

# Persistence and the First-Person Perspective

---

*Dilip Ninan*

University of St Andrews

## 1. Introduction

A familiar divide in the theory of personal identity over time is between those who think personal identity is essentially a matter of psychological continuity and those who think it is essentially a matter of bodily (or biological) continuity. But these theorists share a common opponent: the philosopher who thinks personal identity is a “further fact,” over and above facts about psychological and bodily continuity. Parfit (1984) calls the view that personal identity is essentially a matter of psychological or bodily continuity “the Complex View,” and he calls the view that personal identity is a further fact, over and above the continuity facts, “the Simple View.”

The Simple View is not popular in contemporary philosophy, but it is not completely without appeal.<sup>1</sup> As several philosophers have noted, the Simple View looks compelling when one thinks about one’s own persistence from the first-person point of view or “from the inside.”<sup>2</sup> But in spite of its first-person appeal, the Simple View has few adherents,

For helpful comments and discussion, thanks to Rachael Briggs, Alex Byrne, Richard Holton, Vann McGee, Sarah Moss, Joel Pust, Judith Jarvis Thomson, Stephen Yablo, Seth Yalcin, Zoltán Gendler Szabó, and to audiences at MIT, University College London, and Arché (St Andrews). I’m especially grateful to Caspar Hare, Robert Stalnaker, and an anonymous referee for the *Philosophical Review* for helping to clarify key aspects of this essay.

1. Nor completely without defenders. See, for example, Chisholm (1969), Madell (1981), Swinburne (1984), Lowe (1996, 41–44), and Merricks (1998). Earlier defenders of the Simple View include Butler (1736) and Reid (2002 [1785]).

2. See, for example, Chisholm 1969; Madell 1981; Johnston 1987; Nagel 1986, chap. 3; Blackburn 1997; and Baker 2000, 132ff.

largely because it has certain undesirable theoretical consequences: it conflicts with the widely held doctrine of *physicalism* (the view that everything supervenes on the physical), and it posits hidden facts about our persistence, facts that cannot even in principle be discovered by well-placed observers. The aim of this essay is to suggest a way of respecting our first-person judgments about our persistence without having to accept these consequences.

I begin in §2 by giving a more precise statement of the Simple and Complex Views. In §3, I set out the central dilemma of the essay, first noting the first-person case for the Simple View and then showing how the Simple View conflicts with physicalism and requires us to posit hidden facts. In §4, I present an account of first-person imagining and its relation to possibility: the basic idea is that first-person imaginings have *centered worlds* contents in the sense of Lewis [1979] 1983, and thus serve as evidence for *centered possibilities*, a type of possibility analyzed in terms of centered worlds (Lewis [1983a] 1999; Lewis 1986, §4.4). In §5, I use this account—along with a claim about the ontology of persistence—to show how our first-person judgments about our persistence can be reconciled with physicalism and with the claim that there are no hidden persistence facts. I close by considering an objection to the resulting account.

## 2. Simple and Complex

The Complex View says that personal identity is essentially a matter of some kind of physical or psychological continuity. So we can think of the Complex View as a long disjunction, one that says that either the psychological continuity theory is true or the bodily continuity theory is true or . . . and so on for each plausible, informative view of personal identity. A standard way of formulating a continuity theory is to say something like this:

### Schema 1

For all possible worlds  $w$  and all stages  $x$  and  $y$  in  $w$ ,  $x$  and  $y$  are stages of the same person in  $w$  iff  $x$  bears  $R$  to  $y$  in  $w$ .

*Stages* are understood by four-dimensionalists to be human-shaped objects that exist only for a brief moment; entities like you and me are aggregates of such stages, and so we are extended in time just as we are extended in space (Lewis 1976). Three-dimensionalists can understand a stage as a pair consisting of a continuing object and a time at which the

object exists (Sider 1999).<sup>3</sup> In the above schema, 'R' can be replaced by 'the relation of psychological continuity' to yield the psychological continuity theory or by 'the relation of bodily continuity' to yield the bodily continuity theory, and so on.

Although this may seem like a reasonable way of setting things up, it has the unfortunate consequence of ignoring certain views of how we persist. *Animalists* think that we have the persistence conditions of human organisms or human animals, so that our persistence through time is essentially a matter of biological continuity.<sup>4</sup> Animalists often claim that we are essentially human organisms and only contingently persons (Olson 2007, 45–46). On this view, even if we replaced 'R' in **Schema 1** with 'the relation of biological continuity', we would not have fully characterized our persistence conditions for the resulting account would tell us only what it takes for us to survive *and* remain persons not what it takes for us to survive simpliciter.<sup>5</sup> But it is the latter question that is the central one: what are the necessary and sufficient conditions for my survival, whether or not I will survive as a person? Now it might be that animalists are wrong and that we are in fact essentially persons. But since the Complex View is supposed to be neutral between various different views of our persistence conditions, our definition of the Complex View should not exclude animalism from the outset. For this reason, we might be better off formulating continuity views as follows:

**Schema 2**

For all  $x \in S$  and all possible worlds  $w$  in which  $x$  exists, and all stages  $y$  and  $z$  in  $w$ , if  $y$  is a stage of  $x$  in  $w$ , then  $z$  is a stage of  $x$  in  $w$  iff  $z$  bears  $R$  to  $y$  in  $w$ .

The set  $S$  is intended to include you, me, and all the other uncontroversial cases of human persons, actual and nonactual. All theories of our persistence have something to say about the persistence conditions of these entities; we can leave it open whether or not these entities are essentially persons or essentially human organisms or essentially something else altogether.

3. For an introduction to, and thorough discussion of, the debate between three-dimensionalists and four-dimensionalists, see Sider 2001b.

4. Defenders of animalism include Snowdon (1990), van Inwagen (1990), and Olson (1997). See also Olson 2007.

5. Indeed, the theory that results from such a substitution is false according to most versions of animalism since animalists usually maintain that  $x$  and  $y$  can be biologically continuous with each other without  $x$  being the same person as  $y$ .

Another caveat: I don't think the Complex View rules out the idea that there are cases in which it's indeterminate whether or not someone should be described as having survived some vicissitude. For example, it may be that when psychological continuity comes apart from bodily continuity, it's indeterminate what happens to the initial individual, whether he or she goes where his or her body goes or where his or her psychology goes (Sider 2001a). So in addition to the instances of **Schema 2** we get by replacing '*R*' with a description of some continuity relation, I want our long disjunction to include views that allow for (a significant amount of) indeterminacy.

Call our long disjunction '**Disjunction**'. The Complex View is the view that **Disjunction** is true. Anyone who believes one of the disjuncts—animalists, psychological continuity theorists, and so on—is an adherent of the Complex View. The Simple View entails that **Disjunction** is false.

Note that **Disjunction** entails a supervenience thesis:<sup>6</sup> it entails that facts about our persistence supervene on facts about continuity in this sense:

### Supervenience

For all  $x, x' \in S$ , and for all possible worlds  $w, w'$  such that  $x$  exists in  $w$  and  $x'$  exists in  $w'$ , and for all pairs of stages  $\langle y, z \rangle$  in  $w$ ,  $\langle y', z' \rangle$  in  $w'$ : if  $y$  is a stage of  $x$  in  $w$ , and  $y'$  is a stage of  $x'$  in  $w'$ , and  $y$  and  $z$  in  $w$  and  $y'$  and  $z'$  in  $w'$  are the same with respect to continuity, then  $z$  is a stage of  $x$  in  $w$  iff  $z'$  is a stage of  $x'$  in  $w'$ .<sup>7</sup>

6. To see that **Disjunction** entails **Supervenience**, suppose that an arbitrary disjunct of **Disjunction** is true but that **Supervenience** is false. An arbitrary disjunct:

For all  $x \in S$  and all possible worlds  $w$  in which  $x$  exists, and all stages  $y$  and  $z$  in  $w$ , if  $y$  is a stage of  $x$  in  $w$ , then  $z$  is a stage of  $x$  in  $w$  iff  $z$  bears  $R$  to  $y$  in  $w$ .

Since we're supposing that **Supervenience** is false, there must be individuals  $x$  and  $x'$ , and worlds  $w$  and  $w'$ , and pairs of stages  $\langle y, z \rangle$  in  $w$  and  $\langle y', z' \rangle$  in  $w'$  such that:  $y$  is a stage of  $x$  in  $w$ , and  $y'$  is a stage of  $x'$  in  $w'$ , and  $y$  and  $z$  and  $y'$  and  $z'$  are the same with respect to continuity, but  $z$  is a stage of  $x$  in  $w$  while  $z'$  is not a stage of  $x'$  in  $w'$ . Now since  $y$  and  $z$  and  $y'$  and  $z'$  are the same with respect to continuity,  $z$  bears  $R$  to  $y$  in  $w$  iff  $z'$  bears  $R$  to  $y'$  in  $w'$  (I'm assuming that  $R$  is symmetric). There are two cases: First case:  $z$  bears  $R$  to  $y$  in  $w$  and  $z'$  bears  $R$  to  $y'$  in  $w'$ . So we have that  $y'$  is a stage of  $x'$  in  $w'$  and  $z'$  bears  $R$  to  $y'$  in  $w'$ . But then given our disjunct, it follows that  $z'$  is a stage of  $x'$  in  $w'$ , which contradicts the supposition. Second case:  $z$  doesn't bear  $R$  to  $y$  in  $w$ , and  $z'$  doesn't bear  $R$  to  $y'$  in  $w'$ . But then since  $y$  is a stage of  $x$  in  $w$ , and  $z$  doesn't bear  $R$  to  $y$  in  $w$ , it follows from our disjunct that  $z$  isn't a stage of  $x$  in  $w$ , and that again contradicts our supposition. Thus any arbitrary disjunct of **Disjunction** entails **Supervenience**, and so all of them do, and so **Disjunction** does.

7. **Supervenience** would need to be qualified if there were any cases in which it is indeterminate whether or not two stages are stages of the same individual. This could be done

Two pairs of possible stages  $\langle x, y \rangle$  in  $w$  and  $\langle x', y' \rangle$  in  $w'$  are *the same with respect to continuity* iff:

- $x$  is psychologically continuous with  $y$  in  $w$  iff  $x'$  is psychologically continuous with  $y'$  in  $w'$ ;
- $x$  is bodily continuous with  $y$  in  $w$  iff  $x'$  is bodily continuous with  $y'$  in  $w'$ ;
- $x$  is “brain continuous” with  $y$  in  $w$  iff  $x'$  is “brain continuous” with  $y'$  in  $w'$ ;
- and so on for each relevant continuity relation (including “non-branching” versions of the above relations).<sup>8</sup>

The basic idea behind **Supervenience** is that any two possible situations that share the same continuity facts share the same persistence facts. There can be no difference in how you and I persist without that difference being grounded in the relevant continuity facts.

If **Supervenience** is true, its truth is not brute; if it is true, it is true because one of the disjuncts of **Disjunction** is true (remember: **Disjunction** includes as disjuncts *all* plausible informative theories of our persistence). So we needn't consider a view which denies **Disjunction** but endorses **Supervenience**. I take the Complex View to be committed to **Disjunction**, and therefore committed to **Supervenience**. And I take the Simple View to be committed to the denial of **Supervenience**, and so to the denial of **Disjunction**.

### 3. The Dilemma

#### 3.1. The Case for the Simple View

Here's something I like about the Simple View: it seems true when I consider my own persistence through time from the inside or from the first-person point of view. I'm not alone in finding this aspect of the Simple View appealing. A recurring theme in the literature on our persistence is that, when one considers one's own persistence from the first-person point of view, it seems as though one's persistence is a basic and fundamental fact,

---

by changing the consequent of the above conditional to: '( $z$  is a stage of  $x$  in  $w$ , and  $z'$  is a stage of  $x'$  in  $w'$ ) or ( $z$  is not a stage of  $x$  in  $w$ , and  $z'$  is not a stage of  $x'$  in  $w'$ ) or (it is indeterminate whether  $z$  is a stage of  $x$  in  $w$ , and it is indeterminate whether  $z'$  is a stage of  $x'$  in  $w'$ )'.

8. Given a continuity relation  $R$ , the corresponding nonbranching relation is the one that holds between stages  $x$  and  $y$  iff (i)  $x$  bears  $R$  to  $y$ , and (ii) there is no stage  $z$  ( $z \neq x$  and  $z \neq y$ ) such that one and only one of  $x$  and  $y$  bears  $R$  to  $z$ .

one that is metaphysically independent of facts about physical and psychological continuity.<sup>9</sup>

For example, Blackburn (1997, 181) notes that when one imagines the well-known fission thought experiment from the first-person point of view, one can see three possibilities for one's survival, all of which are consistent with the same continuity facts (see also Madell 1981, 127–31 and Baker 2000, 137). To see this, let us first remind ourselves of the details of the fission case. In that case, someone's brain is divided, and his or her left hemisphere is then transplanted into a waiting body, and his or her right hemisphere is transplanted into a different waiting body. The two resulting individuals bear a high degree of psychological continuity to the initial individual, so much so that, were it not for the other individual, any psychological continuity theorist would happily concede that the initial individual had survived the operation. If we imagine that the left-hemisphere individual is to awake in a bright green room and the right-hemisphere individual is to awake in a bright red room, then Blackburn's point is that when I imagine from the inside undergoing fission, it seems as if there are three ways things could go: either I awake to see green, or I awake to see red, or I do not awake at all. That is, I can imagine from the inside undergoing fission and waking up in the green room, and I can imagine undergoing fission and waking up in the red room, and I can imagine undergoing fission and failing to wake up at all. But in all three cases, the continuity facts would be exactly the same, and so Blackburn's observation suggests that facts about my survival are further facts, over and above the facts about physical and psychological continuity.

More precisely, a problem for **Supervenience** arises as follows. If there is a case in which I go left and see green, there is a possible world  $w$  in which a postfission stage  $z$  of the individual in the green room is a stage of me. If there is a case in which I go right and see red, there is a possible world  $w'$  in which a postfission stage  $z'$  of the individual in the green room is *not* a stage of me (since I end up in the red room). But consider  $y$ , a prefission stage of me in  $w$ , and  $y'$ , a prefission stage of me in  $w'$ .  $y$  and  $z$  in  $w$  and  $y'$  and  $z'$  in  $w'$  are the same with respect to continuity, and yet  $z$  is a stage of me in  $w$  even though  $z'$  is not a stage of me in  $w'$ . So Blackburn's observation seems to give us a counterexample to **Supervenience**.

Blackburn's observation tells directly against **Supervenience**, and so indirectly against **Disjunction** since the latter entails the former. But

9. See the references in footnote 2.

several authors have pointed to considerations that tell directly against **Disjunction** itself. Consider, for example, this passage from Thomas Nagel (1986, 33):

When I consider my own individual life from inside, it seems that my existence in the future or the past—the existence of the same ‘I’ as this one—depends on nothing but itself. . . . My nature then appears to be at least conceptually independent not only of bodily continuity but of all other subjective mental conditions, such as memory and psychological similarity. It can seem, in this frame of mind, that whether a past or future mental state is mine or not is a fact not analyzable in terms of any relations of continuity, psychological or physical, between that state and my present state.

And consider this one from Mark Johnston (1987, 70):

We can imagine many sorts of cases that seem to involve one’s ceasing to be associated with a particular human body and human personality. These cases are particularly compelling when imagined “from the inside.” So I am to imagine undergoing a radical change in my form . . . and perhaps concurrently a wild change in my psychology. There seems to be nothing internally incoherent about such imaginings.

In essence, both Nagel and Johnston are pointing to the fact that, for each disjunct *d* in **Disjunction**, I can imagine from the inside a scenario in which I survive some vicissitude without enjoying the continuity that *d* says is necessary for my survival. The fission case is already one in which one imagines surviving without one’s particular body, but Nagel and Johnston both note that one can also imagine surviving the destruction of one’s psychology, perhaps by undergoing a nefarious medical procedure that results in one’s waking up with an entirely new set of apparent memories, beliefs, desires, character traits, and so on. As Nagel and Johnston say, such cases are most compelling when imagined from the inside, or from the first-person point of view. The observation that one can imagine undergoing a complete psychological transformation—a “brain zap”—was also made by Chisholm 1969 and Williams 1970. Roderick Chisholm, who was explicitly arguing for the Simple View, noted that in order to appreciate the possibility that someone might survive such a vicissitude, “it may be necessary to imagine that the person in question is oneself” (Chisholm 1969, 103–4).

Philosophers inclined toward the idea that our persistence is a matter of psychological continuity have not, as far as I am aware, rejected the possibility of such first-person imaginings. Instead, they suggest that if we subjected an individual *A* to a complete brain zap, the individual *B* that

resulted would not be identical to *A* (see, for example, Shoemaker 1984, 87 and Perry 2002, 115–18). This may be correct as a claim about what we should say about such a case “from the outside,” that is, as a claim about what to say were we to observe someone undergoing a brain zap, but it says nothing about the claim that surviving a brain zap seems possible from the inside.

In what follows, I will, for the sake of simplicity, focus on Blackburn’s observation that, when I look at fission from the first-person point of view, it seems that there are three ways things could go. This observation gives us an apparent counterexample to **Supervenience**, and thus illustrates the general theme that the Simple View looks true when we look at hypothetical cases from the first-person point of view.

### 3.2. *Against the Simple View: Physicalism*

The Simple View takes seriously the possibility judgments I’m inclined to make when I imagine cases from the first-person point of view. That’s something I like about the Simple View. But here’s something I don’t like about it: it conflicts with *physicalism*, the doctrine that everything supervenes on the physical.

Physicalism is often characterized as the thesis that any two possible worlds that are physically just alike are alike simpliciter.<sup>10</sup> One might think that the denial of **Supervenience** entails the denial of physicalism so defined. But, strictly speaking, it doesn’t. For to deny **Supervenience** is to deny that facts about our persistence supervene on the *local* continuity facts. But one can deny this while maintaining physicalism if one holds that facts about our persistence supervene on the *totality* of the physical facts. That is, one might hold that two situations *s* and *s'* could be exactly alike in their continuity facts and yet differ on the persistence facts because there is *some* physical difference between the possible worlds in which *s* and *s'* are located—perhaps *s*’s world contains an extra fundamental particle floating around in some remote region of space.

This is a possible position but not a plausible one. Consider the Simple View’s treatment of fission, which is motivated by Blackburn’s observation that when I imagine facing fission from the first-person point of view, I can imagine things going in three different ways. In taking this to

10. Some authors prefer to restrict the domain of quantification to worlds “which contain the same basic laws and ingredients as our world” (Jackson 1998, 13). See Jackson 1998, chap. 1 and Stoljar 2005 for some discussion of the subtleties involved in formulating physicalism.



challenge **Supervenience**, all three cases must be ones in which the relevant continuity facts are exactly the same. But it's simply not the case that the only way to accomplish this imaginative feat is to hold that there is *some* physical difference between the three imagined situations; on the contrary, the normal way of imagining the case is to hold *everything*—including every physical fact—fixed, while varying only the facts about what happens to oneself. And if that's right, then Blackburn's observation seems to conflict with physicalism for it would seem to suggest that there are three possible worlds  $w$ ,  $w'$ ,  $w''$  that are physically identical but differ on what happens to me after fission: in  $w$ , I end up in the green room; in  $w'$ , I end up in the red room; and in  $w''$ , I don't exist after fission. These three worlds are physically just alike, but not alike simpliciter, and so would seem to constitute a counterexample to physicalism.<sup>11</sup>

I find the Simple View's incompatibility with physicalism unattractive, and I'm not alone in this. Indeed, I suspect this is one of the main reasons the Simple View has so few advocates in contemporary philosophy. A number of theorists explicitly reject the Simple View on the grounds that it contradicts physicalism (or "materialism"). For example, Eric Olson (1997, 3–4) begins his book-length defense of animalism with the assumption of materialism, thus ruling out the Simple View from the get-go. Sydney Shoemaker (1984, 71) similarly takes materialism to be a constraint on a theory of our persistence:

An account of personal identity . . . ought also to cohere with the rest of what we know about the world. In my own view, this last requirement means that an account of personal identity ought to be compatible with a naturalistic, or materialistic, account of mind.

Olson and Shoemaker are not unusual in treating physicalism as a background assumption. In contemporary philosophy, physicalism is essentially the default view; it is *departures* from physicalism that are usually taken to require extensive motivation. So if, like Olson, Shoemaker and

11. One could try to avoid this argument by claiming that physicalism is really the view that any two worlds that are physically alike are *qualitatively* alike and then claiming that the only differences between these three worlds are nonqualitative ("haecceitist") differences. As far as I can see, this move will work only if one also accepts a generous ontology of persistence, of the sort I discuss in §5. Although it's beyond the scope of the present essay to explore this interesting option in detail, let me offer two reasons for not going this route: first, I'm not sure that the combination of a generous ontology of persistence plus haecceitism is ultimately sustainable; and second, unlike the view developed in this essay, the haecceitist view would seem to lack the resources to explain why it is that the three possibilities in question are most easily seen from the first-person point of view.

many other philosophers, you find physicalism attractive, you will find yourself with a reason to dislike the Simple View.

### 3.3. *Against the Simple View: Hidden Facts*

Here is another thing I don't like about the Simple View: it posits hidden facts about our persistence. Again I'm not alone in objecting to this feature of the Simple View. Johnston (1987, 63, 74), for example, takes it to be a constraint on an account of our persistence that it be compatible with a non-mysterious account of how we can come to know the persistence facts, and he rejects the Simple View on the ground that it violates this constraint.<sup>12</sup>

To see what's at issue, consider the fission case again. Suppose that you and I are watching someone—call him 'Sam'—undergo fission. Prior to fission, we are permitted to study his body, brain, and psychology very closely. We run tests, perform brain scans, and ask him lots of detailed questions. In sum: we have access to as much information about Sam as we like. And when the operation occurs, we are permitted to monitor it very closely. And after the operation, we are again given access to all the relevant information about the two resulting individuals, the one with the left hemisphere in the green room, and the one with the right hemisphere in the red room. We learn what their bodies, brains, and psychologies are like, and how their bodies, brains, and psychologies are related to Sam's. We know that both are conscious, we know what they are thinking, we know "what it's like" for them. In sum: we now have all the information that could possibly be relevant for deciding what happened to Sam.

The Simple View tells us that there are three ways things could have gone for Sam: either he ended up with the left hemisphere in the green room, or with the right hemisphere in the red room, or he failed to survive. But it seems that we cannot possibly come to know which possibility has obtained. For no matter which possibility has obtained, all of our relevant evidence will be exactly the same. No matter which possibility has obtained, *all* the facts about these three individuals-at-times—all the facts about their bodies, brains, and psychologies, and how they are related—will be the same. But since these are the only facts relevant for deciding whether two individuals-at-times are the same or not, all of the evidence relevant to the question of what happened to Sam will be exactly the same in all three cases. It's thus hard to see how we could ever come to know what happened to Sam.

12. See also Shoemaker 1963, 12 and Shoemaker 1984, 123–24.

That the Simple View posits hidden facts isn't a decisive objection to it—there can be unknowable truths. But it is, I think, an unattractive feature of the view: other things being equal, a view that doesn't posit hidden persistence facts is preferable to one that does. One reason for this is that it just doesn't seem like there are hidden persistence facts: when I imagine observing every last detail of a fission case, it just seems false to me that there is a hidden fact about the case, a possibility whose obtaining or failure to obtain is one we could not even in principle detect. When I imagine examining these cases very closely, learning every fact about the relevant bodies, brains, and psychologies, it seems to me that I know everything that there is to be known about the case. The hypothesis that there is a further, completely undetectable fact here just seems unwarranted.

Another reason for disliking this feature of the Simple View emerges when we compare it to the Complex View. For the Complex View offers us a simple and straightforward account of how we come to know facts about our persistence: if we know the continuity facts—which are ordinary, externally observable facts about bodies, brains, and psychologies—then we can come to know the persistence facts if we know how they supervene on the continuity facts. In contrast, it's far from clear what the Simple theorist can say about how we come to know facts about our persistence. Since the Simple View denies **Supervenience**, there will have to be cases—like the fission case—in which we know all the continuity facts but cannot even in principle come to know the persistence facts. So it is hard to see how the Simple theorist will be able to give an account of how we come to know the persistence facts that is as simple and straightforward as the Complex theorist's account. While perhaps not a decisive consideration, this does seem like a serious cost of accepting the Simple View, a cost that adherents of the Complex View don't have to pay.

### *3.4. Summary*

Like Nagel and Blackburn and others, I find something appealing about the Simple View for it takes seriously how things seem to me from the first-person point of view. When I imagine cases from the inside, it seems that facts about my persistence are not fixed by facts about physical and psychological continuity.

But there are two things I don't like about the Simple View: first, it conflicts with physicalism, the attractive idea that everything supervenes on the physical, and, second, it has the consequence that I could know all

the relevant physical and psychological continuity facts in a given situation but still be in the dark about the persistence facts.

I'd like it if there was a view that: (i) respected my judgments about how things look from the first-person point of view, but (ii) didn't conflict with physicalism, and (iii) didn't posit hidden facts. That is, I'd like a view that allowed me to avail of the benefits of the Simple View without incurring its main costs. The rest of this essay is devoted to developing such a view.

#### **4. Imagination and Possibility**

The main conflict we've been discussing is over the number of possibilities for the persistence facts compatible with a given specification of the continuity/physical facts. From the inside, it seems that there are multiple possibilities for my persistence compatible with a given specification of the continuity facts and also multiple possibilities for my persistence compatible with a given specification of the physical facts. But physicalism requires there to be only one possibility for my persistence with a given specification of the physical facts. And unless there is only one possibility for the persistence facts compatible with a given specification of the continuity facts, it is hard to see how hidden persistence facts can be avoided, given that the continuity facts constitute the only evidence relevant to determining the persistence facts.

So far we've been assuming that all the possibilities in question—including the ones we see from the inside—are possible worlds. But I want to show that there is some motivation for questioning this assumption; in particular, there is some motivation for questioning the assumption that the possibilities we see from the inside are possible worlds. But if they're not possible worlds, what are they? The answer to that question emerges rather naturally once we get clearer on what it is to imagine something from the inside.

##### *4.1. Imagination*

The distinction between imagining from the inside and imagining from the outside is often made in the philosophical literature, though it's not entirely clear whether or not everyone has precisely the same distinction in mind.<sup>13</sup> The distinction we're interested in is between imagining a

13. The distinction (or something like it) appears, in different guises, in a variety of philosophical discussions. See, for example, Williams [1966] 1973; Nagel 1974, n. 1;

hypothetical case in which some events befall some individuals, any of whom may or may not be identical to you or me, on the one hand, and imagining *being* one of the participants in a hypothetical case, on the other (as when one imagines *undergoing* fission). The distinction is rough but clear enough for present purposes.

I want to suggest that it is useful to think of this distinction as a distinction between two kinds of content. To see what I have in mind, think first about imagining from the outside. Pretend for a moment that I have superpowers of imagining and that I can imagine a scenario in *maximally specific* detail—that is, suppose I can imagine a whole world, fixing every last detail of the imagined world. Suppose I imagine from the outside a situation in which the New England Patriots win the Super Bowl over the New York Giants, and I imaginatively fix every possible detail of this situation. How could we characterize the *content* of this imagining? Since the content of my imagining specifies in complete detail a certain possible situation, a way things could be, it is natural to represent it by a certain possible world: the possible world that is exactly as my imagining specifies the world to be (there is only one, since my imagining is maximally specific). Call that world ‘*w*’: *w* represents the content of my imagining from the outside that the Patriots win the Super Bowl.

Continue to pretend I have superpowers of imagining. Suppose now I undertake a different imaginative task: I imagine *exactly* the same scenario described above—the very same game, the very same world—but I imagine that scenario from the point of view of one of its participants. Suppose, for example, that I imagine being Tom Brady, the quarterback of the New England Patriots, in the scenario described above. I imagine completing touchdown passes, leaving the field with victory in hand, and falling into the arms of my supermodel girlfriend.

How could we represent the content of *this* imagining? Here’s an idea: represent the content as the *pair* consisting of the possible world *w* and Tom Brady, that is  $\langle w, \text{Tom Brady} \rangle$ . That *w* is the first member of the pair represents the fact that the *objective scenario* I’m imagining corresponds to the possible world *w*, and so is the very same objective scenario I earlier imagined from the outside. That Tom Brady is the second member of the pair represents the fact that I am looking at *w* from Brady’s point of view: I am imagining being Tom Brady in that objective scenario.

---

Wiggins 1976, 143; Wiggins 1980, 153–54; Peacocke 1985; Reynolds 1989; Walton 1990, 28–35; Shoemaker 1994; Velleman 1996; Hill 1997; Martin 2002; Higginbotham 2003; Recanati n.d.; Recanati 2007; and Ninan 2008.

Note that I could imagine the very same scenario from another point of view: I could imagine being Eli Manning, the quarterback for the New York Giants. I imagine being sacked by the Patriots' front line, losing the game, and then being pilloried in the *New York Post*. The content of this imagining could be represented as  $\langle w, \text{Eli Manning} \rangle$ , representing the fact that I am again imagining the very same objective scenario ( $w$ ) as I did in the two imaginings described above, but this time I am viewing things from Manning's point of view.

So this account enables us to represent the fact that in all three cases I imagine the same objective scenario: in the first case, the content of the imagining is simply the possible world  $w$ , reflecting the fact that I'm imagining that scenario from the outside and so from no one's point of view; in the second case, the content of my imagining is  $\langle w, \text{Tom Brady} \rangle$ , representing the fact that I'm viewing that same scenario from Brady's point of view; in the third case, the content is  $\langle w, \text{Eli Manning} \rangle$ , representing the fact that I'm viewing that scenario from Manning's point of view.

A note on terminology: a pair consisting of a possible world and one of its inhabitants is a *centered world*. I'll have more to say about centered worlds shortly, but first let me extend the present account to deal with the fact that I don't have superpowers of imagining. Since my imaginative abilities are limited, I never imagine a possible situation in full detail; my imaginings are never maximally specific. When I imagine the Patriots' winning the Super Bowl, my imagining is silent on ever so many details. I do not imaginatively "settle" whether the game is being played on a Friday or on a Sunday, nor whether there is an even or an odd number of people in the crowd, nor whether the Patriots' coach is wearing a wristwatch or not, and so on. We can capture the lack of specificity of my imagining by saying that its content is a *set* of possible worlds, rather than a single world. All the worlds in the set (the set of worlds compatible with what I imagine) agree on certain details—for example, all are worlds in which the Patriots win the Super Bowl over the Giants. But they will disagree on details that my imagining doesn't settle—in the present case, some of the worlds in the set are worlds in which the game is played on a Friday, some are ones in which it is played on a Sunday, some are ones in which there is an even number of people in the crowd, some are ones in which there is an odd number of people in the crowd, and so on. Similar considerations suggest we should treat the content of an imagining from the inside as a set of centered worlds, rather than as a single centered world.

So our present account says that the content of an imagining from the outside is a set of possible worlds, or a *possible worlds proposition*, while

the content of an imagining from the inside is a set of centered worlds, or a *centered proposition*. Centered propositions (or *properties*) are, of course, what David Lewis ([1979] 1983) uses in his account of *de se* and *de dicto* attitudes. A *de se* thought is an ‘I’-thought, a thought about oneself when one thinks of oneself in the first-person way; a thought that is not *de se* is *de dicto*. It is not surprising that the inside-outside distinction should have some close connection to the *de se-de dicto* distinction since both are distinctions between first- and third-person mental states. Focusing on the attitudes of belief and desire, Lewis argued that the content of a *de se* attitude should be represented as a centered proposition since it cannot be represented as a possible worlds proposition. I won’t rehearse Lewis’s arguments for this here; instead, I want to show how recent work on the semantics of *de se* attitude ascriptions provides further motivation for our account of imagination. But before I do that, let me make two further remarks about Lewis’s theory of the *de se*.

First, I should note that, as I use the term ‘centered world’, centered worlds can be centered on things that persist over time. Lewis, on the other hand, took centered worlds to be centered on *stages* and also took stages to be the bearers of psychological attitudes (see Lewis [1979] 1983, 143–44 and the postscript to Lewis 1976 in Lewis 1983b). This decision was tied to his treatment of *de nunc* attitudes, or attitudes about what time it is. We could instead treat *de nunc* attitudes by taking centered worlds to be world-time-individual triples (Egan 2004, 64–65). That’s what we would do if we were interested in theorizing about the *de nunc*; since we’re not, we will omit the temporal coordinate and continue to take centered worlds to be world-individual pairs.

Second, a remark about possible worlds propositions. One of the central themes of Lewis [1979] 1983 is that we can use sets of centered worlds to do all the work sets of possible worlds can do and more (though Lewis [1995] 1999, 320, later refers to this as a “cheap trick.”) Any content that could be characterized as a set of possible worlds  $\{w : \phi(w)\}$  could instead be characterized as the set of centered worlds  $\{\langle w, x \rangle : \phi(w)\}$ . The distinction between *de se* and *de dicto* contents can then be made in terms of *types* of sets of centered worlds. Following Egan 2006, 107, we can say that a centered proposition  $p$  is *boring* iff for all worlds  $w$  and inhabitants  $x$  and  $y$  of  $w$ ,  $\langle w, x \rangle$  is in  $p$  iff  $\langle w, y \rangle$  is too. Since a boring centered proposition doesn’t distinguish between worldmates, it is essentially equivalent to a set of possible worlds, and so a *de dicto* content can be represented by a

boring centered proposition. An *interesting* centered proposition is one that isn't boring; a *de se* content is an interesting centered proposition.<sup>14</sup>

We initially stated our account of imagining in terms of possible worlds versus centered propositions, but we could instead follow Lewis and formulate it in terms of interesting versus boring centered propositions. Or we could remain neutral on the issue, by formulating our account of the inside-outside distinction as follows:

#### **Uncentered Imagination**

The content of an imagining from the outside is an uncentered content, something that determines a possible worlds proposition.

#### **Centered Imagination**

The content of an imagining from the inside is a centered content, something that determines an interesting centered proposition.<sup>15</sup>

A boring centered proposition  $\{(w, x) : \phi(w)\}$  determines the possible worlds proposition  $\{w : \phi(w)\}$ , so if the content of an imagining from the outside is a boring centered proposition, **Uncentered Imagination** will be true. The important point is that the content of an imagining from the outside is something that places a condition on a possible world, whereas the content of an imagining from the inside is something that only places a condition on a centered world. A boring centered proposition  $\{(w, x) : \phi(w)\}$  places a condition on a possible world that is satisfied just in case  $\phi(w)$ . But it makes no sense to ask what condition an interesting

14. Lewis's claim that all attitude contents can be characterized as sets of centered worlds has been challenged recently by Nolan (2006), who points out that Lewis's official account is incompatible with the existence of certain *de se* desires (for example, the wish that one had never existed). A simple (if somewhat artificial) patch is available if we follow Lewis [1983a] 1999, 398, and say that the domain of each possible world contains a "null individual," denoted by ' $\star$ '. The content of the wish that one had never existed can then be identified with the set of centered worlds centered on the null individual. Alternatively, one could adopt a metaphysical view like the one defended in Williamson 2002 and then claim that the wish that one had never existed is really the wish that one had been nonconcrete, in which case the content of the wish will be the set of centered worlds in which the center is nonconcrete (Turner forthcoming, n. 4).

15. Another advantage of this way of formulating our account of the inside-outside distinction is that it could be endorsed by those who reject coarse-grained propositions in favor of some type of structured entity, so long as they were willing to think of the content of an imagining from the outside as something that determines a possible worlds proposition (such as a structured proposition) and of the content of an imagining from the inside as something that determines an interesting centered proposition (such as a structured property) (see Lewis [1979] 1983, 150).



centered proposition places on a possible world, for an interesting centered proposition cannot properly be understood to be true or false at a *possible* world; it is true or false only relative to a *centered* world.

Semantic theories of *de se* ascription typically understand the *de se de dicto* distinction in terms of the boring-interesting centered proposition distinction. So our linguistic argument for **Uncentered** and **Centered Imagination** is, in the first instance, an argument for the view that the content of an imagining from the outside is a boring centered proposition and the content of an imagining from the inside is an interesting centered proposition.

The linguistic argument begins with an observation about the data. When philosophers write about the distinction between imagining from the inside and imagining from the outside, they often emphasize that it is the distinction between imagining *doing* or *experiencing* something versus imagining *that* something is the case. For example, consider this passage from Kendall Walton (1990, 29):

Imagining from the inside is . . . a form of self-imagining characteristically described as imagining *doing* or *experiencing* something (or *being* a certain way), as opposed to merely imagining *that* one does or experiences something or possesses a certain property.

Suppose I imagine from the inside playing for the Patriots and winning the Super Bowl. Walton's observation suggests that I might use (1) to report my imagining:

1. I imagined winning the Super Bowl.

If I imagined from the inside winning the Super Bowl, then (1) (as uttered by me) would be true. As I will presently argue, an independently motivated semantic theory tells us that if (1) (as uttered by me) is true, then the content of my imagining is the interesting centered proposition in which the center wins the Super Bowl. So putting these two thoughts together, we get: if I imagine from the inside winning the Super Bowl, then the content of my imagining is the interesting centered proposition in which the center wins the Super Bowl, which is an instance of **Centered Imagination**.

To argue that the truth of (1) entails that the content of my imagining is the interesting centered proposition in which the center wins the Super Bowl, we start with a claim about the syntactic structure of (1). According to standard syntactic theory, the subject of (1)'s lower clause is

subject-control 'PRO', a phonologically null pronoun.<sup>16</sup> So the real structure of (1) is:

I imagined PRO winning the Super Bowl.

What is the correct interpretation of 'PRO'? An observation due to Morgan 1970 is that subject-control 'PRO' gives rise to unambiguously *de se* readings. To see this, first consider these two hypothetical scenarios:

#### **Election 1**

John is drinking whisky and watching TV. He is watching the speeches of various candidates in the upcoming election. He is impressed by one candidate in particular, a wise-looking bearded man. John comes to think that the bearded candidate will win. The bearded candidate is none other than John himself, but because he's so intoxicated, John fails to realize this. In fact, he is rather pessimistic about his own prospects and thinks to himself, *I'm not going to win.*

#### **Election 2**

While contemplating his prospects in the upcoming election, John (sober this time) thinks to himself, *I'm sure to win.*

In **Election 2**, John has a *de se* expectation that he is going to win; in **Election 1**, he doesn't.

Now consider this sentence:

2. John expects (PRO) to win.

We note that the subject of the embedded clause of (2) is again subject-control 'PRO', at least according to standard syntactic theories. Note also that (2) is true in **Election 2** but false in **Election 1**. This suggests that (2) must be read *de se*: it is only true if John thinks to himself, *I'm going to win.* It's not enough for him simply to be the *F* and to expect that the *F* is going to win. (Note the contrast with the finite clause ascription 'John expects that he will win', which is arguably true in both **Election 1** and **2**.)

There are a few different semantic theories based on Lewis's account of the *de se* that are designed to account for the fact that (2) can be used only to report a *de se* expectation. All the theories in question assume that 'expects' expresses quantification over the centered worlds compatible with what the subject expects, and all treat the complement

16. Not every gerundive phrase has 'PRO' as its subject; but standard diagnostics (for example, idiom and dummy-'it' tests) reveal that 'imagines' is a subject-control rather than a raising verb.

clause ‘PRO to win’ as expressing an interesting centered proposition. Consider, for example, the approach to these issues taken in Anand and Nevins 2004 and von Stechow 2005.<sup>17</sup> In these systems, the semantic value of an expression is given relative to a context and an index. Ignoring tense and temporal expressions, we can take a context  $c$  to be a world-individual pair  $\langle w_c, x_c \rangle$ , and we can also take an index  $i$  to be a world-individual pair  $\langle w_i, x_i \rangle$ . The crucial feature of the semantics is the semantic value of ‘PRO’, which we take to be the individual coordinate of the index:

$$[[\text{‘PRO’}]]^{c,i} = x_i.$$

(The double brackets ‘[[ ]]’ denote a three-place interpretation function that takes expression-context-index triples to extensions.)

With these assumptions in place, we have the following semantic value (intension) for ‘PRO to win’:

$$\begin{aligned} \lambda i. [[\text{‘PRO to win’}]]^{c,i} \\ = \lambda i. x_i \text{ wins in } w_i. \end{aligned}$$

This is a function from centered worlds to truth values, the characteristic function of a set of centered worlds. Such functions are a notationally equivalent way of representing centered propositions. And in this particular case, the function represents an *interesting* centered proposition, a *de se* content. So (2) is true iff the content of John’s expectation is the interesting centered proposition in which the center wins, that is, iff John expects *de se* that he will win. Since John lacks this *de se* expectation in **Election 1**, the sentence is correctly predicted to be false in that situation. Since he has this *de se* expectation in **Election 2**, the sentence is correctly predicted to be true in that situation.

That’s the motivation for our semantic theory. Note what it tells us about (1), repeated here:

1. I imagined (PRO) winning the Super Bowl.

The embedded clause of (1) is ‘PRO winning the Super Bowl’, and so the intension of this clause will be an interesting centered proposition:

$$\begin{aligned} \lambda i. [[\text{‘PRO winning the Super Bowl’}]]^{c,i} \\ = \lambda i. x_i \text{ wins the Super Bowl in } w_i. \end{aligned}$$

17. See Schlenker 2003; von Stechow n.d.; and von Stechow 2003 for a related but slightly different take on these issues.

So our semantic theory entails that (1) (as uttered by me) is true iff the content of my imagining is the interesting centered proposition in which the center wins the Super Bowl. Our observation about the data was that if I imagine from the inside winning the Super Bowl, then (1) (as uttered by me) is true. Putting these claims together gives us the result that if I imagine from the inside winning the Super Bowl, then the content of my imagining is the interesting centered proposition in which the center wins the Super Bowl, an instance of **Centered Imagination**.

What about **Uncentered Imagination**? Suppose I imagine from the outside that the Patriots win the Super Bowl. How would I report that imagining? By saying something like (3):

3. I imagined that the Patriots won the Super Bowl.

If I imagined from the outside that the Patriots won, (3) would be true. But notice also that if I imagine from the inside being a member of the Patriots and winning the Super Bowl, (3) would still be true (though perhaps misleading because I could have said something stronger). All the truth of (3) requires is that the objective scenario I imagine be one in which the Patriots win the Super Bowl. But it apparently places no restriction on *how* I imagine that objective scenario.

This data can be predicted if we assume **Uncentered** and **Centered Imagination** and the semantic theory we've been developing. On our semantic theory, the intension of the complement clause will simply be the boring centered proposition in which the Patriots win the Super Bowl:

$$\begin{aligned} & \lambda i. [[\text{'the Patriots won the Super Bowl'}]]^{c,i} \\ & = \lambda i. \text{the Patriots won the Super Bowl in } w_i. \end{aligned}$$

Given the appropriate lexical entry for 'imagines', (3) will be true just in case all the centered worlds  $\langle w, x \rangle$  compatible with what I imagine are such that the Patriots win the Super Bowl in  $w$ .

This truth condition doesn't specify whether the content of my imagining is a boring or an interesting centered proposition: it simply says all the centered worlds  $\langle w, x \rangle$  compatible with what I imagine must be such that the Patriots win the Super Bowl in  $w$ . According to our account, if I imagine from the outside that the Patriots win the Super Bowl, then the content of my imagining is a boring centered proposition in which the Patriots win the Super Bowl. But then all the centered worlds  $\langle w, x \rangle$  compatible with what I imagine will be such that the Patriots win the Super Bowl in  $w$ , and so the truth condition of (3) will be satisfied. And according to

our account, if I imagine from the inside playing for the Patriots and winning the Super Bowl, then the content of my imagining will be an interesting centered proposition in which the center plays for the Patriots and wins the Super Bowl. But then it will still be true that every centered world  $\langle w, x \rangle$  compatible with what I imagine will be such that the Patriots win the Super Bowl in  $w$ , and so the truth condition of (3) will again be satisfied. So, together with our semantics, our account of the inside-outside distinction predicts that (3) can be made true either by imagining from the outside that the Patriots win the Super Bowl or by imagining from the inside playing for the Patriots and winning the Super Bowl. This is the result we want since it seems intuitively correct that the sentence can be made true in either of these two ways.<sup>18</sup>

That completes my case for **Centered** and **Uncentered Imagination**.

#### 4.2. Possibility

While there is much discussion in the philosophical literature about the relationship between imaginability and possibility,<sup>19</sup> there is relatively little on what (if any) the metaphysical consequences of the inside-outside distinction are.<sup>20</sup> But given our account of this distinction, an intriguing idea is that each type of imagining is a guide to a different kind of

18. One might worry whether a parallel argument would show that our semantics predicts that (1) could be true while the content of my imagining is a boring centered proposition, so long as every centered world compatible with what I imagine is contained in the interesting centered proposition in which the center wins the Super Bowl (that is, so long as the content of my imagining is a *subset* of the interesting centered proposition in which the center wins the Super Bowl). But no such argument can be run, for while some interesting centered propositions do have (nonempty) boring subsets, the centered proposition in which the center wins the Super Bowl is *superinteresting*: it is interesting and has no (nonempty) boring subsets. The proof of this invokes the null individual  $\star$  discussed in footnote 14. Take any arbitrary (nonempty) subset  $q$  of  $p = \{\langle w, x \rangle : x \text{ wins the Super Bowl in } w\}$ . To show that  $q$  is interesting we need to show that there is a world  $w$  with inhabitants  $x$  and  $y$  such that  $\langle w, x \rangle$  is in  $q$  and  $\langle w, y \rangle$  is not in  $q$ . Let  $\langle w', x' \rangle$  be any centered world in  $q$ . Recall that  $\star$  is in the domain of each world, and so is in the domain of  $w'$ . But  $\langle w', \star \rangle$  is not in  $q$ , because  $\star$  does not win the Super Bowl in any world, and so does not win it in  $w'$ . So  $q$  is interesting. Since  $q$  was an arbitrary subset of  $p$ , it follows that every subset of  $p$  is interesting.

19. For some recent discussion, see the articles in Szabó Gendler and Hawthorne 2002.

20. Though Wiggins (1980, 153–54) asks, “Is my imagining being Moses, or an elephant, equivalent to my imagining (the impossible state of affairs of) *my* being Moses or an elephant?” Assuming his answer is “no,” Wiggins seems to be suggesting that when I imagine being Moses or an elephant, I am imagining something possible.

possibility. Perhaps just as there are centered and uncentered *contents*, there are centered and uncentered *possibilities*.

Suppose that, normally, imagining  $p$  provides evidence that  $p$  is possible, no matter what kind of content  $p$  is. When I imagine  $p$  from the third-person perspective,  $p$  is an uncentered content, something that has a truth value at a possible world. So one might take this imagining as evidence that there is a *possible world* accessible from the actual world at which  $p$  is true. Here,  $p$  is an uncentered possibility, a possible way for *the world* to be. But when I imagine  $p$  from the first-person perspective,  $p$  is a centered content, something that has a truth value only at a centered world. So one might take this imagining as evidence that there is a *centered world* accessible from  $\langle$ actual world, me $\rangle$  at which  $p$  is true.<sup>21</sup> Here,  $p$  is a centered possibility, a possible way for *me* to be. From the outside, I see the ways the world could be; from the inside, I see the ways I could be.

Following Lewis [1983a] 1999; 1986, §4.4, let us suppose that uncentered and centered possibilities are both genuine species of possibility.<sup>22</sup> But adding to Lewis, let us suppose that outside imagining is a guide to uncentered possibility and that inside imagining is a guide to centered possibility. Here is the basic idea:

#### **Uncentered Guide**

Imagining from the outside is a guide to uncentered possibility. If I can imagine an uncentered content  $p$  and am inclined to judge that  $p$  is possible on that basis, then I have evidence that there is a possible world  $w$  accessible from the actual world such that  $p$  is true at  $w$ .<sup>23</sup>

#### **Centered Guide**

Imagining from the inside is a guide to centered possibility. If I can imagine a centered content  $p$  and am inclined to judge that  $p$  is possible on that basis, then I have evidence that there is a centered world  $\langle w, x \rangle$  accessible from  $\langle$ actual world, me $\rangle$  such that  $p$  is true at  $\langle w, x \rangle$ .

21. The accessibility relation is just a place-holder at the moment; we'll discuss what it means shortly.

22. See also Hazen 1979.

23. Many philosophers believe: (i) that every possible world is metaphysically accessible from every other possible world, and (ii) that 'possible world' and 'metaphysically possible world' are coextensive. As a result, these philosophers do not need to mention accessibility relations in their analyses of modal claims. But some theorists reject (i) and/or (ii), in which case the reference to a metaphysical accessibility relation is not redundant. See, for example, Salmon 1998 and the references therein.

Now this idea is interesting only if these two types of possibility can come apart, if some claims of the form ‘I could have been *F*’ are true in one sense but not in the other. One type of possibility claim that could be analyzed as a true centered possibility claim but that would be false when analyzed as an uncentered possibility claim would be a claim like ‘I could have been Fred’, as uttered by someone other than Fred. Here’s how we would analyze that claim as a centered possibility claim:

‘I could have been Fred’ is true as uttered by  $x_c$  in  $w_c$  iff there is a centered world  $\langle w, x \rangle$  accessible from  $\langle w_c, x_c \rangle$  such that  $x$  is Fred in  $w$ .<sup>24</sup>

On the standard possible worlds analysis, that claim would be false when uttered by anyone other than Fred. Given that I am not Fred and given the necessity of distinctness, there is no possible world in which I am Fred, and so, in the uncentered sense, I could not have been Fred.

But as several philosophers have noted, the idea that (as each of us would put it) I could have been someone else is not without intuitive support:

My being TN (or whoever in fact I am) seems accidental. . . . So far as what I am essentially is concerned, it seems as if I just *happen* to be the publicly identifiable person TN—as if what I really am, this conscious subject, might just as well view the world from the perspective of a different person. (Nagel 1986, 60)

Here am I, there goes poor Fred; there but for the grace of God go I; how lucky I am to be me, not him. Where there is luck there must be contingency. I am contemplating the possibility of my being poor Fred, and rejoicing that it is unrealized. (Lewis [1983a] 1999, 395)

‘I might have been somebody else’ is a very primitive and very real thought; and it tends to carry with it an idea that one knows what it would be like for this ‘I’ to look out on a different world, from a different body, and still be the same ‘I’. (Williams [1966] 1973, 40)

24. Our apparatus is consistent with but does not require the thought that centered possibility is an essentially *first-personal* kind of possibility. That is, the above analysis of centered possibility is compatible with but does not require the claim that centered possibilities are only expressible with the first-person singular pronoun ‘I’. This can be seen by looking at how we could analyze ‘Dilip could have been Fred’ as a centered possibility claim:

‘Dilip could have been Fred’ is true at  $w_c$  iff there is a centered world  $\langle w, x \rangle$  accessible from  $\langle w_c, \text{Dilip} \rangle$  such that  $x$  is Fred in  $w$ .

See the the appendix to Ninan 2008 for some discussion of the semantics of centered possibility claims.

Now you don't have to accept this idea in order to accept our picture of imagination and possibility. But if you're inclined to accept it—or at least to think it's coherent—then our account gives you a way of making sense of this thought. But the important point here is simply that centered and uncentered possibility might come apart in interesting ways.

One worry about our account of centered possibility is that it does not explain the notion of centered possibility in independently understood terms. What is the relevant accessibility relation here? Metaphysical accessibility, which relates a possible individual to each of the possible individuals he or she could have been? If that is the only answer we can give, then the analysis is to some degree circular. But in this, the centered worlds analysis is not much different from the possible worlds analysis. Possible worlds semantics analyzes:

It is possible that Aristotle could have died as a child

as:

There is a possible world (accessible from the actual world) in which Aristotle died as a child.

But if you ask most advocates of possible worlds what a possible world is, you will be told that it is a “way things could be” or a maximally specific property the universe could instantiate. But this just explains the notion of a possible world in terms of our original modal idiom. In neither case—centered or uncentered—does the analysis in question explain the relevant modal claims in independently understood terms. But the analyses do help to clarify the structure of the relevant modal claims.<sup>25</sup>

A different sort of question about the accessibility relation concerns not its *analysis* but its *extension*: which centered worlds are accessible from me? Could I have had origins different from the ones I in fact have? Could I have been Napoleon? Could I have been a poached egg? I do not know the answer to these questions. But I don't think that our entitlement to appeal to the notion of centered possibility stands or falls with our ability to precisely delineate the extension of the accessibility relation. Again, it is instructive to compare this situation with the case of possible worlds. Few philosophers would think they know exactly what the space of possible worlds is like. Are there possible worlds where Adam has all the qualitative properties Noah in fact has, and vice versa? Are there possible

25. At any rate, this is the line most actualists take. Lewis, on the other hand, aims to reduce the modal to the nonmodal, and so would reject this as a characterization of his possible worlds analysis of modality. I side with the actualists here.



worlds physically identical to this one, but in which no one is conscious? Are there possible worlds broadly similar to this one but in which backward causation takes place?

I do not know the answers to these questions. For some of them, I do not even know how to go about answering them. Despite that, I think I do know certain possible worlds claims: there is a possible world in which my sister is a schoolteacher, there are no possible worlds in which Barack Obama is a nonhuman robot, and no two possible worlds have exactly the same physical facts but different moral facts (the moral supervenes on the physical). The legitimacy of using the possible worlds apparatus to elucidate *particular* modal claims that one accepts doesn't require one to answer every question of the form, "But is there a possible world in which such-and-such happens?" Similarly, unless one has an antecedent reason to be skeptical of the notion of centered possibility, it seems legitimate to use the centered possibility apparatus to elucidate particular centered possibility claims even if one cannot answer every question of the form, "Is there a centered world accessible from you centered on a being like *this*?"

But let me say this: if we think of imagining from the inside as our basic way of representing centered possibilities, then it is natural to assume that accessibility is a relation between centered worlds that are centered on things that have *perspectives* or *points of view*.<sup>26</sup> This might be a necessary condition on the accessibility relation: any centered world  $\langle w, x \rangle$  accessible from me must be such that  $x$  has a "perspective" in  $w$  (whatever that amounts to). Whether this is also a sufficient condition is a question we can leave unanswered for present purposes. This condition is, of course, seriously underspecified insofar as we lack an account of what it takes to have a perspective. But, at the very least, this view would seem to rule out the possibility that there are centered worlds accessible from me that are centered on poached eggs or bedposts. This view of accessibility thus differs from Lewis's construal, according to which every possible object is accessible from every other (Lewis 1986, 239–43).

Another question about this proposal is whether centered possibilities are meant to be *conceptual* or *metaphysical* possibilities. This is a much larger question than I can hope to answer here, not least because there is no consensus on how to draw the conceptual-metaphysical distinction, nor on whether there is even any distinction here to be drawn. My primary claim is that the possibility judgments we arrive at when we imagine from the inside can and should be analyzed in terms of quantification over

26. Thanks to Seth Yalcin for discussion here.

centered worlds. I'm less concerned with the question of whether or not these possibility judgments ought to be labeled as 'conceptual' or 'metaphysical', in part because I'm uncertain of what exactly these labels mean.

However, if one adopts one particularly sophisticated account of the conceptual-metaphysical distinction—*epistemic two-dimensionalism* (Chalmers 1996, 2002a, 2002b; Jackson 1998)—then it would be natural to treat centered possibilities as conceptual possibilities and uncentered possibilities as metaphysical possibilities. For two-dimensionalists analyze conceptual possibility in terms of centered worlds and metaphysical possibility in terms of possible worlds. Readers sympathetic to two-dimensionalism may wish to adopt this interpretation of centered and uncentered possibility. But note that our proposal doesn't *require* this two-dimensional interpretation, nor does Lewis arrive at his proposal via two-dimensionalism.

There may be other worries about the idea of centered possibility beyond those discussed above. But I think it is legitimate to postpone questions about the ultimate tenability of this proposal, at least for the moment. I think this is legitimate because whether or not we should adopt this proposal depends in part on whether doing so would make other philosophical problems more tractable. And in the remainder of the essay I shall argue that this hypothesis helps us to reconcile claims about our persistence that are compelling from the first-person point of view with the principal objections to the Simple View.

## 5. The Centered View

Let us call the view we are developing 'the Centered View'. We now have in place one of the major components of the Centered View: the account of imagination and possibility just discussed. As I shall develop it here, the second main component of the Centered View is a suitably generous ontology of persistence. Once this second component is in place, I'll show how the Centered View helps us to reconcile our first-person judgments about our persistence with physicalism and with a ban on hidden persistence facts.

While I think the most natural way to develop the Centered View involves using a generous ontology of persistence, this commitment might not be strictly necessary. So after showing how to formulate the Centered View using a generous ontology, I'll say something briefly about the possibility of formulating it with less expansive metaphysical resources.

5.1. *Ontology*

In this subsection, I'll describe the ontology of persistence that we will use in formulating the Centered View; in the next subsection, I'll demonstrate what role that ontology plays in our resolution of the dilemma.

Let me begin by looking at what four-dimensionalism says about the ontology of the fission case. According to four-dimensionalism, in any fission case there are three salient entities, all of which completely spatially coincide prior to fission. That is, there are three entities that share the same birth-to-fission temporal part. One of these three entities survives fission and ends up with the left hemisphere; another survives fission and ends up with the right hemisphere; the third fails to survive fission (so this third entity is a proper temporal part of the other two).

As we're developing it, the Centered View uses an ontology that countenances all three of these entities. So the Centered View could be developed using the resources of four-dimensionalism. But it doesn't require an ontology of persistence as abundant as four-dimensionalism.<sup>27</sup> The Centered View could also be developed using a version of three-dimensionalism like the one Sider calls 'promiscuous endurantism', which says:

In the vicinity of every person [there is] a plurality of coincident entities, which share the same momentary properties but differ in their persistence conditions. In my vicinity, there is a psychological-person, a body-person, and perhaps other entities corresponding to other criteria of personal identity. (Sider 2001a, 193)

Promiscuous endurantism is the view that, for every plausible account of our persistence conditions, there is, in my vicinity, a being whose persistence conditions are described by that account. For every continuity relation  $R$  (bodily continuity, nonbranching psychological continuity, branching psychological continuity, and so forth), there is in my vicinity a being whose stages are maximally  $R$ -interrelated.<sup>28</sup> (By 'in my vicinity', I mean that all of these beings presently completely spatially coincide with me.) Like four-dimensionalism, promiscuous endurantism holds that three distinct entities coincide prior to fission, two surviving, one expiring: the two survivors' stages are maximally interrelated by the branching psychological continuity relation, and the third individual's stages are maximally

27. Thanks to Caspar Hare for helping me to see this point.

28. An individual  $x$ 's stages are *maximally interrelated* by a relation  $R$  iff any two of  $x$ 's stages are  $R$ -related, and there is no stage  $y$  that is  $R$ -related to one of  $x$ 's stages and yet not itself a stage of  $x$ .

interrelated both by the nonbranching psychological continuity relation and by the bodily continuity relation.

### 5.2. *Supervenience and Physicalism*

As we're presently formulating it, the Centered View combines the account of imagination and possibility developed in §4 with a generous ontology of persistence. It now remains to be shown how the Centered View reconciles our first-person judgments about our persistence with physicalism and a ban on hidden persistence facts. We start with physicalism.

The argument that our first-person judgments require us to give up physicalism is very similar to the reasoning that led us from Blackburn's observation about the fission case to the denial of **Supervenience**, and we shall discuss that piece of reasoning first. Blackburn's observation can be put by saying that, when I imagine fission from the inside, the following possibility claims all appear to be true:

4. (a) I could undergo fission and survive with the left hemisphere.
- (b) I could undergo fission and survive with the right hemisphere.
- (c) I could undergo fission and fail to survive.

We interpreted these claims as possible worlds claims; we said that these claims could all be true only if the corresponding claims in (5) were true:

5. (a) There is a possible world in which I undergo fission and survive with the left hemisphere.
- (b) There is a possible world in which I undergo fission and survive with the right hemisphere.
- (c) There is a possible world in which I undergo fission and fail to survive.

If (5a) is true, then there is a possible world  $w$  in which I exist such that: if  $y$  is a prefission stage of me in  $w$ , and  $z$  is a postfission stage of the individual with the left hemisphere in  $w$ , then  $z$  is a stage of me in  $w$ . And if (5b) is true, then there is a possible world  $w'$  in which I exist such that: if  $y'$  is a prefission stage of me in  $w'$ , and  $z'$  is a postfission stage of the individual with the left hemisphere in  $w'$ , then  $z'$  is not a stage of me in  $w'$ . Since  $y$  and  $z$  in  $w$  and  $y'$  and  $z'$  in  $w'$  are the same with respect to continuity, this gives us a counterexample to **Supervenience**.

That was the reasoning that led us from the joint truth of the claims in (4) to the falsity of **Supervenience**. But in light of our new account of imagination and possibility, it should be clear where this reasoning goes

wrong. Blackburn's observation is that the claims in (4) look true when we imagine fission from the inside. But then **Centered Guide** tells us that what we ought to conclude is that the claims in (4) are true on their *centered* readings:

6. (a) There is a centered world  $\langle w, x \rangle$  accessible from  $\langle$ actual world, me $\rangle$  such that  $x$  undergoes fission and survives with the left hemisphere in  $w$ .
- (b) There is a centered world  $\langle w, x \rangle$  accessible from  $\langle$ actual world, me $\rangle$  such that  $x$  undergoes fission and survives with the right hemisphere in  $w$ .
- (c) There is a centered world  $\langle w, x \rangle$  accessible from  $\langle$ actual world, me $\rangle$  such that  $x$  undergoes fission and fails to survive in  $w$ .

But if what Blackburn's observation really supports is the joint truth of the claims in (6), then the argument against **Supervenience** cannot be run. For that argument depended on their being possible worlds  $w$  and  $w'$  that share the same local continuity facts but differ on the facts concerning my persistence. But no such argument can be run using the centered possibility claims in (6) since those claims simply say nothing about *possible worlds* at which I exist; they make claims only about what sorts of centered worlds are accessible from me. So according to the Centered View, Blackburn's observation is actually compatible with **Supervenience**. The first-person considerations we took to support the Simple View do not, in light of **Centered Guide**, in fact support that view.

Note the role that our generous ontology of persistence is playing here. One of our aims is to vindicate our first-person judgments about the fission case, which is to vindicate the claims in (4). We are attempting to do this by showing that they are true on their centered readings, which are given in (6). But in order for the claims in (6) to be true, there have to be three centered worlds meeting these conditions: there has to be a centered world centered on something that undergoes fission and survives with the left hemisphere (in order for (6a) to be true), there has to be one centered on something that undergoes fission and survives with the right hemisphere (for (6b) to be true), and there has to be one centered on something that undergoes fission and fails to survive (for (6c) to be true). Since a centered world is just a pair of a possible world and an object in the domain of that world, this means that there must be, somewhere in the space of possible objects, an object that undergoes fission and goes left, and one that undergoes fission and goes right, and one that undergoes fission and goes nowhere. This requirement is most

naturally met by adopting a suitably generous ontology. We need the objects provided by a generous ontology in order to center our centered worlds on them.

Having seen how the Centered View can be developed in the context of a generous ontology of persistence, let me say something about how this commitment might be avoided by those who favor less abundant ontologies. Suppose you think that there are no coincident entities in a fission case: when someone undergoes fission, he or she ceases to exist and two new individuals come into existence. Still, we may be able to analyze claims like (4a) ('I could undergo fission and survive with the left hemisphere') as true centered possibility claims if we alter our conception of what a centered world is. Suppose that instead of defining a centered world as a pair consisting of a possible world and an object in that world, we define a centered world as a pair of a possible world  $w$  and a (possibly partial) *function* from times to objects in  $w$ . We could then say that the possibility claim (4a) is true on its centered reading iff there is a centered world  $\langle w, f \rangle$  accessible from me such that there exists times  $t$  and  $t'$  such that  $f(t)$  is the individual who undergoes fission in  $w$  and  $f(t')$  is the individual with the left hemisphere after fission in  $w$ . Given this sparse ontology,  $f(t)$  will not be identical to  $f(t')$ —I am one object at  $t$  and another at  $t'$ .

This way of developing the Centered View is worth keeping in mind, especially if you are inclined toward less expansive ontologies of persistence. But for the rest of the essay, I will work with the generous ontology version of the Centered View, partly because I think such ontologies can be independently motivated (see Sider 201b for a defense) and partly because I have yet to think through the details of the "world-function pair" proposal carefully.<sup>29</sup>

On to physicalism. Blackburn's observation is that, from the first-person point of view, it seems that there are three ways things could go in the fission case even if all the continuity facts were the same. As we noted earlier, the plausibility of this observation doesn't require one to hold that there are physical differences between the three possibilities: it seems that the three situations could be physically just alike, differing only on the issue of what happens to me after fission. This leads to a conflict with physicalism *if* we adopt the Simple View's assumption that the possibilities in

29. It's interesting to note that Lewis (1976, n. 4) is officially neutral on whether 'aggregates of stages' should be taken to denote mereological sums of stages or something "less metaphysical," such as functions from times to individuals.

question are possible worlds. For then we have three possible worlds that are physically just alike but differ on my postfission spatial location, and so we have three possible worlds that are physically just alike but not alike simpliciter.

But according to **Centered Guide**, the imaginative acts in question support centered possibility claims not uncentered ones. The three possibilities in question are not three possible worlds that are physically just alike but differ on my postfission spatial location, but three *centered worlds* accessible from me. These three centered worlds,  $\langle w, x \rangle$ ,  $\langle w', x' \rangle$ , and  $\langle w'', x'' \rangle$ , are located in physically identical possible worlds, but they differ on the postfission spatial location of the center. Since  $w$ ,  $w'$ , and  $w''$  are physically just alike, and since we're supposing that physicalism is true, it follows that these three possible worlds are alike simpliciter, in the sense that every uncentered claim (possible worlds proposition) that is true at one is true at the others. But the relevant centered worlds differ on what happens to the center after fission: one,  $\langle w, x \rangle$ , is centered on something that undergoes fission and survives with the left hemisphere, another,  $\langle w', x' \rangle$ , is centered on something that undergoes fission and survives with the right hemisphere, and a third,  $\langle w'', x'' \rangle$ , is centered on something that undergoes fission and fails to survive. So we can have three distinct possibilities for my survival that agree on the same totality of physical facts without having to give up physicalism. The Centered View, unlike the Simple View, vindicates our first-person intuitions about the fission case without forcing us to reject physicalism.

### *5.3. Hidden Facts*

Our other complaint about the Simple View was that it posited hidden persistence facts or possibilities whose obtaining or failure to obtain cannot be detected, not even by an observer who is given access to all the relevant physical and psychological information. This seems wrong: when I imagine learning all the physical and psychological facts about a fission case, it seems that I know everything there is to know about the case—there just doesn't seem to be a further, completely undetectable fact here. In addition, this feature of the Simple View means that the Simple theorist cannot offer the simple and straightforward account of how we come to know facts about our persistence available to the Complex theorist.

Let's think about this problem in more detail. According to **Centered Imagination**, when I imagine observing Sam undergo fission, I imagine a set of centered worlds  $\langle w, x \rangle$  in which  $x$  observes Sam undergoing

fission in  $w$ . Now if the Simple View were true, there would be three types of centered worlds in that set:

- Centered worlds  $\langle w, x \rangle$  such that  $x$  is observing Sam undergo fission in  $w$ , and Sam survives with the left hemisphere in  $w$ .
- Centered worlds  $\langle w, x \rangle$  such that  $x$  is observing Sam undergo fission in  $w$ , and Sam survives with the right hemisphere in  $w$ .
- Centered worlds  $\langle w, x \rangle$  such that  $x$  is observing Sam undergo fission in  $w$ , and Sam fails to survive fission in  $w$ .

So if the Simple View is true, the set of centered worlds compatible with what I'm imagining can be partitioned into three cells according to what happens to Sam after fission. Each cell represents a possibility for Sam's survival. But note that the relevant continuity facts are the same in each possibility: Sam is related continuity-wise to the two postfission individuals in exactly the same way in all three possibilities. Thus, even if I had access to all the relevant continuity information—all the information about how these three individuals-at-times are related in terms of psychological and physical continuity—my information wouldn't distinguish between these three possibilities: it wouldn't tell me which of the three cells of the partition I was located in. It seems then that I couldn't come to know which possibility had obtained, even though I would have access to all the relevant information.

That's how things would be if the Simple View were true. But what if the Centered View were true? Does the Centered View posit hidden facts? To see that it doesn't, it's important to see that the Centered View is compatible with the main Complex View accounts of fission. In thinking about how this is so, it will be important to keep in mind the fact that a possibility claim can be true in the centered sense but not in the uncentered sense, and vice versa. Recall our earlier example: it may be true in the centered sense but false in the uncentered sense that I could have been Fred. There may be a centered world accessible from me centered on Fred even though no possible world accessible from the actual world is one in which I am identical to Fred. Similarly, the claims in (4) may all be true in the centered sense even while one or more of them is false in the uncentered sense:

4. (a) I could undergo fission and survive with the left hemisphere.
- (b) I could undergo fission and survive with the right hemisphere.
- (c) I could undergo fission and fail to survive.

And according to the most prominent Complex View accounts of fission, at least one of those claims is false when read as an uncentered possibility



claim. There are two main Complex View accounts of fission. On one of them, every possible world in which I undergo fission is one in which I fail to survive. This might be so either because animalism or the bodily continuity theory is true or because the nonbranching psychological continuity theory is true. The other main Complex View response to fission is Lewis's, according to which any possible world in which I undergo fission is a world in which I have two counterparts who coincide prior to fission and then split off afterward, one surviving with the left hemisphere, the other with the right.

If Lewis is right, then (4c) is false when understood as an uncentered possibility claim. If the others are right, (4a) and (4b) are both false when understood as uncentered possibility claims. So both accounts hold that at least one of the claims in (4) is false when understood as an uncentered possibility claim. But this is consistent with all three being true when read as centered possibility claims, just as the claim 'I could have been Fred' may be false on its uncentered reading but true on its centered reading. So the Centered View treatment of fission is compatible with both Complex View treatments of fission.

This will help explain why the Centered View doesn't entail that there are hidden persistence facts. When I imagine observing Sam undergoing fission, I'm imagining a set of centered worlds  $\langle w, x \rangle$  in which  $x$  is observing Sam undergoing fission in  $w$ . So each centered world  $\langle w, x \rangle$  is located in a possible world  $w$  in which Sam undergoes fission. What is a possible world in which Sam undergoes fission like? If Lewis is wrong and the others are right, then each such possible world is one in which Sam doesn't survive fission. Thus, every centered world compatible with what I imagine is one in which Sam doesn't survive fission since Sam doesn't survive fission in a centered world  $\langle w, x \rangle$  iff Sam doesn't survive fission in the possible world  $w$ . So there aren't multiple possible outcomes here for Sam. There's only one way things can go: Sam doesn't survive. So there are no hidden persistence facts here: if I know all the relevant continuity facts, and I know how Sam's persistence supervenes, then I will know what happens when Sam undergoes fission.

Similarly, if Lewis is right and the others are wrong, then every centered world compatible with what I imagine is one in which Sam has two counterparts, both of whom survive fission. But, again, there are not multiple outcomes here for Sam: every world in which he undergoes fission is one in which he has two counterparts who coincide initially and then split off after fission. So again, there are no hidden persistence facts here: if I know all the relevant continuity facts, and I know that Lewis is right about

fission, then I will know what happens when Sam undergoes fission: I'll know that there are "two Sams" both of whom survive fission.

So unlike the Simple View, the Centered View doesn't posit hidden persistence facts. So unlike the Simple View, the Centered View can avail itself of the same simple and straightforward account of how we come to know facts about our persistence offered by the Complex View.<sup>30</sup>

#### 5.4. Multiple Outcomes and the Third-Person Perspective

Now that we've seen the virtues of the Centered View—it respects our judgments about what seems possible from the first-person point of view and yet doesn't conflict with physicalism or posit hidden persistence facts—I want to close by considering an objection to it.

*Objection:* You've argued that the joint truth of the possibility judgments we're inclined to make when we imagine fission from the inside is compatible with physicalism and a ban on hidden facts. But one can also generate the intuition of "multiple fission outcomes" by imagining the fission case *from the outside*. Consider, for example, the following two cases:

##### **Fission 1**

Someone—call him 'Sam'—is about to undergo fission. He is frightened and doesn't know what will happen. But the next day he awakes—to his great relief—with the left hemisphere in the green room. Later, he meets the individual who awoke with the right hemisphere in the red room.

##### **Fission 2**

Someone—call him 'Sam'—is about to undergo fission. He is frightened and doesn't know what will happen. The next day he doesn't awake. Two new individuals wake up the next day for the first time: one awakes with the left hemisphere in the green room, the other with the right hemisphere in the red room.

**Fission 1** and **2** are described in a way that encourages us to imagine them from the outside. Yet the judgment that **Fission 1** and **2** are both genuine possibilities is as compelling as the corresponding first-person judgments, (4a) and (4c). There is simply no first-person/third-person

30. There is another Complex View account of fission that I haven't discussed: one which says that it's indeterminate what happens to Sam in the fission case. An adequate discussion of this option would take us too far afield into the topic of indeterminacy, but I am confident that admitting that possibility would make no significant difference to the general point that the Centered View doesn't posit hidden persistence facts.

asymmetry to be found here. But that means that the original dilemma simply reemerges, albeit in a slightly different form. For given **Uncentered Guide**, our inclination to judge **Fission 1** and **2** possible will lead us back to the Simple View, which will in turn lead us back into conflict with physicalism and the ban on hidden persistence facts. So you haven't really solved the problem with which you began.

*Reply:* The claim that there is no significant asymmetry between the possibility judgments we're inclined to make when we imagine fission from the inside and the ones we're inclined to make when we imagine it from the outside is controversial. As we noted in §3.1, many philosophers have claimed that when we imagine cases from the first-person perspective, facts about our persistence seem to be further facts, over and above the continuity facts. If the authors who make this point (Nagel, Blackburn, Chisholm, and so on) had thought that there was no important asymmetry between our first-person judgments and our third-person ones, then their emphasis on the *first-person* point of view would be extremely puzzling. Blackburn, for example, must think that the case for multiple fission outcomes is more compelling when we imagine fission from the inside than it is when we imagine it from the outside—why else would he go to such lengths to emphasize the role of the first-person perspective in generating the relevant intuitions?

Furthermore, advocates of the Simple View—who one might expect to agree with the objector on this matter if anyone did—have not generally held that the Simple View looks true from the outside. For example, a major theme of Madell's book-length defense of the Simple View (1981) is that the previous literature on our persistence is distorted by an overemphasis on the third-person point of view. He thinks the deliverances of the two perspectives conflict, but unlike the rest of the literature, he thinks the first-person perspective should prevail:

Repeatedly in the contemporary literature on personal identity . . . the third-person viewpoint is taken to reveal the truth and the only truth. My argument will be that, far from this being the case, the third-person viewpoint misses absolutely fundamental truths. It is not just that we have to recognize a conflict between third-person and first-person viewpoints. What we also have to recognize is that in this conflict the first-person viewpoint must prevail. (Madell 1981, 22)

And specifically in connection with fission, Madell (1981, 129) suggests that the idea that there are multiple possible outcomes in the fission case is one that “we can *only* make sense of from the first-person viewpoint”

(emphasis added). So the objector's view that our inclination to make the relevant possibility judgments doesn't depend on the perspective from which we imagine the case is controversial, apparently at odds with the judgments of many other philosophers who have thought about the matter, including even those who have sought to defend the Simple View.

But the objector might simply deny that Nagel, Blackburn, Madell, and others are right in thinking that there is any interesting first-person/third-person asymmetry here. If he or she does, I'm not sure what could be said to change his or her mind about what seems right *as a matter of intuition*. So what else can we say to the objector?

In a sense, the arguments in this essay constitute a reason for the objector to treat his or her first-person judgments differently from his or her third-person ones. From the objector's point of view, both sets of judgments are initially compelling, and both seem to conflict with two other desiderata on a theory of our persistence (compatibility with physicalism and the ban on hidden facts). But what we've shown is that while a plausible account of imagination and possibility allows us to reconcile our first-person judgments with these other theoretical commitments, that same account does *not* allow us to reconcile the objector's *third-person* judgments with those commitments. Assuming the objector doesn't want to give up physicalism or accept hidden persistence facts, this seems to give the objector a reason to treat the two sets of judgments differently, rejecting the third-person ones, while retaining the first-person ones, in spite of the fact that he or she finds both sets initially compelling.

The methodological stance taken here is a familiar one: possibility judgments made on the basis of imagining a hypothetical case are defeasible, not ones we must retain come what may. Sometimes we have theoretical reasons to reject possibility judgments that we are initially inclined to accept. Thus, the objector can read this essay as providing an argument, based on theoretical considerations, for rejecting his or her (controversial) third-person judgments, while retaining his or her (more widely shared) first-person judgments.

To Blackburn, myself, and others, the idea that there are three possible fission outcomes is, as a matter of intuition, more compelling from the first-person point of view than it is from the third-person point of view. The Centered View respects our "asymmetry" intuition since it endorses the joint truth of the relevant centered possibility claims but rejects the joint truth of the relevant uncentered possibility claims. To the objector, the case for multiple fission outcomes is as intuitively compelling from the third-person point of view as it is from the first-person point of

view. The Centered View doesn't respect the objector's "symmetry" intuition; instead, it provides him or her with a reason for rejecting it.

## References

- Anand, Pranav, and Andrew Nevins. 2004. "Shifty Operators in Changing Contexts." In *Proceedings from Semantics and Linguistic Theory* 14, ed. Robert B. Young. Ithaca, NY: CLC Publications, Cornell University.
- Baker, Lynne Rudder. 2000. *Persons and Bodies*. New York: Cambridge University Press.
- Blackburn, Simon. 1997. "Has Kant Refuted Parfit?" In *Reading Parfit*, ed. Jonathan Dancy, 180–201. Oxford: Blackwell.
- Butler, Joseph. 1736. *Of Personal Identity*. First appendix to *The Analogy of Religion*. Reprinted in Perry 1975.
- Chalmers, David. 1996. *The Conscious Mind*. New York: Oxford University Press.
- . 2002a. "Does Conceivability Entail Possibility?" In *Conceivability and Possibility*, ed. Tamar Szabó Gendler and John Hawthorne, 145–200. New York: Oxford University Press.
- . 2002b. "On Sense and Intension." *Philosophical Perspectives* 16: 135–82.
- Chisholm, Roderick. 1969. "The Loose and Popular and the Strict and Philosophical Senses of Identity." In *Perception and Personal Identity*, ed. Norman Care and Robert Grimm, 82–106. Cleveland: Press of Case Western Reserve University.
- Egan, Andy. 2004. "Second-Order Predication and the Metaphysics of Properties." *Australasian Journal of Philosophy* 82: 48–66.
- . 2006. "Secondary Qualities and Self-Location." *Philosophy and Phenomenological Research* 72: 97–119.
- von Fintel, Kai. 2005. "Lecture Notes for Pragmatics in Linguistic Theory." Notes for class taught at the LSA Summer Institute, MIT.
- Hazen, Allen. 1979. "Counterpart-Theoretic Semantics for Modal Logic." *Journal of Philosophy* 76: 319–38.
- Higginbotham, James. 2003. "Remembering, Imagining, and the First Person." In *Epistemology of Language*, ed. Alex Barber, 496–533. Oxford: Oxford University Press.
- Hill, Christopher. 1997. "Imaginability, Conceivability, Possibility and the Mind-Body Problem." *Philosophical Studies* 87: 61–85.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford: Clarendon Press.
- Johnston, Mark. 1987. "Human Beings." *Journal of Philosophy* 84: 59–83.
- Lewis, David K. 1976. "Survival and Identity." In *The Identities of Persons*, ed. Amelie O. Rorty, 17–40. Berkeley: University of California Press. Reprinted in Lewis 1983b, 55–77.

- . [1979] 1983. "Attitudes De Dicto and De Se." *Philosophical Review* 88: 513–43. Reprinted in Lewis 1983b, 133–59. Page references are to the 1983 reprint.
- . [1983a] 1999. "Individuation by Acquaintance and by Stipulation." *Philosophical Review* 92: 3–32. Reprinted in Lewis 1999, 373–402. Page references are to the 1999 reprint.
- . 1983b. *Philosophical Papers*, Vol. 1. New York: Oxford University Press.
- . 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- . [1995] 1999. "Reduction of Mind." In *A Companion to Philosophy of Mind*, ed. Samuel Guttenplan, 412–31. Oxford: Basil Blackwell. Reprinted in Lewis 1999, 291–324. Page references are to the 1999 reprint.
- . 1999. *Papers in Epistemology and Metaphysics*. Cambridge: Cambridge University Press.
- Lowe, E. J. 1996. *Subjects of Experience*. Cambridge: Cambridge University Press.
- Madell, Geoffrey. 1981. *The Identity of the Self*. Edinburgh: Edinburgh University Press.
- Martin, M. G. F. 2002. "The Transparency of Experience." *Mind and Language* 17: 376–425.
- Merricks, Trenton. 1998. "There Are No Criteria of Identity over Time." *Noûs* 32: 106–24.
- Morgan, Jerry. 1970. "On the Criterion of NP Deletion." CLS 6.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *Philosophical Review* 83: 435–50.
- . 1986. *The View from Nowhere*. New York: Oxford University Press.
- Ninan, Dilip. 2008. "Imagination, Content, and the Self." PhD diss., MIT.
- Nolan, Daniel. 2006. "Selfless Desires." *Philosophy and Phenomenological Research* 73: 665–79.
- Olson, Eric T. 1997. *The Human Animal*. New York: Oxford University Press.
- . 2007. *What Are We?* New York: Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Peacocke, Christopher. 1985. "Imagination, Experience, and Possibility: A Berkeleyan View Defended." In *Essays on Berkeley*, ed. John Foster and Howard Robinson, 19–35. Oxford: Clarendon Press.
- Perry, John, ed. 1975. *Personal Identity*. Berkeley: University of California Press.
- . 2002. *Identity, Personal Identity, and the Self*. Indianapolis: Hackett Publishing Company.
- Recanati, François. 2007. *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford: Oxford University Press.
- . n.d. "Imagining *de se*." Unpublished manuscript.
- Reid, Thomas. 2002 [1785]. *Essays on the Intellectual Powers of Man*. University Park: Pennsylvania State University Press.
- Reynolds, Steven. 1989. "Imagining Oneself to Be Another." *Noûs* 23: 615–33.

*Persistence and the First-Person Perspective*

- Salmon, Nathan. 1998. "The Logic of What Might Have Been." *Philosophical Review* 98: 3–34.
- Schlenker, Philippe. 2003. "A Plea for Monsters." *Linguistics and Philosophy* 26: 29–120.
- Shoemaker, Sydney. 1963. *Self-Knowledge and Self-Identity*. Ithaca, NY: Cornell University Press.
- . 1984. "Personal Identity: A Materialist's Account." In *Personal Identity*, ed. Richard Swinburne and Sydney Shoemaker, 67–132. Oxford: Blackwell.
- . 1994. "The First-Person Perspective." *Proceedings and Addresses of the American Philosophical Association* 68: 7–22. Reprinted in Shoemaker 1996, 157–75.
- . 1996. *The First-Person Perspective and Other Essays*. New York: Cambridge University Press.
- Sider, Theodore. 1999. "Global Supervenience and Identity across Time and Worlds." *Philosophy and Phenomenological Research* 59: 913–37.
- . 2001a. "Criteria of Personal Identity and the Limits of Conceptual Analysis." *Philosophical Perspectives* 15: 189–209.
- . 2001b. *Four-Dimensionalism: An Ontology of Persistence and Time*. New York: Oxford University Press.
- Snowdon, Paul F. 1990. "Persons, Animals, and Ourselves." In *The Person and the Human Mind*, ed. C. Gill, 71–86. Oxford: Clarendon Press.
- von Stechow, Armin. 2003. "Feature Deletion under Semantic Binding: Tense, Person, and Mood under Verbal Quantifiers." In *Proceedings of NELS* 33, ed. Makoto Kadawaki and Shigeto Kawahara, 379–404. Charleston: BookSurge Publishing.
- . n.d. "Binding by Verbs: Tense, Person, and Mood under Attitudes." Unpublished manuscript.
- Stoljar, Daniel. 2005. "Physicalism." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta.
- Swinburne, Richard. 1984. "Personal Identity: The Dualist Theory." In *Personal Identity*, ed. Richard Swinburne and Sydney Shoemaker, 1–66. Oxford: Blackwell.
- Szabó Gendler, Tamar, and John Hawthorne, eds. 2002. *Conceivability and Possibility*. New York: Oxford University Press.
- Turner, Jason. Forthcoming. "Fitting Attitudes *De Dicto* and *De Se*." *Noûs*.
- Van Inwagen, Peter. 1990. *Material Beings*. Ithaca, NY: Cornell University Press.
- Velleman, David. 1996. "Self to Self." *Philosophical Review* 105: 39–76.
- Walton, Kendall. 1990. *Mimesis as Make-Believe*. Cambridge, MA: Harvard University Press.
- Wiggins, David. 1976. "Locke, Butler, and the Stream of Consciousness." In *The Identities of Persons*, ed. Amelie O. Rorty, 139–73. Berkeley: University of California Press.
- . 1980. *Sameness and Substance*, 1st ed. Oxford: Blackwell.

- Williams, Bernard. [1966] 1973. "Imagination and the Self." *Proceedings of the British Academy* 52: 105–24. Reprinted in Williams 1973, 26–45. Page references are to the 1973 reprint.
- . 1970. "The Self and the Future." *Philosophical Review* 79: 161–80. Reprinted in Williams 1973, 46–63.
- . 1973. *Problems of the Self: Philosophical Papers 1956–1972*. Cambridge: Cambridge University Press.
- Williamson, Timothy. 2002. "Necessary Existents." In *Logic, Thought and Language*, ed. A. O'Hear, 233–51. Cambridge: Cambridge University Press.