Potential Traits and Personal Identity

Julian De Freitas & George A. Alvarez

Harvard University

Header: Potential Traits and Personal Identity

Version: 01/22/20

Total word count: 10,078

Abstract (117 words)

Previous work suggests that an important factor that influences whether or not we personally identify with an imagined version of ourselves is whether it has a mind. But is personal identity simply a reflection of whether an agent currently has a mind, or does it also track whether it has the *potential* to have a mind in the future? In 6 experiments we find that participants are willing to identify with agents who are completely mindless—such as their imagined fetus or future self who suffers serious brain damage—provided they believe that these agents have this potential. We suggest that feelings of personal identity track not only mindhood but potential mindhood, and that in some cases potential wins out.

Keywords: Lay theories, identity, mind perception, consciousness, potential

A striking feature of mental life is that we are not limited to our immediate sensory experience, but can also travel within our own minds, imagining remembered, anticipated, or hypothetical versions of ourselves, and putting ourselves in others' shoes. This ability plays a critical role in decision-making, navigation, and social cognition (Ames, Jenkins, Banaji, & Mitchell, 2008; Buckner & Carroll, 2007). Yet how do we determine whether an imagined version of ourselves—say, someone who is much younger or deteriorated—is *us*?

A core idea within psychology and philosophy is that we solve this self-correspondence problem by determining whether the imagined version of ourselves would continue to possess a mind (Blok, Newman, Behr, & Rips, 2001; Hood, Gjersoe, & Bloom, 2012; Locke, 1689; Nichols & Bruno, 2010; Parfit, 1984). To reach this conclusion, previous work has tested how the discontinuity in specific mental contents or traits over time — e.g., due to aging, life-changing events, or even experimentation by mad scientists — interrupts intuitions about personal identity (Bartels & Rips, 2010; De Freitas, Cikara, Grossmann, & Schlegel, 2017; Tobia, 2015). When people are asked to imagine hypothetical situations in which significant deteriorations occur to a person's mental characteristics, as in a Phineas Gage-type tragedy (Tobia, 2015), they are inclined to say that the person is no longer the same person they used to be. Further, people do not simply treat all parts of the mind as equally important for personal identity. Certain mental traits, such as moral ones, are viewed as particularly important, whereas others, like perceptual ones, are not (Strohming & Nichols, 2014), and people believe that the morally *good* traits are most essential of all (De Freitas, Cikara, et al., 2017; De Freitas, Cikara, Grossmann, & Schlegel, 2018; Newman, De Freitas, & Knobe, 2015). These intuitions do not simply reflect a general optimism or the individualistic ideals of a Western population, since people make similar judgments whether they are misanthropes, judge outgroup members, or live in an interdependent culture (De Freitas & Cikara, 2018; De Freitas,

Sarkissian, et al., 2018), and these beliefs are observed in young children (Lefebvre & Krettenauer, 2020).

A way of summarizing this previous work is that whether a person at t_0 is viewed as personally continuous with a person at t_1 depends on the continuity of particular mental contents and traits between them: when their fundamental moral goodness, memories, or other core beliefs are altered or erased, they are no longer recognized (by themselves or by others) as the same person.

The Mind Only hypothesis

Here we test a logical corollary of this work: if personal identity depends on the continuity of particular mental contents or traits (henceforth “mind”), people should not identify with imagined selves who have no mind at all. In philosophy, this is referred to as the *psychological view* of personal identity (Locke, 1689; Parfit, 1984; Shoemaker, 2008). It states that what matters is the literal presence of sophisticated mental states, and predicts that people are unwilling to ascribe personal identity to imagined versions of themselves that lack these mental states. Since there are not many people without minds walking about, philosophy exploring these ideas tends to focus on imagined fetuses and patients in a persistent vegetative state, both of whom lack the capacity for mental states¹. Subscribers to the psychological view, then, would predict that people are equally *unwilling* to ascribe personal identity to these agents, since both agents lack minds; people should say that they did not exist when they were a fetus, and would not exist if they entered into a persistent vegetative state (PVS). Despite a substantial amount of philosophical theorizing on personal identity based on these two agents (Shoemaker, 2008), no work on personal identity has empirically contrasted them.

For our purposes, we will refer to this view as the *Mind Only Hypothesis*, since it predicts a clear limit on commonsense personal identity: People do not identify with imagined selves that

¹ In reality, a minority of PVS patients has been found capable of some mental capacities (Jennett & Plum, 1972).

currently lack a mind. Note that we are not referring here to pedestrian examples of ‘mindlessness’, as in daydreaming or sleeping, but to the biological inability to execute even basic mental functions like recognition, thinking, and memory. We want to emphasize that the Mind Only Hypothesis strikes us as a very reasonable view. After all, if personal identity depends on the continuity of mind, then how can we say that we are continuous with an entity that lacks one? If we believe that the presence of a mind defines us, then so long as we imagine an agent without a mind, we should think that it could not be us; our sense of a continuing identity should be disrupted. It would seem there is no way to logically escape this view.

The Mind Potential hypothesis

Yet we might reach an alternative view if we think not only about whether the agent *currently* has a mind, but about whether it has the underlying *potential* to have one. Under this view, when people decide whether they are psychologically continuous with an imagined agent, they are not exclusively taking into account whether the imagined agent currently has a mind, but also whether the agent may develop one in the *future*.

Returning to the case of the fetus and vegetative patient, we see that this view makes quite different predictions from the Mind Only hypothesis. Even though a fetus does not have a mind, it has the potential to develop one, so people should identify with it. A vegetative patient also has no mind, yet the brain damage it has suffered renders it incapable of ever developing one in the future, so people should not identify with it. Thus, this hypothesis predicts that people should be more likely to identify with the fetus than the vegetative patient. This strikes us as a surprising prediction, since a fetus has never had a full-fledged mind, whereas some vegetative patients have had a mind for most of their lives. Therefore, the crucial difference would have to be that the fetus’ *potential* mind is more complete than the vegetative patient’s current or potential mind. Put differently, the

fetus represents the beginning of a temporal process that ends in a full-fledged mind, whereas a vegetative patient represents the endpoint of a process that results in no mind.

Since this view states that what matters is not just the presence of a mind but the potential to develop one, we refer to it as the *Mind Potential Hypothesis*. It, too, makes a clear prediction: when reasoning about agents without minds, people are willing to identify with them so long as they have the potential to develop a mind.

The Same-Mind Potential hypothesis

Here it may also help to be more precise about what *kind* of mind the agent will need to have the potential to develop. Unlike the Mind Potential view, which states that the agent should simply have the potential to develop *a* mind, the *Same-Mind Potential Hypothesis* has a more specific requirement: the agent has to have the potential to develop *my* mind. Building on previous work showing that people do not view all parts of their mind as equally essential to their personal identity (De Freitas, Cikara, et al., 2017; Strohminger, Knobe, & Newman, 2017), we will loosely define ‘my mind’ as those mental traits that one values most about oneself (whatever they may be; for more detail on this matter, see the studies reviewed above). In Sum, the Same-Mind Potential view predicts that people should identify with a mindless agent if it has the potential to develop the same mind that they consider theirs, rather than just any mind whatsoever.

The Current Studies

To arbitrate between the Mind Only, Mind Potential, and Same-Mind Potential hypotheses, we start by taking inspiration from existing philosophical work, asking people whether they identify with imagined versions of themselves that lack a mind— a fetus, and a vegetative patient (*Experiment 1*). Counter to the Mind Only hypothesis, we find that people preferentially identify with the fetus. Since an important assumption of our design is that people view these agents as mindless, *Experiment 2* measures the extent to which people attribute various mental capacities to these agents, and finds

that, if anything, they ascribe more mind to the vegetative patient than the fetus. *Experiment 3* then marshals more direct evidence for the role of potential, by showing that people would be less likely to identify with their imagined fetus in the counterfactual case where it was not destined to develop a valued trait (vs. unvalued trait). We then find that these effects also operate at an implicit level, as measured by a so-called *self-reference* effect in a speeded categorization task involving the fetus vs. PVS (*Experiment 4*). Finally, *Experiments 5a-b* find that the effect generalizes to other agents who are ascribed no mind, and suggest that personal identity specifically tracks whether an imagined agent has the potential to develop the Same-Mind that one currently has, rather than any mind whatsoever.

Experiment 1. Imagining agents without minds

This experiment explores whether people identify with two agents without minds— a fetus and a patient in a PVS. Above and beyond the existing concepts that people have of these agents, we took several steps to ensure that participants understood that the agents were both equally without minds. Specifically, we stipulated: [the agent] “has no mental contents, memories, or capacities. It has no capacity for conscious awareness or reasoning and is not sentient”. We also only analyzed data from participants who correctly selected the option “It has no mental contents” in a comprehension question about this information.

If participants identified to the same extent with both agents, then this would support the Mind Only hypothesis, but if they preferentially identified with the fetus then this would support the Mind Potential and Same-Mind Potential hypotheses.

Method

One hundred and one participants were recruited to take part in a 5-min study for \$0.50 compensation, and 17 of them were excluded for failing any of the attention or comprehension checks, leading to a final sample of 84 participants (mean age = 42.1 years; 43% female). In the

absence of any previous studies that had used the paradigm created for the present experiments, Experiment 1 began with the heuristic assumption that the resulting effect size would be comparable to a recent study on personal identity (De Freitas & Cikara, 2018), yet we recruited half this number because our own intuitions about the vignettes convinced us that the effects would be particularly robust (we also replicated the effect in pilot studies). Participants in all studies were recruited from the online crowdsourcing platform, Amazon's Mechanical Turk. Such participants are more diverse than typical lab samples, yet provide responses that are equally reliable (Buhrmester, Kwang, & Gosling, 2011). We screened participant ids to ensure that no one took part in more than one experiment across the experiments in this project. All studies were approved by the Harvard Internal Review Board.

At the beginning of all experiments, participants answered an attention check (*How many fatal heart attacks have you had?*). Additionally, at the end of all experiments they answered two-multiple choice comprehension questions about the type of agent that they were asked about and whether it had mental capacities (*1. According to the story, which of the following is referred to? (i) A victim of a serious accident, (ii) A fetus, (iii) A very old dog; 2. According to the story, which is true of the thing referred to? (i) It has many thoughts and desires, (ii) It has no mental contents, (iii) None of the above.*). Participants had to pass all attention and comprehension checks in order to be included in the analyses. Finally, they answered basic demographics questions: age, gender, native language, and level of science and philosophy education (*1=None at all, 9=A very large amount*). All data and code are available upon request.

Participants then read about either a fetus or vegetative patient (presented between-subjects). Importantly, we used the same wording to stipulate that both agents had no mental capacities (the vegetative patient variant is presented in squared parentheses):

Once there was a two-month-old fetus in your mother's womb that eventually developed into the person that you are today [Imagine that you get into a very serious car accident that severely damages your brain].

At two-months [After the injury] the body has a brain, but it has no mental contents, memories, or capacities. It has no capacity for conscious awareness or reasoning, and is not sentient. Only some of its vital functions work, e.g., circulation, digestion, immune system.

As the main identity measure, participants were then asked to rate the extent to which they agreed with the following statement ($1=Definitely\ disagree$, $5=Neutral$, $9=Definitely\ agree$): *Although my body continued to exist from the fetus to the adult [from before the accident to after the accident], I did not exist when it was a fetus [I no longer exist after the accident].*

Results and Discussion

In line with the Mind Potential and Same-Mind Potential views, which predict that people will identify with a mindless agent so long as it has the potential to develop a mind of some kind, people were more inclined to agree that they did not exist as the vegetative patient ($M = 6.82$, $SD = 2.56$) than the fetus ($M = 4.44$, $SD = 3.21$; $t(82) = 3.79$, $p < .001$, $d = 0.83$; Figure 1A). Since both agents were described as equally lacking a mind and we only analyzed participants who remembered this detail, these results contradict the Mind Only hypothesis, which predicts that people do not personally identify with an imagined self that does not currently possess a mind.

Figure 1. Results from Experiments 1 and 2. Participants' mean scores (reverse-coded) on (A) a measure of personal identification with their imagined fetus and a vegetative patient, and on (B) 14 mind ascription items that rated the extent to which each agent had the capacity for pain, experience, agency, and consciousness. Participants identified more with the fetus, even though they ascribed more mind to the PVS (patient in a persistent vegetative state). Error bars depict standard error of the mean.

Experiment 2. Measuring mind ascription

A natural alternative account of the effects from Experiment 1 is that people did not identify more with the fetus because it has the *potential* to develop a mind, but because they viewed it as *already* possessing more mind. If this is the case, then the identity results from Experiment 1 actually support the Mind Only hypothesis, since they track different levels of existing, rather than future, mind.

One common way of determining whether people are ascribing existing mental traits to an agent is to simply ask them. Previous studies have asked people to rate the extent to which they think agents have mental contents/traits like *experiences* (e.g., desires, feelings, emotions, pleasure, fear, hunger, personality), *pain* specifically, *agency* (e.g., planning, self-control, memory, emotion recognition, moral understanding, communication, influencing situations), and *consciousness* of themselves and the world (H. M. Gray, Gray, & Wegner, 2007; Ward, Olsen, & Wegner, 2013).

In fact, some of these studies have even measured mind attribution to fetuses and patients in a persistent vegetative state. For instance, Gray et al. (2007) asked participants to consider the following agents:

7 week fetus. At 7 weeks, a human fetus is almost half an inch long—roughly the size of a raspberry.

...and:

Gerald Schiff. Gerald Schiff has been in a persistent vegetative state (PVS) for the past six months. Although he has severe brain damage--Gerald does not appear to communicate with others or make purposeful movements--his basic bodily functions (such as breathing, sleeping, and circulation) are preserved.

One relevant finding from this previous work is that, even though the vegetative patient was described as having brain damage, it was ascribed slightly more agency (yet equal amounts of experience) than the fetus. We note that if such an effect also applied in the current studies, then it would have worked *against* the Mind Potential and Same-Mind Potential hypotheses for which we found support, thereby providing even stronger evidence against the Mind Only hypothesis.

Another finding worth highlighting from previous work is that, on average, participants do not pick the lowest point on the mind attribution scale, even when judging fetuses and vegetative patients. For instance, one study focusing on mind ascription to vegetative patients found the following mean levels of mind ascription on a 7-point Likert scale (1 being the lowest point): ability to feel pain (3.09), capacity for experience (2.75), capacity for agency (1.96), endorsement of conscious awareness (2.31) (Ward et al., 2013). Those who frequently work with Likert Scales know that it is exceedingly rare for *all* participants to pick the extreme values of the scale. Even so, might these findings suggest that fetuses and vegetative patients are not good examples of mindless agents? Suggesting otherwise, Gray et al. (2011) found that vegetative patients are ascribed even less mind than dead people. Therefore, even if people ascribe some (low) amount of mind to vegetative patients and fetuses on average, this is still as little mind as we could hope to see attributed. Also, note that our own vignettes took several additional steps to stipulate that the agents were equally without the mental capacities described above, and we only analyzed data from participants who remembered this fact.

At the same time, we cannot simply assume that these manipulations were effective nor can we lean on previous mind attribution findings, given that our own vignettes differed in important

ways from those studies: first, we emphasized that the agents were *equally* without mental capacities, and second, we asked participants to imagine that the bodies of these agents were continuous with their *own* bodies, i.e., we measured first-person rather than third-person ascriptions. Thus, before we can conclude that the identity results from Experiment 1 contradict the Mind Only hypothesis, we need to also measure mind ascriptions to these agents. To this end, Experiment 2 asked participants to answer several mind ascription measures about the particular fetus and vegetative patient that participants were asked to imagine in Experiment 1. We predicted that, in line with previous work and our own stipulations about these agents' mental capacities, people would ascribe low and equal levels of mind to both agents.

Method

We again recruited one hundred and one participants to take part in a 5-min study for \$0.50 compensation, and 49 of them were excluded for failing the attention or comprehension checks, leading to a final sample of 52 participants (mean age = 33.9 years; 25% female). We expect that the large number of exclusions was because participants were asked the comprehension question about mental capacities after having just answered the main items about whether each agent had several mental contents/traits, suggesting to them that it might be ok to say that the agents had some level of mental traits— even if they encoded the fact that the story specified that the agent did not. Importantly, we included the comprehension questions at the end of the experiment, to ensure that they did not affect participants' responses on the main mind attribution items for this experiment, and we also analyze the data without excluding participants based on this question.

After reading the same vignettes from Experiment 1, participants were asked to *Please rate the extent to which you agree with the following statements about this fetus [victim]: It...*, followed by several mind attribution items about the agent, presented in randomized order to each participant (all items are

listed in the Appendix). Each measure was presented as a statement, and participants responded on a scale ranging from 1 (*Strongly disagree*) to 9 (*Strongly agree*).

Results and Discussion

We analyzed the attribution items as in previous work, by averaging together items related to the same category of mental capacity (Ward et al., 2013). Figure 1B shows that participants attributed relatively low levels of mind to both agents, with all but one rating falling below the midpoint of the scale. Independent t-tests confirmed that there were no differences in how much experience, agency, or consciousness each agent was ascribed (p s = .545, .899, and .792 respectively). However, participants ascribed *more* capacity for pain to the vegetative patient ($M = 5.35$, $SD = 2.80$) than the fetus ($M = 3.77$, $SD = 2.70$; $t(50) = 2.07$, $p = .044$, $d = 0.57$). We also find these effects when we do not exclude participants based on the comprehension question about mental states: the vegetative patient is ascribed a higher ability to feel pain than the fetus ($N = 66$, $p = .002$), but there is no difference in how much of the other mental contents/traits the agents are ascribed.

Given the identity results of Experiment 1, the Mind Only hypothesis predicts that in the current experiment participants should also ascribe more mental capacities to the fetus than the vegetative patient. Yet if anything we found the opposite result: they ascribed *more* mental capacity to the vegetative patient. Thus, despite ascribing more existing mental capacities to the vegetative patient, participants still identified more with the fetus (Figure 1). This result suggests that mental ascriptions worked *against* the Mind Potential and Same-Mind Potential hypotheses, reinforcing the existing evidence against the Mind Only hypothesis.

This finding also helps settle potential concerns about the fact that both agents are still ascribed some low levels of mind. If mind is attributed to some extent, the Mind Only hypothesis also predicts that identity ascriptions will track any relative changes in mind attribution, yet we found no evidence for this prediction. If anything, when the relative amount of mind increased from one

agent to the other, we saw no corresponding increase in identity ascription; in fact, the increase went in the opposite direction.

In light of Experiments 1 and 2, we can now begin to consider why people identify more with the fetus than the vegetative patient, if not due to existing mind. We suggest that it is because the fetus has the potential to have a full-fledged mind. Specifically, even though the vegetative patient has had a mind for most of its life, and even though it presently has more mind than the fetus, the fetus will develop a mind that is more capable than the mind which the vegetative patient has or will have—the fetus has a better potential mind. Experiment 3 explores this hypothesis more directly.

Experiment 3: The role of potential

The findings from Experiments 1 and 2 are consistent with the Mind Potential and Same-Mind Potential hypotheses, but now we would like to test if there is direct evidence for the role of potential. Experiment 3 explores a ‘causal’ manipulation in which people are asked to imagine a counterfactual scenario in which their fetus experienced a genetic mutation that prevented it from developing the one trait that they value most in themselves. We contrast this case with one in which the mutation affects a more superficial trait about themselves that they should value less. Since in both cases the fetus is identical except for how the genetic mutation affects its future, any difference in personal identification between these conditions would suggest that personal identity tracks changes in potential traits.

Another benefit of this design is that it holds the type of agent constant. So far we have been contrasting fetuses and vegetative patients, for three reasons; we have wanted to: (i) consider more than only one mindless agent, and these agents happen to be the only ones that occur naturally; (ii) explore the role of potential, and these agents differ on this dimension; and (iii) connect our empirical findings to existing theoretical work on personal identity in philosophy, which has focused

on these agents (Degrazia, 2002, 2005; Marquis, 1989; McInerney, 1990; Olson, 1997a, 1997b; Schechtman, 2014; Shoemaker, 2008). For instance, the so-called *fetus problem* (Olson, 1997a) discusses precisely the sorts of issues that we investigated empirically in Experiments 1-2.

Yet one concern is that fetuses and vegetative patients differ on various dimensions other than just whether they have the potential to develop a mind. A fetus transforms into an adult gradually, building memories, values, and other hallmarks of personal identity over time. In contrast, an adult would transform into a vegetative patient suddenly (in these studies, after a “very serious car accident”), losing these hallmarks of personal identity in a flash. The fetus and vegetative patient also vary relative to the current self on several dimensions, including: their direction of trajectory (gain vs. loss), temporal direction (past vs. future), hypotheticality (everyone has been a fetus, whereas most people will never become a vegetative patient), valence (ostensibly neutral vs. negative), and potentially more. For our purposes, however, we would ideally like to hold all of these factors constant, in order to isolate the role of potential.

Therefore, it would be informative to test whether the effects still hold when we hold the agent constant, or when we compare agents other than a fetus which differ only in their potential. Experiment 3 takes the first approach, and we take the second approach in Experiments 5a-b.

Method

We again recruited one hundred and one participants to take part in a 5-min study for \$0.50 compensation, and 54 of them were excluded for failing any of the checks, leading to a final sample of 47 participants (mean age = 34.1 years; 34% female).

We intended to contrast participants’ intuitions about cases in which their fetus lost either a valued or unvalued trait. In order to ensure that participants were thinking concretely about their valued trait, we asked them to *Think about one particular aspect of yourself that you value more than any other. In one sentence, please explain what it is.* They could then type their answer in a provided textbox.

We then presented participants with the same scenario about the fetus from Experiment 1, in which they were asked whether they once existed as the fetus. At the end of the scenario, we asked them to imagine a genetic mutation that would cause it to lose either a valued or unvalued trait, presented between-subjects (the valued trait is indicated in squared parentheses): *Imagine that this fetus experienced a genetic mutation that completely prevented it from ever developing a tolerance to lactose [the one aspect of you that you value most], while everything else about it was preserved.* We refer to these conditions as Consequential Change and Superficial Change, respectively.

Since the identity question asked about a counterfactual that had the potential to be confusing (i.e., how people would turn out if their fetus had experienced a mutation), we simplified the identity question to *Do you think that this fetus would still be you?* Participants responded on a scale ranging from 1 (*No, definitely not me*) to 9 (*Yes, definitely me*).

Results and Discussion

Those in the Superficial Change condition were more inclined to agree that they would still exist as the fetus ($M = 6.00$, $SD = 2.59$) than those in the Consequential Change condition ($M = 4.29$, $SD = 2.99$; $t(45) = 2.09$, $p = .042$, $d = 0.61$). We also find this result if we do not exclude participants based on the comprehension check about mental states ($N = 54$, $p = .003$).

Since in both cases the genetic mutation would only have affected how the fetus turned out in the future, these findings provide direct evidence for the role of potential in identity judgments. Note, however, that these results do not yet arbitrate between the Mind Potential and Same-Mind Potential hypotheses, since losing the potential for a valued trait is consistent with both, i.e., there are fewer valued traits, but also fewer mental traits in general. We arbitrate between these accounts in Experiment 5a.

Next, since this experiment asked people to describe their most valued trait, we also decided to analyze these descriptions. Doing so is informative for two reasons: First, although existing work

predicts that these traits will tend to be about the mind rather than about the body or other factors entirely (Blok et al., 2001; Hood et al., 2012; Locke, 1689; Nichols & Bruno, 2010; Parfit, 1984), we wanted to confirm that this was indeed the case: were the potential traits that won out over existing mental traits still mental traits themselves? Second, to our knowledge all existing studies of personal identity except for one (Chen, Urminsky, & Bartels, 2016) have used pre-constructed items. One concern with this sort of approach is that findings will be limited to the imaginations of the researchers, who might leave out items that are nonetheless important for personal identity. Therefore, we also wanted to test whether the factors which people spontaneously generate in the current situation cohere with those highlighted in previous work.

To answer these questions, we first screened the sentences to determine what categories of traits they fell under. We then asked two independent coders, blind to both condition and the experiment's hypotheses, to categorize the sentences into one of the following pre-determined categories: *moral* (*values that relate to how they treat others*), *intelligence* (*abilities that they have that go beyond personality traits, involving critical reasoning, creativity, etc.*), *personality* (*personal traits and characteristics that they naturally have*), *physical* (*physical traits*), or *NA* (*misinterpreted the question*). Inter-coder reliability for the full set of sentences was high, with coders selecting the same category for 40 of the 47 responses, yielding a Cohen's kappa of .80. For analysis, we excluded the 7 cases in which coders chose different categories, as well as the 6 cases in which coders agreed that the participants' responses were not applicable, yielding 34 final responses for analysis. This is a standard way to code and exclude sentences in free-response data (De Freitas, DeScioli, Nemirow, Massenkoff, & Pinker, 2017).

Consistent with previous work (Chen et al., 2016), a large proportion of self-generated sentences referred to moral (38.2%) and personality traits (17.6%), although we also found a large proportion which referred to intelligence traits (41.1%). Only a tiny proportion referred to physical

traits (3%). Together, these results support two main conclusions. First, they suggest that intelligence traits are an important category for personal identity that has largely been neglected in existing work, which has tended to focus on morality (De Freitas, Cikara, et al., 2018; Strohminger et al., 2017). Second, people must have identified more with a fetus than a vegetative patient because they viewed it as possessing, at a potential level, precisely those mental traits that it lacked at a literal level— traits like intelligence, morality, and personality. Therefore, mind ascription is still relevant to personal identity, it's just that it can be driven by potential (inferred) mind, rather than just by existing (literal) mind.

Experiment 4: Implicit personal identification

Now that we have suggested an important role for potential in personal identity, we can ask what exactly is the psychological nature and consequences of these effects? One likely possibility is that the effects reflect how people think about personal identity at a rational, System 2 level (Kahneman, 2011), but do not tell us anything about how they represent personal identity at a more implicit, System 1 level. In line with this idea, previous work finds that people care less about future selves that are temporally distant (Parfit, 1984; Pronin & Ross, 2006), perhaps because their representations of these selves are less detailed (Wakslak, Nussbaum, Liberman, & Trope, 2008). Perhaps people have even less detailed representations of completely mindless agents that they neither remember (fetus) nor have personally been (vegetative patient), making them incapable of identifying with these agents at an implicit level— even if they say that they do at an explicit level.

One way to determine whether participants identify with an agent at an implicit level is to test whether they exhibit certain behavioral effects that typically occur when people associate an item with their current selves. Specifically, we can test whether participants exhibit a *self-reference* effect, in which they perform better when categorizing that item vs. others (Rogers, Kuiper, & Kirker, 1977). Typically, self-reference effects occur for items which people have been asked to

associate with their current selves. As an example, one previous study asked participants to associate their self, a friend, and a stranger with simple geometric shapes such as a triangle, circle, and square (the reason for using simple geometric shapes is that participants are unlikely to have pre-existing associations with them). After learning these associations, participants were serially shown each shape for just 100ms, paired with either its correct label or an incorrect one, and they had to rapidly (within 1s) say whether each pairing was a match or mismatch. Even though what a shape signified should not have affected performance on the task (which had the same requirement regardless of what a shape signified) participants were more accurate for the shape that they associated with their self than with others, exhibiting a self-reference effect (Sui, He, & Humphreys, 2012). This performance benefit is thought to arise because self-relevant items are processed by a dedicated network for self relevance (Sui, Rotshtein, & Humphreys, 2013), which couples with attention and memory enhance processing of the stimulus (Sui & Humphreys, 2015; Sui et al., 2013).

Relevant to the current studies, recent work suggests that this effect can also be used to diagnose which item is associated with the self in less obvious cases, as when determining which of two imagined selves one personally associates with (De Freitas, Rips, & Alvarez, 2019). So long as a participant performs better for one imagined self vs. another, we may conclude that the performance benefit occurred because one of the items was preferentially represented as the self. Also relevant, when multiple imagined selves are pitted against each other within the same paradigm, then people only exhibit a self-reference effect for one of imagined selves at a time, suggesting that there is a limit on how many imagined selves can be cognitively represented as *the* self concurrently (De Freitas et al., 2019).

Applying these insights to the current question, the self-reference effect can be used to diagnose whether participants implicitly identify with the fetus or the vegetative patient. Given that the current agents are completely mindless, unlike the imagined agents used in previous work, it is

not obvious whether we should find a preference. If we don't, then this suggests that people preferentially identify with the fetus at an explicit level only—perhaps because their representation of it lacks the necessary detail to compete for self-representation at an implicit level.

Method

To determine the sample size, we first replicated the original self-reference effect for one's *current self* versus a friend or stranger, using forty one participants and all the same exclusion measures that we planned to employ in the current study. Thus, we recruited forty one participants to take part in a 20-min study for \$2.00 compensation, of which 17 were excluded for failing any of the following checks: an attention check, two comprehension checks about the order of vignettes and mental capacities of the agents (*1. According to the story, which of the following is referred to? A fetus, and then a victim of a serious accident., A victim of a serious accident, and then a fetus, None of the above.; 2. According to the story, which is true of what was referred to? They had many thoughts and desires., They had no mental contents., None of the above.*), or a comprehension check shown after the main perceptual task (*Which, if any, of the following images did you observe in the sequence of images you saw? A. Two cars crashing, B. A building on fire, C. An ambulance, D. Different shapes, E. None of the above*). Participants were also excluded for scoring lower than 55% accuracy overall (50%=chance). Finally, we excluded individual trials with reaction times faster than 100ms (people jumping the gun, or holding down a key). These stringent exclusions led to a final sample of 24 participants (mean age = 31.2 years; 49% female), which is consistent with previous exclusions using the same criteria for this paradigm on Mechanical Turk (De Freitas et al., 2019).

As in Experiment 1, we asked each participant to make identity judgments about the fetus and the vegetative patient presented in the prior studies, except this time the two agents were presented within-subjects, with the order of presentation was counterbalanced between-subjects. We then asked participants to associate these agents with different geometric shapes (e.g., fetus =

triangle, victim = square). We also asked them to associate a third shape with a stranger ‘named John’, so that the overall task would be challenging. After they learned these associations, they completed the same sorts of speeded matching trials described above (the paradigm is displayed in Figure 2A). Importantly, their task was the same regardless of what the shape represented, so if they performed better for a given agent, then this would suggest that the shape was implicitly associated with the self.

During *learning trials*, participants completed 27 untimed trials, in which the different shapes ($190 \times 190\text{px}$) were sequentially presented in a randomized order on top of a white fixation cross (14pt), which itself was centered on a $800\text{px} \times 554\text{px}$ gray background. Participants said whether the shape represented the fetus, vegetative patient, or stranger by clicking one of three numbered keys (‘1’, ‘2’, or ‘3’) that were matched with the words *fetus*, *victim*, and *stranger* respectively (19pt). Participants received feedback on each trial, then once they passed 21 learning trials they proceeded to the practice trials (else learning trials repeated). Note that since we cannot control for monitor size in an online experiment, we have reported all stimulus dimensions in pixel values and font sizes.

During the *practice trials*, participants saw either the correct shape-label pair (matched condition), or a shape paired with an incorrect label (mismatched condition). Each trial consisted of the following sequential events: a fixation cross (500ms), a pairing of a shape and a label presented 107px below it (90ms), a blank screen during which participants pressed ‘n’ for match or ‘m’ for mismatch (‘m’) (900ms), performance feedback that read *Correct!*, *Error!* or *Too slow!* (500ms), and an inter-trial blank screen (100ms). Participants had to pass half of the 18 practice trials in order to proceed to the main trials (else the practice trials repeated). *Main trials* were the same as the practice ones, except there were 300 trials total, consisting of 150 match and 150 mismatch trials, with shape-label pairs fully counterbalanced and presented in a randomized order within-subjects.

Figure 2. The paradigm (A) and results (B) for Experiment 4, plotted for each agent condition (left column) and for each participant as a set of ordered performance differences (fetus – vegetative patient; right column). PVS = patient in a persistent vegetative state. Error bars depict standard error of the mean. $**p < 0.01$.

Results and Discussion

Performance levels in the matching and mismatching conditions for a given agent's shape were combined, in order to form a measure of d' (Green & Swets, 1966). Since we did not instruct participants to respond as quickly as possible, we did not analyze reaction time. Participants performed better for the fetus ($M = 2.43$, $SD = 0.96$) than the vegetative patient ($M = 1.92$, $SD = 0.88$; $t(23) = 3.69$, $p = .001$, $d = 1.06$) and stranger ($M = 1.79$, $SD = 0.98$; $t(23) = 3.58$, $p = .002$, $d = 1.03$), whereas they performed indistinguishably for the vegetative patient and the stranger ($t(23) = 0.79$, $p = .436$, $d = 0.23$; Figure 2B).

Adding to Experiments 1-3, these results suggest that participants preferentially identified with the fetus at *both* explicit and implicit levels, as evidenced by the self-reference effect. Since self-reference effects are thought to be supported by a limited neural mechanism that can only represent one self at a time (De Freitas et al., 2019; Sui et al., 2013), the current results suggest that this same mechanism was used to represent the fetus as the self. Further, given that participants had to first understand the meaning of the vignette before these associations could be made, these results also suggest fresh information processed at an explicit level can drive self-referential behavior at an implicit level, in a top-down fashion.

Experiment 5a. Generalizing to other agents

So far we have found that ascriptions of personal identity are affected by an agent's potential, even if the agent does not currently have a mind. We find this both when contrasting a fetus and vegetative patient (Experiments 1-2), and two fetuses with different potentials (Experiment 3). Yet one possibility is that these studies only tell us about the concept of fetuses, for which potential might be especially integral and salient, rather than about the effect of potential on personal identity more generally. Therefore, in order to know whether these effects are more general, we should replicate them in agents other than just fetuses.

In this vein, the current study had four aims: (i) to test for a generalization of the effects to other agents, while (ii) controlling for the absence of a mind across agents, (iii) including a direct manipulation of role, and additionally (iii) adjudicating between the Mind Potential and Same-Mind Potential hypotheses. To do this, we asked participants to imagine that they were in a car accident that either: damaged their brain permanently (No-Mind Potential); damaged their brain but that surgeons would be able to reinstate their *same* mental contents after one year (Same-Mind Potential); or wiped out their mental contents (while leaving their brain otherwise undamaged) but that surgeons would be able to reinstate *different* mental contents after one year (Different-Mind Potential)². Both of the final conditions involve regaining a mind, but only the Same-Mind Potential condition involves regaining the same mental traits that one used to have. If people value those mental traits more than just any mental traits, then they should identify most with this agent.

In order to ensure that each agent was viewed as lacking a mind, we took three steps. First, as in the previous studies, we used similar language across conditions to stipulate that the agents were without minds, and only analyzed participants who correctly answered a comprehension about both this fact and about what kinds of mental traits the agent would/would not gain in the future. Second, to provide a strong test, notice that we deliberately stacked the deck against the Same-Mind Potential Hypothesis by stipulating that the agent in this condition suffered brain damage, whereas the agent in the Different-Mind Potential condition suffered no brain damage but only had their mental contents wiped out. We expected that this would lead people to ascribe relatively more existing mind to the agent in the Different-Mind Potential than Same-Mind Potential condition (Experiment 5b tests this prediction directly). Third, in our identity question we emphasized that we were asking people to make an identity judgment about the agent *immediately after the accident*, and even underlined this part of the question. In sum, depending on what differences we found, we

² The agent in the Different-Mind Potential condition is loosely inspired from the protagonist in the movie *Eternal Sunshine of the Spotless Mind*.

expected to be in a position to make secure conclusions about whether the effects were due to existing vs. potential mindhood.

Methods

Aiming for the same sample size per condition as in Experiments 1-3, we recruited one hundred and forty eight participants to take part in a 5-min study for \$0.50 compensation, and 49 of them were excluded for failing any of the checks, leading to a final sample of 99 participants (mean age = 35.7 years; 45% female). In addition to the same comprehension questions used in Experiments 1-3, participants had to correctly answer a question about the agent's future (*According to the story, which will be true of the brain after one year of surgeries? (i) It will not function properly, (ii) It will have the same mental contents as what you had before the accident, (iii) It will have different mental contents than what you had before the accident*).

Participants were then given one of three possible endings (presented between-subjects). In the No-Mind Potential and Same-Mind Potential conditions, they were shown similar information, except that the agent in the Same-Mind Potential condition was described as having the potential to recover the same mental contents in one year (this condition is presented in squared parentheses):

Imagine that you get into a very serious car accident that completely damages your brain.

After the injury, the body has a brain, but it has no mental contents, memories, or capacities. Only some of its vital functions work, e.g., circulation, digestion, immune system.

The brain is unable to develop or learn any mental capacities.

Brain experts will conduct a series of careful micro surgeries on the brain over a yearlong period. They determine that after the last surgery at the end of the year, the brain will still not function properly [will be fully functioning again, and all of its contents, memories, and capacities will be identical to the ones you had before the accident].

The Different-Mind Potential condition was similar, except that the agent was described as having its mental contents wiped out rather than damaged, and as having the potential to recover different mental contents in one year:

Imagine that you get into a very serious car accident that completely wipes out the contents of your brain without otherwise damaging it.

After the injury, the body has a brain, but it has no mental contents, memories, or capacities. Only some of its vital functions work, e.g., circulation, digestion, immune system.

The brain is able to develop and learn all mental capacities, but it must start from scratch like an infant.

Brain experts will conduct a series of careful micro surgeries on the brain over a yearlong period. They determine that after the last surgery at the end of the year, the brain will be fully functioning again, but all of its contents, memories, and capacities will be different from the ones you had before the accident.

Participants then answered the same identity question as in Experiment 1, although we first emphasized that they should *Please rate the extent to which you agree with the following statement about the body immediately after the accident:* Finally, they answered the comprehension and demographic questions.

Results

A 3-way (Condition: No-Mind Potential, Same-Mind Potential, Different-Mind Potential) univariate ANOVA yielded the predicted main effect of condition $F(2, 96) = 5.58, p = .005, \eta_p^2 = .104$ (Figure 3A). Follow up t-tests revealed that participants were less inclined to agree that they would cease to exist if they were the victim in the Same-Mind Potential Condition ($M = 5.34, SD = 2.53$) than the No-Mind Potential ($M = 6.95, SD = 2.30; t(64) = 2.69, p = .009, d = .67$) and Different-Mind Potential Conditions ($M = 7.30, SD = 2.52; t(60) = 3.05, p = .003, d = .78$). There was no significant difference between the No-Mind Potential and Different-Mind Potential conditions ($t(68) = 0.62, p = .537, d = .15$). These results suggest that people are willing to currently identify with a mind-damaged accident victim, so long as it has the potential to reinstate the same mind that they currently have— even if this will only happen in a year.

Figure 3. Results from Experiments 5a and 5b. Participants' mean scores on (A) a measure of personal identification (reverse-coded) with a mindless agent who in one year will have either no mind, a different mind, or the same mind, and on (B) 14 mind ascription items that rated the extent to which each of these agents has the capacity for pain, experience, agency, and consciousness. Participants identified more with the agent who would have the same mind, even though they ascribed equal or more mind to the other agents. Error bars depict standard error of the mean.

Experiment 5b. Teasing apart potential from current mind

As for Experiment 1, we cannot simply assume that participants ascribed equal levels of mind to each of the agents, so Experiment 5b asked people to answer the same mind ascription measures from Experiment 2, but with respect to the agents introduced in Experiment 5a. We also tested our prediction that, despite identifying more strongly with the agent in the Same-Mind Potential than Different-Mind Potential condition, participants would ascribe more existing mind to the agent in the latter case, given that its mental contents were 'merely' wiped out rather than

damaged. As in Experiment 5a, we emphasized that participants should make judgments about the mental capacity of the agent *immediately after the accident* while underlying this portion of the question.

Methods

We recruited one hundred and fifty one participants to take part in a 5-min study for \$0.50 compensation, and 45 of them were excluded for failing any of the checks, leading to a final sample of 106 participants (mean age = 36.3 years; 46% female).

The design of the experiment was similar to Experiment 5a, except that after reading the vignettes participants were asked to *Please rate the extent to which you agree with the following statements about the body immediately after the accident. It...*, then they answered the same mind ascription items from Experiment 2.

Results

Figure 3B shows that participants attributed relatively low levels of mind to all agents, with some exceptions for the agent in the Different-Mind Potential condition. We ran separate 3-way univariate ANOVAs for each of the four categories of mind ascription, and found significant results for each one: pain ($F(2, 103) = 7.31, p = .001, \eta_p^2 = .124$), experience ($F(2, 103) = 12.91, p < .001, \eta_p^2 = .200$), agency, ($F(2, 103) = 4.28, p = .017, \eta_p^2 = .077$) and consciousness ($F(2, 103) = 6.46, p = .002, \eta_p^2 = .111$).

Table 1 summarizes the follow up t-test comparisons between conditions for each of mental capacities. As expected, for all mental capacities there were no differences in mind ascription between the agents in the No-Mind Potential and Same-Mind Potential conditions, likely because both were described as suffering brain damage. As predicted, however, the agent in the Different-Mind Potential Condition was ascribed more mind than the other agents. Specifically, relative to the No-Mind Potential agent it was ascribed more mind on every measure, and relative to the Same-

Mind Potential agent it was ascribed more mind on all measures except for agency (for which there was no significant difference).

Thus, we see that participants in Experiment 5a ascribed the most identity to the agent in the Same-Mind Potential Condition, even though they also ascribed it either equal or lower mind than the other agents. These findings strongly suggest that ascriptions of identity tracked potential, not existing, mind. Together with the previous experiments, they also suggest that the effect of potential is not specific to a particular kind of agent, but generalizes across different cases— whether contrasting fetuses and vegetative patients (Experiments 1-2 and 4), fetuses with different potentials (Experiment 3), or accident victims with different potentials (Experiment 5a-b).

Table 1. Levels of mind ascription to each of the agents in Experiment 5b

Condition Contrast	<i>M (SD) vs. M (SD)</i>	<i>t</i>	<i>d</i>
Pain			
No Mind v Same Mind	5.39 (2.54) v 5.58 (2.88)	0.30	0.07
Different Mind v Same Mind	7.53 (1.93) v 5.58 (2.88)	3.28**	0.77
No Mind v Different Mind	5.39 (2.54) v 7.53 (1.93)	3.81***	0.95
Experience			
No Mind v Same Mind	4.40 (0.11) v 4.41 (0.14)	0.33	0.08
Different Mind v Same Mind	4.54 (0.12) v 4.41 (0.14)	4.18***	0.99
No Mind v Different Mind	4.40 (0.11) v 4.54 (0.12)	4.89***	0.08
Agency			
No Mind v Same Mind	2.30 (1.34) v 2.84 (2.29)	1.19	0.28
Different Mind v Same Mind	3.71 (2.07) v 2.84 (2.29)	1.70	0.40
No Mind v Different Mind	2.30 (1.34) v 3.71 (2.07)	3.28**	0.81
Consciousness			
No Mind v Same Mind	3.24 (2.07) v 3.20 (2.44)	0.09	0.02
Different Mind v Same Mind	5.03 (2.61) v 3.20 (2.44)	3.09**	0.73
No Mind v Different Mind	3.24 (2.07) v 5.03 (2.61)	3.07**	0.76

Note. The results are for independent-sample t-statistics. Sample sizes for each of the conditions were: No-Mind Potential: 33, Different-Mind Potential: 32, Same-Mind Potential: 41. ** $p < .01$, *** $p < .001$.

Discussion

Does an agent need to have a mind in order for people to personally identify with it? We find that the answer is no, so long as the agent has the potential to develop the same mind in the

future. Experiment 1 found that people were more inclined to personally identify with a fetus than a vegetative patient, even though both agents were ascribed almost no mind—if anything, the vegetative patient was ascribed more (Experiment 2). Experiment 3 found direct evidence that these effects were driven by inferences about potential, and Experiment 4 uncovered that the effect also operated at an implicit level. Finally, Experiments 5a and 5b established that the effects generalize to mindless agents other than fetuses, and that personal identity specifically tracks the potential to develop the same mind that one currently has, not just any mind whatsoever.

The Role of a Mind

One could argue that these studies have little to do with everyday personal identity and mental ascription, because we do not typically reason about mindless agents. Yet that would be missing the point of testing intuitions about mindless agents, which was to try to hold mind completely constant. Our study could have tested intuitions about agents who do have minds, while tracking whether changes in mind ascription affect personal identity, but this would have been a weaker test than measuring identity in the near absence of a mind.

Next, although people's judgments tracked potential traits, Experiment 3 found that these traits still related to the mind, such as morality, intelligence, and personality. Therefore, mindhood of some form drives personal identity, it is just that the current effects were driven by *inferred* minds; in some sense, potential minds can substitute for absent ones. How common is it for potential to fully win out over existing mind ascription? We suspect that whenever people reason about an agent's future and infer that its future mind will be more full-fledged than its existing mind, then personal identity will side with potential.

The current effects are also relevant to morally-valenced policies concerning agents who are vulnerable because they are unconscious or minimally conscious. Much empirical work on this issue has tended to focus on the existing mental capacities of these agents (K. Gray et al., 2011; Ward et

al., 2013). Yet the current studies suggest that, in line with some strands of philosophical theorizing (Marquis, 1989), another factor at play may be whether the agent has a potential mind. Future work should investigate whether the current intuitions about personal identity affect moral attitudes toward abortion and assisted death.

The Role of Potential

The current work contributes to existing work showing the effects of potential—including on evaluations of events/individuals in the present. For instance, people have a tendency to be optimistic about future events in the world (Weinstein, 1980), and to over-estimate the happiness they will experience for future positive events (Wilson & Gilbert, 2003). Even more surprising, people prefer individuals with potential over those with proven performance (Massey & Thaler, 2013; Tormala, Jia, & Norton, 2012), and prefer current performances by individuals with potential even though they could watch superior performances by others (Poehlman & Newman, 2014). The current studies show that the influence of potential even extends to a notion as fundamental as whether a given person is currently us.

But were people identifying with the present agents (“that’s me”), or with the future agents who would realize their potential (“that would become me”)? Here three points are relevant. First, across experiments we emphasized that we were soliciting identity judgments for the agents before their potentials were realized, e.g., asking participants how much they agreed that “I did not exist when it was a fetus” (Experiment 1), or stipulating that they should make the judgment “about the body immediately after the accident” (Experiment 5a). Second, it can’t be that people were answering these questions for the future agents, otherwise they would have also ascribed high levels of mind to these agents, yet they didn’t. Third, we found an implicit self-reference effect for the fetus in Experiment 2, which indicated that participants associated it with their current selves.

Together, these points suggest that information about an agent's potential determined whether people identified with the agent in its existing state.

How commonly does potential affect ascriptions of personal identity? In practice, we suspect that it depends in part on the temporal *perspective* that one takes of an imagined self. In many cases, our perspective may be zoomed in on the agent's current moment in time, yet sometimes we may take a more zoomed-out view in which we see the agent as an event over time. In some cases, this zoomed out view may be more salient, as for fetuses (for whom potential looms large); in other cases, we may come to take on this perspective for other reasons, e.g., as when deciding how long to provide life support to a brain-damaged patient. In line with this thinking, one study found that even though potential affects valuations of goods consumed in the present, this effect is attenuated when people focus more on the present than the future (Poehlman & Newman, 2014).

Since participants were willing to identify with a brain-damaged version of themselves that would recover only in a year's time, one implication of the current studies is that the intervening period between the existing time of the agent and its projected end is less important than the end itself. Some previous work suggests that event endings may be special, e.g., when evaluating the moral character of others, people weight more heavily the behaviors at the end of the individual's life (Newman, Lockhart, & Keil, 2010). A related question is the extent to which such effects can be retrospective. For instance, people might identify with their fetus because they are viewing it from the current vantage point in which they have a fully-functioning mind, but would they identify less with their fetus when asked to imagine that their life ended in a PVS?

Origin of the Effect

Why is personal identity affected by potential traits? We see three overlapping possibilities: (i) it is a by-product of how the self is represented, (ii) it is a by-product of how potential is represented, and (iii) it is adaptive and rational. For (i), we know that personal identity tracks existing

mind ascription (Blok et al., 2001; Hood et al., 2012; Locke, 1689; Nichols & Bruno, 2010; Parfit, 1984), so we need to explain why this effect extends to potential minds that are inferred. For (ii), we know that there are various effects of potential in other domains of reasoning (Massey & Thaler, 2013; Poehlman & Newman, 2014; Tormala et al., 2012), so we need to explain why it extends to personal identity. For (iii), reasoning about potential is typically considered adaptive, because it indicates that people value *anticipated* payoffs, which incentivizes them to strive toward items/goals so as to actually obtain these payoffs. Similarly, valuing potential seems rational provided that decision-making is considered within an extended temporal context, rather than just the present—although you cannot experience the payoff now, you will eventually after striving for it.

Thus, the more pressing question may be whether identifying with potential minds is a useless mashup of these other abilities, or serves some additional use. We cannot be sure, but the current effects do seem useful. Humans have self-interested reasons for wanting to thrive, and part of thriving requires imagining and charting out their own, unique futures. Doing so requires making good decisions for their future selves, and tracking how those decisions will affect them specifically. Personal identity based on potential traits seems to satisfy these needs, guiding choices at various future scales (e.g., “I’m not going through with the surgery tomorrow, because if it impacts my personality then I’ll lose what make me *me*”; or “If I start investing all my time in programming now, what kind of person will I be in five years?”). Perhaps the current effects are a symptom of how people intuitively ‘feel out’ their potential trajectories, helping them determine which ones to actually live out.

Acknowledgments

Thanks to Samuel G. B. Johnson, Lance Rips, Fiery Cushman, David Shoemaker, and Joshua Knobe for helpful feedback, and Patcharapon Thammathorn and Michelle Zhao for coding sentences.

Appendix: Mind Attribution Items in Experiment 2, Taken from Ward (2013)

Please rate the extent to which you agree with the following statements about this fetus [victim]:

It...[1 = Strongly disagree, 5 = Neutral, 9 = Strongly agree].

[Pain item]

can experience pain.

[Experience items]

has a personality.

can experience desire.

can experience feelings.

can experience emotions.

can experience pleasure.

can experience hunger.

can experience fear.

[Agency items]

is able to plan its actions.

is able to control its actions.

Potential Traits and Personal Identity

remember the events of its life.

can understand the thoughts and emotions of those around it.

can understand right from wrong.

can influence the outcome of situations.

can communicate its thoughts and feelings to others.

[Consciousness]

is conscious of itself.

is conscious of the people and world around it.

References

- Ames, D. L., Jenkins, A. C., Banaji, M. R., & Mitchell, J. P. (2008). Taking another person's perspective increases self-referential neural processing. *Psychological Science*, 19(7), 642-644.
- Bartels, D. M., & Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, 139(1), 49-69.
- Blok, S., Newman, G., Behr, J., & Rips, L. J. (2001). *Inferences about personal identity*. Paper presented at the Proceedings of the twenty-third annual conference of the cognitive science society, Edinburgh, Scotland.
- Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11(2), 49-57.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1), 3-5.
- Chen, S. Y., Urminsky, O., & Bartels, D. M. (2016). Beliefs about the causal structure of the self-concept determine which changes disrupt personal identity. *Psychological Science*, 27(10), 1398-1406.
- De Freitas, J., & Cikara, M. (2018). Deep down my enemy is good: Thinking about the true self reduces intergroup bias. *Journal of Experimental Social Psychology*, 74, 307-316.
- De Freitas, J., Cikara, M., Grossmann, I., & Schlegel, R. (2017). Origins of the belief in good true selves. *Trends in Cognitive Sciences*, 21(9), 634-636.
- De Freitas, J., Cikara, M., Grossmann, I., & Schlegel, R. (2018). Moral goodness is the essence of personal identity. *Trends in Cognitive Sciences*, 22(9), 739-740.
- De Freitas, J., DeScioli, P., Nemirow, J., Massenkoff, M., & Pinker, S. (2017). Kill or die: Moral judgment alters linguistic coding of causality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(8), 1173.
- De Freitas, J., Rips, L. J., & Alvarez, G. A. (2019). The limit of personal identity. *Under review*.
- De Freitas, J., Sarkissian, H., Newman, G. E., Grossmann, I., De Brigard, F., Luco, A., & Knobe, J. (2018). Consistent belief in a good true self in misanthropes and three interdependent cultures. *Cognitive Science*, 42, 134-160.
- Degrazia, D. (2002). *Are we essentially persons? Olson, Baker, and a reply*. Paper presented at the The Philosophical Forum.
- DeGrazia, D. (2005). *Human identity and bioethics*. Cambridge: Cambridge University Press.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *science*, 315(5812), 619-619.
- Gray, K., Knickman, T. A., & Wegner, D. M. (2011). More dead than dead: Perceptions of persons in the persistent vegetative state. *Cognition*, 121(2), 275-280.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Hood, B., Gjerse, N. L., & Bloom, P. (2012). Do children think that duplicating the body also duplicates the mind? *Cognition*, 125(3), 466-474.
- Jennett, B., & Plum, F. (1972). Persistent vegetative state after brain damage: a syndrome in search of a name. *The Lancet*, 299(7753), 734-737.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Lefebvre, J. P., & Krettenauer, T. (2020). Is the true self truly moral? Identity intuitions across domains of sociomoral reasoning and age. *Journal of Experimental Child Psychology*, 192, 104769.
- Locke, J. (1689). *An essay concerning human understanding* (K. Winkler Ed.). Indianapolis: Hackett Publishing.
- Marquis, D. (1989). Why abortion is immoral. *The Journal of Philosophy*, 86(4), 183-202.

- Massey, C., & Thaler, R. H. (2013). The loser's curse: Decision making and market efficiency in the National Football League draft. *Management Science*, 59(7), 1479-1495.
- McInerney, P. K. (1990). Does a fetus already have a future-like-ours? *The Journal of Philosophy*, 87(5), 264-268.
- Newman, G. E., De Freitas, J., & Knobe, J. (2015). Beliefs about the true self explain asymmetries based on moral judgment. *Cognitive Science*, 39(1), 96-125.
- Newman, G. E., Lockhart, K. L., & Keil, F. C. (2010). "End-of-life" biases in moral evaluations of others. *Cognition*, 115(2), 343-349.
- Nichols, S., & Bruno, M. (2010). Intuitions about personal identity: An empirical study. *Philosophical Psychology*, 23(3), 293-312.
- Olson, E. T. (1997a). *The human animal: Personal identity without psychology*. Oxford: Oxford University Press.
- Olson, E. T. (1997b). Was I ever a fetus? *Philosophy and Phenomenological Research*, 57(1), 95-110.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Poehlman, T. A., & Newman, G. E. (2014). Potential: The valuation of imagined future achievement. *Cognition*, 130(1), 134-139.
- Pronin, E., & Ross, L. (2006). Temporal differences in trait self-ascription: when the self is seen as an other. *Journal of Personality and Social Psychology*, 90(2), 197.
- Rogers, T. B., Kuiper, N. A., & Kirker, W. S. (1977). Self-reference and the encoding of personal information. *Journal of Personality and Social Psychology*, 35(9), 677.
- Schechtman, M. (2014). *Staying alive: Personal identity, practical concerns, and the unity of a life*. Oxford: Oxford University Press.
- Shoemaker, D. (2008). *Personal identity and ethics: A brief introduction*. Toronto: Broadview Press.
- Strohminger, N., Knobe, J., & Newman, G. (2017). The true self: A psychological concept distinct from the self. *Perspectives on Psychological Science*, 12(4), 551-560.
- Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159-171.
- Sui, J., He, X., & Humphreys, G. W. (2012). Perceptual effects of social salience: evidence from self-prioritization effects on perceptual matching. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5), 1105.
- Sui, J., & Humphreys, G. W. (2015). The integrative self: How self-reference integrates perception and memory. *Trends in Cognitive Sciences*, 19(12), 719-728.
- Sui, J., Rotshtein, P., & Humphreys, G. W. (2013). Coupling social attention to the self forms a network for personal significance. *Proceedings of the National Academy of Sciences*, 110(19), 7607-7612.
- Tobia, K. P. (2015). Personal identity and the Phineas Gage effect. *Analysis*, 75(3), 396-405.
- Tormala, Z. L., Jia, J. S., & Norton, M. I. (2012). The preference for potential. *Journal of Personality and Social Psychology*, 103(4), 567.
- Wakslak, C. J., Nussbaum, S., Liberman, N., & Trope, Y. (2008). Representations of the self in the near and distant future. *Journal of Personality and Social Psychology*, 95(4), 757.
- Ward, A. F., Olsen, A. S., & Wegner, D. M. (2013). The harm-made mind: Observing victimization augments attribution of minds to vegetative patients, robots, and the dead. *Psychological Science*, 24(8), 1437-1445.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5), 806.
- Wilson, T. D., & Gilbert, D. T. (2003). Affective forecasting. *Advances in experimental social psychology*, 35(35), 345-411.