

©American Psychological Association, 2024. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. The final article is available, upon publication, at: [10.1037/xge0001711](https://doi.org/10.1037/xge0001711)

Is Personal Identity Intransitive?

Julian De Freitas¹ and Lance J. Rips²

¹Marketing Unit, Harvard Business School

²Department of Psychology, Northwestern University

Abstract

There has been a call for a potentially revolutionary change to our existing understanding of the psychological concept of personal identity. Apparently, people can psychologically represent people, including themselves, as multiple individuals at the same time. Here we ask whether the intransitive *judgments* found in these studies truly reflect the operation of an intransitive *concept* of personal identity. We manipulate several factors that arbitrate between transitivity and intransitivity and find most support for transitivity: in contrast to prior work, most participants do not make intransitive judgments when there is any reason to favor one individual over another. People change which single individual they personally identify with, depending on which individual competes more strongly or weakly for identity, rather than identifying with both individuals. Even when two individuals are identical and therefore both entitled to be the same person, we find that people make more transitive judgments once they understand the practical commitments of their responses (Experiment 4), and report not being able to actually imagine two perspectives simultaneously when reasoning about the scenario (Experiment 5). In short, we suggest that while people may make intransitive judgments, these do not reflect that they psychologically represent identity in an intransitive manner.

Keywords: Identity; Persistence; Transitivity; Mental simulation

Public Significance Statement

This work investigates whether people think that an individual person can be the same, identical person as another co-existing person. Using a thought experiment in which participants are asked to imagine that a single person is duplicated, we find that people believe a single person can be identical only to themselves. Likewise, leveraging a thought experiment in which participants imagine that they are teleported into two places simultaneously, we find that participants do not actually experience two co-existing consciousnesses when imagining the scenario but only one. The results suggest that we psychologically represent people, including ourselves, as strictly singular individuals. The findings have fundamental implications for how we make various practical decisions about ourselves and others, such as who is entitled to certain privileges like loved ones or belongings, as well as for how we think about new technology that can increasingly ‘clone’ human likeness—suggesting we will not view these clones as identical to ourselves.

Everywhere, people change. Despite these changes, we are skilled at keeping track of ourselves and others, as when determining whether the 8-year-old bucktooth in the photobook is the same person as the 30-year-old academic holding the photobook. More generally, we can determine whether we are continuous with past, future, or hypothetical versions of ourselves that we encounter, as in some representational form (e.g., in the perceptual form of a photobook, or a name written on an old trophy board) or imagined form (e.g., as when judging whether we would still be the same person if we had Alzheimer's disease or locked-in syndrome). We will refer to judgments about whether an entity at t_0 is the same as the one at t_1 , as judgments of *identity persistence*, or *persistence* for short. Persistence judgments about humans specifically are typically referred to as *personal identity* (Hume, 1740; Lewis, 1976; Locke, 1689; Parfit, 1984; Perry, 1972).

In psychology, persistence judgments about individuals are thought to reflect the operation of a cognitive representation allocated to representing the single individual. As such, one would expect these representations to abide by the basic logical properties of the identity relation, which say that identity is reflexive, symmetric, and transitive (Church, 1956; Mendelson, 2009). Specifically, if x_i , x_j , and x_k are individuals, then:

Principle 1a. $x_i = x_i$ (reflexivity)

Principle 1b. If $x_i = x_j$ then $x_j = x_i$ (symmetry)

Principle 1c. If $x_i = x_j$ and $x_j = x_k$, then $x_i = x_k$ (transitivity)

In this paper, we focus on the transitivity criterion, which effectively says that an entity at t_0 cannot be identical to two different entities at t_1 . As an example, if there are two people in my living room, a 30-year old academic (x_j) and a 30-year old entertainer (x_k), then it is easy to

appreciate that they are not the same person ($x_j \neq x_k$), since each of them can go his or her own way, think separate thoughts, and so on. Now let's say we assert that the 8-year-old bucktooth in the photobook is the same person as *both* the 30-year-old academic ($x_i = x_j$) and the 30-year-old entertainer ($x_i = x_k$). If this is true, then the 30-year-old academic and entertainer must also be the same person ($x_j = x_k$), given the transitivity criterion. Yet, we initially observed that they are not the same person ($x_j \neq x_k$), and so our assertion has violated the transitivity rule.¹

Despite the fact that the transitivity criterion for numerical identity is coherent, some empirical studies have found that people make seemingly-intransitive judgments about the persistence of both non-human entities—such as icebergs, tigers and hamsters (Hood et al., 2012; Rips, 2011; Rips et al., 2006)—and about the personal identity of themselves and others (Weaver & Turri, 2018; Woike et al., 2020). The findings have led some to claim that the psychological concept of identity is itself intransitive, i.e., that people have an underlying, intuitive conviction that personal identity is intransitive (Weaver & Turri, 2018; Woike et al., 2020).

If these authors (and others we review below) are correct that the mind has an intransitive representation of personal identity, this revision to our understanding of identity would be revolutionary, overturning basic assumptions in intuitive physics and psychology. Moreover, it would overturn much of our coherent cognitive experience, which seems to depend on tracking singular individuals over time in a transitive manner. In the photobook example, for instance, it

¹ Sometimes people use “identical” to refer to similarity, as in, *the chair in my office is identical to the one in my home* Alrenga (2006). Similarity need not satisfy the transitivity criterion, though, since one object can be similar to a second, and the second can be similar to a third, without the first being similar to the third. In this work, we are not asking whether two entities at two spatial or temporal points merely have similar properties, but whether they are the very same, numerical individual. In this sense, numerical identity is the relation that everything has to itself and to nothing else.

would be tremendously surprising if the academic thought that he was simultaneously identical with the entertainer, since this would presumably affect the academic's thoughts and behavior dramatically, such as thinking that the entertainer was entitled to his belongings, lovers, and even thoughts and wrongdoings. Before accepting this revolution, therefore, we think it is worth scrutinizing the case for both transitive and intransitive identity carefully.

Before doing so below, however, we want to quickly distinguish between questions about the constitution of the self (which may have many aspects) and questions about personal identity (how the mind represents people, including oneself, over time). Both empirical and philosophical research on the constitution of the self have led to the recognition of the divided self (Dennett, 2014; Kurzban, 2012; Nagel, 1971). This perspective says that the mind is composed of multiple, parallel processes that function independently and can even be at odds with one another. Empirical evidence for this view includes observations of phenomena such as self-deception, lack of self-control, subliminal influences, automatic behaviors, conflicting moral responses, confabulation, and other related phenomena. Yet even though people may have multiple ways of presenting and experiencing the self, the possibility of these selves is consistent with the presence of a single person, a person who can vote only once in an election, count as only one person in the US census, be the only person in a study, etc. In the current work, we ask whether they also psychologically represent people in this singular way that is consistent with transitivity, or in a more pluralistic way that is consistent with intransitivity.

The Case for Transitive Identity

Several findings in psychology suggest that people employ a transitive concept of personal identity. Studies involving both children and adults find that, even when there are multiple contenders for identity, people are sensitive to which contender preserves what they

deem to be essential (as opposed to superficial) features of the original (Blok et al., 2001; De Freitas et al., 2018; Keil, 1989; Nichols & Bruno, 2010). For instance, developmental studies find that children believe that proper nouns refer to individuals that preserve the same physical form despite changes to surface characteristics, e.g., that Mr. Red is still Mr. Red even if he is painted green (Hall et al., 2003). Likewise, children believe that proper nouns refer to unique individuals rather than to all similar individuals (Sorrentino, 2001). They also behave in a way that suggests they think objects sum and subtract in a manner that obeys the transitivity criterion, e.g., that two objects minus one is one, not two (Wynn, 1992).

Other studies have found consistent effects using thought experiments. For example, child and adult participants in one study were asked to judge whether one object can become two in a simplified Ship of Theseus paradigm (Marchak & Hall, 2022), in which a ship's parts are gradually replaced until there are two ships—one made from the ship's old parts, and one made from the new parts that are spatiotemporally continuous with the original ship (Hobbes, 1839–1845; Plutarch, 2001). When participants were asked to pick which of the two ships (old parts or new) is the Ship of Theseus, most participants picked the object with the old parts rather than “both”, suggesting they viewed just one object as continuous with the original ship.

Beyond developmental studies and thought experiments, the notion of transitive identity is supported by studies in perceptual psychology. When confronted with a simple ‘fission’ display of one object splitting into two, both adults and infants have a difficult time maintaining a coherent object representation, as evidenced by a reduction in the object-specific preview benefit in adults (Mitroff et al., 2004) and an inability to recognize or remember the correct number of objects in infants (Cherries et al., 2008). These findings suggest that the mind cannot psychologically represent one object as being equivalent to two individuals, but is perceptually

sensitive to cues—such as the closure of an object by a single object boundary—that indicate the spatiotemporal persistence of a singular object over time (Cherries et al., 2009). Likewise, adults find it much more challenging to track multiple objects that disintegrate into particles than they do multiple objects which maintain their spatiotemporally cohesive object boundaries (vanMarle & Scholl, 2003).

Even seeming violations of transitivity in perception, such as apparent motion (seeing two closely-timed flashes as a single flash that moves from a first point to a second) (Ullman, 1979) can be explained as the perceptual system's attempt to use informative cues like spatiotemporal priority in order to construct transitive percepts of single, persisting individuals (Flombaum et al., 2009), rather than as intransitive perceptual operations. Another related example is object persistence through occlusion, in which the object that leaves occlusion is irresistibly perceived as the same object that entered occlusion, provided it disoccludes within a certain spatiotemporal window relative to the occlusion event (De Freitas et al., 2016; Scholl & Pylyshyn, 1999). Because spatiotemporal continuity is a highly informative indicator that you are looking at the same object as it becomes occluded, the visual system appears to have mechanisms built into it for unifying independent, spatiotemporally proximal percepts into that of a single, persisting individual.

The Case for Intransitive Identity

Despite these previous findings suggesting transitivity, other work involving thought experiments in cognitive psychology and philosophy has argued that people may have an intransitive concept of identity. Most of this work asks participants to consider thought experiments in which there are multiple contenders for identity, as in (i) the already mentioned Ship of Theseus problem (Hobbes, 1839–1845; Plutarch, 2001), (ii) brain swapping scenarios

wherein the brains of two individuals are transplanted into each other's bodies and participants are asked whether the original person's body or brain is the continuer of the original, and (iii) fission cases, in which one person splits into two (Parfit, 1984; Perry, 1972).

In work using the Ship of Theseus thought experiment, one study found that roughly 50% of adult participants pick each option, i.e., ship with the old parts and new parts (Cakar & Hohenberger, 2015; Rose et al., 2020). These findings are consistent with an intransitive concept of identity, provided we interpret the choices as reflecting a belief that both ships are identical to the original, although we note that this study did not provide an explicit "both" option. In fact, another study using this thought experiment found that some participants interpret phrases like "Theseus's ship" as more of a description (i.e., a ship owned by Theseus) that could logically apply to more than one entity, rather than as a proper name that can apply to only a single individual (Marchak & Hall, 2019), suggesting that the term is ambiguous unless its intended meaning is clarified.

More support for intransitive identity comes from thought experiments in which brains are swapped between bodies. One study invited participants to imagine that the brain of Jane is transplanted into the body of Anne, and vice versa (Protzko et al., 2023). When participants were asked which of Jane's mind or body was allowed to use the gym membership that Jane originally paid for, there was a trimodal pattern—some picked the body, others the mind, and most picked 'both', consistent with intransitivity (Protzko et al., 2023). Another study asked participants to imagine a future in which someone with traumatic brain damage named Tom was able to heal his brain by receiving a donation of brain tissue from another deceased individual (Finlay & Starmans, 2022). The study found that most participants only agreed that Tom's identity had

changed when he received 100% of the donor's brain tissue, but not even when he received 75%, perhaps suggesting that they thought identity resides in both the mind and the body.

An intransitive concept of identity is also supported by studies that employ fission thought experiments, in which an object, animal, or person splits into two, and then participants are asked which of the resulting copies is the original (Parfit, 1984). Unlike much of the above work, these studies have tended to provide explicit "both" and "neither" options, perhaps because the identical properties of the copies make the "both" option more salient to experimenters. Surprisingly, as many as *half* of participants in one study selected "both" when asked which of two pieces of a split iceberg was continuous with the original iceberg (Rips, 2011), and another study found similar results when participants considered an animal that split into two copies (Rips et al., 2006).

In the context of personal identity, another fission study asked participants to imagine a protagonist named Derek, who is copied and then simultaneously beamed to two planets using a Star Trek-like teletransporter, such that one copy hugs his mother on Mars and the other copy hugs his wife on Venus (Weaver & Turri, 2018). An impressive 70% of participants chose the option, *Derek hugged his wife and Derek hugged his mother*, instead of the other available options (*Derek hugged his wife but someone else hugged his mother*; *Derek hugged his mother but someone else hugged his wife*; or *Someone else hugged both Derek's wife and his mother*). Along similar lines, another study asked people to imagine that they would be split into two copies, then asked how they would distribute a sum of \$100,000 between the copies (Woike et al., 2020). Most participants split the money evenly, rather than giving all the money to just one of the two copies, seemingly violating the transitivity criterion.

Finally, beyond work squarely focused on personal identity, analogous claims implying an intransitive concept of personal identity have been made in other literatures. For example, one study on self-localization concludes: “healthy humans can bi-locate in two different bodies at the same time...participants experience a sense of being split in two selves” (Aymerich-Franch et al., 2016, pp. 105-106), and another concludes, “self-location may...be present in both the gamer and avatar simultaneously... Both perspectives appeared therefore to be processed at the same time” (Furlanetto et al., 2013, pp. 3-4). Similar claims have also been made in philosophical arguments for ‘pluralistic’ selves (Shoemaker, 2016; Shoemaker, 2007; Tierney, 2020; Tierney et al., 2014; Williams, 1970).

A Second Look at ‘Intransitive’ Personal Identity

Before concluding that these recent studies revise our understanding of the personal identity concept in favor of intransitivity, we need to ask whether these judgments are truly based on an underlying, intuitive conviction that identity is intransitive. An alternative possibility is that participants in these studies are not interpreting the questions as intended. Another alternative is that, while they endorse intransitivity, they are still unable to represent this notion psychologically—much as people are incapable of intuitively comprehending the idea of four or more dimensions, while appreciating that these dimensions can still exist. If these possibilities are true, then this would mean that existing findings of intransitive judgments do not provide sufficient basis for claims like, “ordinary judgments about personal identity are not constrained by the one-person-one-place” (Weaver & Turri, 2018, p. 236) or that “many participants identified the original with both continuers simultaneously, violating the transitivity of identity relations” (Woike et al., 2020, p. 1)

The Present Studies

For these reasons, the present studies aim to arbitrate between a transitive vs. intransitive concept of identity. Since some of the strongest arguments for intransitive identity have come from thought experiments involving personal identity, we focus on such cases as well.

Despite often employing hypothetical scenarios, thought experiments have several features suitable for our purposes. First, they can be constructed in a controlled manner to reveal the inferences underlying psychological representation of identity, much as visual illusions and other ‘artificial’ stimuli can be used to understand those psychological processes. Second, although technology today is increasingly enabling people to clone aspects of their likeness, voice, etc., it is not yet at the point where people can create perfect copies of themselves or others. As such, thought experiments provide the means to get at issues of intransitivity by having people consider full-blown replicas.

Whereas the scenarios we ask participants to consider are hypothetical, it is important to underscore that representations of identity are psychologically fundamental and crucial for normal everyday functioning (Kuhlmeier et al., 2004; Spelke et al., 1995). As a recent review concludes, identity persistence has “consequences for attitudes and judgments or decisions, motivation, intentions and behavior, and psychological and physical health” (Sedikides et al., 2023).

In Experiments 1-3, we ask participants to imagine a scenario in which they are duplicated. Since a transitive concept allows only one agent to be identical to the original, it follows that all causally entitled contenders for the identity must compete for this single identity representation. Thus, if one contender competes more strongly, it may surpass the activation threshold for being represented as that identity. Similarly, if one contender competes more

weakly, it may lose to another contender for representation as that identity (Nozick, 1981). To test these predictions within our paradigm, we ask whether participants are more likely to identify with the copy if the original is killed off (thereby weakening how much the original competes for the identity representation) or if they are asked to imagine themselves in the first-person perspective of the copy (thereby strengthening how much the copy competes for the identity representation). Likewise, if people employ a transitive concept of personal identity, then we would also expect those who identify with the copy after the original is killed to *switch back* to identifying with the original if it is revived, since this revival should cause the original to compete once again for the identity representation (Experiment 3). To preview our results, across experiments we find that these manipulations do indeed affect people's judgments, consistent with a transitive concept of personal identity.

Finally, Experiments 4-5 provide the strongest test of the transitive view by considering cases of fission, in which participants are asked to imagine splitting into two copies while their original bodies disappear. Experiment 4 asks whether participants are interpreting the identity question differently than intended, by testing whether they are less likely to make intransitive judgments when asked whether both individuals are entitled to the practical matters of the original (e.g., sharing a bedroom with the original's spouse). Experiment 5 then probes whether they are truly capable of intransitive psychological representations, by investigating whether they imagine fission scenarios by mentally simulating multiple simultaneous perspectives (consistent with intransitivity) or only a single perspective (consistent with transitivity).

All procedures were approved by the Harvard Internal Review Board, and data and code for all studies are publicly available here: https://github.com/Ethical-Intelligence-Lab/self_duplication.

Experiment 1: Do People Personally Identify with Two Co-Existing Selves?

Experiment 1 asked participants to imagine a hypothetical scenario in which they are duplicated. In theory, an intransitive concept of personal identity allows participants who imagine such a scenario to identify with both the original and the duplicate simultaneously, whereas a transitive concept allows them to identify with at most one of these individuals at a time. In line with previous work (Hood et al., 2012), in this scenario we expected that by default people start out preferentially identifying with the original, given their longer shared causal history with it (i.e., until the point of duplication, they have always been the original). Importantly, we also tested additional predictions that fall out of a transitive view of personal identity: participants should be more likely to switch to identifying with the copy if their association with it is strengthened (accomplished by asking participants to imagine themselves from the copy's perspective), or if their association with the original is weakened (accomplished by killing off the original).

Method

Participants. We recruited 377 participants and excluded 27 for failing any of the attention or comprehension checks, yielding a final sample of 350 ($M_{\text{age}} = 36$). 175 participants reported their gender as female, 171 as male, and 4 as other. Participants in all studies were recruited from the online crowdsourcing platform, Amazon's Mechanical Turk, to take part in a 5-min study for \$0.50 compensation. In the absence of any previous studies that had used the paradigm created for the present experiments, we used the heuristic assumption that the resulting effect sizes would be comparable to the last study that the first author published on personal identity (De Freitas & Cikara, 2018): 100 participants per condition.

Participants were assigned to one of four conditions in a 2 (original: dead vs. alive) x 2 (perspective: first vs. third) design. Sample sizes in each condition were: alive-first ($n=93$), alive-third ($n=84$), dead-first ($n=85$), dead-third ($n=88$).

Procedure. Participants read the following vignette:

“Imagine you are living in a future where scientists have figured out how to make a perfect copy of the human body and brain of a person at a given point in his or her life. You are invited into a laboratory, where a perfect copy of you is made. The scientists place the copy of you in front of the original you, and the copy of you looks identical to the original you.”

To investigate the effect of weakening associations with the original, we manipulated the original’s state (alive v. dead), telling them either that “The copy of you then awakens” or “The scientists then kill the original you. Instantly after the original you dies, the copy of you then awakens.” To investigate the effect of strengthening associations with the copy, we manipulated the participant’s perspective of the copy (third v. first), telling them either that “The scientists inspect the copy, and confirm that all of its physical and psychological traits are identical to that of the original you”, or “Here’s what it’s like to wake up as the copy: It feels just like you, but that you somehow switched bodies. You look back at your original living body and can’t believe that you were just there. You no longer feel like you’re there. You feel like you are here, in the copy’s body.”

Participants then answered the identity question, “At the end of this whole procedure, who do you believe is you?”, by choosing from one of four options, “The original you, the copy, neither, both.”

Finally, they answered two comprehension checks about the original’s state (“According to the story, which of the following is true about the original you? 1. The original you remains alive. 2. The scientists kill the original you. 3. The original you runs away”) and the basic

scenario (“According to the story, how many copies of the original you are made? No copies. One copy; Two copies”), followed by basic demographics items on their age and sex (multiple choice, with the options ‘male’, ‘female’, and ‘other’).

Results and Discussion

The proportion of “both” responses varied across conditions, but was not the modal response in any condition, and was dramatically lower than the 70% proportion found in previous work (Weaver & Turri, 2018; Woike et al., 2020). When the original was still alive and participants were given a third-person perspective of the situation (the baseline condition), 87% chose original, and 13% (i.e., the rest) chose both, indicating that people largely identified with just one individual.

Furthermore, our manipulations of whether one individual or the other competed more strongly for identity representation affected participant’s choices. Because participants chose one of four options, we analyzed these effects using multinomial logistic regression (Hosmer Jr et al., 2013). Consistent with transitive personal identity, weakening ties to the original (by killing it off) made participants more likely to identify with the copy (“copy responses” when original is dead: 24% vs. when alive: 15%; $\beta = 1.00$, $SE = 0.32$, $p = .002$), and strengthening ties to the copy by providing its first-hand perspective increased “copy” responses (first-hand: 33%, third-hand: 6%; $\beta = 2.57$, $SE = 0.38$, $p < .001$; Figure 1).

We also found that participants were significantly more likely to identify with both (vs. the original) when given the first-hand perspective of the copy vs. not (“both” responses when given first-hand: 25%; third-hand: 13%; $\beta = 1.45$, $SE = 0.30$, $p < .001$), a result that is consistent with the idea that some participants employed an intransitive concept of identity. An alternative interpretation of this result, however, is that these participants interpreted “both” as meaning

serially (rather than simultaneously) identical, i.e., “I once was the original, but now am the copy”. Although the question wording strove to prevent this interpretation by specifying that it was asking for judgments at the *end* of the duplication procedure, we cannot rule out the possibility that some participants had this interpretation. There were no other statistically significant effects (all $ps > 0.09$; Figure 1).

Finally, as far as we know Experiment 1 is the first time a perspective manipulation has been used to increase *personal* identity for oneself (i.e., whether one feels numerically identical to one person or another) as opposed to empathy with others (Galinsky & Moskowitz, 2000; Richardson et al., 1994). In fact, perspective taking had the biggest effect of our two identity switching manipulations. For this reason, Experiment 2 follows up by asking what it is about this manipulation that swayed participant’s judgments.

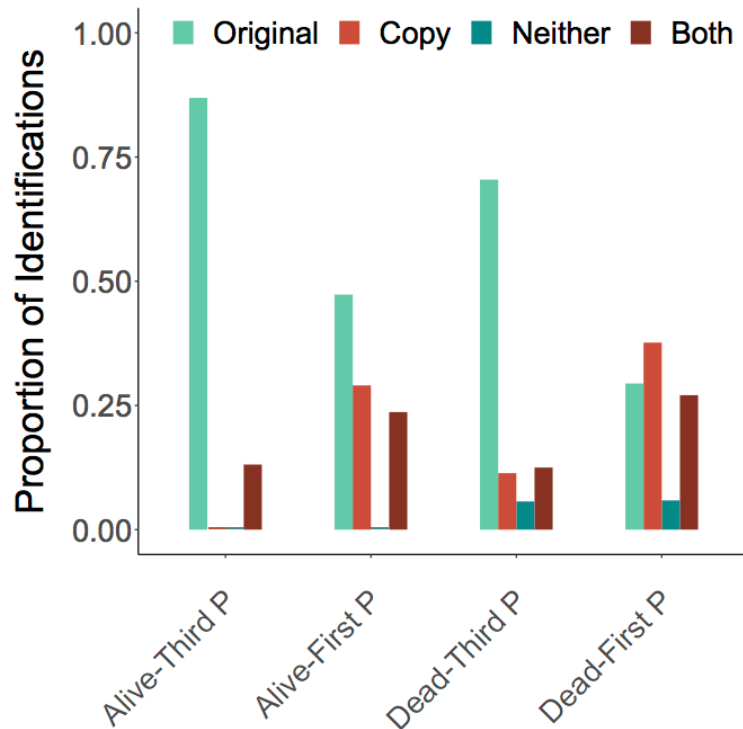


Figure 1. Proportions of participants who chose each option in Experiment 1. Alive/Dead refers to whether the original is alive or dead. First/Third P refers to whether participants take a first or third-person perspective of the copy.

Experiment 2: How Does the Perspective Manipulation Affect Personal Identity?

We see three potential explanations of how the perspective manipulation increased personal identification with the copy in Experiment 1: (1) as the name of the manipulation suggests, imagining the copy's perspective induces one to represent the copy as oneself (Ames et al., 2008), (2) being told that the copy *feels* like the original suggests that the copy has equivalent mental states to the original, or (3) being told that the copy felt as though it had just been where the original was suggests that the copy not only feels like the original but feels *psychologically continuous* with it. These three explanations can be likened to psychologically traversing ever closer to experiencing the other as if from one's own first-person perspective, such that one not only sees the world from the other's perspective (1), but also feels what they feel (2) and appreciates the continuity of experience from the original's location to that of the copy (3). Broadly in line with construal level theory (Libby & Eibach, 2002; Trope & Liberman, 2010), in which people provide more vivid reports when imagining the same activity from a first-person perspective than a more 'distant' third-person one (Libby & Eibach, 2002), we expect that people are more likely to personally identify with the copy as they move from (1) to (3).

Method

Participants. We recruited 291 participants and excluded 29 for failing comprehension checks, yielding a final sample of 262 ($M_{\text{age}} = 33$). 142 participants reported their gender as

female, 117 male, 2 as other, and 1 was unspecified. Participants were assigned to one of three conditions (empathy only, perspective only, and psychological continuity) in a between-subjects design. The sample sizes in these conditions were empathy only ($n=84$), perspective only ($n=99$) and psychological continuity ($n=79$).

Procedure. We tested the effect of each of factors (1)-(3) relative to a control condition that featured no perspective manipulation (the same condition from Experiment 1 in which the original is killed and participants are simply informed that the duplication succeeded, without being told what it is like from the copy's perspective). Thus, this study compares three new conditions to this control condition from Experiment 1.

We manipulated factors (1) to (3) between-subjects: (1) In the *perspective only* condition, instead of reading the whole description of the copy's first-hand perspective from Experiment 1, participants were simply asked to adopt the copy's perspective, using a standard prompt adapted from previous perspective-taking studies (Ames et al., 2008; Galinsky & Moskowitz, 2000): "Imagine for a moment that you are the copy, standing in their shoes and seeing the world through their eyes. Think about how you, as the copy, would experience the event."

(2) In the *feels like original* condition, participants read only the first sentence of the original perspective manipulation (without being asked to imagine being in the shoes of the copy): "Here's what it's like to wake up as the copy: It feels just like you, but that you somehow switched bodies."

(3) In the *psychological continuity* condition, they read the whole perspective manipulation from Experiment 1, including the stipulation that the copy felt as though it had just been where the original is. As in the control condition taken from Experiment 1, all three of these

conditions specified that the original was killed. The comprehension and demographics questions were identical to Experiment 1.

Results and Discussion

Since each condition can be thought of as cumulatively adding a factor to the ones before it, our prediction (in light of Experiment 1) was that each subsequent condition would further increase the choices of “copy” and “both”. Therefore, we ran a multinomial logistic regression with the type of perspective condition (1=control, 2=perspective only, 3=feels like original, 4=continuous with original) coded as an ordinal predictor. We found that as the perspective conditions increased in comprehensiveness, participants were increasingly likely to choose “copy” (control: 11%; perspective only: 26%; feels like original: 22%; continuous with original: 29%; $\beta = .46$, $SE = 0.13$, $p < .001$) and “both” (control: 13%; perspective only: 15%; feels like original: 27%; continuous with original: 32%; $\beta = .56$, $SE = 0.13$, $p < .001$). The perspective manipulations did not affect how much participants chose “neither” ($p = .835$; Figure 2).

While these results do not further arbitrate between transitive and intransitive personal identity, they do reveal that several aspects of the perspective manipulation used in the prior experiments swayed people toward identifying with the copy, including the act of imagining the copy’s perspective, learning that it had the same feelings as the original, and learning that it felt psychologically continuous with the original. These results also provide support for our starting assumption that judgments about one’s own personal identity tap into the same cognitive process normally implicated in social perspective-taking tasks, given that we could use a perspective taking manipulation to influence *personal* identification; to our knowledge, this is the first time that a perspective manipulation has been found to induce this effect.

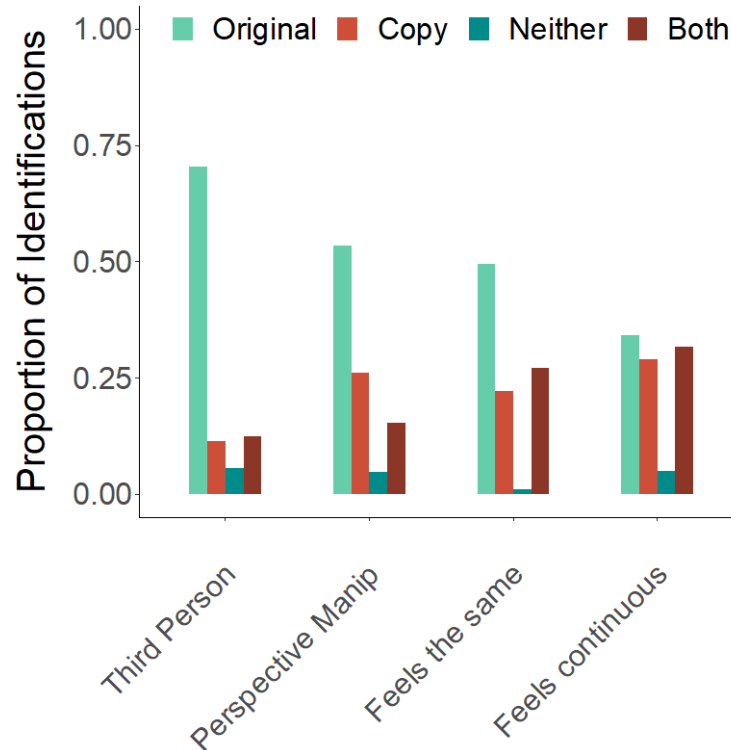


Figure 2. Proportions of participants who chose each option in Experiment 2. Note: “Third Person” refers to the condition from Experiment 1 in which the original is killed and participants are simply informed that the duplication succeeded, without being told what it is like from the copy’s perspective.

Experiment 3: Reviving the Original

A transitive concept of personal identity should also allow those who switch to identifying with the copy when the original dies to *switch back* to identifying with the original if it is *revived*, because reviving the original leads it to newly compete for identity representation. If people employ an intransitive concept of personal identity, however, then reviving the original should make them more likely to identify with “both” but not also with the original.

Method

Participants. We recruited 203 participants and excluded 27 for failing any of the checks, yielding a final sample of 176 ($M_{\text{age}} = 35$). 76 participants reported their gender as female, 99 as male, and 1 as other. Participants took part in two conditions in a within-subject design.

Procedure. We wanted as many participants as possible to start out identifying with the *copy*, so we re-ran the condition from Experiment 1 in which the original is killed and participants are given a full, first-hand description of what it is like from the perspective of the copy (i.e., Experiment 2's 'feels continuous' condition). Otherwise, Experiment 3 employed a similar design to Experiments 1-2. On the first page, participants read the vignette and answered the identity question as before. On the second page, they read that the original was revived ("The scientists decide to bring the original you back to life, by using an advanced electrical stimulation technique. The original awakens, and is functioning perfectly. Therefore, both the original you and the copy are now alive"). After reading this new information, they were asked to answer the identity question again.

Comprehension items were as in Experiment 2, except the third option for the first comprehension check was changed to "The original you is killed then brought back to life", given that this experiment manipulated whether the original was revived after being killed off.

Results and Discussion

We ran a multinomial logistic regression with condition (original dead or revived) as a factor and found that "copy" responses decreased when participants learned that the original was revived (dead: 32% vs. revived: 16%; $\beta = -.89$, $SE = 0.29$, $p = .002$), whereas reviving the original did not affect how many participants chose both or neither ($ps = .946$ and $.494$; Figure 3). Consistent with a transitive concept of personal identity, these results suggest that when the

original was revived participants felt that they had to choose between *either the copy or the original*, leading some participants to choose the original.

Further, we focused on the subset of participants who initially identified with copy and found that after the original was revived 33.93% of them continued to identify with the copy, 1.79% switched to neither, 23.21% switched to original, and 41.07% switched to both. Thus, a sizeable number switched to identifying with the original, in line with transitive personal identity.

With that said, it is notable that a sizeable subset of participants switched to identifying with both, even more than those who switched to identifying with the original. This pattern is consistent with intransitive personal identity. At the same time, intransitive personal identity cannot explain why a sizeable number of participants switched to original.

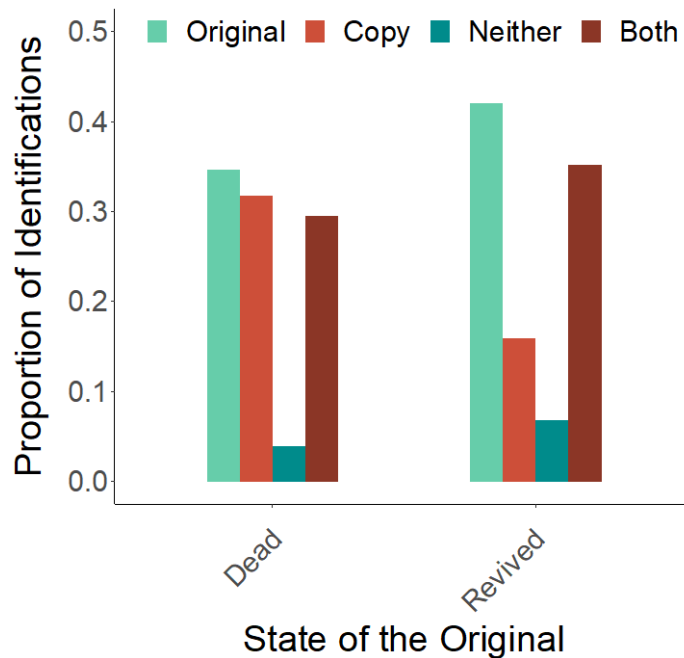


Figure 3. Proportions of participants who chose each option in Experiment 3.

Experiment 4: Fission and Practical Matters

Although so far the findings are consistent with a transitive concept of personal identity, a plausible interpretation of the “intransitive” view is that “both” responses might only be reliably elicited under sufficiently stringent conditions, where the two candidates are not only similar but also neither of them is significantly “better” than the other. Such an argument cannot account for why some participants were more likely to personally identify with “copy” (i.e., the ‘worse’ candidate) when conditions were favorable to it in Experiments 1-3. Why would a change in perspective render the copy better? Even so, fission cases could provide a stronger test of the “intransitive” view, given that the copies are completely identical in historical respects as well (Weaver & Turri, 2018).

One possible explanation of intransitive judgments in fission cases is that participants in prior work were not interpreting questions about identity as intended. Perhaps participants were responding to the superficial similarities between the two candidates rather than thinking through the implications of intransitive responses. Since personal identity is also connected to practical matters like obligations and privileges (e.g., whether one owes a loan, or can be intimate with certain people) (Protzko et al., 2023), here we probe intuitions about identity by testing who is viewed as having the original person’s privileges after a fission event occurs. Specifically, we reason that, while participants might be willing to endorse that both copies in a fission event are Derek, they will feel less sure when it comes to asking which copy is allowed to be intimate with Derek’s loved one. Such a pattern would suggest that “both” choices might not reflect a deeper psychological commitment to intransitive identity. Although it may not be possible to eliminate all responses based on similarity, a focus on substantive matters may at least make this dependence less frequent.

Method

Participants. We recruited 319 participants and excluded 50 for failing a comprehension check, yielding a final sample of 269 ($M_{\text{age}} = 37$). 94 participants reported their gender as female, 174 male, and 1 preferred not to disclose. Participants were assigned to one of two conditions (scenario: identity, privilege) in a between-subjects design. The sample size in each condition was $n=138$ (identity) and $n=131$ (privilege).

Procedure. Participants were shown a scenario closely modeled after Weaver and Turri (2018). The identity and privilege scenarios were identical, except that the privilege scenario included the section in squared parentheses:

“The year is 2450 and human civilization has advanced so far that we could barely comprehend it. Derek is currently at his home. His friend is 5 miles north of his house. His relative is 5 miles south of his house.

Derek enters the Quantum Teletransporter in his house. Then he presses the button. In an instant, the quantum device scans his body and records the exact state of all his cells and brain states.

Instantly, the information travels through an information wormhole to his friend’s house and his relative’s house, where it is perfectly reconstituted in physical form. Instantly, a person steps out of the Teletransporter in Derek’s friend’s house. With a smile, there is a hug and she says, “My friend! I’m so happy to see you!”.

Simultaneously a person steps out of the Teletransporter in Derek’s relative’s house. With a smile, there is a hug and she says, “My relative! I’m so happy to see you!” Derek’s house is now empty.

[Later that evening, the two figures both return to Derek’s house, where Derek’s wife is getting ready to go to bed.]”

In the identity condition, participants answered the question: “At the end of the story, who is Derek”, whereas in the privilege condition they answered, “At the end of the story, who can get into bed with Derek’s wife?” The answer options for both these questions were: “Only the

person from his friend’s house”; “Only the person from his relative’s house”; “Both persons”; and “I am not sure”.

Finally, they answered a comprehension check (“In the story, you considered a scenario involving: Plane travel; Teletransportation; Farming”), followed by demographics questions.

Results and Discussion

We ran a multinomial logistic regression with condition (identity or privilege) as a factor and found that “both” responses decreased when participants answered the practical matters question (identity: 62% vs. privilege: 57%; $\beta = -.79$, $SE = 0.37$, $p = .030$; Figure 4). So, when participants reckoned with Derek’s practical matters, fewer felt that both copies were entitled to the same relationship as the original. At the same time, it is noteworthy that in this fission context “both” was still the modal response in both conditions. Experiment 5 investigates this pattern further.

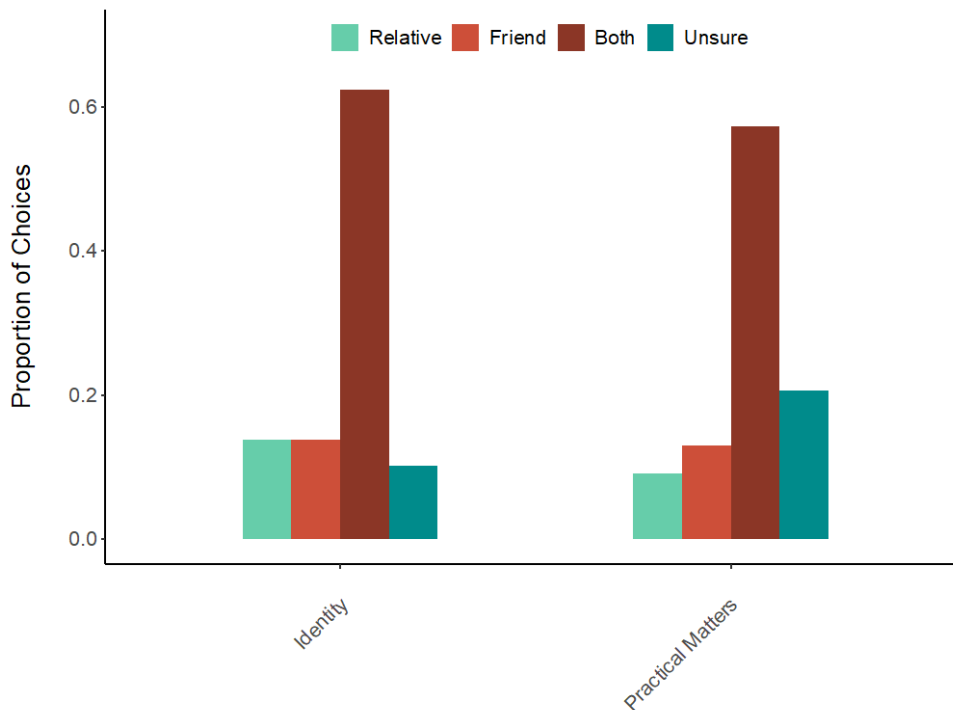


Figure 4. Proportions of participants who chose each option in Experiment 4.

Experiment 5: Do People Employ Single or Multiple Perspectives to Answer Fission Cases?

Experiment 4 finds, consistent with prior work, that a large proportion of participants choose “both”. But do these responses reflect an actual cognitive ability to simultaneously represent two individuals as having the same identity? For instance, in the context of yourself, can you simultaneously represent both copies as ‘you’? Experiment 5 investigates this question by asking participants how they imagine a fission scenario.

Method

Participants. We recruited 299 participants and excluded 91 for failing a comprehension check, yielding a final sample of 208 ($M_{\text{age}} = 31$). 59 participants reported their gender as female, and 149 as male. Participants were assigned to one of two conditions (scenario: basic, perspective) in a between-subjects design. The sample size in each condition was $n=105$ (basic) and $n=103$ (perspective).

Procedure. Participants were shown a scenario similar to that in Experiment 4, except that the protagonist was oneself rather than Derek, and there was no mention of returning to Derek’s wife. In the Basic condition, they answered the extent to which they agreed, “At the end of the story, I exist at both houses at the same time” (0 = Completely Disagree, 100 = Completely agree).

In the perspective condition, they were first asked to “Imagine experiencing both houses simultaneously: what do you see, hear, and smell?” As a means of ensuring they followed these instructions, we asked them to also type what they see, hear, and smell. Thereafter, we probed the nature of their imagined perspective by asking “Which of the following images best captures what you were able to imagine?”. They were shown the images in Figure 5 as options, with the

order of images randomized on the page. The images were accompanied by the following descriptions:

- *Two consciousnesses*: You felt that you were two consciousnesses seeing the world at the same time.
- *Single consciousness, flipping*: You felt that you were a single consciousness flipping back and forth between seeing the world from A and B.
- *Single consciousness, simultaneous*: You felt that you were a single consciousness seeing a single superimposed version of the two scenes.
- *Single consciousness, split screen*: You felt that you were a single consciousness outside of the two scenes, seeing the two scenes side-by-side, like a split-screen.

The first option captures an intransitive perspective, whereas the last three options capture variants of a transitive perspective. Thereafter, they answered the same scaled identity item as the basic condition. Participants in both conditions answered the same comprehension check and demographic items from Experiment 4.

Two consciousnesses



You₁



You₂

Single consciousnesses, flipping



You-at-time₁



You-at-time₂

Single consciousness, simultaneous



You

Single consciousness, split screen



You

Figure 5. Image options presented in Experiment 5. Images generated by OpenAI’s DALL-E model.

Results and Discussion

Consistent with Experiment 4, participants generally agreed they were in ‘both’ places and were significantly more likely to do so in the perspective condition ($M = 78.49$, $SD = 22.68$) than basic condition ($M = 70.54$, $SD = 22.83$), $t(204) = 2.36$, $p = 0.019$, $d = 0.33$. Asking participants in the perspective condition to imagine being both individuals simultaneously likely increased agreement with the “both” item.

However, only 26.2% of participants in the perspective condition said they experienced being two consciousnesses simultaneously. Among the 73.8% who said they experienced being a single consciousness, the most to least common options were: split screen (30.1%), flipping (24.3%), and simultaneous (19.4%).

Furthermore, based on these choices, mean agreement with the “both” rated item was as follows. Two consciousnesses simultaneously: $M = 71.00$, $SD = 21.68$; single consciousness, split screen: $M = 79.97$, $SD = 21.95$; single consciousness, flipping: $M = 81.96$, $SD = 21.10$; single consciousness, simultaneous: $M = 81.95$, $SD = 26.14$. One possible interpretation of this result is that those who employed a single consciousness to imagine the scenario were more likely to identify with each individual separately, leading them to judge that they identified with both individuals—even if they were incapable of representing themselves as two. As an analogue, if one saccades across a scene and sees each object in it, one might feel he is aware of all objects in the scene even if the brain is not capable of being aware of all objects simultaneously.

60% of participants complied with the instructions to type what they see, hear and smell. Of those who did, 100% provided descriptions compatible with the type of perspective they chose. For instance, a participant who said, “In my friend’s house I see my friend cooking. In my relative’s house I see my relative sweeping” chose the “single consciousness, split screen option”.²

The results suggest that, although participants provided intransitive *judgments*, these judgments were not based on an intuitive experience of identity as intransitive.

General Discussion

We started by noting that some studies have claimed that the mind employs an intransitive concept of identity, which makes it psychologically coherent to say that an entity at t_0 is identical to *two* separate individuals at t_1 . If true, this would entail a revolution to our understanding of the psychological representation of identity, overturning assumptions and empirical findings across both high-level cognition (De Freitas et al., 2017; Hall et al., 2003; Sorrentino, 2001) and perception (Cherries et al., 2009; Mitroff et al., 2004; vanMarle & Scholl, 2003). As such, we took a second look at whether seemingly intransitive *judgments* are truly based on an underlying, intuitive conviction that identity is intransitive, by employing manipulations that were designed to dissociate the transitive and intransitive accounts.

Inspired by the transitive account, Experiments 1-3 manipulated several factors that should affect whether one of two contenders for identity representation competes more strongly

² Note that all results show a consistent pattern and remain statistically significant if we repeat the tests with only this 60% subset. Agreement with “both” item in the perspective condition ($M = 79.06$, $SD = 18.81$) versus basic condition ($M = 70.54$, $SD = 25.83$), $t(158) = 2.45$, $p = 0.015$, $d = 0.36$. Choice proportions: two consciousnesses simultaneously (27.4%), split screen (32.3%), flipping (22.6%), and simultaneous (17.7%). Mean agreements with the “both” item based on these choices were: Two consciousnesses simultaneously: $M = 74.88$, $SD = 16.35$; single consciousness, split screen: $M = 78.35$, $SD = 20.45$; single consciousness, flipping: $M = 78.79$, $SD = 21.48$; single consciousness, simultaneous: $M = 87.18$, $SD = 15.25$.

or weakly. We found several results consistent with transitivity: unlike previous studies, the modal response was not “both”, and identity judgments were influenced by factors such as whether the original was killed, whether participants were asked to imagine themselves in the first-person perspective of the copy, and whether the original was revived after being killed.

Experiments 4-5 then provided further tests of a transitive account, by revisiting fission cases in which a person splits into two copies and the original disappears. Experiment 4 found that a smaller proportion of people chose “both” when judging whether the copies are entitled to the original’s privileges, as opposed to when making identity judgments. At the same time, and consistent with prior work, a large proportion still picked “both”. Experiment 5 then found that even when most participants agree they were both individuals, they do not imagine experiencing the world from two perspectives simultaneously, suggesting that their intransitive judgments do not reflect an underlying conviction that personal identity is intransitive.

Why Do People Pick Both?

Why do so many participants choose or agree with “both” in the fission scenario, despite not being able to psychologically represent this reality? The fact that far fewer do so when an original is still present, suggests that participants are inclined to pick just one individual, but might not know who to choose in fission scenarios, wherein both individuals are equal in all causal respects (Nozick, 1981; Rips, 2011). People might pick or agree with “both” because they appreciate that each contender is *causally entitled* to the original—having the qualifications needed to be the original, if the other contender were no longer around (Nozick, 1981; Rips et al., 2006)—even though they may believe that only *one* continuer can be numerically identical to the original, aka the ‘uniqueness’ criterion (Nozick, 1981).

Another possibility is that participants truly believe that they would be both individuals in a fission scenario, even though they realize at a different level the logical constraints on such beliefs. Much as people cannot visualize four dimensions given cognitive limits while still logically appreciating that four or more dimensions may exist, people may make intransitive identity judgments while intuitively being limited to a transitive representation of personal identity. Interestingly, philosophers of identity like Derek Parfit have also logically concluded that it would be correct to say that fission produces two individuals who are equally identical to the original (Parfit, 1984)—at least at the exact moment of duplication, before each individual begins to accumulate new experiences that lead their identities to diverge. The fact that such arguments were considered revolutionary might further underscore the notion that they are not based on how people intuitively represent personal identity.

A related question is why Experiments 1-3 found that “both” was not the modal response and Experiments 4-5 found that the modal response was “both” (though note that Experiment 5 was not forced choice). Part of the reason may be due to the different paradigms of cloning versus fission: in the cloning case the original is still present, and so may compete more strongly for identity, preventing as many participants from picking both. In the fission scenario, in contrast, each individual is qualitatively identical, so there is no reason to favor either a priori. But another reason, suggested by our results in Experiment 5, is that when participants agree with “both” they mean that they identified with each individual in serial, not simultaneously. This interpretation is consistent with our findings that those who used a single perspective were more likely to agree with the “both” statement. Again, this finding underscores that people’s ratings were not based on how they actually represented personal identity.

Does Intransitive Personal Identity Exist?

Experiment 5 found that only 26% of participants report experiencing two perspectives simultaneously. Given that everyone else reported experiencing one perspective, one interpretation is that this 26% of participants is simply mistaken. Another way to disambiguate this issue would be to employ behavioral or neuroscientific methods to determine whether these participants are truly representing multiple selves in parallel, independent of their subjective reports. For example, behavioral work finds that self-relevant items are prioritized in perception and working memory—the so-called self-reference effect (Keyes & Brady, 2010; Knoblich & Flach, 2001; Macrae et al., 2017; Rogers et al., 1977). If this 26% of participants is truly representing multiple selves simultaneously, then they should experience self-reference effects for *both* individuals in the scenario, leading them to perform better overall compared to participants who report experiencing just one perspective (consistent with intransitivity). But if these participants are in fact utilizing a single perspective, then they should perform just as well as other participants overall (consistent with transitivity).

Constraints on Generality

Our studies involved large samples of online participants using a diverse set of paradigms. Thus, we expect the results to generalize to all adults who pass the attention and comprehension checks. Given our theoretical framework, which posits that our effects arise from a fundamental way in which people think about numerical identity, we expect the results to be reproducible with similarly large samples of in-person participants. However, we do not have evidence for such effects. Our studies were limited to cases involving duplicates of persons, so might not generalize to cases involving duplicates of other entities, like objects or animals. We have no

reason to believe that the results depend on other characteristics of the participants, materials, or context.

Transparency and Openness

Data and code for all studies are publicly available here: https://github.com/Ethical-Intelligence-Lab/self_duplication.

Author Note

For helpful feedback, thanks to Nathan Couch, Tomer Ullman, and Steven Pinker. Ideas from this paper have not been previously disseminated. Both Julian De Freitas and Lance J. Rips were responsible for conceptualization, and methodology. Julian conducted the investigation and wrote the original draft, then both authors reviewed and edited 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.
- Aymerich-Franch, L., Petit, D., Ganesh, G., & Kheddar, A. (2016). The second me: Seeing the real body during humanoid robot embodiment produces an illusion of bi-location. *Consciousness and Cognition, 46*, 99-109.

- Blok, S., Newman, G., Behr, J., & Rips, L. J. (2001). Inferences about personal identity. Proceedings of the Cognitive Science Society, Edinburgh, Scotland.
- Cakar, T., & Hohenberger, A. (2015). The Ship of Theseus: Using Mathematical and Computational Models for Predicting Identity Judgments. *Proceedings of the Cognitive Sciences Society*.
- Cherries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2008). Cohesion as a constraint on object persistence in infancy. *Developmental Science*, *11*(3), 427-432.
- Cherries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2009). Do the same principles constrain persisting object representation in infant cognition and adult perception? The cases of continuity and cohesion. In B. Hood & L. Santos (Eds.), *The origins of object knowledge* (pp. 107-134). Oxford University Press.
- Church, A. (1956). *Introduction to mathematical logic* (Vol. 1). Princeton University Press.
- 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., Myers, N. E., & Nobre, A. C. (2016). Tracking the changing feature of a moving object. *Journal of Vision*, *16*(3), 22-22.
- 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.
- De Freitas, J., Tobia, K. P., Newman, G. E., & Knobe, J. (2017). Normative judgments and individual essence. *Cognitive Science*, *41*, 382-402.

- Dennett, D. C. (2014). The self as the center of narrative gravity. In *Self and consciousness* (pp. 111-123). Psychology Press.
- Finlay, M., & Starmans, C. (2022). Not the same same: Distinguishing between similarity and identity in judgments of change. *Cognition*, *218*, 104953.
- Flombaum, J. I., Scholl, B. J., & Santos, L. R. (2009). Spatiotemporal priority as a fundamental principle of object persistence. In B. Hood & L. Santos (Eds.), *The origins of object knowledge* (pp. 135-164). Oxford University Press.
- Furlanetto, T., Bertone, C., & Becchio, C. (2013). The bilocated mind: new perspectives on self-localization and self-identification. *Frontiers in Human Neuroscience*, *7*, 1-6.
- Galinsky, A. D., & Moskowitz, G. B. (2000). Perspective-taking: decreasing stereotype expression, stereotype accessibility, and in-group favoritism. *Journal of Personality and Social Psychology*, *78*(4), 708-724.
- Hall, D. G., Waxman, S. R., Brédart, S., & Nicolay, A. C. (2003). Preschoolers' use of form class cues to learn descriptive proper names. *Child Development*, *74*(5), 1547-1560.
- Hobbes, T. (1839–1845). De Corpore. In W. Molesworth (Ed.), *The English works of Thomas Hobbes* (Vol. 1). John Bohn.
- Hood, B., Gjersoe, N. L., & Bloom, P. (2012). Do children think that duplicating the body also duplicates the mind? *Cognition*, *125*(3), 466-474.
- Hosmer Jr, D. W., Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression*. John Wiley & Sons.
- Hume, D. (1740). *A treatise of human nature*. Courier Corporation.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. MIT Press.

- Keyes, H., & Brady, N. (2010). Self-face recognition is characterized by “bilateral gain” and by faster, more accurate performance which persists when faces are inverted. *Quarterly Journal of Experimental Psychology*, *63*(5), 840-847.
- Knoblich, G., & Flach, R. (2001). Predicting the effects of actions: Interactions of perception and action. *Psychological Science*, *12*(6), 467-472.
- Kuhlmeier, V. A., Bloom, P., & Wynn, K. (2004). Do 5-month-old infants see humans as material objects? *Cognition*, *94*(1), 95-103.
- Kurzban, R. (2012). *Why everyone (else) is a hypocrite: Evolution and the modular mind*. Princeton University Press.
- Lewis, D. (1976). Survival and identity. In A. O. Rorty (Ed.), *The identities of persons* (pp. 17-40). University of California Press.
- Libby, L. K., & Eibach, R. P. (2002). Looking back in time: self-concept change affects visual perspective in autobiographical memory. *Journal of Personality and Social Psychology*, *82*(2), 167-179.
- Locke, J. (1689). *An essay concerning human understanding* (K. Winkler, Ed.). Hackett Publishing.
- Macrae, C. N., Visokomogilski, A., Golubickis, M., Cunningham, W. A., & Sahraie, A. (2017). Self-relevance prioritizes access to visual awareness. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(3), 438-443.
- Marchak, K. A., & Hall, D. G. (2019). Designators, descriptions, and artifact persistence. *Cognition*, *192*, 103999.
- Marchak, K. A., & Hall, D. G. (2022). Children's Understanding of Proper Names and Descriptions. *Journal of Child Language*, *49*(6), 1295-1306.

- Mendelson, E. (2009). *Introduction to mathematical logic*. CRC press.
- Mitroff, S. R., Scholl, B. J., & Wynn, K. (2004). Divide and conquer: How object files adapt when a persisting object splits into two. *Psychological Science*, *15*(6), 420-425.
- Nagel, T. (1971). Brain bisection and the unity of consciousness. *Synthese*, 396-413.
- Nichols, S., & Bruno, M. (2010). Intuitions about personal identity: An empirical study. *Philosophical Psychology*, *23*(3), 293-312.
- Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.
- Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
- Perry, J. (1972). Can the self divide? *The Journal of Philosophy*, *69*(16), 463-488.
- Plutarch. (2001). *Plutarch's Lives (Vol. 1)* (H. Clough, Ed.) [Dryden, J]. Random House LLC. (AD 46 - AD 127)
- Protzko, J., Tobia, K., Strohminger, N., & Schooler, J. W. (2023). Do Obligations Follow the Mind or Body? *Cognitive Science*, *47*(7), e13317.
- Richardson, D. R., Hammock, G. S., Smith, S. M., Gardner, W., & Signo, M. (1994). Empathy as a cognitive inhibitor of interpersonal aggression. *Aggressive behavior*, *20*(4), 275-289.
- Rips, L. J. (2011). Split identity: Intransitive judgments of the identity of objects. *Cognition*, *119*(3), 356-373.
- Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, *113*(1), 1-30.
- 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), 677-688.

- Rose, D., Machery, E., Stich, S., Alai, M., Angelucci, A., Berniūnas, R., Buchtel, E. E., Chatterjee, A., Cheon, H., & Cho, I.-R. (2020). The Ship of Theseus Puzzle. *Oxford Studies in Experimental Philosophy Volume 3*, 3, 158-174.
- Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, 38(2), 259-290.
- Sedikides, C., Hong, E. K., & Wildschut, T. (2023). Self-continuity. *Annual Review of Psychology*, 74, 333-361.
- Shoemaker, D. (2016). Heart of Animalism. *Animalism: New Essays on Persons, Animals, and Identity*, 303.
- Shoemaker, D. W. (2007). Personal identity and practical concerns. *Mind*, 116(462), 317-357.
- Sorrentino, C. M. (2001). Children and adults represent proper names as referring to unique individuals. *Developmental Science*, 4(4), 399-407.
- Spelke, E. S., Kestenbaum, R., Simons, D. J., & Wein, D. (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, 13(2), 113-142.
- Tierney, H. (2020). The Subscript View. *Oxford Studies in Experimental Philosophy Volume 3*, 3, 126.
- Tierney, H., Howard, C., Kumar, V., Kvaran, T., & Nichols, S. (2014). How many of us are there. *Advances in experimental philosophy of mind*, 181.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440-463.
- Ullman, S. (1979). *The interpretation of visual motion*. Massachusetts Inst of Technology Pr.

vanMarle, K., & Scholl, B. (2003). Attentive tracking of objects vs. substances. *Psychological Science, 14*, 498-504.

Weaver, S., & Turri, J. (2018). Personal identity and persisting as many. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy* (Vol. 2, pp. 213-242). Oxford University Press.

Williams, B. (1970). The self and the future. *The Philosophical Review*, 161-180.

Woike, J. K., Collard, P., & Hood, B. (2020). Putting your money where your self is: Connecting dimensions of closeness and theories of personal identity. *PloS one, 15*(2), e0228271.

Wynn, K. (1992). Addition and subtraction by human infants. *Nature, 358*(6389), 749-750.