First-Personal Continuity and the Explanatory Limits of Derivative Identity

Kevin Don | Department of Philosophy, University of Edinburgh

Introduction

Philosophical discussions of personal identity often begin by asking what makes a person at one time identical with a person at another. Analytic theories usually answer this by specifying relations that are taken to settle the issue: psychological continuity, bodily persistence, biological continuity, or the right kind of causal connection. Once the relevant relations obtain, the facts about identity are supposed to be fixed.

In this paper, I will call accounts of this kind *derivative*. I use the term in a deliberately broad and stipulative sense. A theory is derivative when facts about diachronic personal identity are taken to be wholly determined by more basic facts concerning psychological states, bodies, organisms, causal relations, or some combination of these. This does not imply that all derivative theories give the same criterion of identity. Locke, Olson and Parfit plainly do not. What they share, for present purposes, is the thought that nothing further concerning personal identity remains once the relevant underlying facts have been fixed.

There are at least two different questions that can be asked of such accounts. One concerns *numerical identity*: whether a person existing at one time is one and the same individual as a person existing at another. The second concerns the first-personal structure through which a subject anticipates its own future. When I fear tomorrow's pain or look forward to tomorrow's pleasure, the anticipated experience is presented not merely as an experience that will occur to someone appropriately related to me, but as part of my own future. These questions are closely connected, but it does not follow that answering the first automatically explains the second.

The argument of this paper is therefore more limited than an objection to derivative theories as such. I do not argue that psychological or biological accounts give the wrong verdict about numerical identity. Nor do I argue that there must be some additional immaterial entity, Cartesian ego, or metaphysical fact over and above the physical and psychological facts. My claim is instead explanatory: a criterion that settles numerical identity does not, merely by settling it, explain the first-personal character of

anticipation. A theory may correctly determine that a future person is identical with me while leaving further work to be done in explaining how that future is given as mine from within an ongoing conscious life.

Teletransportation makes this distinction especially vivid. The point of the case will not be to assume that a reconstructed individual is merely a successor rather than the original person. That would beg the question against the derivative theorist. Instead, I will grant, for the sake of argument, whatever identity verdict the relevant derivative theory supplies and ask what explanatory resources that verdict provides for understanding first-personal anticipation.

Barry Dainton's work on phenomenal continuity is important here because it shows that the continuity of consciousness has received a detailed treatment within contemporary analytic philosophy. His account also forces an important qualification on any straightforward appeal to a single continuous stream of consciousness: phenomenal continuity can involve difficult cases of division, interruption and branching. I therefore do not offer first-personal continuity as a simple alternative criterion of numerical identity. Rather, I use Dainton together with Bergson's account of duration to identify a feature of conscious persistence that may not be exhausted by the relations used to determine identity.

I begin by clarifying derivative accounts and the explanatory target of the paper. I then reconsider teletransportation without assuming that the reconstructed subject is numerically distinct from the original. I next situate the argument in relation to Dainton's account of phenomenal continuity and Parfit's cases of divided consciousness. I then use Bergson's account of duration to explain why the temporal structure of first-person experience may resist straightforward reconstruction from relations among independently specifiable states. Finally, I consider objections from Parfit and Williams and clarify the limited conclusion that follows.

1. Derivative Identity and the Explanatory Target

Locke provides the classic starting point for psychological approaches to personal identity. For Locke, personal identity follows consciousness rather than sameness of substance, with memory playing a central role in extending consciousness backwards to past experiences (Locke, 1694/1975). A familiar difficulty for a simple memory criterion is circularity. If genuine memory already requires that the remembered experience was originally undergone by the person who now remembers it, then personal identity has been presupposed in determining whether the relevant memory relation exists. Shoemaker's introduction of quasi-memory is

intended to avoid this problem by specifying a memory-like psychological relation without building personal identity into its definition (Shoemaker, 1963). Psychological continuity can then be understood through overlapping chains of appropriately caused psychological connections rather than through direct memory alone.

The phrase *appropriately caused* also requires clarification. Mere psychological resemblance is plainly insufficient. Someone might accidentally possess beliefs or apparent memories resembling mine without thereby being psychologically continuous with me. The relevant later psychological states must therefore arise from the earlier states in the appropriate causal way. Parfit similarly builds the "right kind of cause" into Relation R, his relation of psychological connectedness and continuity (Parfit, 1984, p. 215). For present purposes, the important point is not to settle every dispute about what the right causal relation consists in. It is that psychological continuity theories appeal to facts about psychological states and the causal relations between them in order to determine facts about persistence.

Biological accounts change the criterion rather than this general explanatory structure. Olson's animalism holds that we persist because we are human animals and continue for as long as the relevant organism continues (Olson, 1997). Psychological change, including substantial loss of memory or personality, need not therefore amount to the loss of the individual. Here the relevant identity facts are grounded primarily in biological rather than psychological continuity.

Williams's thought experiments reveal an important pressure on purely psychological accounts. In "The Self and the Future", the same imagined transformation can be presented either as an exchange of psychological characteristics between bodies or as a series of increasingly radical psychological interventions performed on one continuing body (Williams, 1970). The second presentation makes fear of what will happen to the body appear difficult to dislodge. Williams therefore gives us reason to resist the assumption that prudential concern straightforwardly follows psychological continuity.

Parfit moves in a different direction. He argues that personal identity itself is not what fundamentally matters in survival. What matters is psychological connectedness and continuity with the right kind of cause, or Relation R (Parfit, 1984). Strict numerical identity is the one-one relation in which something is identical only with itself. Relation R, by contrast, can in principle take branching forms in which identity cannot. This allows Parfit to claim that what matters in survival may be preserved even where numerical identity is indeterminate or absent.

These theories disagree about what makes someone persist. Yet a distinction can be drawn between three questions. First, what criterion determines whether the person at t_1 is numerically identical with the person at t_2? Second, what facts constitute or ground that identity? Third, what explains the first-personal structure through which the subject at t_1 anticipates experiences at t_2 as its own?

Derivative theories are principally designed to answer the first two questions. Nothing in this paper establishes that their answers are mistaken. The present argument concerns the move from those answers to the third question.

That distinction matters because an identity verdict and an explanation of anticipation are not obviously the same thing. Suppose a theory establishes that the person who will undergo an operation tomorrow is numerically identical with me. This tells us who will undergo the operation. It does not yet analyse the distinctive way in which I now dread that pain as something that will be experienced by me. The first-personal character of that anticipation is not merely an additional proposition concerning the future individual's identity. It is a feature of the present subject's temporal perspective.

To describe this feature as *non-derivative* is therefore not to posit an additional identity fact. The claim is weaker. First-personal continuity is non-derivative in the sense that it is itself part of the phenomenon requiring explanation rather than something that can simply be replaced, without argument, by a statement of the relations used to determine numerical identity.

This distinction also limits what the paper needs to establish. Showing that Locke, Olson or Parfit encounters a difficult case would not demonstrate that every possible derivative criterion must fail. A better criterion may be available. The argument does not depend on denying this possibility. The question is whether discovering such a criterion would, by itself, provide a complete explanation of first-personal anticipation.

2. Teletransportation and First-Personal Anticipation

Consider a familiar teletransportation case. A machine records a person's complete physical and psychological structure, destroys the original body, and reconstructs an individual elsewhere with the same memories, character, intentions and dispositions. For a psychological continuity theory, the relevant relations may be preserved to an extraordinary degree.

It would be question-begging to describe the reconstructed individual simply as a duplicate or successor and then demand an explanation of why I should regard that

individual's experiences as mine. If the psychological theorist holds that the relevant continuity relation constitutes personal identity, then the reconstructed individual is, according to the theory, not merely someone related to me. The individual is me.

The argument can grant this verdict.

Suppose, then, that the pre-transport subject at t_1 and the post-transport subject at t_2 are numerically identical. The question raised here is not whether some additional metaphysical fact must determine whether this verdict is correct. It is what explanatory work the identity criterion performs in accounting for anticipation from t_1.

Before entering the machine, the subject can know every relevant fact specified by the theory. The future subject will remember entering the machine, continue the same projects, possess the same character and stand in the appropriate causal relation to the earlier subject. If these facts settle identity, then the subject can also accept that the person appearing at the destination will be numerically identical with the person entering the machine.

There nevertheless remains a distinction between accepting an identity proposition and describing the first-personal structure of anticipation. The subject does not ordinarily experience anticipation as the detached inference that a person satisfying some identity criterion will later undergo an experience. The anticipated future is presented as the continuation of the subject's own experiential life.

This need not introduce a mysterious further fact. The claim is instead that the phenomenon has not yet been described merely by giving the conditions under which the later subject counts as identical with the earlier one. A derivative theorist can respond that the relevant psychological or biological continuity does not merely establish identity but also explains first-personal anticipation. That is a coherent response. But it requires an argument connecting those relations to the temporal structure of first-person experience. Restating the identity criterion is not yet such an argument.

Teletransportation makes this explanatory distinction unusually visible because it separates two ideas that coincide in ordinary life: preservation of psychological structure and the apparent continuation of an unfolding conscious process. In ordinary survival, memories, intentions, bodily processes and conscious experience develop through an overlapping temporal history. Teletransportation asks us to imagine the relevant structure being reproduced through an unusual causal process. A psychological theory may consistently regard this process as identity-preserving. The present argument does not deny that. The case nevertheless forces us to ask whether

reproducing or preserving the relations that constitute identity is also sufficient to explain the first-personal temporality within which a future is anticipated as one's own.

Williams's work is relevant at this point, although it does not provide the conclusion I am defending. His torture case shows how unstable our judgements can become when bodily and psychological continuities are pulled apart. Presented in one way, psychological transfer encourages us to follow memories and character; presented differently, the prospect of torture encourages concern to remain with the body (Williams, 1970). One possible conclusion is that bodily continuity rather than psychological continuity supplies the correct criterion of personal identity.

Nothing in the present argument rules that out. But a Williams-style bodily criterion would still leave a distinct explanatory question. To say that tomorrow's tortured body is numerically identical with my present body determines who is tortured. The first-personal character of my present fear remains a feature of the way my future is anticipated from within this life. Williams therefore supports the importance of prudential perspective without requiring us to adopt his apparent preference for bodily continuity.

This also clarifies the role of intuition. Teletransportation is not being used on the assumption that reluctance to enter the machine reveals metaphysical truth. Intuitions about unusual cases can be unreliable. The point is instead methodological. The case separates questions that ordinary cases allow us to run together. A theory may tell us that identity survives teletransportation. The remaining philosophical question is what, within that theory or in addition to its identity criterion, explains the first-personal structure of anticipating the post-transport experiences as one's own.

3. Dainton, Phenomenal Continuity, and Fission

The idea that continuity of consciousness deserves an independent analysis is not peculiar to phenomenological philosophy. Dainton's work develops one of the most detailed contemporary accounts of phenomenal unity and continuity. In *The Phenomenal Self*, Dainton asks how experiences occurring at different moments belong to a single stream of consciousness. His account of diachronic phenomenal unity appeals to overlapping specious presents connected through forms of diachronic co-consciousness (Dainton, 2008, pp. 51–74). On this model, the stream is not constructed simply by placing isolated experiences side by side. Successive experiential phases overlap and are internally connected in ways that generate phenomenal continuity.

Dainton's account is directly relevant because it treats phenomenal continuity as philosophically substantive rather than assuming that the unity of a conscious life is automatically explained once facts about personal identity have been specified. His broader C-theory also develops an account of subjects in terms of capacities for consciousness rather than a Cartesian substance (Dainton, 2008). This is important because it shows that taking first-personal or phenomenal continuity seriously does not require positing an immaterial ego.

The position defended here nevertheless differs from Dainton's in an important respect. I am not attempting to provide a complete ontology of subjects, nor do I claim that phenomenal continuity supplies a sufficient criterion of personal identity. Dainton offers a positive theory of the structure and persistence of conscious subjects. My argument is more limited: phenomenal continuity identifies an explanatory target that should not be assumed to disappear once a criterion of numerical identity has been given. Bergson will later provide a different explanation of why this target may resist analysis purely in terms of externally related states.

Parfit's cases of divided consciousness create an immediate difficulty for any stronger claim. In his Physics Exam case, a subject is imagined as temporarily dividing consciousness into two streams in order to pursue two calculations before the streams are reunited. Each divided stream appears continuous with the single stream that preceded the division (Parfit, 1984, pp. 245–248). Similar considerations become more dramatic in cases of permanent fission.

If phenomenal continuity were proposed as a straightforward sufficient criterion of numerical identity, this would be a serious problem. If the original stream is phenomenally continuous with both later streams, strict numerical identity cannot simply follow that continuity, since numerical identity cannot branch in the same way. The lesson should be accepted rather than resisted. First-personal continuity cannot simply replace psychological or biological continuity as another uncomplicated criterion of numerical identity.

But that concession does not undermine the weaker claim of this paper. It instead sharpens the distinction on which the paper depends. Fission shows that questions about phenomenal continuity and questions about strict numerical identity can come apart. There can be meaningful questions about the continuation, division and organisation of conscious perspective even where the language of one-one numerical identity becomes difficult to apply.

Dainton himself takes fission seriously rather than treating it as fatal to a consciousness-based account. He considers whether phenomenal continuity can branch and develops a model involving branching personal time (Dainton, 2008, pp. 364–408). I do not need to adopt his particular solution. What matters here is that the fission problem concerns the topology of conscious continuation rather than showing that phenomenal continuity is philosophically unreal or irrelevant.

Indeed, Parfit's example helps reveal why the present claim must remain explanatory rather than criterial. During the divided phase of the Physics Exam, each stream can possess its own internally unified first-person perspective. Asking what unifies the experiences within each stream is not identical to asking how many numerically distinct persons exist during the experiment. Parfit himself uses the case to challenge attempts to explain experiential unity merely by saying that the experiences belong to the same person. If one person has two mutually inaccessible streams, saying that all of the experiences are that person's experiences does not explain what makes each stream internally unified.

The direction of explanation therefore matters. Personal identity cannot simply be invoked as an all-purpose explanation of phenomenal unity when cases can be constructed in which the unity and identity questions diverge. Equally, phenomenal unity should not simply be promoted into a new criterion of numerical identity. The more modest conclusion is that the structure of conscious continuity constitutes a distinct philosophical problem.

This point also helps distinguish the present view from a simple appeal to the intuition that survival must consist in one unbroken stream of consciousness. That claim is too strong. Fission, sleep, anaesthesia and other interruptions make clear that conscious life is more complicated. What requires explanation is not uninterrupted awareness but the temporally organised first-person perspective through which experiences belong to an unfolding life.

4. Bergson, Duration, and the Structure of Conscious Persistence

Bergson's relevance to this problem does not depend upon treating his famous example of a melody as an argument for personal identity. Taken by itself, an analogy between a melody and a life would establish very little. The important part of Bergson's account lies in the distinction between time as represented and time as lived.

Bergson argues that we often represent time by treating successive moments as if they were objects arranged alongside one another in space. Time becomes homogeneous, divisible and measurable. This representation is indispensable for

practical and scientific purposes, but Bergson argues that it does not reproduce the character of duration as it is experienced (Bergson, 1910). Lived duration is qualitative. Earlier phases of experience are not simply absent points standing externally related to the present. They remain implicated in the way the present is experienced.

The melody illustrates this claim rather than establishing it. A melody is not heard as a collection of independent notes whose identities remain unaffected by their temporal position. The earlier notes are retained within the experience of what follows, so that the same physical note occurring in a different place can have a different experiential significance. What matters is not merely which elements occur but the temporally organised way in which they unfold.

The connection to first-personal continuity can now be made more precisely. If conscious experience has this temporally extended structure, then the subject's present perspective is not simply one independently existing psychological state standing in an external relation to earlier and later states. The past is already involved in the structure of the present through memory, retention, character, expectation and the accumulated organisation of experience. Anticipation similarly projects the present towards a future that is experienced as belonging to the continuation of this unfolding perspective.

Derivative accounts can represent many of the relations involved in this process. Psychological continuity theories can describe memory, intention and causal dependence. Biological theories can describe the persistence of the organism that sustains conscious life. Nothing in Bergson's view requires those descriptions to be false.

The question is whether a description of relations between temporally located states is equivalent to an explanation of the temporal unity from which those states are abstracted.

This is where Bergson's criticism of spatialisation becomes philosophically relevant rather than merely metaphorical. Analysis requires us to distinguish states, properties and relations. In doing so, we can treat a conscious life as though it consisted of a sequence of independently specifiable temporal stages linked in the correct ways. Bergson's claim is that lived duration has a different structure: its phases interpenetrate rather than merely stand next to one another. If this is correct, then an analysis conducted entirely in terms of externally specifiable relations may describe the structure produced by a continuing life while leaving the character of that continuation itself under-described.

Teletransportation brings that possibility into focus. The machine is stipulated to preserve whatever physical and psychological information the relevant derivative theory requires. The psychological theorist can therefore conclude that identity is preserved. Bergson does not give us grounds simply to deny that conclusion. What his account gives us is a reason to distinguish the reproduction or preservation of a structural profile from an explanation of the temporally extended process through which a conscious perspective unfolds.

The derivative theorist may respond that this distinction ultimately makes no metaphysical difference. Perhaps the right causal and psychological relations simply are what the continuation of the process consists in. That possibility remains open. But if this is the response, an argument is required showing why Bergson's distinction between lived temporal unity and its analytical representation does not create an explanatory remainder. It is not enough to observe that all of the states and relations used by the criterion have been preserved, because the point under dispute is whether the phenomenon of first-person temporal continuity is exhausted by such specification.

A further difficulty concerns interruptions in consciousness. If first-personal continuity consists in an uninterrupted unfolding, then dreamless sleep, general anaesthesia or coma appear to destroy the person. That conclusion is implausible, and the phrase should therefore be rejected. The relevant claim concerns *diachronic unity*, not uninterrupted occurrent consciousness.

Dainton provides one possible way of making this distinction. His later development of the C-theory treats subjects in terms of capacities for consciousness, allowing a subject potentially to persist through periods in which consciousness is not being exercised (Dainton, 2008, pp. 311–340). The present paper need not adopt that theory to make the more limited point that gaps in occurrent experience do not straightforwardly entail gaps in the existence of the subject or the destruction of every form of first-personal continuity. A consciousness-based approach therefore need not claim that we cease to exist every night.

The case of a newborn raises a related but different issue. If continuity involves the way past phases are retained within later phases, how can a conscious life begin when there is no earlier phase to retain? The answer is that persistence and origination are different questions. A process need not have existed eternally in order subsequently to possess diachronic continuity. Its initial experiential phase does not owe its existence to continuity with an earlier phase; continuity becomes relevant as later phases emerge from and incorporate what precedes them. A complete theory of persons would need a fuller account of when a conscious subject begins to exist. The present argument does

not provide one, just as it does not provide a complete criterion of personal identity. Its subject is persistence and anticipation once a first-person perspective exists.

These limiting cases therefore constrain the account without defeating it. First-personal continuity should not be understood as a permanently unbroken chain of conscious episodes. It is better understood as the temporally organised continuation of a perspective whose experiential phases are internally related in ways that are not obviously captured merely by assigning them to the same numerically identified person.

5. Parfit, Prudential Concern, and the Limits of the Claim

The strongest reductionist response is Parfit's. If all of the physical and psychological facts are fixed, Parfit denies that there is a further fact about personal identity waiting to be discovered. Relation R captures what matters in survival, and the demand for something more may reflect a mistaken picture of persons as separately existing entities (Parfit, 1984).

Nothing in the present paper establishes the existence of the further fact Parfit rejects. This point is important. To ask for an explanation of first-personal anticipation is not to assume that an additional metaphysical object must travel from one moment to another. Nor does the persistence of a first-person question prove that there must be some hidden answer of the kind reductionism denies.

Parfit can also challenge the appeal to prudential concern more directly. Perhaps our ordinary concern for our own future should not be treated as evidence about the metaphysics of identity at all. On Parfit's view, recognising that identity is not what matters can alter our practical attitudes. His famous description of escaping the "glass tunnel" presents reductionism not as failing to capture ordinary self-concern but as giving us reason to loosen its grip (Parfit, 1984, p. 281). Metaphysics may therefore correct prudential concern rather than being constrained by it.

This objection succeeds against a stronger version of my argument. We cannot infer simply from the fact that I fear my future pain differently from another person's pain that there must be a deep metaphysical fact explaining this difference. Ordinary prudential concern is not an infallible guide to ontology.

But two questions again need to be separated. One is *normative*: how much concern should I rationally have for my own future rather than the future of others? Parfit's reductionism may have revisionary consequences for that question. The other is *explanatory*: what is the structure of the first-personal anticipation that we presently

have, and how is it related to diachronic consciousness? The fact that an attitude ought to change does not eliminate the need to understand what that attitude is or how it is possible.

Parfit could answer that psychological continuity provides both explanations. Because future psychological states are appropriately connected with present ones, the future is experienced and anticipated as one's own. This is a substantive answer, and the present paper does not show that no such account can succeed. What it does show is that the answer requires more than pointing to the existence of Relation R. The connection between those relations and the first-personal temporality of anticipation must itself be explained.

A second objection concerns the scope of the argument. Even if Locke, Olson and Parfit fail to provide such an explanation, why conclude that the derivative method is itself limited? Perhaps the correct derivative account simply has not yet been formulated.

This objection should also be accepted insofar as it targets a universal claim. The examples considered here do not establish that every possible derivative account must fail. The conclusion is deliberately weaker. Existing derivative criteria should not be regarded as explanatorily complete merely because they settle the identity question. A future account might show how the relevant biological, psychological or causal facts also constitute first-personal temporal continuity. If so, it would answer the challenge raised here rather than evade it.

The disagreement is therefore not best understood as a choice between mysterious first-person metaphysics and respectable third-person analysis. Nor is it simply a contest between intuition and theory. It concerns what a successful explanation of persistence is required to explain.

For a derivative account, one central task is to identify the relations in virtue of which a person persists. The present argument adds another demand: if the account is intended to explain survival as it matters from within a conscious life, it must connect those relations to the temporally structured first-person perspective in which a future is anticipated as one's own.

Bergson gives us reason to suspect that this connection cannot simply be assumed because analytical representation tends to reconstruct temporal processes as relations among distinguishable states. Dainton gives us reason to take phenomenal continuity seriously while also showing why it cannot simply function as a replacement criterion for numerical identity. Parfit's fission cases demonstrate that phenomenal continuity

and strict identity can diverge. Williams shows that prudential anticipation does not straightforwardly follow psychological continuity.

Taken together, these considerations do not overthrow derivative accounts. They show why there is an explanatory question that such accounts must answer.

Conclusion

Dominant theories of personal identity attempt to determine the conditions under which a person persists over time. Psychological, biological and reductionist theories give importantly different answers, but each seeks to settle identity through more basic facts about persons, organisms, psychological states or causal relations.

I have argued that the success of such a criterion should be distinguished from a complete explanation of first-personal anticipation. Teletransportation makes the distinction visible, but it does not establish that the reconstructed subject is merely a successor. The argument can grant the derivative theorist's identity verdict. What remains is the question of how the relations that determine identity explain the distinctive temporal perspective from which a subject anticipates future experience as its own.

Dainton's account of phenomenal continuity demonstrates both the importance and the complexity of this question. Parfit's fission cases show why continuity of consciousness cannot simply be substituted for existing identity criteria: phenomenal continuity may branch where numerical identity cannot. Yet precisely this divergence shows that phenomenal unity and numerical identity are not straightforwardly the same explanatory problem.

Bergson helps explain why the distinction arises. If conscious life is characterised by duration, its temporal phases are not merely discrete states standing in external relations. Earlier experience is incorporated into the organisation of the present, while anticipation belongs to the movement through which a perspective projects itself towards its future. Analytical descriptions of psychological and biological relations may capture indispensable facts about this process without thereby exhausting its first-personal structure.

This conclusion is deliberately limited. I have not shown that derivative theories are false, that teletransportation necessarily destroys the original person, or that first-personal continuity constitutes a further metaphysical fact. Nor have I supplied an alternative sufficient criterion of personal identity. The claim is instead that settling numerical identity and explaining first-personal continuity are distinguishable

philosophical tasks. Derivative theories may succeed at the former. If they are also to explain survival as it is anticipated from within a conscious life, the latter remains work they must still do.

Reference List

Bergson, H. (1910) *Time and Free Will: An Essay on the Immediate Data of Consciousness*. Translated by F.L. Pogson. London: George Allen & Unwin.

Dainton, B. (2000) *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

Dainton, B. (2008) *The Phenomenal Self*. Oxford: Oxford University Press.

Locke, J. (1975) *An Essay Concerning Human Understanding*. Edited by P.H. Nidditch. Oxford: Oxford University Press. (Original work published 1694).

Olson, E.T. (1997) *The Human Animal: Personal Identity Without Psychology*. Oxford: Oxford University Press.

Parfit, D. (1984) *Reasons and Persons*. Oxford: Clarendon Press.

Shoemaker, S. (1963) *Self-Knowledge and Self-Identity*. Ithaca, NY: Cornell University Press.

Williams, B. (1970) 'The self and the future', *The Philosophical Review*, 79(2), pp. 161–180.