

The Persistence of Self

Tristan Reiniers

Faculty Mentor: Mason Cash

.....

ABSTRACT: This paper primarily addresses Barry Dainton and Tim Bayne’s article, “Consciousness as a Guide to Personal Persistence.” In this article, Dainton and Bayne reject psychological continuity in favor of phenomenal continuity as a criterion for personal persistence. They define phenomenal continuity as the kind of connection between a person’s experiences that obtains when those experiences are components of a unified stream of consciousness. I summarize Dainton and Bayne’s position and defend them in bringing attention to phenomenal continuity as an important factor in personal persistence. However, I argue that they go too far in holding that complete loss of psychological continuity is survivable. I make a distinction between an individual’s self and person, arguing that phenomenal continuity is a sufficient condition for the persistence of a person but not that person’s self — only psychological continuity can enable a self to persist through time. Since it is our selves that we should really be concerned about, we should not be willing to settle solely for continuity among our experiences. Still, while phenomenal continuity isn’t a sufficient condition for the persistence of a self-conscious entity, it is a necessary one, and so in the end I construe the concept of phenomenal continuity as an important addition to Lockean accounts of personal persistence.

..... *Republication not permitted without written consent of the author.*

In their paper, “Consciousness as a Guide to Personal Persistence,” Barry Dainton and Tim Bayne explicate a Lockean account of personal identity, in that it is framed in terms of mental states and capacities. However, it diverges from more mainstream accounts by concentrating on phenomenal continuity, as opposed to psychological continuity of the type described by Locke (241) or more recently Parfit (206), as a criterion for the persistence of personal identity over time (Dainton and Bayne 549).¹ Phenomenal continuity refers to a relationship between experiences; that is, when experiences are members of unified streams of consciousness of the kind we usually enjoy, they are related by phenomenal continuity (Dainton and Bayne 549). Dainton and Bayne view this shift of perspective as obligatory for Lockean methodology by certain thought experiments, particularly a longstanding dilemma described by Bernard Williams, discussed below (549). I seek to prove in this paper that Dainton and Bayne’s analysis is correct to the extent that phenomenal continuity is indeed an indispensable concept in accounting for the persistence of a person over time. However, their first and second theses — which, taken together, hold that phenomenal continuity is sufficient to preserve a person over time and that the loss of psychological continuity can be survived (559) — require adjustment to yield more consistently compelling and intuitive conclusions regarding certain imaginary cases. I submit that persisting, self-conscious entities (such as the average human being) should be understood as constituted by a person and a self. To introduce this idea very briefly here, I define the person as the subject who experiences the contents of a stream of consciousness; the person persists as long as the entity senses a relationship of temporal continuity between his/her successive phenomenal contents. In contrast, I define the self (roughly) as the psychology of the person. I amend Dainton and Bayne’s second thesis in light of this distinction to make it more consistent with our intuitions regarding our persistence conditions.

In “The Self and the Future,” Bernard Williams poses three imaginary cases that are only slightly dissimilar from one another. The first involves a situation that could be construed as a body-switch: two human beings (hereafter referred to as “subjects”), A and B, are hooked up to a brain-state transfer device which imprints A’s body with information from B’s brain and vice versa (161). Prior to the brain-state transfer, A and B are told that one of the two entities that results from the procedure will be given a large sum of money and the other will be tortured;

A and B must choose (prior to the procedure) which resultant subject — the A-body-subject or the B-body-subject — receives which treatment. Suppose A privately tells the experimenter that he chooses the B-body-subject to get the money and B requests the opposite circumstances. The switch is then thrown and the brain states are transferred. The experimenter, realizing he cannot fulfill both requests, arbitrarily accords with B’s wishes, giving the money to the A-body-subject and torturing the B-body-subject. Williams depicts the reactions of each ensuing subject: the B-body-subject (since he has A’s memories) complains that he did not choose this outcome and also that he chose in the way he did precisely because he did not want the unpleasant things to happen to him. Meanwhile, the A-body-person will be gratified that he received the money and will be vindicated that he chose wisely. Williams writes:

These facts make a strong case for saying that the experimenter has brought it about that B did in the outcome get what he wanted and A did not. It is therefore a strong case for saying that the B-body-person really is A, and the A-body-person really is B; and therefore for saying that the process of the experiment really is that of changing bodies. . . . This seems to show that to care about what happens to me in the future is not necessarily to care about what happens to this body. (164)

Following this line of reasoning, Williams holds that this scenario demonstrates that one should identify oneself with one’s memories (167).

The second imaginary case is similar, but designed to pull the reader’s intuitions in the opposite direction, towards the idea that “[one’s] undergoing physical pain in the future is not excluded by any psychological state [one] may be in at the time” (169) and thus towards a more bodily criterion for personal survival.² To this effect, the second case includes only A and the experimenter.³ The experimenter tells A, always using the second-person pronoun, that A will be tortured after having impressions of a past, which exactly fits the past of another currently living person, installed into his brain (167, 168). In this case (in contrast to the first, above) A becomes fearful. Why would he fear the torture of his body in this case but not the previous one? Consider the two key differences between Williams’s first and second cases. First, the

A continues to exist after the procedure, psychologically altered but still capable of experiencing pain. To this effect, Williams portrays A's inner monologue thusly:

I can at least conceive the possibility, if not the concrete reality, of going completely mad, and thinking perhaps that I am George IV or somebody; and being told that something like that was going to happen to me would have no tendency to reduce the terror of being told authoritatively that I was going to be tortured. (168)

Interpreting the procedure as a brainwashing seems very tenable and renders it at least plausible that A ought to exhibit self-concern for the one who will be tortured. Unlike the first case (in which the B-body-person came to bear A's memories), the A-body-person has no competitor (impostor?) seeking to claim the name of "A" for him/herself.

While on the subject of determining which labels to assign to entities, it is worth mentioning that the second difference between Williams's two cases is that the experimenter uses only the second-person pronoun in the second case; that is, A is told "the torture will happen to you." This choice of pronoun makes it difficult not to sympathize with A's self-concern for the resultant person who will be tortured. Williams contends that it is not obvious whether the experimenter misrepresents the situation in speaking that way.

Williams's third imaginary case further frustrates an attempt to use any sort of memory criterion to define the persistence of a person. The third "case" is composed of six different scenarios, each involving a psychological alteration being performed on A, followed by the A-body-person's torture. The psychological alterations depicted in the scenarios proceed piecemeal from the commonplace (amnesia) to the outcome described in Williams's first imaginary case (in which A's and B's mental information are reciprocally exchanged). More specifically, Williams depicts a first scenario in which A is tortured after an operation that causes total amnesia, a second scenario in which this torture is preceded by amnesia and the inducement of certain changes in A's character, a third that is the same as the second except that illusory memory beliefs are also induced in A, a fourth in which these memories are modeled after another actual person (B), a fifth in which the result of the fourth is accomplished

by putting the information into A directly from the brain of B (leaving B the same as before), and a sixth that is essentially the same as Williams's first imaginary case — A's and B's mental information are reciprocally exchanged (172). Clearly, the difference between each successive psychological alteration depicted in the above situations is only slight and incremental (perhaps even merely superficial), leading to the idea that the differences among the scenarios are differences of degree and not of kind. Williams believes that A would have straightforward reasons for fearing the pain of torture if he foreknew that his prospect would be that of situation number one (inducement of amnesia followed by torture). Williams is confident that most would agree with this belief (as am I). Given that A has reasons for fearing the pain of torture in the first situation, and given that each situation features various psychological alterations that differ only incrementally and by degrees from one another, Williams argues that A has reasons for fearing torture if he foreknew that his prospect were any of the above situations. He seeks to show that it is at least plausible that there are cases in which "one's fears can extend to future pain whatever psychological changes precede it" (180). Still, this arguably tacit endorsement of a more-or-less bodily criterion for the survival of persons is reluctant and unsure (180).⁴ So in the end, Williams's paper forces ambivalence on the reader, since his imaginary experiments are ambiguous with respect to whether A has survived (176, 180). When, upon reading any of his three imaginary cases, readers intuitively favor one interpretation (e.g., that A has survived a procedure or not), Williams implicitly makes the point that this confidence is a result of how the facts are presented (such as when his first case was neatly arranged to favor an "entities switching bodies" interpretation) and less a result of the facts themselves (Dainton and Bayne 550). As Dainton and Bayne write, "Given an appropriate narrative context, the stipulation that a brain-state transfer device shifts a person from one body to another can seem as natural and plausible as the stipulation that it merely affects a drastic form of brainwashing on a subject who remains in their original body" (552).

Williams's inducement of a lack of faith in imaginary cases when they are used as tools in philosophical argumentation amounts to a critical blow to Lockean accounts of personal identity because, as Dainton and Bayne point out, Lockeans generally rely on imaginary cases to justify their arguments (549). In what sense do Lockeans depend on imaginary cases? Lockean accounts of personal identity are characterized by the fact that

Self depends on consciousness, not on substance. Self is that conscious thinking thing,—whatever substance made up of, (whether spiritual or material, simple or compounded, it matters not) — which is sensible or conscious of pleasure and pain, capable of happiness and misery, and so is concerned for itself, as far as that consciousness extends. . . . [U]pon separation of [one’s] little finger, should [one’s] consciousness go along with the little finger, and leave the rest of the body, it is evident the little finger would be the person, the same person; and self then would have nothing to do with the rest of the body. (245)

In questions of personal persistence, Locke uses memory as the decisive factor: “Absolute oblivion separates what is thus forgotten from the person, but not from the man” (246). (By “person,” Locke means “a forensic term” which “is the name for this self” [249].)

So far as we know, consciousness and bodies are inseparable in practice. As a result, it does not matter in practice whether one determines personal persistence in terms of diachronic bodily continuity or diachronic psychological continuity. This irrelevance means that in order to justify and illustrate his points, Locke (and for that matter, modern-day philosophers sympathetic to his approach) uses imaginary cases in which the two continuities (that is, diachronic bodily and psychological continuity) diverge. For example, Locke relies on posing situations like the following:

Could we suppose two distinct incommunicable consciousnesses acting in the same body, the one constantly by day, the other by night, and, on the other side, the same consciousness, acting by intervals, two distinct bodies: I ask, in the first case, whether the day and night-man would not be two as distinct persons as Socrates and Plato? And whether, in the second case, there would not be one person in two distinct bodies, as much as one man is the same in two distinct clothings? (248)

Lockeans’ recourse to hypothetical scenarios like those above is what philosophers have in mind when they write that Lockean accounts of personal identity rely

on imaginary cases to justify their arguments; Lockeans’ reliance on this rhetorical technique is what leads Dainton and Bayne to note that effective attacks against the use of imaginary cases as tools for justifying arguments serve as effective attacks against Lockean accounts of personal identity in general.

So how do Dainton and Bayne discharge the threat and make sense of Williams’s argument? They argue that Williams’s scenarios are only able to induce contradictory responses in the reader because no mention is made of phenomenal continuity in any scenario (550). They hold that, when a distinction is made between phenomenal and psychological continuity, no ambiguity can exist (550). What is phenomenal continuity? The “base ingredients” are not dispositional states, as in psychological continuity;⁵ rather, they are phenomenal states (i.e., experiences). Phenomenal states are unified, in both the synchronic and diachronic case, by phenomenal connectedness (553). Synchronic phenomenal connectedness is manifested in the “togetherness” or “unity-within-consciousness” of our experience of the various contents of an average conscious state, such as current conscious thoughts, bodily sensations, perceptions, etc. (553, 554). Phenomenal connectedness also holds diachronically in that “each brief phase of a stream of consciousness is experienced as flowing into the next” (554).⁶ Furthermore, experiences at different times which are not phenomenally connected can be phenomenally continuous if they are linked by an overlapping chain of direct phenomenal connections (554), which introduces transitivity to the relation of phenomenal connectedness. This distinction between phenomenal connectedness and phenomenal continuity mirrors the distinction Parfit makes between psychological continuity and connectedness (Parfit 206) that effectively addresses the kind of transitivity objection Reid aims at Locke’s theory of identity (Reid 248, 249).⁷ Dainton and Bayne hold that phenomenal continuity is different from psychological continuity, however, in that psychological continuity is a causal relationship, whereas phenomenal continuity is an experiential one (549). They acknowledge some room for argument on this particular point (though they personally find implausible the type of reductive analysis necessary to make such an argument against their position [555]), but they contend that those favoring an opposing argument must at least acknowledge that phenomenal and psychological continuity are distinct since the former cannot be reductively analyzed in terms of beliefs or memories (554). Dainton and Bayne’s explication of their concept of phenomenal continuity sets up their answer to Williams: they say that the outcomes of his imaginary cases depend on whether phenomenal continuity

is preserved for A and B throughout any of the memory-swapping/brainwashing procedures that Williams portrays, then A and B follow their flow of experience (559), regardless of any changes in their psychology. As Dainton and Bayne write,

Williams's scenarios should be baffling; we are unsure what to make of them because we are left in the dark about what really matters from the point of view of one's continued existence. The readers fill in the details about phenomenal continuity for themselves, and the narrative structure of [Williams's scenarios] leads them to fill in the relevant details in different ways [for each scenario]. But . . . clarifying the fate of a subject's stream of consciousness also removes any doubt about the fate of the subject themselves: they invariably follow the flow of experience. (559)

Thus, they resolve Williams's puzzle, something which they contend that neither physical nor standard psychological accounts of personal identity can do (570).

In their paper, Dainton and Bayne make three particularly important claims that I will reproduce mostly verbatim below, in the interest of brevity, and because I will be referring to them in the subsequent argument:

The Inseparability Thesis: self and phenomenal continuity cannot come apart; all the experiences in a single (non-branching) stream of consciousness are co-personal. (557) (T1) [Phenomenal continuity, either by itself or combined with psychological continuity] is person-preserving, even across changes in brain and body. (559) (T2) Loss of psychological continuity is survivable. (559)

I take issue with T1 and T2. When they are applied to certain imaginary cases in an effort to determine whether a person has survived some procedure, they can yield very counterintuitive results that are not at all compelling.

Before posing such a case, it is necessary to go through some details of Dainton and Bayne's paper in order to defend some inferences I will draw. To this end, consider that phenomenal continuity is sustained by a stream of consciousness. Dainton and Bayne ask, "How simple can a stream of consciousness get before it leaves its owner

behind? How primitive can the contents of a stream become before it ceases to support one of us?" (560). One can imagine a situation — meditative trance, for example — in which one's stream of consciousness excludes all the sensations with which we are familiar except an awareness of one's breathing. Extreme cases like this suggest that the contents of a stream of consciousness must meet some minimum level of sophistication in order to constitute a stream of consciousness in the first place. Dainton and Bayne suggest placing this minimum at a low threshold, proposing that we could survive on a very rudimentary consciousness, perhaps like that which people have in the womb, with no cognitive sophistication at all (560, 561). (While Dainton and Bayne acknowledge some room for debate here regarding whether we experience rudimentary consciousness of this sort prior to birth, I agree that one's phenomenal continuity, at least, is sustained by such a simple kind of consciousness in the womb.⁸ As they say, why shouldn't a person persist as long as his/her stream of consciousness flows on, regardless of how primitive its character becomes? [561])⁹ I hold that at least this minimally sophisticated type of content must exist in all persons' streams of consciousness. This line of reasoning follows from Dainton and Bayne's above conclusions: the one thing all people have in common is their personhood, and the minimally sufficient criterion for the persistence of personhood is the continuing presence of the type of minimally sophisticated content described above. In other words, to qualify as a person in the first place, one must experience a stream of consciousness that is at least at the minimum level of sophistication. Any more sophisticated type of content enjoyed by a person's stream of consciousness must exist in addition to this minimum type of consciousness. This notion — that all people's streams of consciousness have some minimally sophisticated content in common — is in keeping with any sensible notion of what an increase in the sophistication of consciousness must entail; for example, if a basic bodily feeling (say, awareness of the passage of time) is present at the minimal level of sophistication of conscious experience, how is one to conceive that this faculty could be absent at a more sophisticated level of conscious experience? Whatever people can experience at consciousness's lowest threshold of primitiveness must be indispensable to personhood and therefore exist (at least tacitly) in all people.

Now that it is understood that all people's streams of consciousness share a certain content in common, I will pose a thought experiment somewhat similar to one of Dainton and Bayne's, in which a person's phenomenal and psychological continuities come apart, via the use of a

streamal diverter (which diverts a person's stream of consciousness instantaneously from one brain to another, but does not transfer any of his/her psychology [556]). This experiment follows: I -- Call me Pat -- will have my stream of consciousness instantly diverted such that it becomes sustained by another's body. (In other words, I will come to experience events as they are rendered by this other person's nervous system.) The other person's stream of consciousness will not undergo any such transfer, and so will be annihilated during the procedure. Remember, my psychology will not make the transfer since a brain-state transfer device is not being used in tandem with the streamal diverter; as a result, none of my memories, beliefs, desires, and other psychological states will survive. The person whose brain my stream of consciousness will be diverted into has no cognitive sophistication at all, and his/her brain can support only a stream of consciousness of the minimum sophistication required for him/her to be considered a person in the first place. This other person can be imagined to be a fetus in the womb if the reader subscribes to the notion, as Dainton and Bayne seem to (561), that such a creature is in fact a person; if not, the reader can substitute whatever he/she considers to be the most rudimentary form of person imaginable.

What would actually happen in such a procedure? According to Dainton and Bayne, I would survive, since their first and second theses (T1 and T2) assert that phenomenal continuity is all that is required to preserve a person and the loss of psychological continuity can be survived (559). But is this really a sufficiently accurate description of what occurs in the procedure? I do not think so. Consider what actually happens — everything that has ever characterized my existence as a unique human (my psychology) is annihilated. So what is left? Only that minimally sophisticated stream of consciousness, that bare aspect of personhood which is common to all people everywhere. To say, because of this minimum degree of personal persistence, that “I” survived the procedure is to equate my self entirely and solely with the single, bare quality of personhood itself, and nothing else — no further unique characteristics. Following the procedure, I am only that which all people sustain in themselves — the bare capacity for awareness. I find it extremely counterintuitive to describe me simply as “surviving” in this transformation. Such an account of the events is not descriptive enough. This procedure kills the “Pat” in me — every quality I had ascribed to myself as Pat ceases to characterize me anymore. My memories as Pat are irretrievably lost. If the entity into which my stream of consciousness is diverted is a fetus,

this fetus will eventually grow more cognitive capacity and the contents of my stream of consciousness will be characterized by increasing sophistication according to the fetus's development of memories, personality characteristics, and other dispositional states. But this new psychology will be completely different from that which I had previously possessed. In light of this complete psychological reformatting, I would see no reason to identify my current self with this ensuing entity; in fact, if I were told that the ensuing individual were going to be, say, tortured, I would not be any more concerned than if I were told the same of any random stranger.¹⁰

It seems to me that Dainton and Bayne, in order to allow for a more thorough, accurate, and intuitive account of the above thought experiment, should amend their argument slightly. I submit that persisting, self-conscious entities (such as the average human being) should be understood as necessarily composed of two parts, a self and a person. I define the person as the subject who experiences the contents of a stream of consciousness. (The person persists as long as phenomenal continuity obtains in his/her case.) I define the self as all the cognitive processes that are potentially realizable (in the brain) and that stand consistently ready to be triggered in specific types of situations so as to impinge upon the person's stream of consciousness in (more or less) dependable, regular ways. (However, I intend this definition to exclude potentially realizable cognitive processes that are standing available to impinge upon consciousness solely because of the spatiotemporal orientation of the person's body with respect to its surroundings. For example, if you stare at a painting in a museum, the cognitive process by which the painting can be depended on to remain in the visual field should not count as a constituent of your self, because such a process is available to impinge upon consciousness in a dependable way solely because of the orientation of your body in its environment. In other words, potentially realizable cognitive processes that are standing available to impinge upon consciousness solely as a result of the immediate operation of the physiological methods of perception, such as sight, hearing, taste, smell, touch, equilibrioception, thermoception, proprioception, nociception, etc., are not constituents of the self.)

An example of such a cognitive process that stands ready to be triggered in the right sort of situations can be illustrated by the following scenario. Consider a

pedestrian and an experienced skateboarder traversing a particular stretch of sidewalk riddled with cracks. Looking down, both are confronted with the same visual field. However, the skateboarder (perhaps unconsciously) distinguishes more features in the scene than the pedestrian does since the skateboarder must quickly classify the various cracks according to which are passable and which are impassable with respect to the diameter of the skateboard's wheels. This discriminatory ability and its tendency to be engaged (consciously or not) in the right sorts of situations (such as when the skateboarder is traversing uneven surfaces) should count as a constituent of the skateboarder's self. Another example of a constituent of a self might be a particular memory, insofar as the memory is rendered by a cognitive process and can be depended on to be accessible in the right sorts of situations (such as when a person consciously tries to recall it). Personality traits (if they are defined as dispositions to act a particular way in a given situation) are a third example, since they seem reducible, again, to potentially realizable cognitive processes that stand consistently ready to be triggered in specific types of situations.

All these constituents of the self might be considered objective in the sense that they exist (for the most part) whether or not the person acknowledges them. For example, one can only be so successful in consciously repressing memories — a particular memory can remain available to present itself in conscious awareness whether or not one chooses to recall it. By virtue of their objective existence, these constituents of the self might be contrasted with the person's narrative self (as defined by Dennett) which is subjective in the sense that the narrative self is the result of an act of free interpretation of the objective self and is amenable to almost complete revision at any time. In the account being developed here, the person weaves together all the constituents of his/her objective self (memories, tendencies, capacities, etc.) into a subjective whole, the narrative self, in order to construct a life story which makes sense of his/her existence for him/herself and others.

The continuous availability of all of these (objectively existent) cognitive processes for contributing to a self-narrative and impinging upon consciousness (and the fact that a person can depend on their activation in critical circumstances, such that the skateboarder can approach sidewalk cracks with confidence) afford the person a sense of continuity above and beyond mere

phenomenal continuity. It seems to me undeniable that the preservation of this further continuity — I'll label it "psychological continuity" because it seems analogous to that familiar term — is a large part of what self-conscious entities (such as the average human being) have in mind when they express their desire to persist over time.

Having concluded my redefinitions of "self" and "person," I submit that survival admits of degrees — full survival requires the preservation of both the person and the constituents of the self, but preservation of only the person constitutes partial survival (as in the above thought experiment with the fetus) and is preferable to death. (I think that most human beings would prefer amnesia to death, for example.)

This distinction between person and self increases the descriptive power of Dainton and Bayne's phenomenalist account while maintaining all of the progress their position has already made in disambiguating the aforementioned threatening imaginary cases. For instance, my definitions, like Dainton and Bayne's, allow for the argument that Williams's scenarios induce contradictory responses in the reader only because no mention is made of whether phenomenal continuity is preserved in any of the cases. In fact, it seems that my account disambiguates the outcome of each scenario more than Dainton and Bayne's does. For instance, if we were to assume in Williams's first scenario that the subjects were conscious throughout the brain-state transfer, then, by my account, each subject undergoes a partial death: the persons would persist through the brain-state transfer but each person's old self would "die" as a new one replaced it. Such an account does justice to the fear many people would likely experience at the prospect of undergoing such a procedure. In contrast, Dainton and Bayne's account would simply stipulate that each subject survives because their person survives, even though their self does not.

Despite such differences, my distinction does not radically alter Dainton and Bayne's conclusions. As described below, most of their theses remain true following my redefinition of terms, although some are made to mean something slightly different. Still, where their meanings are altered it seems that Dainton and Bayne's theses are brought into line with a more reasonable concept of what it means to survive.

For example, the first clause of the Inseparability Thesis would remain true according to my interpretation of terms: self and phenomenal continuity still cannot come apart, for if one's stream of consciousness ceased to flow, then his/her person would die, since the minimally sufficient condition that must be maintained for a person's persistence (as reasoned above) is that he/she be a conscious being.¹¹ It follows that in the absence of his/her person, his/her self cannot exist, since any self is defined in terms of the availability of streamal content (that is, content of a stream of consciousness) to a person. If the person ceases to exist, the constituents of the self are unavailable to anyone, and therefore cease to constitute a self.¹²

The second clause of the Inseparability Thesis is similarly unthreatened: The notion that all the experiences in a single (non-branching) stream of consciousness are co-personal is entirely consistent with my definition of a person.

T1 is also left intact. Phenomenal continuity holds whenever a given temporal "snapshot" of the contents of a stream of consciousness is experienced as flowing into the next temporal snapshot, and so on. It must therefore be true that a stream of consciousness must exist in any instance in which phenomenal continuity holds. The continued existence of a stream of consciousness obviously necessitates a subject of consciousness, which by my definition is a person. Therefore, it remains true that phenomenal continuity, either by itself or combined with psychological continuity, is person-preserving. And if phenomenal continuity can be maintained across changes of brain and body, then so can the existence of a person as I have defined it.

It turns out that the foremost effect that my distinction between selves and persons has on Dainton and Bayne's argument is that it requires that T2 be clarified. T2 reads, "Loss of psychological continuity is survivable" (559). Recall that I define a persisting, self-conscious entity as composed of a self and a person, in which the self is the person's psychology, so a loss of psychological continuity implies the death of the self and is therefore not compatible with the notion of full survival. With this definition in mind, I propose that T2 be amended to "loss of psychological continuity is survivable by one's person, but not one's self."¹³ This amended version of T2 echoes the point made above, that persistence admits of degrees. Full survival requires the preservation of both

the person and self, but preservation of only the person constitutes partial survival.

My amended readings of Dainton and Bayne's theses provide a much more descriptive and intuitive account of the imaginary fetal "streamal diverter" case I depicted above. According to the amended framework of analysis, that experiment results in the following events coming to pass: the fetus's person dies, taking the fetus's self with it; my self dies when my old psychology becomes unavailable to any person, but my person survives in the body of the fetus and eventually sustains a new self as the fetus matures and accumulates experiences. Is this not a more descriptive, realistic account than that which Dainton and Bayne would have to offer, saying simply that the fetus dies and I live?

In closing, Dainton and Bayne, in explicating the concept of phenomenal continuity, have provided a tool which seems indispensable in constructing a truly robust Lockean theory of personal identity. They are correct in saying that adjusting the primacy of psychological continuity downward in favor of phenomenal continuity when considering questions of personal persistence is obligatory for Lockeans — imaginary cases such as Bernard Williams's scenarios are otherwise intractable.¹⁴ Nevertheless, they are too quick to dismiss the importance of psychological continuity. After all, most people care strongly about what happens to them. What exactly is it that a person cares about when he/she cares about him/herself? I (and probably most others) understand the concept of "self" to refer to those qualities that make a person distinct from others. For example, what does it mean to be, say, an individual by the name of Pat? It means remembering a litany of events from Pat's perspective, having Pat's set of attitudes and opinions regarding a range of topics; it means having Pat's particular set of goals and intentions; it means caring intimately about what happens to Pat. The fact that one's attachment to one's self consists largely in a desire to maintain the characteristics that make one unique is why phenomenal continuity is an insufficient (albeit necessary) condition for the survival of a self; because, in its most rudimentary form (as in the above fetal example) the maintenance merely of phenomenal continuity does not guarantee that what makes a person distinct from others — the person's psychology — is kept intact. Surely, we must identify our selves as more than merely minimally sophisticated streams of consciousness devoid of psychological content if we are to construct a theory of the persistence of beings over time that does justice to

our intuitions in every situation, however hypothetical or physically impossible such situations may be.

NOTES

¹ While I write that Locke accounts for personal persistence in terms of psychological continuity, he might more technically be said to endorse a psychological connectedness criterion for the persistence of personal identity over time since his notion of the identity relation does not allow for transitivity (as Reid points out [248, 249]).

² When Williams writes that “[one’s] undergoing physical pain in the future is not excluded by any psychological state [one] may be in at the time” (169), he notes the exception of certain psychological states, such as unconsciousness, which preclude the experience of pain (169).

³ Actually, Williams does not frame the second case in terms of A and the experimenter; instead, he uses the first-person perspective. For example, he writes “Someone in whose power I am tells me that I am going to be tortured tomorrow” (167). Since this is the only one of Williams’s cases thus depicted, I describe it in the third-person for the sake of comparison.

⁴ That is, Williams waters down his conclusion by preceding it with the word “perhaps” (180) and acknowledging that he risks being incorrect.

⁵ Components of one’s psychology are dispositional states in that they manifest themselves as tendencies to believe, remember, opine, etc., in certain ways. To hold a given belief, for example, is a dispositional state — it consists in being disposed to act a certain way in response to certain stimuli. If I believe it is wrong to litter, this belief manifests itself as a disposition to harangue those who litter (and probably refrain from littering myself).

⁶ The amount of time over which diachronic phenomenal connectedness can obtain is very short, lasting perhaps about one second (Dainton and Bayne 554).

⁷ Reid takes issue with Locke’s notion that “as far as [one’s] consciousness can be extended backwards to any past action or thought, so far reaches the identity of that person” (Locke 241). Reid points out that this doctrine implies a contradiction. Assume some person, A, currently remembers the time when he/she performed a certain action, B; furthermore, at the time when he/she did B, he/she remembered performing an even earlier action, C, but now A’s memory of C is lost. Logically (if identity is a transitive relation), A should

be the same as the one who did C, but Locke would have to hold that this is not so, since A’s consciousness does not currently extend so far backward.

⁸ But not one’s “self” — I will develop this particular point below.

⁹ To avoid confusion: Neither Dainton and Bayne nor I hold that phenomenal continuity can be maintained over periods of unconsciousness such as dreamless sleep, anesthesia, etc. They acknowledge that this creates a problem for their theory (do we die every time we fall asleep?); they label it the “bridge problem” and offer some solutions at the end of their piece (562).

¹⁰ I do not mean to imply that the degree of my concern corresponds with the degree of my moral indignation. I would submit that one can feel an equal amount of moral indignation given two distinct situations and nevertheless be more concerned about one than the other, insofar as one commands more attention than the other (for example, by being torture that happens to me).

¹¹ To clarify, when I say “the first clause of the Inseparability Thesis” I mean its first grammatical clause.

¹² It might be objected that in certain situations where a human being is unconscious (severing phenomenal continuity), it is often the case that the constituents of the self are preserved (in the configuration of the brain) until the human being wakes up. This preservation of one’s psychological elements through periods of unconsciousness could be taken to demonstrate that my amended reading of the Inseparability Thesis is false, that phenomenal continuity and self are in fact separable. My only disagreement with this objection to my interpretation of the amended Inseparability Thesis would be somewhat pettily terminological. It is true that the self could be said to lie dormant during periods of unconsciousness, like dreamless sleep, and then make itself available to the person immediately upon waking. However, it is still the case that at any given moment in which unconsciousness obtains, the self is not available to any conscious subject. Since I define the self in terms of such availability, it seems inappropriate to say that there is a self during periods in which unconsciousness obtains. (Such a statement would amount to saying that the constituents of the self are available to consciousness during unconsciousness.) Instead, it would be more clear to describe a situation of dreamless sleep as one in which the physical system responsible for rendering the self retains its functionality, although the self is not present because there is no person to which it can be made available.

¹³ The alternatives (that it is survivable by both person and self, or by self alone) are disallowed since they are not consistent with the way I defined my terms.

¹⁴ In my terms, perhaps I ought to say that downplaying psychological continuity in favor of phenomenal continuity is obligatory when considering the question of “person-self persistence” — not just personal persistence — in order to capture a whole individual in the language I use (as opposed to just his/her person).

WORKS CITED

Dainton, Barry, and Tim Bayne. “Consciousness as a Guide to Personal Persistence.” *Australasian Journal of Philosophy* 83.4 (2005): 549 – 71.

Dennett, Daniel. “Why Everyone Is a Novelist.” *The Place of Mind*. Ed. Brian Cooney. Stamford, CT: Wadsworth/Thomson Learning, 2000. 466-74.

Locke, John. *An Essay Concerning Human Understanding*. 1690: 241 – 49. NetLibrary, Inc. 1990s. 16 July 2008 <<http://www.netlibrary.com/Details.aspx>>.

Parfit, Derek. *Reasons and Persons*. New York: Oxford University Press, 1984.

Reid, Thomas. *Essays on the Intellectual Powers of Man*. Ed. James Walker. New York: J.C. Derby, 1855: 248 – 49. Google Books. 16 July 2008 <<http://www.google.com/books?id=PxAQAAAAYAAJ>>.

Williams, Bernard. “The Self and the Future.” *Philosophical Review* 79.2 (1970): 161 – 80.