism, with the corresponding synchronic and diachronic identity conditions. Spelling out precisely what these identity conditions are is a non-trivial task, but will involve details relating to body shape and size, chemical structure and constitution, genetics, and so forth. A C-system is a collection of capacities for unified states and streams of consciousness. These are obviously very different kinds of thing. Although the experiential capacities of human subjects are (very probably) grounded in their brains, as we have already noted, there are easily conceivable circumstances in which subjects and their organisms can go their separate ways. The bite of a vampire would terminate the life of the human animal I am currently housed in, but it would not put an end to *me*. Alternatively, I might be reduced to the condition of some envatted consciousness-sustaining brain-parts. I could not continue to exist in the absence of my biological organism if I were numerically identical with it.

So much seems clear, but questions remain. If I am not numerically identical with my human animal or organism, just how am I related to it? There remains a story to tell about this. Moreover, my non-identity with my organism is itself a source of difficulties. My organism may be an animal rather than a subject of experience, but it has an impressively wide range of properties. It can walk, run, sleep and digest, but it can also read and write, talk, think thoughts, remember its past, feel emotions and perceive the world around it. By virtue of having a brain, my organism has a mind, and by virtue of having a mind my organism has an array of cognitive and experiential capabilities. Indeed, it looks to have exactly the same cognitive and experiential capabilities that I have! But if so, we are confronted with a worrying problem. If I am a thinking, experiencing being, and my organism is a thinking, experiencing being, then if I and my organism are distinct things we have *two* thinking experiencing beings. Since the same applies to every human subject, the world's

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

population of thinking, experiencing beings is double what we have previously assumed. And while this is bad enough, there is worse to come. What makes me so sure that the thinker of *these* thoughts is the subject, rather than the organism? If I can't rule this out, how can I rule out the possibility that I am the organism after all, rather than the subject?

In recent years this objection to variants of the Lockean approach to selfhood has been forged into a potent weapon by a number of philosophers, among whom Paul Snowdon and Eric Olson have been particularly prominent.[9] Much-discussed in the literature, it goes by a variety of labels, e.g. the 'too many thinkers problem' or the 'too many subjects issue'. Whatever label we choose, the problem *is* potentially a serious one. Any account of the self which leads to a doubling of the number of experien*cers* without a corresponding doubling in the number of experien*ces*, is in very serious trouble.

There are a number of ways of countering this overpopulation menace, but my preferred solution has long been a mereological one. On this view, subjects, in the form of C-systems, are related to their human organisms as parts to wholes. This is a very natural position to adopt; after all, there is a lot more to a typical human organism than a collection of experiential capacities. It also allows us to respond to the overpopulation problem as follows:

> Can human organisms have thoughts and experiences? In one sense it's true to say that they can, but in another it's not. Organisms have experiences in the way that a car can tell the time: by having a *proper part*—in the form of a clock on the dashboard, say—that can tell the time. The vast bulk of a car (its engine, its chassis, the seats and wheels and so forth) is entirely irrelevant to its time-keeping capacity, which is wholly located in one small part, the clock. To register this difference we can say that the whole car possesses the capacity to tell the time in a *secondary* way, the clock possesses this capacity in a *primary* way. Analogously, a human organism only possess experiential capacities in a secondary way. The primary or non-derivative possessor of these capacities is the subject of experience, in the form of a C-system, which is in turn located in the organisms brain.

[9] Paul Snowdon, 'Persons, Animals and Ourselves', in C. Gill (ed.), *The Person and the Human Mind* (Cambridge: Cambridge University Press, 1990), 83–107. Eric Olson, *The Human Animal* (Oxford: Oxford University Press, 1997); 'An Argument for Animalism', in J. Barresi and R. Martin (eds), *Personal Identity* (Oxford: Blackwell, 2003).

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

If we take this line, the overpopulation problem is dissolved at a stroke. Recognizing that human organisms possess experiential capacities in a secondary way does not increase the number of primary possessors of these capacities. There remains just one primary possessor – and hence subject – per organism. Nor does it increase the number of experiential capacities or experiences. Human organisms possess the capacity to metabolize alcohol in a secondary way, by virtue of possessing a part – the liver – which possesses this capacity in a primary way. There is just *one* metabolic capacity per animal, possessed in two different ways.[10]

6.

In his recent paper, Derek Parfit has espoused a solution to the overpopulation problem along just these lines, and for (largely) these reasons.[11] More recently still, Olson has discussed this alleged solution. He calls it *thinking-subject minimalism*, or just 'minimalism' for short, characterizing it as the view that an organism cannot think 'because it has a smaller part that *really* thinks'.[12] Olson goes on to argue that although this view may seem to have its attractions, when examined more closely it turns out to be deeply troubled.

One of the problems to which Olson draws our attention runs thus. According to minimalism, a being thinks or experiences in the primary non-derivative sense only if all of its constituent parts are directly involved in its thinking and experiencing. This is why, after all, an entire organism only possesses experiences or experiential capacities in a secondary or derivative manner. But can we draw, in a principled way, a boundary between those parts of an organism that

[10] For a fuller exposition see *The Phenomenal Self* , §7.7; I also suggest a solution along these lines in my 'Review of Eric Olson's *The Human Animal*, in *Mind* **107**, (1998), 679–82. Galen Strawson adopts the same position in 'The Self and the SESMET', *Journal of Consciousness Studies*, **6**:4, (1999), 99–135.

[11] D. Parfit, 'We Are Not Human Beings', *Philosophy*, **87**:1 (2012), 5–28.

[12] E. Olson, 'The Nature of People', in S. Luper (ed.) *The Cambridge Companion to Life and Death* (Cambridge University Press 2014, 30–46); also his 'On Parfit's View That We Are Not Human Beings', in A. O'Hear (ed.), *Mind, Self and Person* (Cambridge University Press, forthcoming)

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

are directly involved in experiencing, and those which are not? Consider the case of walking or eating, and imagine what it would be like trying to specify which of the atoms in your body are 'directly' involved in these activities. How would you go about it? Some parts of your body are *more* involved with walking than others (e.g. your feet), but even here there are some parts of your feet (e.g. your toenails) that aren't obviously involved at all. So in these cases, Olson suggests, the idea that some parts of a person are 'directly involved' in their walking is simply confused: there are no such parts. Moreover, he argues, precisely the same applies in the case of thinking or experiencing.

Transferring these considerations into a C-theoretical key, is Olson right in claiming that it is impossible, even in principle, to specify which parts of an organism are directly responsible for the capacities for consciousness it possesses? I cannot see that he is. Olson is right that it isn't easy to specify which parts of an organism are specially relevant to its capacity for walking. The ability to walk (in a normal fashion) involves host of different factors: muscular and skeletal conditions must be met, but a functioning nervous system and brain (or some of it) are required too. In assessing the minimalism approach to subjects of experience a better analogy would be functions or capacities that are localized in a single organ, such as the ability to pump blood round an organism's body (which our hearts accomplish), or extract oxygen from the atmosphere (our lungs are good at this). The experiential capacities relevant to subjecthood are, after all, localized in a similar fashion in our brains – if they weren't, we wouldn't be able to enjoy a subjectively normal life while reduced to a (functioning) brain-in-a-vat.

There is no reason why things must stop there. As neuroscience progresses, we may well learn more about precisely which neural systems are the minimal causal bases for the different kinds of experience we can enjoy. These minimal bases will no doubt overlap considerably, but not completely. It is already clear, for example, that different parts of the brain are involved in vision, hearing, taste and smell. Moreover, if certain aspects of the brain's functioning are destined to remain mysterious there is every reason to think that there *exist* minimal causal bases for different experiential capacities, even if neuroscience proves incapable of discovering their exact contours. We should also bear in mind that the task of isolating these minimal causal bases may be easier for some subjects than others. Subjects whose experiential capacities are grounded in digital computers, for instance – if there be such – will no doubt prove less troublesome in this regard.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

7.

Olson finds minimalism problematic for a second reason. Let's suppose – as Olson is prepared to grant, for the sake of argument – that our different mental capacities have distinct (but partially overlapping) causal bases. The neural systems which constitute the causal base for your capacity to remember faces will be different from your capacity to see faces, or hear the sound of a trumpet, feel sensations of warmth on your skin or think conscious thoughts. Each of these systems is a distinct physical thing, directly responsible for a different form of experience. In the context of the C-theory, embracing minimalism amounts to equating subjects with only those parts of an entire organism which are directly involved in the production of experience. If we follow this minimalist maxim, Olson argues, it looks very much as though a human organism sustains not *one* subject, but *many*. There will be the subject that consists of the neural system that is responsible for your ability to remember faces, another that consists the system responsible for your ability to hear trumpet-sounds, another for the system responsible for your ability to feel warmth, and so on, and on. As Olson observes, the result is a disaster: 'What we take to be a person able to perform all sorts of mental operations is really many beings, each able to perform only one. If, as Locke said, a person is by definition both intelligent and self-conscious, then there are no people, and we do not exist.'[13]

Subjects of experience (as I am construing them here) are not Lockean persons, since they do not essentially possess high-level cognitive abilities required for full self-consciousness in Locke's sense. Nonetheless, if each our experiential capacities turns out to constitute a distinct subject, in the way Olson suggests, the C-theory would be in very serious trouble. In fact, this is not the case.

As I have already indicated, the C-theory is based on a simple but compelling criterion for allocating experiential capacities to individual subjects. At any given time, experiential capacities which have the ability to produce co-conscious experiences belong to the same subject, and those that lack this ability belong to different subjects. Accordingly, if at any given time the neural systems in your brain that are responsible for your ability to see faces, hear trumpet-sounds and feel sensations of warmth are such that when active and producing experience, the experiences they produce all belong to a single unified conscious state, then these systems – or rather, the experiential capacities they possess – all belong to the same

[13] 'On Parfit's View That We Are Not Human Beings', this volume.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

C-system and hence the same subject, in this case you. In contrast, if when exercised simultaneously the experiences these capacities produce do not exist within the same unified stream of consciousness, they do *not* all belong to the same C-system or subject.

For Olson's argument to get off the ground, it would have to be plausible to suppose that group of experiential capacities could be such that (a) they all belong to a single C-system, and (b) each capacity in the group also constitutes a distinct subject in its own right. But as far as I can see, it is not remotely plausible to suppose this could be the case. By virtue of being able to produce co-conscious experiences when active together, experiential capacities which belong to the same C-system are manifestly *not* distinct and separate subjects in their own right. The ability to contribute experiences to the same stream of consciousness is surely the defining trait of capacities which are unambiguously constituent part of a *single* subject's mind. The situation is altogether different in the case of experiential capacities which, when active simultaneously, produce experiences which are not co-conscious with any other experiences, and which lack the ability to contribute to unified conscious states which include the contributions of other experiential capacities. These entirely isolated capacities might very well constitute distinct subjects in their own right.[14] But isolated capacities such as these are also very different from capacities which are integral parts of C-systems – i.e., capacities which, when active together, can and do contribute to fully unified conscious states.

There is a more general lesson here. When considering the circumstances in which one or more experiential capacities constitute a single subject it isn't sufficient to consider solely the forms of experience *these* capacities taken in isolation can produce when active; one must also consider their relationships with *other* experiential

[14] In *The Phenomenal Self* (§8.1–8.3) I broach the question of how simple a subject of experience could possibly be. Could there be a subject with a mind so primitive that it possesses the capacity for just a single simple form of sensation? An issue anticipated by Hume: 'We can conceive a thinking being to have either many or a few perceptions. Suppose the mind to be reduc'd even below the life of an oyster. Suppose it to have only one perception, as of thirst or hunger. Consider it in that situation' (*A Treatise of Human Nature*, (ed.) Selby-Bigge, Oxford: Clarendon Press, 1739/1978: 634) I argue that such minds may well be (logically) possible. If so, then the entirely isolated experiential capacities envisaged by Olson *would* constitute subjects of experience in their own right – or at least, they would if (as we are assuming) each produces an entirely separate stream of consciousness.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

capacities. Provided minimalists do not forget this, they have nothing to fear from Olson's objection.

8.

Olson's third main criticism of the minimalist approach tacks in a different direction. In rejecting the claim that we are animals minimalists avail themselves of a solution to the overpopulation problem, but they also incur a debt: just what are we, if we are not organisms? One option is to say that we are identical with our brains. However, as Olson points out, holding that we *are* a biological organ has highly counterintuitive consequences. After your (bodily) death, your brain might be pickled in formaldehyde and preserved, completely lifeless, in a jar for decades. If you are your brain, then that's *you* in pickling-jar. Needless to say, this is not a very plausible result, particularly for anyone who believe that continuing to have some form of mental continuity is essential to their continuing to exist, and most minimalists fall into this category. However, all is not lost. There is an alternative view, which Olson suggests might prove appealing to minimalists: we identify ourselves not with brains *per se*, but with brains which possess the capacity to sustain a mental life of some specified kind or level. Olson calls this the 'functioning brain' view.

In effect, the functioning brain view is a variant of the 'S-generator' doctrine I considered earlier. It is also vulnerable to the objections I leveled against that doctrine: a rapid replacement of parts will be fatal to a functioning brain, but not to the subject it sustains, provided the brain(s) capacities for experience remain intact, and capable of producing unified states and streams of consciousness. Olson puts his finger on a different problem. If we say we are functioning brains, we are confronted with an awkward fact: our *ordinary brains* haven't gone away. As biological organs, with the persistence conditions of such, our ordinary brains are capable of doing things that functioning brains can't do, such as survive being kept lifeless for years in a pickling jar – or more simply, continuing to exist without being able to sustain any sort of mental life. But equally, under normal circumstances, when normally connected to healthy human bodies, our ordinary brains are capable of doing everything that a functional brain can do. Just as functional brains are capable of sustaining a conscious mental life, so too are ordinary brains. So if functional brains are subjects, by virtue of this ability, why aren't ordinary brains? We are confronted once more with the overpopulation problem.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

For advocates of the functional brain view this is a very real difficulty. However, consisting as they do of nothing more than capacities for experience, C-systems are a quite different kind of entity. From the perspective of the C-theory, subjects – in the form of C-systems – are nothing more that *dispositional properties* of both ordinary brains and functional brains. More specifically, C-systems are properties of the minimal causal basis of a brain's experiential capacities mentioned earlier. Although there is much yet to be discovered about the neural correlates of consciousness, given what we do know it seems very unlikely that our *entire* brains are involved in the production of consciousness – damage to certain brain regions has no effect on our experiential capacities at all. We can thus appeal once again to a version of the mereological solution to the threatened doubling of subjects: both ordinary and functional brains possess experiential capacities in a derivative or secondary way, whereas the primary possessor is the C-system itself. Since there is just *one* primary possessor of experiential capacities in the picture, the overpopulation problem simply doesn't arise.[15]

9.

The thesis that we are subjects of experience, where subjects are entities whose defining trait is their capacity to have conscious experiences may be an appealing one, but there is troublesome scenario, of a rather exotic kind, which suggests that it is fatally flawed, or so Tim Bayne and Mark Johnston have both argued.

In his *Surviving Death* (2010), Johnston envisages the following sequence of events. First of all, we are to suppose that a subject whose experience has hitherto been qualitatively quite similar to your own, becomes convinced – correctly, it turns out – that their experience is not veridical, but hallucinatory. What's more the beings who are in control have the power to provide the subject

[15] The issue of precisely how C-systems are related to brains is bound up with several very general metaphysical questions concerning (a) the relationship between objects and their properties, and (b) the nature of dispositional properties. To illustrate, for anyone inclined to identify dispositional properties with their physical bases it will be an option to maintain that in the human case, experiential capacities are identical with certain neural systems. On the resulting view, C-systems will almost certainly be *proper parts* of human brains, rather than properties of them. But irrespective of whether C-systems are parts of brains, or properties of them, it remains the case that they are *less* than entire brains, and so not identical with them.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

with any form of experience they choose. (In fact, it is by exercising this power that they manage to convince the subject that their experience is in fact hallucinatory.) This subject then receives the following message from the controllers:[16]

> We have had the ability to create artificial conscious beings for centuries, and we enjoy experimenting on our creations in all sorts of ways. However, despite being in possession of vast psychophysical knowledge, in all these centuries we have never managed to achieve the goal of one physical realizer of consciousness to one stream of consciousness. What we know of psychophysics tells us that physical realizers always produce streams of consciousness (qualitatively identical) in batches of seven. What's more, our best psychophysics strongly suggests that this problem cannot be overcome. So there is in fact no single physical system that is especially causally responsible for any one of the seven distinct streams of consciousness in which these words are now appearing.

Johnston goes on to argue that if you were to find yourself in these circumstances, there is only one conclusion that you should draw: contrary to what you have doubtless assumed thus far, you do not in fact exist. In the envisaged circumstances, there is no *you* thinking your thoughts. Or at least, this is the conclusion you should draw if you are inclined to equate subjects with physical systems possessing the capacity for consciousness.

Bayne introduces an imaginary case along the same lines, one which involves single physical system producing several simultaneous streams of consciousness. He goes on to spell out precisely why this 'multiplicity scenario' (as I will call it) poses a problem for anyone who wants to equate selves or subjects with 'the underlying substrate that is responsible for generating the stream of consciousness – the *machinery* in which consciousness is grounded.'[17] If we know anything about subjects it is that they are the kind of thing which possess a single stream of consciousness. In other words, at any one time a subject who is conscious will have just *one* synchronically unified state of awareness. Hence the problem. The multiplicity scenario shows us that there is no *a priori* guarantee 'that a single consciousness-generating mechanism will produce

[16] Here I paraphrase a passage from Johnston's *Surviving Death* (New Jersey: Princeton University Press), 143.

[17] T. Bayne, *The Unity of Consciousness* (Oxford: Oxford University Press, 2010), 288.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

only one stream of consciousness' at a given time.[18] If consciousness-generating mechanisms can produce several streams of consciousness at once they cannot *be* subjects, they're just the wrong kind of thing.

Although the C-theory identifies subjects with systems consisting of *capacities* for consciousness, not consciousness-generating mechanisms, it is difficult to see that this suffices to circumvent the problem. In our own case at least, C-systems are (very probably) grounded in our brains. If a single brain can produce multiple streams of consciousness at once, what are we to make of the situation? Does the brain possess one C-system or several? If the latter, then – by the lights of the C-theory – this brain sustains several subjects, and since each of these subjects possesses a single stream of consciousness there is no problem. But if there is just a single C-system involved we are in trouble, for Bayne is right: it is difficult to see how a single subject could have multiple streams of consciousness at any one time. In these circumstances it looks very much as though we would be obliged to say that the brain in question does not sustain a subject at all, even though it possesses experiential capacities, and is in fact generating several streams of consciousness. So we are left confronted with this question: just how many C-systems *does* such a brain possess?

10.

Bayne and Johnston's multiplicity scenario leaves the C-theory in an awkward position, for it gives rise to questions for which there are no easy or obvious answers. These questions only need be addressed if the scenario itself is really possible. But is it? The envisaged state of affairs is certainly very exotic indeed, and I am by no means certain that it is logically possible. But equally, I am not certain that it isn't.

So far as I can see, we would have to know a good deal more about the way in which conscious experience is related to the physical world before we would be in a position to answer this question with any confidence. There are certainly some positions on this relationship which leave the multiplicity scenario looking very dubious. Suppose consciousness is wholly material in nature, and that our ordinary conscious states extend through almost the entire volume of our brains – as might be the case, for example, if different forms of experience are constituted by different modes of excitation in some physical field generated by neural activity. In this case, supposing that a single

[18] Ibid., 288.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

brain could produce seven streams of consciousness would be as absurd as supposing that a single pint class could contain seven pints of water.

But on other views of the nature of consciousness, the multiplicity scenario looks a good deal less problematic. Suppose, for example, that a form of property dualism is true. Experiences are non-physical particulars that are nomologically correlated with certain patterns of physical activity – the sort of view which David Chalmers has sympathetically explored on various occasions.[19] In our world (or so we generally assume), the psychophysical laws are such that human brains produce just one stream of consciousness, but in other possible worlds the psychophysical laws may be such that the brains that are physically indistinguishable from ours produce two streams of consciousness, or seven, or seven hundred. So far as I can see, there is nothing metaphysically incoherent in these multiplicious states of affairs. Since the additional streams of consciousness are non-physical, they do not exist in physical space, and so are not competing for room – the situation is very *un*like squeezing several pints of water into a pint glass. The fact that these experiences lack a spatial location might itself be a source of trouble. If, as I have been assuming, all of the multiple streams of consciousness produced by a given brain over a given interval are qualitatively indistinguishable, what individuates them? What makes it the case that there are seven streams (say), rather than just one? However, provided that the property dualist takes experiences to be basic particulars, there is no difficulty here either. The identity and individuality of an experience, and hence of the streams of consciousness they compose, would be as autonomous and secure as that of any material object – or even the entire physical universe.

There may well be other possibilities compatible with the multiplicity scenario. Perhaps consciousness is wholly physical in nature, but in some physically possible universes it exploits the 'hidden' spatial dimensions that string theorists believe to exist. In which case, a three-dimensional brain might sustain multiple streams of consciousness which are located in these additional spatial dimensions. I am not suggesting that this is very likely. The lesson, rather, is that when it comes to the relationship between the physical and phenomenal realms, at present we are confident of very little. So much so that we cannot, with any confidence, deem the multiplicity scenario to be metaphysically impossible.

[19] Such as his *The Conscious Mind* (Oxford: Oxford University Press, 1996).

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

11.

So we are left with the task of working out how to make best sense of Johnston and Bayne's hypothetical state of affairs. For anyone who takes the view that subjects of experience are *identical* with their streams of consciousness – Galen Strawson defends this line – there is no difficulty here.[20] If your brain is producing seven streams of consciousness, it is sustaining seven subjects of experience, and you are one of these seven. But the C-theory I have been defending does not equate subjects with streams of consciousness. Instead, it equates them with the clusters of capacities for consciousness. We are thus returned to the question posed earlier. If a brain is producing seven streams of consciousness, does it possess seven C-systems, or just one? If the latter, the brain in question is sustaining a single subject, one possessing seven streams of consciousness. If the former, the brain is managing to sustain seven subjects, each with their own stream of consciousness. The idea that a single subject could have several entirely distinct streams of consciousness at the same time looks to be problematic: as Bayne rightly stresses, subjects are the kind of thing that have just one unified consciousness at any given time. So the situation would be easier – for the would-be C-theorist at any rate – if we can say that in the envisaged circumstances the brain is sustaining seven C-systems, rather than one. But is such a state of affairs really intelligible? Before attempting to answer this question it will help if we take a closer look at the properties from which C-systems are built.

An experiential capacity is a particular species of dispositional property. A typical physical object – a glass vase, say – has a great many properties of this sort. If filled with water, it will not leak (provided it is oriented the right way); if dropped, it will shatter (provided the surface it falls on is sufficiently strong and rigid), and so on. Objects have the dispositions they do because of their particular natures, or more precisely because of their causally relevant features. If an object O possesses the dispositional property D to produce an effect of type E under conditions C, it is because O possesses some feature F that is causally relevant to occurrence of E. In the not inconsiderable literature on the nature of dispositional properties, feature F is known as the 'ground' or 'base' of the disposition D. This is all

[20] G. Strawson, *Selves* (Oxford: Oxford University Press, 2009). Strawson, notoriously, also holds that subjects only last a fraction of a second, as measured by clock time – but this is because he believes that our streams of consciousness are much briefer than most of us tend to think.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

relatively uncontroversial.[21] However, when it comes to the precise relationship between dispositions and their bases, there is a significant divergence of views, one that is very relevant to how the multiplicity scenario should be interpreted.

One view, associated with Armstrong (1969) and Mellor (2000) is that dispositions and their bases are identical: D and F are simply one and the same property.[22] Taking this stance has its advantages. If we adopt it, there is no mystery as to how dispositions can be causally efficacious. A particular kind of berry is *poisonous* because some of the damaging effects some of the chemicals in it have on the human body, and these same chemicals constitute the disposition's base. Also, since we don't have to recognize dispositional properties over and above the base-properties, the Armstrong-Mellor view is parsimonious. However, in the eyes of some it is also problematic. Many different foods and liquids possess the property of being poisonous-to-humans, and in many cases the chemicals involved are themselves different, e.g. with holly berries the toxic agent theobromine, in yew berries the prime culprit is taxine. If we hold that dispositions and their bases are identical, then the property *poisonous* doesn't exist. What exists are the (very) many different specific chemical compounds that produce unpleasant or fatal effects in people when ingested.

To circumvent this unwelcome (to some) consequence, Prior, Pargetter and Jackson (1982) argued that it is a mistake to identify dispositional properties with their bases.[23] Instead, dispositions should be viewed as second-order properties of *having a first-order property which plays a causal role of a particular sort*. On this view, something has the property of being poisonous if it possesses a first-order property (containing the compound potassium cyanide, say) which has the relevant causal consequences (ingestion proves troublesome to humans). When construed in this 'functionalist' way, a dispositional property *has* a causal base, but it is not identical

[21] It is not *completely* uncontroversial. There are those – causal structuralists, for example – who hold that there can be capacities or dispositions which are not grounded in properties of a non-dispositional sort. Whether 'bare dispositions' are possible is an interesting question, but as far as I can see it is largely irrelevant to the present issue – for reasons which will emerge.

[22] D. Armstrong, 'Dispositions are Causes', *Analysis* **30** (1969), 23–6; H. Mellor, 'The Semantics and Ontology of Dispositions', *Mind*, **109** (2000), 757–780).

[23] E.M. Prior, R. Pargetter, F. Jackson, 'Three Theses about Dispositions', *American Philosophical Quartely*, **19** (1982), 251–257.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

with this base. This approach has its advantages, but it too has its critics. Heil (2003), for example, complains that it renders dispositions themselves causally impotent, since all the causal work is done by the first-order properties.[24] Another option, defended by Mumford is to opt for token-identities between particular instances of dispositions and instances of their bases.[25] There are many ways of being poisonous, and so many different bases, but each instance of this dispositional property is identical with some particular base.

The debate on the nature of dispositions is very much alive, and this brief survey does not do justice to its intricacies, but we have what is needed for present purposes. The question we were considering was this: if a particular brain sustains seven streams of consciousness at the same time, how many C-systems are involved, seven, or just one? Since there is just one brain involved, and this brain is in all non-experiential respects similar to an ordinary brain which sustains just a single consciousness (or so we are supposing), the causal base is going to be precisely the same in both cases. Appreciating the significance of this will be easier if we focus on a simple example.

Let's begin with a single experiential capacity, call it P1, whose causal base F is such that it produces an single token E-type experience when triggered in circumstances of type C. Let's now suppose that the psycho-physical laws are different, and that when F is a triggered in the same way it produces not one, but two token E-type experiences. In this case, can we hold that F is grounding two distinct experiential capacities, P1 and P2, each for E-type effects? If so, then there is no obstacle to our supposing that an otherwise ordinary brain which is sustaining seven simultaneous streams of consciousness is also sustaining seven C-systems.

Here the different views on the nature of dispositions in general enter the scene. If we follow Armstrong, Mellor and Mumford in holding that dispositions are identical with their causal bases, in either type- or token-specific ways, we immediately encounter a problem. The hypothesis that P1 and P2 are numerically distinct powers is simply incoherent. Since P1 and P2 are identical with the same (token) base, on this view of dispositions they must be identical with one another. If this view of dispositions is correct, a single brain can sustain just one C-system at a time.

For proponents of the higher-order functionalist account, the situation is less cut-and-dried. Since the causal base F we are considering

[24] J. Heil, *From an Ontological Point of View*, (Oxford: Oxford University Press, 2003).

[25] S. Mumford, *Dispositions* (Oxford: Oxford University Press, 1998).

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

can produce two E-type experiences, rather than just one, there is a case for holding that it grounds two numerically distinct capacities for single E-type experiences. Since on this view dispositional properties are not identical with their bases, the fact that these capacities – if they existed – would have a common base does not entail that the capacities themselves must be identical.

However, before concluding that the multiple C-system interpretation of the multiplicity scenario is viable, at least if we construe dispositions in the Prior-Pargetter-Jackson way, there a hurdle to be overcome. Although we can interpret the case as involving two capacities, each for an E-type experience, we could also interpret it as involving a single capacity for *two* E-type experiences. Since each of these hypotheses fits all the facts equally well, it is difficult to see why we should prefer the two capacity interpretation simply because it makes life easier for the C-theory. Certainly, the facts of the matter do not seem to *compel* us to adopt the two capacity interpretation. Indeed, the fact that the one- capacity interpretation is more economical might be a reason for preferring it over the rival two-capacity view.

In the light of these considerations, the multiple capacity – and hence multiple C-system – interpretation looks to be on very shaky ground. Firstly, because it is only intelligible at all on the 2nd-order view of dispositions, a view which many find problematic. Second, because even if the 2nd-order view proves to be correct, the issue of whether or not a brain is sustaining a single subject or many looks to be grounded on nothing more solid than a *decision* to describe matters in one way rather than another, a decision that is in no way compelled by the underlying facts. Many will find it objectionable to suppose that the issue of how many distinct subjects of experience a given brain is supporting at a given time could be conventional in this sort of way – I certainly do.

12.

Since the alternative has proved unsatisfactory, it looks as though we should accept that if a brain is producing seven streams of consciousness it possesses just one C-system. This means, of course, that if we equate subjects of experience with C-systems, there is just a single subject present, one which is enjoying seven streams of consciousness at once.

Since this can easily seem completely absurd, one option open to the C-theorist is to subscribe to a more restricted position. Yes,

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

subjects are C-systems, but only when the C-system in question can produce a single stream of consciousness. C-systems which can produce multiple streams of consciousness do not constitute subjects. Imposing this restriction has a clear rationale – subjects with multiple streams of consciousness *are* bizarre – and it provides the C-theorist with an effective counter to the Bayne-Johnston multiplicity objection: just because *some* possible C-systems are not subjects it doesn't necessarily follow that *no* C-systems are subjects. There is no obstacle at all to taking well-behaved single-stream C-systems to be subjects.

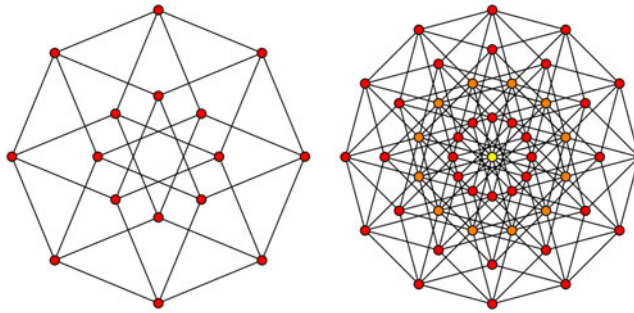
But while this response is adequate, it may not be the best. Another, more radical option is available. We are all familiar with ordinary cubes. They possess six square sides of equal size, which meet at right-angles and extend through three spatial dimensions. For most of human history, the notion that there could be *hypercubes* possessing more than six sides, and extending through more than three dimensions, would have seemed absurd. Yet thanks to the efforts of mathematicians over the past couple of centuries, we know that 4, 5, 6, 7 … *n* dimensional cubes are perfectly possible, mathematically, and the properties have been successful explored. Since our imaginations are confined to three dimensions, we cannot mentally picture such entities, but this does not prevent our recognizing that they are possible. Analogously, perhaps *hyperselves* are possible: subjects of experience possessing more than a single stream of consciousness at a time. Since we cannot our imagine more than a single consciousness at a time – doing so is incompatible with the unity of our streams of consciousness – we cannot imagine (fully) what it would like to be such a subject, but we already know that reality is not confined to what we are able to mentally picture. A single C-system with the ability to generate multiple streams would clearly qualify as a hyperself, provided of course that we continue to equate C-systems and subjects of experience.

There is, of course, a problem that is still to be overcome. As Bayne points out, it is plausible to suppose that a subject of experience is the kind of being whose consciousness is fully unified at any given time. This constraint on the mental life of subjects can be formally stated like this:

> *Subject-Unity Principle*: at any one time, the experiences of a single subject *S* are all synchronically co-conscious, and necessarily so.

A hyperself, if it were to exist, would clearly fail to conform to the Subject-Unit Principle. Although some of its experiences at any

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press



The shadows cast by 4 and 6 dimensional hypercubes onto a 2 dimensional surface.

given time are co-conscious (those existing within individual streams of consciousness) many of its experiences are very definitely not co-conscious (those found in different streams). However, deeming hyperselves to be impossible would be premature.

Why does the Subject-Unity Principle seem plausible? The answer, I take it, is that (a) under normal circumstances our own streams of consciousness are fully unified, as far as we can tell, and (b) we assume that we are not hyper-selves, and have only one stream of consciousness at a time. Anyone wishing to defend the possibility of hyperselves might be tempted to put pressure on (b). Can you really be certain that you are not in fact a hyperself? Since the multiple stream of consciousness of such a subject are experientially self-contained, the qualitative character of each stream is just as it would be if the other streams did not exist. But while this is true, there is of course no reason to believe that our brains *are* producing multiple streams at any one time, so there is no reason to take this exotic possibility very seriously. Moreover, and more importantly, anyone who believes the notion of a hyperself to be incoherent will follow Bayne and Johnston, and argue that if our brains are in fact producing multiple streams of consciousness, then these streams do not belong to a single subject – in effect, there is no subject on the scene at all. The Subject-Unity Principle demands that we interpret the situation thus.

Have we reached a stand-off? I think not, for the Subject-Unity Principle itself is not beyond question. This is because there are easily conceivable circumstances in which it fails to hold at all. In science fiction tales – not to mention philosophical thought-experiments – featuring time travel, it is by no means unusual to find tales which involve a time-travelling protagonist who meets up with

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

their earlier (or later) self. In the 2012 movie *Loopers*, the complex plot revolves around precisely this possibility. The eponymous loopers are assassins who hide their victims bodies in the past, and when young they are hired to kill (unknowingly) their older and no-longer-quite so-deadly selves: we thus find Bruce Willis spending much of the movie trying to avoid the efforts of his younger self to kill him.[26] Now, when a time traveller spends a period co-existing with an earlier or later version of themselves, we have a single person, at different stages of their life, existing twice-over at the same time. Although Young-Willis and Old-Willis are one and the same person, during those periods when they co-exist, they each have their own entirely distinct and separate streams of consciousness. When Young-Willis punches Old-Willis, the pain in the jaw felt by the latter is not felt by the former; when Old-Willis retaliates and hits back – hard – the same applies, but in reverse. We thus have a situation in which it is not the case that all the experiences of a single subject at a given time are all mutually co-conscious.

Opinions differ on the issue of whether or not time travel is really a genuine logical or nomological possibility. (Complicating the issue still further, travel through some forms of time – e.g. a block universe – may be easier than in others – e.g. a presentist universe – and there is nothing approaching agreement as to which of these competing views of the nature of time is correct.) For present purposes, however, I am not sure this matters much. The fact that the relevant time travel scenarios are so easily conceivable surely suggests that the Subject-Unity Principle cannot true, or at least, not without emendation or qualification.

As for the way in which this constraint needs qualifying, we need look no further than the literature on time travel. Suppose a time traveler spends an hour in their time machine travelling a thousand years into the past. Their trip might seem paradoxical: the traveler arrives both an hour *after* setting off, and a thousand years *before* setting off. But there is really nothing paradoxical here at all. As Lewis argued, if we are to make sense of time travel we need to recognize a distinction between two forms of time: ordinary objective world-time, and the

[26] The complexity of some of the time travel scenarios devised by philosophers easily match the complexity of those devised by most Hollywood scriptwriters. In his 'Who was Dr Who's Father?' MacBeath defends the logical possibility of a case (devised by Jonathan Harrison) which implies 'that a man can be his own father, that he can eat himself (that is, eat himself up), and that he can die before he is born.' See MacBeath 'Who Was Dr Who's Father?', *Synthese* **51** (1982), 397.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

personal or *subjective* time of the time traveler.[27] It is with respect to ordinary objective time that the traveler arrives a thousand years before setting off, and with respect to their subjective time that they arrive an hour later.

Lewis himself took subjective time to be determined by physical, psychological and causal factors: 'First come infantile states. Last come senile ones. Memories accumulate. Food digests. Wristwatch hands move. The assignment of coordinates that yields this match is the time travelers personal time'.[28] There is nothing wrong in allowing factors such as these determine someone's subjective time if these same factors also determine that person's persistence through time. Unless subjective time and personal identity are guaranteed to march in step circumstances might arise in which episodes in the person's life are not assigned any location in their subjective time at all (or the wrong one). Since Lewis believed that personal identity is determined primarily by psychological and causal factors, it was natural for him to characterize subjective time in the way he did. But if, as I have been arguing, we should trace our existence through time in terms of experiential continuity – both actual and potential – rather than physical or psychological continuities, we will need to think about subjective time in a different way. According to the C-theory, we are essentially subjects of experience, and a subject remains in existence for a given interval of time if throughout that interval it is either continually conscious or retains the capacity for consciousness. If such a subject can survive dramatic ruptures in psychological and/or physical continuity – as we saw at the outset, there are reasons for thinking this is the case – if subjective time were determined solely by psychological and/or physical factors then such ruptures would terminate the subject's subjective time, even though the subject continues to exist.

Fortunately, we needn't look very far for an alternative way of generating subjective temporal orderings. At the heart of the C-theory is a simple idea: over short periods all and only those experiential capacities which have the ability to contribute to unified states and streams of consciousness belong to the same subject. As a moment or two's introspection reveals, a typical stream of consciousness is experientially continuous: each brief phase is experienced as flowing seamlessly into the next. (Think of what it is like to hear a sound *continuing on*, or the sky *remaining resolutely blue* as you stare

[27] D. Lewis, 'The Paradoxes of Time Travel', *American Philosophical Quarterly* **13** (1976), 145–152.
[28] Ibid., 70

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

up at it.) Given this, we can say of two experiential capacities C1 and C2, that C1 is *subjectively earlier* than C2 if the experiences C1 can produce would be experienced as flowing into the experiences C2 can produce. C1 and C2 are *subjectively simultaneous* only if the experiences they can produce would be experienced together as synchronically co-conscious. In this sort of way the relationships which bind experiences into unified streams – i.e. co-consciousness, in its synchronic and diachronic forms – can be used to define a subjective temporal ordering for a collection of experiential capacities (or phases of such) distributed through ordinary objective time.[29]

According to the Subject-Unity Principle, in the form given above, at any given time the experiences of a single subject are necessarily all mutually co-conscious – in effect, a subject is confined to just one stream of consciousness at a time. However, while on the face of it the Principle is a very appealing one, it runs into trouble: there are easily conceivable time travel scenarios in which a single subject enjoys two (or more) streams of consciousness at a given time. Having acknowledged the distinction between objective and subjective time, there is a natural and appealing way of relieving this tension. Let's return to the case of the time-travelling Young-Willis and Old-Willis, who are engaged in a fist-fight with one another during one of the periods during which they coexist. At any moment during this interval the experiences of Young-Willis and Old-Willis are simultaneous with respect to ordinary objective time, but non-simultaneous in their respective subjective times. Given the way we are now thinking of subjective times, the experiences of the two Willis' would have to be co-conscious if they were to be subjectively simultaneous, but evidently they aren't.

There is a more general lesson to be drawn from this case, and it's a simple one. The problematic constraint on the relationship between subjects, time and experience should be reformulated along these lines:

> *Subject-Unity Principle*: at any given moment in the **subjective time** of a subject *S*, the experiences of *S* are all mutually synchronically co-conscious, and necessarily so.

[29] For two experiences E1 and E2 to be synchronically co-conscious it is not enough for them simply to be experienced simultaneously, i.e. be experienced at the same moment of ordinary objective time. E1 and E2 are only synchronically co-conscious if they are *experienced together*, i.e. unified in the same distinctive way as the current contents of your stream of consciousness.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

Bayne, it will be recalled, insists that there is a constitutive *a priori* connection between subjecthood and the unity of consciousness: a subject is the kind of thing whose nature is such that at any one time its consciousness is fully unified. I too find this plausible. This connection is not abandoned in the revised Subject-Unity Principle, it is simply relativized to locations in subjective time, rather than objective time. What's more, if we conceive of subjective time in the way I have suggested, we now have a fuller understanding of why this constitutive connection exists in the first place. Since subjective simultaneity is *defined* in terms of synchronic co-consciousness, it simply makes no sense to suppose that the experiences a subject is having at a given location in its personal time could be anything other than synchronically co-conscious. Experiences which aren't related in this must either belong to a different subject altogether, or to the same subject, but at different locations in its subjective time. The *a priori* link between phenomenal unity and subjecthood is thus not only retained, it is explained.

13.

This is all very well, you may by this point be thinking, but hyperselves are not time travelers, they are beings of a seemingly paradoxical sort. From the vantage point of the C-theory, a hyperself looks very much as though it constitutes a single subject. The C-theory equate selves with subjects, and subjects with materially (or immaterially) grounded capacities for experience, in the form of C-systems. Since in the case of a hyperself there is just one material (or immaterial) base – a single brain-like system, say – as we saw earlier, there is a strong case for holding that in these circumstances there is just a single C-system involved, and hence a single subject. However, there's a problem: how can a single subject possess more than one stream of consciousness at a time? Travelling through time to co-exist with one's earlier or later self is one way of achieving this, but time travel is of no assistance in the case of hyperselves, whose multiple streams of consciousness exist at the same location in their personal time as well as ordinary world-time.

While this is true, it is also not the end of the story. If a seven-streamed hyperself suffers a headache at precisely noon, the pain will exist in each of its seven streams of consciousness. Do these seven pains exist simultaneously? In one sense, yes: so far as world-time is concerned, they all occur at noon. However, from the perspective of this hyperself's subjective time the situation is altogether

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

different. Since each pain occurs within a distinct stream of consciousness, they are not synchronically co-conscious with one another, and hence are not related in the way experiences have to be related if they are to be subjectively simultaneous. What goes for these pains goes for *all* the experiences in the seven separate streams of consciousness: these experiences are simultaneous with experiences occurring in the same stream, but not simultaneous with any of the experiences occurring in the other streams.

Given these facts, a way of interpreting the envisaged state of affairs almost imposes itself: a hyperself is a single subject, but one whose experience unfolds in multiple *parallel* dimensions of subjective time. Here 'parallel' time-dimensions are those where (a) every distinct moment in each dimension occurs before or after every other moment, (b) no moments in different dimensions are temporally related in any way to any moments in any of the others. If Lewis' concrete possible worlds exist, the events in each such world unfold in a different parallel objective temporal dimension. If a hyperself were to exist, the experiences in each its streams of consciousness unfold in a different parallel subjective temporal dimension. As far as I can see, these scenarios are equally possible, metaphysically speaking.

There is an important difference, however. Lewis' worlds are entirely isolated from one another. The experiences which occur in a hyperself's various streams of consciousness at a given (objective) time are experientially isolated from one another, but nonetheless intimately related by virtue of their belonging to a common subject of experience. But we already know how *this* is possible. It is possible because the experiences are being produced by a single C-system, one possessing a single material (or immaterial) dispositional base.

There is a second point of divergence. Lewis' parallel universes each constitutes (or is located within) a different objective time-system. A hyperself's multiple parallel subjective temporal dimensions all exist within a single objective time-system. However, this does not mean that the experiences in these streams at a given objective moment all *really* exist simultaneously. Viewed from the vantage point of one temporal framework they do exist simultaneously, viewed from the perspective of another – perfectly valid – temporal framework they don't. In the time travel case, the distinction between objective and subjective times forces itself upon us because genuinely temporal relations which cannot be accommodated in objective time alone. If you are about to travel a thousand years into the past, there is a legitimate sense in which the experiences you have after you arrive (one subjective hour from now) will occur in your future, even though they are also located in the objective past. In

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press

this sort of case, we are compelled to recognize the existence of subjective time because objective time cannot satisfactorily accommodate the temporal relations which exist between the different phases of a time traveller's life. The situation is the reverse of this in the case of hyperselves. Here we are led to acknowledge subjective temporal dimensions because events that are clearly simultaneous from the objective temporal perspective are clearly *not* simultaneous from the subjective perspective. The events in question, are of course, the experiences a hyperself would enjoy, if one were to exist.[30]

University of Liverpool
bdainton@liverpool.ac.uk

[30] For comments on earlier drafts and discussion of the ideas therein, my thanks to Thomas Jacobi, Galen Strawson and Tom Winfield.

https://doi.org/10.1017/S1358246115000065 Published online by Cambridge University Press