

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Inferences About Personal Identity

Permalink

<https://escholarship.org/uc/item/15t6g47n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Authors

Blok, Sergey
Newman, George
Behr, Jennifer
et al.

Publication Date

2001

Peer reviewed

Inferences About Personal Identity

Sergey Blok (s-blok@northwestern.edu)

George Newman (g-newman@northwestern.edu)

Jennifer Behr (jennifer.behr@yale.edu)

Lance J. Rips (rips@northwestern.edu)

Department of Psychology, Northwestern University,
Evanston, Illinois 60208 USA

Abstract

We investigate the features people use in making inferences about continuity of individual persons. Using a transformation paradigm, we show that people weigh both continuity of the brain and continuity of mental content. Across experiments, we document instances in which participants are more likely to assert individual continuity than continuity of personhood. We discuss these results in terms of a hierarchical view of concepts and philosophical work on personal identity.

Introduction

People are sensitive to the effects that transformations have on membership in basic-level categories (e.g., Gelman & Wellman, 1991; Keil, 1989; Rips, 1989). For example, Rips (1989) asked participants to read stories about creatures of one category (e.g., birds) that came to resemble those of another (e.g., insects). If the transformation was due to accidental factors, participants believed that the creature remained a member of the original category. If the transformation was part of normal maturation, however, participants judged the creature a member of the second category. In general, transformations that alter an object's core properties also tend to change the object's category, while transformations that alter an object's surface or incidental properties tend to preserve its category.

Despite the relatively large number of studies that address questions of category membership continuity, there have been few studies addressing reasoning about individual continuity (see Hall, 1998, Johnson, 1990 and Liittschwager, 1994, for exceptions in the developmental literature). The central question is how we decide that a particular individual – for example, my dog Fido – is the same individual (still Fido) across transformations. This question is potentially different from the one about membership – whether this individual is still a dog.

We investigate here two issues concerning reasoning about individuals. First, we explore the kinds of features people use in judging continuity of identity. Second, we contrast the ways in which people reason about class membership and about identity continuity.

Features of Person Identity

What properties does a person at time t_1 need to share with one at time t_2 in order for that individual to be *the same* at both temporal markers? In making such judgments, people may be *phenomenalists*, relying on continuity of appearance.

In a preliminary experiment, we created stories that varied the type of transformation that a hypothetical target person undergoes. One set of participants -- the Plastic Surgery group -- read a scenario about Jim, a male accountant, who receives plastic surgery to alter his appearance cosmetically to resemble that of Marsha, a female actress. Another set of participants -- the Brain Transplant group -- read a similar story in which Jim's brain is replaced with that of Marsha. After reading the story, both groups supplied judgments of Jim's *identity change* – whether the individual was still Jim or had become Marsha after surgery. Results indicated that a greater proportion of participants in the Brain Transplant group believed Jim's identity had changed than in the Plastic Surgery group (45% and 15%, respectively, $\chi^2(1, 39) = 4.29, p < .05$).

These results suggest that changes in appearance are ordinarily not enough to warrant change in identity. The finding parallels earlier studies of natural kinds that show that people tend to reject mere appearance as evidence for category membership when appearance conflicts with deeper properties of the category in question. It has often been suggested that natural kinds' hidden, causally central properties are used for categorization and induction (Ahn, 1998; Gelman & Hirschfeld, 1999), while surface features such as appearance are used for similarity judgments (Keil, 1989; Rips, 1989). But although our results hint that our participants are not folk-phenomenalists, they leave open the question of what criteria they do use to assess personal identity.

The question of criteria for identity is one of the oldest in metaphysics. Writers on the *physicalist* side (e.g., Aristotle, Wiggins) argue that continuity of the body, or, more importantly, the brain, is critical to identity. According to this view, a person P_2 at time t_2 is the same person as P_1 at t_1 if

P_2 has the brain of P_1 ¹. Philosophers arguing from a *functionalist* position (e.g., Locke) propose that what matters for identity is not the physical brain, but rather the mental content – the person’s unique memories, habits, and personality.

If people are folk-physicalists, then a brain transplant that does not preserve the mental content of the original person should be judged to be as person-preserving as a transplant that retains the original memories. We test this hypothesis in Experiment 1. In Experiment 2 we test the alternative hypothesis that people are folk-functionalists.

Individuals and Hierarchies

A prevalent assumption in the cognitive-psychology literature on categories is that individuals inherit properties of the categories to which they belong. For example, if Fido is a dog, then properties of dogs are true of Fido. (However, see Sloman, 1998, for some exceptions.)

A number of philosophical positions also imply that judgments of identity (Fido vs. Rover) and category membership (dog vs. cat) are related. According to these accounts, criteria of identity for an object (whether phenomenal, physical, or functional) are given by membership in a category to which the object belongs (e.g., Geach, 1962; Gupta, 1980; Wiggins, 1980). If a dog, Fido, is somehow transformed so that it is no longer a dog, it must be true that it is no longer the same individual. In some of these theories (Geach, 1962), different categories to which Fido belongs (e.g., dog vs. pet) yield different criteria of identity, whereas in others (Wiggins, 1980) there can be only one set of criteria. In the experiments that follow, we show that people’s reasoning about continuity of identity need not follow *any* obvious category. Instead, participants sometimes rely on distinct sets of features when reasoning about continuity of identity and continuity of category membership.

Experiment 1: Memories and Causality

One goal of Experiment 1 was to determine whether participants perceive continuity of memories as necessary for identity. For the purposes of this paper, we consider “memories” to be unique sets of personal mental representations.

We presented stories that manipulate whether a target person undergoing a brain transplant retains or loses memories in the process. If what matters for identity is continuity of physical parts, such as the brain, then participants should perceive brain transplants that retain memories and those that lose memories as equally conducive to sameness of identity. By contrast, if continuity of memories is essential to continuity of identity, then a brain transplant that retains memories should be more likely to elicit perceptions of sameness than a transplant that does not retain memories.

We also varied whether the memories in question could affect the person’s behavior. Because the more essential features of concepts may be those that are causally central

(Ahn, 1998), we expected memories that have causal efficacy to be more individual-preserving than memories that could not cause behavior.

One methodological limitation of the preliminary study discussed in the introduction is that participants were never queried about whether the target person was still a member of the same category after the transformation (e.g., whether he was still a person, as opposed to still being Jim). We were therefore not able to determine whether participants’ judgments of identity change correlated with their judgments of change in category membership – a correlation that we would expect given the philosophical theories cited earlier. We address this question in Experiment 1 by asking participants to judge the extent to which the post-transformation individual is still a member of his original categories. The category most likely to confer identity on our target individual is the category PERSON itself. However, to determine whether PERSON has special status in this regard, we contrast it with other possible categories, in this case occupation (ACCOUNTANT) and gender (MALE).

Method

Thirty-eight Northwestern University undergraduates read a science fiction story about Jim, a male accountant undergoing a lifesaving brain transplant. Specifically, Jim’s brain was transplanted into a robot body. In a between-groups design, we varied whether Jim’s memories were the same or different after the transformation. In addition, half the participants in each group answered questions about a situation in which the memories could cause behavior in the robot, and half answered questions about a situation in which the memories could not cause behavior. The full story appears in Figure 1.

The Preserved Memory group read a version of the scenario in which the robot received an *unaltered* version of Jim’s brain. The Altered Memory group read a version of the story in which the memories were significantly *altered* during the transformation process.

After reading the story, participants rated their confidence in a number of statements relating to Jim’s continuity. Causal efficacy -- whether the memories were able to affect behavior -- was manipulated by varying whether the questions that followed related to events on Monday (when the robot was off) or Wednesday (when the robot was functional). For example, in the Low Causal Efficacy condition, a probe statement read:

On Monday, before the scientists switch it on, the robot is Jim.

0—1—2—3—4—5—6—7—8—9
<i>strongly disagree</i> <i>strongly agree</i>

¹ A qualified physicalist position need not require strict sameness of matter, just that there be an unbroken chain of intermediate states between the matter that makes up the body now and the matter that made it up in the past.

The Transplant

Jim is an accountant living in Chicago. One day, he is severely injured in a tragic car accident. His only chance for survival is participation in an advanced medical experiment. Jim agrees.

A team of scientists remove his brain and carefully place it in a highly sophisticated cybernetic body (robot). The robot is powered by electricity. The scientists connect the brain to the robot controls. Though all the right connections between the robot and the brain have been made, the scientists cannot “plug” the robot in because they are waiting for a power adapter they have ordered.

On Monday, the scientists come in to work and the power adapter still has not arrived. While they wait, the scientists scan the brain inside the robot and note that [THE MEMORIES / NO MEMORIES] in it are the same as those that were in the brain before the operation.

Finally, on Wednesday, the power adapter arrives and the scientists turn on the robot. The robot appears human-like in its behavior. The robot has senses and can move and talk. Again, the scientists scan the brain inside the robot and find that [THE MEMORIES / NO MEMORIES] in it are the same as those that were in the brain before the operation.

Figure 1: Stimulus story for Experiment 1 showing the memory manipulation.

By contrast, in the High Causal Efficacy condition, participants evaluated the statement:

On Wednesday, after the scientists switch it on, the robot is Jim.

0—1—2—3—4—5—6—7—8—9
strongly *strongly*
disagree *agree*

Each participant provided judgments about four kinds of continuity: **individual continuity** (“... the robot is Jim”), **personhood continuity** (“... the robot is a person”), **gender continuity** (“... the robot is a male”), and **occupational continuity** (“... the robot is an accountant”). Question order was fully counterbalanced across participants.

Results

Individual Continuity Table 1 summarizes the results of this experiment. As expected, continuity of memory was important for identity. Participants gave higher ratings when the transplanted brain retained the old memories than when it did not (mean ratings on a 0-9 scale were 5.28 and 1.70, respectively; $F(1, 34) = 21.21, p < .001$). The capacity of memories to cause behavior also had an effect on continuity. When the robot was described as being “off,” continuity judgments were lower than when the robot was “on” ($M = 2.15$ and 4.83 , respectively; $F(1, 34) = 11.83, p < .001$). The interaction between Memory and Causal Efficacy was reli-

	Memories				
	Altered		Preserved		
	Continuity	Low CE	High CE	Low CE	High CE
Individual	1.30	2.10	3.00	7.56	
Personhood	1.90	3.60	2.00	4.56	
Gender	3.80	4.20	2.56	7.00	
Occupation	1.30	1.60	2.78	6.78	

Table 1: Mean continuity ratings (Experiment 1) as a function of memory continuity and causal efficacy (CE).

able, $F(1, 34) = 5.82, p < .05$. Causal efficacy had a larger effect when Jim’s memories were retained rather than altered, suggesting that when Jim’s old memories were no longer present, it did not matter for identity per se whether the altered memories were able to cause behavior.

Personhood Continuity A similar analysis of personhood ratings showed that, in contrast to identity findings, there was no main effect of Memory, suggesting that continuity of memory was not important for continuity of personhood, $F(1, 34) < 1$. However, there was a main effect of Causal Efficacy, $F(1, 34) = 5.45, p < .05$. Participants answering questions about a post-transformation robot in the “on” state were more likely to judge the robot as being a person than those who responded to queries about an “off” robot ($M = 4.08$ and 1.95 , respectively). The interaction between Memory and Causal Efficacy was not reliable, $F(1, 34) < 1$.

Responses to the personhood question were predictive of identity judgments, but the magnitude of the effect was secondary to that of memory continuity. When memory continuity, causal efficacy and personhood were simultaneously regressed onto identity judgments, the standardized regression coefficient for personhood ($\beta = 0.30$) was relatively small compared to that for memory ($\beta = 0.51$). Adjusted R^2 for the overall model was .56.

Gender and Occupation Continuity Gender continuity (“Still a male?”) did not fit the pattern of identity judgments. In contrast to identity results, there was no main effect of memories, $F(1,34) < 1$. People’s continuity ratings in altered and preserved memory conditions tended towards the middle of the scale ($M = 4.00$ and 4.7 , respectively), indicating that people were not confident about the right way to think about gender continuity. Causal efficacy, however, did play a role in these judgments. When the brain was able to cause behaviors, the object was rated more likely to retain gender than when the brain was unable to cause action ($M = 5.60$ and 3.18 , respectively; $F(1, 34) = 6.62, p < .05$). Also, the effect of causal efficacy was greater when memories were retained, as suggested by a reliable interaction between memory and causal efficacy, $F(1, 34) = 4.62, p < .05$.

By contrast, occupation ratings (“Still an accountant?”) did appear to follow the pattern of the identity ratings. An analysis of variance similar to the one performed on identity continuity judgments revealed parallel effects. We found a main effect of memory continuity such that a brain transplant that preserved memories was also more likely to pre-

serve occupation, $F(1, 34) = 22.35, p < .001$. Also, when the brain had an effect on behavior, occupation was more likely to be preserved than if the brain had no effect, $F(1, 34) = 9.33, p < .01$. Furthermore, as in the case of the identity ratings, there was an interaction between memory continuity and causal efficacy such that causal efficacy was more important for occupation continuity when memories were preserved, $F(1, 34) = 6.91, p < .05$.

Discussion

This experiment provided evidence for the role of memory continuity in perceived identity. Participants who read about a memory-preserving transplant gave higher individual continuity ratings than did participants who read about a transplant that did not preserve memories. This supports the widely held view in the philosophical literature that personal mental representations are central to individual identity. Moreover, our participants granted the highest level of identity continuity to a transplant if these memories had the capacity to cause behavior.

Perhaps the most striking finding is the relative independence between judgments of identity continuity and personhood continuity. People's reasoning about continuity of identity does not appear tightly bound to sameness of membership in normal categories. In fact, participants in some conditions were more likely to agree that the individual was still Jim after the transformation than that he/it was still a person. Specifically, in the condition judged optimal for individual continuity (Preserved Memories, High Causal Efficacy), participants gave a high mean rating of 7.56 when asked if the individual was the same, but a much lower rating of 4.56 when asked if it was still a person, $t(8) = 2.63, p < .05$.

We used gender and occupation categories as foils for personhood. As expected, we found only a poor fit between gender continuity and identity. Occupational continuity fared better in terms of reflecting identity judgments, though it seems likely that individual identity was driving occupation identity rather than the reverse. On intuitive grounds, occupational categories are hardly viable as granting identity to individuals. I can cease being a student, without any significant loss of identity. We revisit the issue of individual and occupation identity in Experiment 2.

In general, it seems possible that our participants used different criteria to judge identity, gender, and occupation. Most importantly, people decided about identity and personhood using different criteria. While the critical property of identity appears to be continuity of memories, personhood may depend more heavily on typical properties of persons, such as having a human body and engaging in human behaviors.

Experiment 2 pursues this issue, asking whether perceived identity continuity can be maintained through a transformation that does not preserve any of the physical parts of the original person.

Experiment 2: Necessary Features of Identity

Experiment 1 showed that a brain without the right memories does not guarantee identity of individuals. It is still

possible, however, that the brain may be a necessary but not sufficient property of individuals. In this case, memories would have to be transmitted in the physical stuff in which they arose. By contrast, if memories are the "software" that is merely running on the brain "hardware," then it is conceivable that physical brains are not even necessary for identity – any computationally adequate device would do. This is the issue that separates physicalists and functionalists. Will people infer individual continuity even if the original person's memories are "implemented" on a machine that is not the original physical brain?

If the answer is "yes," then we may conclude that people's beliefs about identity are relatively unconstrained, allowing for identity to be preserved through a wide range of fairly extreme transformations. Such a folk-functional position is at least intuitively sensible. For example, body cells die and regenerate multiple times throughout the lifespan. It seems odd to consider such physical changes as threatening to personal identity. The competing folk physicalist theory sees brain tissue as at least necessary for identity.

Method

To address the question of physicalism versus functionalism, we modified the brain-transplant scenario from Experiment 1 to include a condition in which the memories in Jim's brain are copied onto a computer designed to control the robot (Computer Copy condition). The story for this condition appears in Figure 2. We also ran a replication of the Brain Transplant scenario from Experiment 1 (Brain Transplant condition) without the passages relating to causal efficacy.

The second factor in the design was whether the memories in the brain (computer) were altered or preserved. This design thus generated four scenarios, which we gave to separate groups of participants. After reading the scenario, participants answered the same set of questions as in Experiment 1. Judgments of individual, personhood, gender, and occupational continuity were made on a 10-point scale.

Jim is an accountant living in Chicago. One day, he is severely injured in a tragic car accident. His only chance for survival is participation in an advanced medical experiment. Jim agrees.

A team of scientists copy the memories in his brain onto a state-of-the-art computer. The computer is placed in a highly sophisticated cybernetic body (robot). All the right connections between the robot and the computer have been made, and the computer is able to control the robot. The scientists scan the computer and note that [*NONE OF*] the memories in it are the same as those that were in the brain before the operation.

When the scientists turn on the robot, the robot appears to be human-like in its behavior. It has senses and can move and talk.

Figure 2: Stimulus story for the computer copy condition in Experiment 2.

	Memories			
	Altered		Preserved	
Continuity	Computer Copy	Brain Transplant	Computer Copy	Brain Transplant
Individual	0.89	1.53	1.97	5.27
Personhood	1.11	2.83	2.08	2.69
Gender	2.23	4.77	4.34	4.58
Occupation	3.17	0.91	5.78	4.16

Table 2. Mean continuity ratings (Experiment 2) as a function of memory continuity and transplant type.

Questions were presented in two different random orders across participants. Sixty-four Northwestern University undergraduate students took part in the study.

Results

Table 2 presents a summary of the results, which appear to favor a folk-physicalist over a folk-functional position.

Individual Continuity Participants in this experiment were more likely to think that the post-transplant individual was still Jim if the transplant included Jim's brain than if it merely included Jim's memories. A Brain Transplant elicited higher continuity ratings than a Computer Copy, $F(1, 56) = 17.95, p < .001$. As in Experiment 1, there was also an effect of the memories themselves. Participants who read about a transformation that preserved memories gave higher continuity ratings than those reading about a transformation that altered memories, $F(1, 56) = 26.81, p < .001$. Most importantly, however, there was also an interaction between transformation type and memory factors: Preserved memories facilitated continuity to a greater extent when the transformation was a Brain Transplant than when it was a Computer Copy, $F(1, 56) = 8.17, p < .01$. There were no reliable effects of question order.

Personhood Continuity As in the case of individual continuity, participants who read the Brain Transplant scenario viewed the robot as more likely to be a person than people who read about a Computer Copy, $F(1, 56) = 4.13, p < .05$. However, there was no reliable effect of memory continuity on personhood judgments: Participants were about as likely to think that the robot was a person whether or not the memories were the same as Jim's, $F(1, 56) < 1$.

As this result suggests, personhood continuity ratings did not fully predict judgments of identity continuity. That is, participants were not simply basing their identity judgments (Still Jim?) on whether they believed the object in question is still a person. We tested this claim as we did in the previous study, by running a simultaneous regression with personhood ratings, transplant type, memory continuity, and the interaction between them as predictors of identity responses. The pattern of regression weights closely resembled those in Experiment 1. Memory continuity ($\beta = .46$) was a better predictor of identity judgments than personhood ($\beta = .32$). Adjusted R^2 for the overall model was .54.

Gender and Occupation Continuity A similar analysis of variance on gender continuity ratings revealed no reliable effects. There was a trend for readers of the Brain Trans-

plant story to assert a higher level of gender continuity than participants reading a Computer Copy story ($M = 4.67$ and 3.29 , respectively; $F(1, 56) = 3.33, p = .07, n.s.$)

Occupational continuity revealed a main effect of transformation type, $F(1, 56) = 8.25, p < .01$. It is important to note, however, that the direction of the effect was reversed relative to the identity findings. A Computer Copy was more convincing than a Brain Transplant in allowing Jim to retain the status of an accountant ($M = 4.48$ and 2.53 , respectively). As in the previous study, there was also a main effect of memory continuity, $F(1, 56) = 18.71, p < .001$. Continuity of memory positively predicted retention of an occupation.

Discussion

These data support the hypothesis that people's naïve construal of individual identity is roughly compatible with folk-physicalism. Continuity of the physical brain had an effect on continuity that went beyond that of functionally equivalent brain content. As in the previous experiment, our data speak against the possibility that this was due entirely to people's beliefs about personhood. While our participants did indicate that a computer copy was less of a person than a Brain Transplant, a regression analysis showed that personhood ratings were not as good a predictor of identity judgments as memory continuity.

We have also replicated the results of Experiment 1 showing that when continuity conditions were optimal (Brain Transplant, Preserved Memories), identity ratings were higher than personhood judgments ($M = 5.27$ and 2.69 , respectively; $t(15) = 4.22, p < .01$).

Gender and occupation continuity were not good candidates for granting identity to individuals. Both gender and occupation ratings exhibited a poor fit with identity judgments. Occupation continuity actually exhibited a reverse pattern on the criteria people used for continuity judgments. In a number of conditions people were more certain about the continuity of an occupation than they were about individual continuity. For example, in the Computer Copy -- Preserved Memories condition, participants were more likely to say that the object in question is still an accountant than they were to assert that Jim is still in existence ($M = 5.78, 1.97$, respectively; $t(15) = 3.81, p < .01$).

One potential limitation of the current study is that the Computer Copy story was ambiguous as to the fate of Jim's original brain. If our participants assigned identity status to whichever object inherits the original brain, and they believed that Jim's brain survived the accident (even if damaged), then we would expect low continuity ratings in the

Computer Copy condition because Jim's brain is a better candidate for being Jim than the computer containing his memories. This kind of view is proposed by Williams (1973; see Nozick, 1985 for a reply).

While the data can not rule out this possibility, a free response questionnaire administered after the experiment showed that none of our participants explicitly considered Jim's original brain as a factor in their continuity judgments. Furthermore, a related study (in preparation), addressing the question of two possible continuers, showed that people are relatively insensitive to the existence of an identity competitor, preferring to base their judgments on sameness of substance.

Summary and Conclusions

In these studies, we explored the set of features people consider important to personal identity. We showed that when people reason about identity continuity, they take into account continuity of the physical brain and its causally efficacious mental content.² People are not phenomenalists, in that appearance is not a necessary feature of an individual's continuity. People are also not unconstrained functionalists, in that they do not assign full continuity if an object only implements a person's unique mental content.

What can a description of a folk theory tell us about the way people form and use concepts? A common assumption is that an individual's identity conditions are given by one or more of the categories to which it belongs. While this provides a convenient way to link categories and individuals, our data show that people do not always use the same sets of characteristics in deciding continuity of an individual and continuity of membership in even its most obvious superordinate category. We have documented instances in which an individual who is viewed as having ceased to be a person is still seen as the same individual.

This finding presents a challenge to the theory that identity conditions are dictated by the superordinate category (or sortal concept) to which that individual belongs. This theory incorrectly predicts that any doubt about proper membership in the category should be reflected in doubt about the individual's survival.

Whether this finding is restricted to reasoning about persons or can be generalized to a wider range of objects remains to be seen. Although relatively minor changes to objects can often cause them to change basic category membership, their individual continuity appears to be much more rigid. Keil (1989) used the example of a coffee pot that was

reshaped as a birdfeeder to show that changes in intended function cause shifts in basic-level categorization for artifacts. Despite this change at the category level, however, the object is likely to be judged as the same individual as the one before the transformation.

Finally, reasoning about individuals may turn out to be fundamentally different than reasoning about categories. Individuation often takes into account the history of an object in a way that category membership does not. Whether Jim is still a person after a transformation may depend on whether the causal forces responsible for personhood are still intact. Whether Jim is still Jim, however, may depend on the trajectory of his parts across time.

Acknowledgments

This research was supported by NSF grant SES-9907414. The authors would like to thank three anonymous reviewers for their comments.

References

- Ahn, W. (1998). Why are different features central for natural kinds and artifacts? *Cognition*, 69, 135-178.
- Geach, P. T. (1962). *Reference and generality*. Ithaca, NY: Cornell University Press.
- Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understanding of the nonobvious. *Cognition*, 38, 213-244.
- Gelman, S. A., & Hirschfeld, L. A. (1999). How biological is essentialism? In D. L. Medin & S. Atran (Eds.), *Folkbiology*. Cambridge, MA: MIT Press.
- Gupta, A. (1980). *The logic of common nouns*. New Haven, CT: Yale University Press.
- Hall, D. G. (1998). Continuity and the persistence of objects. *Cognitive Psychology*, 37, 28-59.
- Johnson, C. (1990). If you had my brain, where would I be? Children's understanding of the brain and identity. *Child Development*, 61, 962-972.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Liittschwager, J. C. (1994). *Children's reasoning about identity across transformations*. Doctoral dissertation, Department of Psychology, Stanford University, Stanford CA.
- Nozick, R. (1981). *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge, England: Cambridge University Press.
- Sloman, S. A. (1998). Categorical inference is not a tree. *Cognitive Psychology*, 35, 1-33.
- Wiggins, D. (1980). *Sameness and substance*. Cambridge MA: Harvard University Press.
- Williams, Bernard (1973): *Problems of the Self*. Cambridge, England: Cambridge University Press.

² A strict physicalist position may be question-begging. If being the same individual depends on having the same physical material, how do we decide about whether physical material is the same? As philosophical theories, both the physicalist and functionalist approaches have some important deficiencies, particularly with respect to possible circularity. However, our purpose here is simply to see whether either theory approximates the reasoning of untrained participants. We leave to further research the question of what would happen if our lay-physicalists were confronted with difficulties for their view.