

Forthcoming in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*. Continuum Press.

How Many of Us are There?*

Hannah Tierney, Chris Howard, Victor Kumar, Trevor Kvaran, & Shaun Nichols

University of Arizona

In trying to chart the contours of our folk conceptions, philosophy often proceeds with an assumption of monism. One attempts to provide a single account of the notion of free will, reference, or the self. The assumption of monism provides an important constraint for theory building. And it is a sensible starting assumption. Monistic views often provide the simplest explanations with the fewest commitments. However, it's possible that for some philosophically interesting notions, people operate with multiple different notions. Experimental philosophy provides new tools for exploring such folk pluralism.

We will argue that in the case of personal identity, monism does not capture folk commitments concerning personal identity. Many of our identity-related practical concerns seem to be grounded in distinct views of what is involved in personal identity. Furthermore, both empirical evidence and philosophical thought experiments indicate that judgments about personal identity are regimented by two (or more) different criteria.¹ Of course, just because people are pluralists about personal identity doesn't mean that this is a sustainable position. In the second half of the paper, we will consider reasons for thinking that the folk commitment to pluralism should be rejected or overhauled. We will offer a tentative case in favor of a pluralist philosophical view about personal identity.

1. Folk Pluralism about Personal Identity

Williams provides one of the most famous sets of thought experiments in contemporary philosophy (1970). He first gives a case involving two persons, A and B. A is about to have all of his psychological characteristics transferred into B's body, and simultaneously, B's psychological characteristics will be transferred to A's body. Before the transfers, A and B are told that one of the resulting persons will be tortured. Now, Williams asks, before the transfer, should A want the torture to be administered to his current body or to B's current body? Suppose A asks for the torture to be administered to the A-body and B asks for the torture to be administered to the B-body. Given the transfer of psychological characteristics, the person in the B-body will remember A's request. After running through various possible requests and outcomes, Williams writes, "all the results suggest that the only rational thing to do, confronted

*We'd like to thank Mike Bruno, Wesley Buckwalter, Joshua Knobe, Carolina Sartorio, and Justin Sytsma for comments and discussion on an earlier draft of this paper. We'd also like to thank David Shoemaker for extensive discussion about the paper.

¹ In this essay, our point is to defend pluralism about personal identity. To that end, we focus on two different criteria: psychological and biological. But it could very well be the case that a dual-criteria pluralism isn't plural enough. There might be additional criteria for personal identity that guide judgment in some cases.

with such an experiment, would be to identify oneself with one's memories, and so forth, and not with one's body" (1970, 167). Williams thus suggests that this thought experiment indicates that what matters for personal identity is one's *psychological characteristics*. Call this the 'trait view' of personal identity.

Next, Williams has us consider the same case described slightly differently. Someone tells me that I am going to be tortured tomorrow, but before the torture, all of my distinctive psychological traits will be removed, and false memories will be inserted. After all that, the torture will be administered. What reaction should I have to the torture?

Fear, surely, would ... be the proper reaction: and not because one did not know what was going to happen, but because in one vital respect one did know what was going to happen – torture, which one can indeed expect to happen to oneself, and to be preceded by certain mental derangements as well. (1970, 168)

On Williams' view, in this case, despite the annihilation of one's psychological characteristics, it seems like one will still be present to feel the pain. Williams maintains that our intuitions about this case conform to a *biological* approach to personal identity. Although the distinctive psychological characteristics are eradicated, the biological organism persists.²

Williams' two cases seem to support starkly different views of personal identity. Sider notes that one natural reaction to these cases is to think that people operate with two notions of persistence for persons:

It appears that we are capable of having either of two intuitions about the case.... A natural explanation is that ordinary thought contains two concepts of persisting persons, each responsible for a separate set of intuitions, neither of which is *our* canonical conception to the exclusion of the other... (Sider 2001, 198)

In the following section, we will argue that empirical research indicates that people really do have two different ways of thinking about personal identity, one in terms of psychological traits and another that conforms more closely to a biological criterion.³

² In keeping with contemporary discussions, we will consider the biological continuity theory as the alternative to the psychological trait theory (Olson 1997; Sider 2001; Shoemaker forthcoming). It's not obvious that the lay intuitions that run against psychological trait theories can be neatly captured by the biological theory. After all, most people believe in survival after death, and that obviously can't be accommodated by the biological theory. It's not clear that beliefs about post-mortem survival neatly conform to the trait theory either, though. In some contexts, intuitions about post-mortem survival might conform best to the soul theory of personal identity (e.g. Reid 1785). We will have to set those complex issues aside for present purposes.

³ While we talk about conditions of personal identity in this paper, most everything we say can be reframed in terms of conditions for 'survival' (see e.g. Parfit 1971, 8; Rovane 1990, 356). Parfit (1971, 1984) argues that fission destroys identity, but does not destroy what matters for survival, and goes on to suggest that what is important in survival is *not* identity but rather *connectedness* (cf. Tierney in prep). Parfit defines psychological connectedness, as "the holding of particular direct psychological connections" (1984, 206; 1971). And this notion is then used to define *psychological continuity*: "the holding of overlapping chains of strong connectedness." As Parfit notes, connectedness is not an identity relation since it isn't transitive (1984, 206ff.) Two individuals are psychologically connected if they share direct psychological relations, and this relation can obviously come in degrees (see also Shoemaker 1996, 320). In contrast, psychological continuity is an all-or-nothing relation that *is* transitive; it holds between two

1.1. Empirical case for pluralism about personal identity

Survey studies

In survey studies, people deploy a trait conception of self in some contexts but not others. For instance, Rips and colleagues presented participants with a scenario in which Jack's brain was placed in a robot, and they varied whether the brain retained Jack's memories. Participants tended to judge that it was still *Jack* only when the brain retained Jack's memories (Rips et al. 2007). In another study, participants were asked directly to list what they regarded as essential for continued existence. Participants were told:

One problem that philosophers wonder about is what makes a person the *same person* from one time to another. For instance, what is required for some person in the future to be the *same person* as you? What do you think is required for that?

Over 70% of the participants explicitly mentioned psychological factors like memory or personality traits as necessary for persistence. All participants were then pressed with a specific question about the necessity of memory: "In order for some person in the future to be *you*, that person doesn't need to have any of your memories". Over 80% disagreed with the claim that they could be identical to someone in the future who didn't have any of their memories (Nichols & Bruno 2011).

Although some surveys trigger responses that fit with a trait conception of self, other studies point to something different. In an adaptation of the key Williams case, participants were told to imagine that they were about to undergo a surgery that would permanently eliminate their distinctive mental states (including thoughts, memories, and personality traits). Next they were told that after the surgery the doctors have to administer a series of painful shots. They were then asked to indicate agreement or disagreement with the sentence "*you* will feel the pain". Participants overwhelmingly agreed with this statement, suggesting that they are thinking of the self in a way that persists despite the complete disruption of the trait-self (Nichols & Bruno 2011).

Manipulation studies

The previous studies elicit people's opinions on simple surveys, and these studies indicate that different kinds of questions elicit responses that conform to different ways of thinking about the persistence of self. A different experimental technique starts by manipulating how people think about the self and then explores how this affects judgment and decision-making. Prompted by Parfit's work on personal identity (Parfit 1971; 1984), Dan Bartels and colleagues have manipulated how participants think about *psychological connectedness*, which Parfit defines as "as 'the holding of particular direct psychological connections'" (1984, 206). In several studies, Bartels and Urminsky (2011) manipulated how people think about psychological connectedness with prompts like the following (high-connectedness in brackets):

individuals if chains of strong psychological connectedness obtain between them (Parfit 1984, 207). According to Parfit, "Of these two general relations, connectedness is more important both in theory and in practice," (1984, 206). Indeed, some of the studies we discuss below focus on the relationship between beliefs about connectedness and judgments about identity-related practical concerns. However, when we discuss the psychological factor in judgments about identity, we intend to be neutral about whether the relevant factor is connectedness or continuity.

Day-to-day life events change appreciably after college graduation, but what changes the most [*least*] between graduation and life after college is the person's core identity. . . . The characteristics that make you the person you are . . . are likely to change radically around the time of graduation [*are established early in life and fixed by the end of adolescence*]. . . Several studies conducted with young adults before and after college graduation have found large fluctuations in these important characteristics [*have shown that the traits that make up your personal identity remain remarkably stable*]. (Bartels & Urminsky 2011, 185)

The first point to make is that these manipulations do affect people's judgments about connectedness of self. When asked to rate on a scale from 0 to 100 (0= "I will be completely different in the future"; 100= "I will be exactly the same in the future"), people in the low-connectedness condition give significantly lower ratings than baseline, and people in the high-connectedness condition give significantly higher ratings than baseline (Bartels & Urminsky 2011, Bartels & Rips 2010, Bartels, Kvaran, & Nichols forthcoming).

So, the manipulation works to affect people's immediate judgments about the connectedness of the self. Bartels and Urminsky (2011) assigned participants either to the high-connectedness or low-connectedness condition, and then presented them with a temporal discounting task; in this task, participants had to decide whether to forego a smaller reward now in favor of a larger reward later. Bartels & Urminsky found that the manipulation affected how 'patient' people were, i.e., how willing they were to wait for the larger reward. Those who were given the low connectedness manipulation were more impatient – more likely to take the smaller reward sooner rather than wait longer for the larger reward, as compared to those who read the high connectedness manipulation.⁴ Subsequent research showed that people given the low-connectedness manipulation are also more charitable with future funds than those given the high-connectedness manipulation (Bartels, Kvaran & Nichols under review). Thus, leading people to think that the self changes a lot seems to make them less concerned with the interests of the future self, and relatively more concerned with the interests of others.

We extend this work by exploring whether manipulating people's beliefs about psychological connectedness also affects their judgments about punishment. Half of the participants were in the low-connectedness condition and half were in the high-connectedness condition. Each of these groups was also split into past (1-year) and present (1-week) conditions (yielding a 2x2 design). Following the connectedness manipulation, participants were instructed to: "Imagine that you cheated on an exam a week [*year*] ago". Participants were then asked if their cheating were discovered now, how much punishment would be appropriate (ranging from none at all [1] to the maximum punishment allowed by the school [6]). We expected that if people are led to believe that there is low psychological connectedness, this shouldn't affect their judgments about the allotment of punishment for a very recent offense. Thus, we predicted no difference between high and low connectedness for participants asked about an offense committed a week ago. But for those asked about an offense committed a year ago, we predicted

⁴ Bartels & Urminsky also got this effect when they simply measured beliefs about connectedness (study 5). They had participants indicate how connected they believed they would be with their future self. Three weeks later, they engaged in a temporal discounting task (without being told that this was related to the activity three weeks prior). Bartels & Urminsky found that people who regarded themselves as less connected in part 1 were more impatient on the task in part 2.

that those who thought there was low psychological connectedness would regard themselves as deserving less punishment. That is exactly what we found (See chart 1).⁵ Intuitively, the idea is that if I think that I'm psychologically very different from my past self, then I will regard my current self as less deserving of punishment.

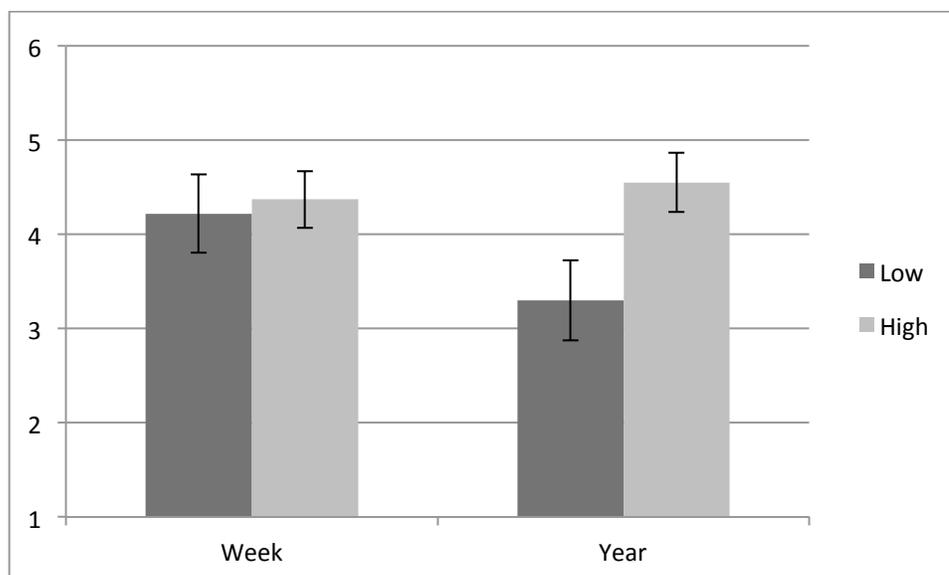


Chart 1: Connectedness and deserving punishment
(bars represent two standard errors of the mean)

Thus far, we've seen that manipulating beliefs about connectedness affects practical concerns about economic decisions and punishment. Next, we wanted to explore a Williams-style case. Recall that in Williams' key thought experiment, he asks whether you would fear pain that was scheduled to occur after massive psychological changes. Philosophers have tended to follow Williams in saying that one would fear the pain. On this basis, we predicted that anxiety about future pain would *not* be affected by the connectedness manipulation. That is, leading people to think that the self actually changes a lot won't affect their anxiety about future pain. As before, participants were either presented with the high or low connectedness manipulation and then given the following instructions:

Imagine that you have to get a root canal in 1 week [1 year]. To what extent would you be anxious about the root canal right now?

While anxiety was greatly affected by whether the pain would occur in a week or a year, connectedness had no significant effect on reported anxiety (see chart 2).⁶

⁵ Consistent with expectations, the analysis revealed main effects of timing (week vs. year) and connectedness ($F(1,181) = 14.83$ and 4.01 , $ps < .05$) and most importantly, the predicted interaction ($F(1,181)=9.17$, $p<.01$). $N=182$, 57 female. All participants were recruited from Amazon mechanical turk.

⁶ There was a main effect of timing [week vs. year] ($F(1, 131)=17.1$, $p<.001$), but was no effect of connectedness ($F(1,131) = .134$, $p=.714$) and no interaction ($F(1,131)=.266$, $p=.607$). $N=132$, 61 female. All participants were recruited from Amazon mechanical turk.

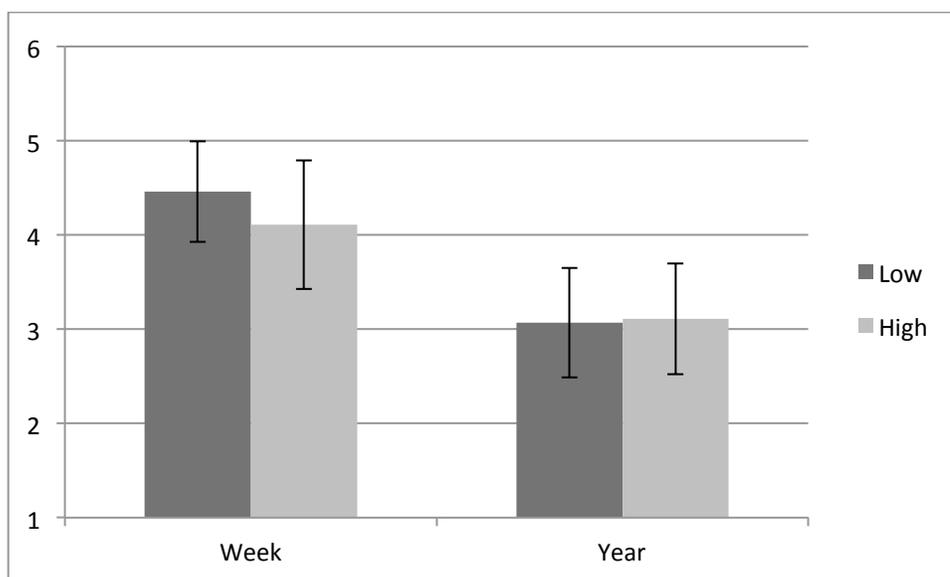


Chart 2: Connectedness and anxiety about future pain
(bars represent two standard errors of the mean)

These manipulation studies use a very different technique than the survey studies, but once again indicate that people operate with two different conceptions of self. When people are led to think that their future self won't be very connected with their current self, this affects their economic decisions and punishment judgments. But when it comes to contemplating future pain, manipulating beliefs about connectedness seems to have no effect. Changing the way people think about the self *qua* collection of traits does not impact their feelings about future pain. This, of course, conforms with the results of the simple survey involving a Williams-style case (Nichols & Bruno 2011).

Senses of Self

Survey and manipulation studies both indicate that people's judgments conform to different criteria for personal identity in different contexts. Recent work on memory provides a partial explanation for this. There are apparently two very different senses of self implicated in long term memory. One of these senses of self is stored in *semantic memory* and presents the self as a collection of psychological traits. A series of studies with brain-damaged patients shows that this trait sense of self is preserved across massive psychological damage, including profound retrograde and anterograde amnesia (e.g. Klein et al. 2004; Tulving 2003). Even patients who can't remember a single episode from their past still have knowledge of their personality characteristics. Given that they have knowledge of their traits despite a complete absence of episodic memory, these cases suggest that semantic memory holds the trait conception of the self.

While semantic memory stores a representation of self as a collection of traits, *episodic memory* delivers a sense of self that is not tied to a trait set (Nichols forthcoming). When I recall an event from my childhood, the representation of the self – that this event happened to *me* at age 12 – is not disrupted by the fact that the child's traits differ radically from my current traits. This

holds even for people with significant brain damage. Despite profound neurological dysfunction, when brain damaged patients recall experiences from their childhood, they give every impression of identifying with that distant and qualitatively different individual (see, e.g. Hilts 1995; Skloot 2004). Episodic memory is constructed such that it can likely produce a sense of identity with some person in the past even if there is *no* psychological continuity with the past person. That's just how episodic memory is built.⁷

The prevailing view in contemporary memory theory is that imagining future experience (“episodic foresight”) recruits some of the same mechanisms as recollecting past experience (e.g. Hayne et al. 2011). It's plausible that one commonality between episodic recollection and episodic prospection is in how the self is presented. When I remember an experience from childhood, the trait changes that I've undergone do not interfere with remembering the experience from the first person perspective; similarly, when I imagine sitting in my rocking chair as an octogenarian, the fact that I will have very different traits does not interfere with imagining rocking back and forth. Just as episodic retrospection generates a sense of identity across trait changes, so too in imagining the future, episodic prospection seems to allow us to project ourselves into new experiences without any attention to trait differences. This point applies to the Williams-style cases. It's easy for me to imagine feeling pain in the future without giving a moment's consideration to the stability of my psychological traits. This explains why in survey studies, people tend to say that even after a complete psychological transformation, *they* would feel the pain. This also explains why getting people to think that the self changes dramatically has no effect on their anxiety about future root canals. When imagining future pain, the trait-conception of self takes a backseat.

1.2. Pluralism and practical concerns

In the previous section, we argued that both survey studies and manipulation studies support a pluralist approach to characterizing folk views regarding personal identity. Moreover, work in memory research provides independent reason to think that the self gets represented in two very different ways. One of these representations presents the self as a set of psychological traits. The other representation is decidedly not defined in terms of psychological traits and fits with a biological or “animalist” conception of self. These two different representations issue different identity criteria.

In a series of important articles, David Shoemaker argues that no single account of personal identity is adequate for the plurality of our practical concerns (Shoemaker 2007, 2011, forthcoming). Thus, Shoemaker is a natural ally for our pluralist agenda. Despite this “pluralism of the practical” (Shoemaker forthcoming 23), Shoemaker does not embrace a pluralistic approach to the criteria for personal identity. Rather, he argues that the concept of *ownership* can ground all of our identity-related practical concerns, without an appeal to personal identity at all. In the following section, we first present Shoemaker's argument for the “pluralism of the practical” and then argue, *contra* Shoemaker, that pluralistic judgments cannot be fully explained with reference to the ownership relation he posits.

⁷ While the conception of self delivered by episodic memory often parallels the biological conception, it is not a perfect fit. For instance, some people report episodic memories of past lives, i.e., memories of events that they think they experienced as *other* biological organisms (see, e.g., Spanos et al. 1991). As mentioned above (footnote 2), we will set aside these complications for the purpose of our paper.

Shoemaker on the “Pluralism of the Practical”

Shoemaker argues that our practical concerns, though commonly thought to track the psychological continuity of persons, don’t track a single notion of personal identity (Shoemaker, 23).⁸ Shoemaker analyzes several identity-related practical concerns and argues that these concerns cannot be captured by a singular, monistic view of personal identity.

According to Shoemaker, social treatment often tracks biological, rather than psychological, continuity:

Consider someone who, due to some traumatic brain injury, undergoes radical psychological discontinuity. She will still be treated as the owner of the pre-transformation-person’s car and other property, and she will also be treated as the spouse of the pre-transformation-person’s spouse, the daughter of her parent, and so forth. (24)

To be sure, we don’t think it would be acceptable for the state to seize the property of people suffering from dementia, even for the demented with no heirs. And we don’t disapprove of the love expectant mothers feel for their fetuses or children feel for their elderly and vegetative parents. In these cases, social treatment depends on biological, not psychological, continuity (see also Schechtman 2010, pp. 275-276). Shoemaker also notes that compensation tracks biological continuity in certain situations:

Suppose Johann suddenly enters a fugue state. Call the radically psychologically discontinuous “fuguer” Sebastian. Suppose that I had broken Johann’s wrist prior to the fugue state but that I now have the medical equipment and expertise to completely heal it and, indeed, make it stronger than before (i.e. to “rejuvenate” it). When I rejuvenate the wrist I broke, it is Sebastian’s. Does what I have done count as compensation? It certainly seems so, despite the psychological discontinuity between Sebastian and Johann. This is because the kind of burden I attempted to rectify was to Johann’s animal self... (Shoemaker, 30)

Despite the lack of psychological continuity between Johann and Sebastian, an act of compensation remains possible. In this context, Shoemaker argues that compensation tracks biological continuity as opposed to psychological continuity.

Of course, both social treatment and compensation can also track psychological continuity. If a child’s entire psychological profile was transferred from her body to a new body, while the old body was destroyed, her mother would presumably treat this “new” individual as her own child, despite the lack of biological continuity. And, if an individual had caused some psychic trauma to the child before the transfer, but rectified the trauma after the transfer occurred, the child would still surely be compensated, again, despite the lack of biological continuity. In this way, it looks as though these identity-related practical concerns fail to track a singular, monistic criterion of identity, but rather follow two distinct criteria in different contexts.

Identity v. ownership

Shoemaker’s cases seem to provide yet further reason to think that judgments about personal identity are pluralistic. However, Shoemaker does not think that this plurality of the practical supports a pluralist theory of personal identity. Indeed, he explicitly rejects pluralism about personal identity (p. 31). Rather, he argues that the concept of *ownership* grounds the plurality of

⁸ In this section, we write primarily in terms of psychological continuity because Shoemaker sees his target as those who defend continuity, as opposed to connectedness, criteria. But as we noted above (footnote 3), we wish to remain neutral between these psychological accounts.

practical concerns, without implicating the concept of identity at all. While it is beyond the scope of this paper to analyze Shoemaker's ownership relation in full, we are not convinced that this relation can capture the relevant set of practical judgments without appealing to identity. To illustrate this, we will analyze two of the practical concerns Shoemaker argues are best grounded in terms of the ownership relation, and raise some worries for each.

According to Shoemaker, ascriptions of responsibility can be explained in terms of ownership. We hold agents responsible for their actions because they are *theirs*, not because these agents are identical to the agents who performed them. While it may seem as though ownership and personal identity are intimately connected, Shoemaker notes that "ownership is a relation between agent and action, whereas identity is a relation some individual bears to itself" (Shoemaker, 25). Though we do not wish to argue that these relations are identical, there is good reason to think that judgments of identity are critical to judgments of ownership. For example, part of what makes a previously performed action *mine* is that I am the same person as the agent who performed the action. Indeed, a common way to deny ownership of (and responsibility for) a given action is to argue that you are not the same person as the agent who performed it.

To prise apart ownership from identity, Shoemaker argues that an agent can own an action without being identical to the agent who performed it. Shoemaker presents two cases in which ownership explains attributions of responsibility yet identity does not obtain.

Shoemaker first points to joint acts, in which two or more agents both own and are responsible for an action, though they are not identical to one another. For example, though Jones and Smith both perform a bank robbery, and are equally responsible for this joint action, neither of them is identical to the "joint agent" that produced the action. But it doesn't follow that identity is irrelevant to these judgments. First, personal identity theorists can admit that multiple agents are capable of being responsible for joint actions without arguing that these agents must be identical to the joint agent. It is enough that an individual is identical to *one* of the people who took part in performing the joint action to establish a judgment of responsibility. More importantly, the judgment that Jones and Smith *own* the action seems to depend on a presupposition of identity. After all, if Jones were not identical to one of the people who helped perform the bank robbery, then we would have good reason to think that Jones does not own the action, and is thus not responsible for it.

Shoemaker then produces cases in which an agent is responsible for an act despite not having performed the act in any way. If a parent tells her child to break a valuable vase, then the parent would be responsible for the child's action even though he is not identical to the child. It seems odd to argue that one could *own* a child's action, even if one is in fact responsible. But in any case, an identity theorist need not argue that an individual is responsible only for those actions he or she performed. Indeed, even if one grants that the ownership relation obtains, it's plausible that it does so in part because the identity relation also obtains. That is, it's because the person who told the child to break the vase was identical to you that we hold you responsible for the action.

Though Shoemaker argues that pluralist intuitions about practical concerns are grounded in terms of ownership, and not in terms of identity, it is far from clear that we can neglect identity in explicating practical concerns. For judgments of ownership seem to depend crucially on judgments about personal identity. Furthermore, the empirical work presented in section 1.1 provides additional reason to think that assessments of identity really do play a critical role in the formation of judgments regarding our practical interests. The survey tasks explicitly ask about identity (Rips et al. 2006; Nichols & Bruno 2010). The manipulation tasks explicitly manipulate

beliefs about identity, and those manipulations affect practical decisions (see also Bartels & Urminsky 2011; Bartels, Kvaran, & Nichols forthcoming). And the fact that memory systems generate two different senses of the self provides an independent explanation for why people might operate with multiple criteria for identity – the different senses of self ground different identity criteria.

Shoemaker has a further reason to resist pluralism about identity– he maintains that pluralism about personal *identity* is incoherent: “we cannot be pluralists about numerical identity; we can only be pluralists about ownership”. One worry is that if there are plural criteria for personal identity, there can be cases that “require that I am both identical with, and not identical with, some past or future individual” (31). More generally, the idea that there are plural criteria for personal identity threatens to generate a problem of too many thinkers. As I deliberate about metaphysics, is the thinker the animal organism or the collection of psychological traits? Are *both* the organism and the psychology engaged in deliberation?

Even if there is some deep incoherence in pluralism about personal identity, it’s important to recognize the appropriate target of this concern. It might well be incoherent to *sustain* a pluralist account of identity. But that wouldn’t show that people don’t *have* a pluralistic conception. As we’ve seen, there is very good reason to think that people do operate as pluralists about personal identity. Of course, if it really is deeply incoherent to be a pluralist about personal identity, that might give us good reason to reject commonsense pluralism. But the goal of this first part of our paper was merely to establish that people are committed to plural criteria of personal identity. In the next part of the paper we examine whether we should abandon those commitments.

2. Philosophical Pluralism about Personal Identity

The fact that individuals’ judgments about practical concerns track varying conceptions of personal identity does not entail that we, as philosophers, should endorse pluralism about personal identity. It could be the case that people are simply mistaken when they make judgments about practical concerns that track biological views of personal identity, for example. Perhaps children are unwarranted in feeling obligated to pay for medical care for their senile parents. Alternatively (or even additionally), it could be the case that people are mistaken when they make judgments about practical concerns that track psychological views of personal identity. Perhaps people are unwarranted in attending to psychological continuity or connectedness in allotting punishment. In this part of the paper, we will explore – in a preliminary and partial way - the philosophical sustainability of pluralism about personal identity.

2.1. The price of monism

Much of philosophy gets its interest precisely because it flows from commonsense conceptions. Accordingly, on many views of metaphysics, folk conceptions define domains (e.g., Lewis 1972; Jackson 1998; Strawson 1959). Jackson defends this view by noting that until we have defined the subject matter, we cannot make progress on metaphysical questions. Jackson uses the analogy of a bounty hunter to make the point. A bounty hunter will “not get very far if they fail to attend to the representational properties of the handbill” (30). “Likewise,” Jackson writes, Metaphysicians will not get very far with questions like: Are there *Ks*? Are *Ks* nothing over and above *Js*? and, Is the *K* way the world is fully determined by the *J* way the

world is? in the absence of some conception of what counts as a *K* and what counts as a *J* (30-1)

How are we to determine the subject matter for philosophical issues like the self and free action, then? Jackson maintains that for many philosophical questions, the most plausible answer is one that appeals to ordinary conceptions:

What ... are the interesting philosophical questions that we are seeking to address when we debate the existence of free action and its compatibility with determinism...? What we are seeking to address is whether free action *according to our ordinary conception* or something close to our ordinary conceptions, exists and is compatible with determinism (31).

If we take this approach to thinking about the subject matter of the metaphysics of personal identity, it provides a presumption in favor of sustaining pluralism about the *subject matter* of personal identity. That is, the fact that we *have* two senses of self and that they regiment judgments and decisions differently in different contexts provides *prima facie* reason to think that when we are doing the philosophy of personal identity, we need to adopt a pluralist approach to the subject matter. In some contexts, the ordinary conception of personal identity is trait based; in other contexts, the ordinary conception of personal identity is biological. Forcing monism on the subject matter would be akin to a bounty hunter searching for one person from the morphed images of two very different people.

It's not just that pluralism fits with what people *say*. As we've seen, pluralism also captures people's *practices*. Our commitments about possessions, inheritance, and familial obligations all conform to the biological conception. To extirpate that conception would be wildly disruptive to our practical concerns. Indeed, it's not clear that it's within the realm of reasonable possibility that we *could* eradicate those kinds of biologically-based practical concerns.

At the same time, people attach enormous importance to their memories and values for personal survival. Furthermore, as we've seen, judgments about punishment are sensitive to psychological continuity and connectedness. This comports well with the fact that some of the pragmatic aims of punishment, e.g., moral education, fail to make sense if our only standard for punishment was a biological criterion. We often punish those who perform illicit actions in order to alter these individuals' psychology such that they recognize the wrongness of their ways, or at least form an aversion to committing future wrongs. Furthermore, as philosophers since Locke have emphasized, what makes an agent responsible for an action seems not simply to be the fact that she is biologically continuous with the agent who performed the action. Rather, we take an agent's psychological characteristics to bear directly on whether she is responsible for an action, and thus whether punishment is appropriate. Thus, our intuitions about the propriety of punishment and some of the most important justifications of punishment cannot be captured in solely biological terms.

2.2. The coherence of pluralism

In debates over personal identity, animalists bring to bear reasons to reject psychological criteria, and their opponents marshal reasons to reject animalism. Obviously we cannot attempt to address the wide range of arguments brought against one theory or the other. Instead, we want to focus on whether there are good reasons to object to the *pluralistic* aspect of commonsense thought about personal identity. One natural concern is that to embrace a pluralistic approach to personal identity is to embrace blatant inconsistencies. For instance, endorsing the pluralistic

approach requires me to say that I both am and am not the same person as a baby born some time ago.

To explore the viability of pluralism, we want to consider how it fares against these kinds of concerns. In the recent literature, pluralism has not been the target of much direct discussion, but the view might seem to be susceptible to one of the most prominent problems raised by animalists: the too-many-thinkers problem. There are many different ways to present the problem. Eric Olson, in his book *What are We?*, sets it up in the following way:

Suppose human animals think in just the way that we do: every thought of yours is a thought on the part of a certain animal. How could that thinking animal be anything other than you? Only if you are one of at least two beings that think your thoughts? (Or maybe you and the animal think numerically different but otherwise identical thoughts. Then you are one of at least two beings thinking exactly similar thoughts.) If you think, and your animal body thinks, and it is not you, then there are two thinkers sitting there and reading this book. (Olson 2007, 35)

According to Olson, this is counterintuitive for several reasons. First, it would be very odd if one body housed several different thinkers who thought exactly the same thoughts, yet only one of these thinkers is *you*. However, this is only odd on a monist's view. A pluralist can happily grant that there are two thinkers: under some circumstances you are the animal; under other circumstances, you are the constellation of psychological traits. Of course, one could still argue that it would be odd for two individuals to be housed in one body. The counterintuitiveness of this is acknowledged by four-dimensionalists themselves (e.g., Lewis 1976, 26). But the pluralist need not be committed to this kind of multiple-occupancy view. For it's plausible that the animal and the psychological traits do not occupy exactly the same spatio-temporal region.

An epistemic worry also arises from the too-many-thinkers problem: if there are two thinkers housed in your body, how do you know which thinker is you?

Worse, you *ought* to wonder which of the two thinkers is you. You may believe that you are the non-animal (because you accept the Psychological Approach, perhaps). But the animal has the same grounds for believing that *it* is a non-animal as you have for supposing that you are. Yet it is mistaken. For all you know, you might be the one making this mistake. If you *were* the animal and not the person, you would never be any the wiser. (Olson 2009, 82)

Again, such a challenge is only a problem for the monist. While proponents of the psychological approach must address the possibility that there is no way, in principle, to distinguish the biological thinker from the psychological thinker, pluralists are not affected by such a worry. On a pluralist view of personal identity, there is no dispute to settle, for both the biological thinker and the psychological thinker are correct—they *are* both you.

While the too-many-thinkers problem may indeed be a problem, it seems only to be a problem for monists about personal identity. Indeed, even the description of this objection strikes us as an attack on monism, not pluralism—you can only have too many thinkers if you think there should only be one.

However, there is another worry closely related to the too-many-thinkers problem that may challenge the metaphysical coherence of pluralism. Shoemaker writes:

...we cannot be pluralists about numerical identity... If the numerical identity we are talking about is identity of individuals like us, then the proposal just given

could require that I am both identical with, and not identical with, some past or future individual. (Shoemaker, 31)

At first glance, Shoemaker's worry seems deeply troubling: how can it be that an agent is both identical and not identical to the same individual? The pluralist view has a natural response to Shoemaker's concern, however. Consider a case in which an individual is biologically continuous or identical with a past individual, but is not psychologically continuous or identical with that same individual. In one sense, it can be said that the individual both is and is not identical with this past individual, but this phraseology is misleading. It is more apt to say that the individual is biologically identical or continuous with a past individual, but not psychologically identical or continuous. Being biologically continuous and psychologically continuous are two very different states of existence—they require different persistence conditions. If we are explicit about which criterion we are using for our persistence claims, Shoemaker's concern is greatly diminished.

None of the foregoing provides a defense of either the psychological or biological approach to identity. Indeed, for all we have said, it might be the case that ultimately, the right view about personal identity is a nihilistic one on which the self is simply an illusion. There might be deep problems with our ordinary conceptions of the self. But the mere fact that we have multiple conceptions does not look to present us with obvious incoherency.

Conclusion: The pull of pluralism

It is a common strategy in experimental philosophy to explain classic philosophical debates—compatibilism versus hard determinism, for example—by isolating different psychological systems that undergird each competing view.⁹ Often, however, this method only succeeds in explaining the debate, not resolving it, for the two competing views, though both grounded in psychological systems, remain incompatible. Interestingly, we may be able to do more when it comes to the debate between biological and psychological approaches to personal identity. When people make judgments about persistence and engage in practical decision-making, different contexts seem to trigger different criteria of personal identity. Both empirical evidence and philosophical thought experiments indicate that judgments about personal identity are regimented by two different criteria, one in terms of psychological traits and one that largely conforms to biological criteria. Only a pluralist conception of the self can capture all of our intuitions about persistence—be it those generated by Williams' famous thought experiment or our everyday attributions of moral responsibility. Furthermore, we've argued that some of the most natural objections against pluralism about personal identity do not threaten the kind of pluralism that seems to be implicit in folk judgments and practices. It is possible to synthesize the biological and psychological approach to personal identity into a coherent pluralist view of the self. The monistic presumption that guides philosophical inquiry might be flatly misplaced when it comes to the self. While the assumption of monism is often a sensible theoretical constraint, a monistic view cannot adequately capture folk conceptions of personal identity. Perhaps pluralism has much more pull than originally thought.

⁹ See Nichols, S. & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions." *Nous*, 41, 663-685.

Works Cited

- Bartels, D., Kvaran, T., and Nichols, S. Under review. Selfless giving.
- Bartels, D. and Rips, L. 2010. Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, 139, 49–69.
- Bartels, D., & Urminsky, O. 2011. On intertemporal selfishness: The perceived instability of identity underlies impatient consumption. *Journal of Consumer Research*, 38, 182-198.
- Hayne, H., Gross, J., McNamee, S., Fitzgibbon, O., & Tustin, K. (2011). Episodic Memory and Episodic Foresight in 3-and 5-Year-Old Children. *Cognitive Development*, 26(4), 343-355.
- Hilts, P. 1995. *Memory's Ghost: The Nature Of Memory And The Strange Tale Of Mr. M*. New York: Simon & Schuster.
- Klein, S., German, T., Cosmides, L., Rami, G. 2004: A Theory of Autobiographical Memory: Necessary Components and Disorders Resulting from Their Loss. *Social Cognition*, 22, pp. 460-490.
- Jackson, F. 1998. *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Lewis, D. 1972. Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy*, 249-258.
- Lewis, D. 1976. Survival and identity. In A. Rorty (ed.) *Identities of Persons*. University of California Press.
- Nichols, S. Forthcoming. The episodic sense of self.
- Nichols, S., & Bruno, M. 2010. Intuitions about personal identity: An empirical study. *Philosophical Psychology*, 23(3), 293-312.
- Nichols, S. & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions." *Nous*, 41, 663-685.
- Olson, E. 1997. *The human animal: Personal identity without psychology*. Oxford University Press.
- Olson, E., 2007. *What are We?* Oxford: Oxford University Press.
- Olson, E., 2009. Personal Identity. in *Science Fiction and Philosophy: From Time Travel to Superintelligence*, S. Schneider. (ed.), Malden, MA: Blackwell.
- Parfit, D. 1971. Personal identity. *The Philosophical Review*, 80(1), 3-27.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Reid, T. 1785. *Essays on the intellectual powers of man*. Edinburgh: J. Bell.
- Rips, L., Blok, S. & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, 113, 1–30.
- Rovane, C. 1990. Branching self-consciousness. *The Philosophical Review*, 99(3), 355-395.
- Shoemaker, D. forthcoming. The stony metaphysical heart of animalism. In Stephan Blatti and Paul Snowdon (eds.) *Animalism*. Oxford: Oxford University Press.
- Shoemaker, D. 1996. Theoretical persons and practical agents. *Philosophy & Public Affairs*, 25(4), 318-332.
- Shoemaker, D. 2007. Personal identity and practical concerns. *Mind*, 116: 316-357.
- Shoemaker, D. 2011. Moral responsibility and the self. In Shaun Gallagher, ed., *Oxford Handbook of the Self*. Oxford: Oxford University Press: pp. 487-520.
- Sider, T. 2001. Criteria of personal identity and the limits of conceptual analysis. *Philosophical Perspectives*, 15, 189-209.
- Skloot, F. 2004. *In the Shadow of Memory*. Lincoln: University of Nebraska Press.

- Spanos, N., Menary, E., Gabora, N., DuBreuil, S., Dewhirst, B. 1991. Secondary identity enactments during hypnotic past-life regression: A sociocognitive perspective. *Journal of Personality and Social Psychology*, 61, 308-320.
- Strawson, P. 1959. *Individuals: An Essay in Descriptive Metaphysics*. New York: Anchor Books.
- Tierney, H. in prep. Breaking Up is Hard to Do: The Failings of Fission Cases.
- Tulving, E. 1993 Self-knowledge of an amnesic individual is represented abstractly. In T. K. Srull & R. S. Wyer (Eds.), *Advances in Social cognition* (Vol. 5, pp.1-49). Hillsdale, N.J.: Erlbaum.
- Williams, B. 1970. The self and the future. *The Philosophical Review*, 79, 161-180.