Out of Sorts? Some Remedies for Theories of Object Concepts: A Reply to Rhemtulla and Xu (2007)

Sergey V. Blok
University of Texas at Austin

George E. Newman
Yale University

Lance J. Rips
Northwestern University

Concepts of individual objects (e.g., a favorite chair or pet) include knowledge that allows people to identify these objects, sometimes after long stretches of time. In an earlier article, the authors set out experimental findings and mathematical modeling to support the view that judgments of identity depend on people's beliefs about the causal connections that unite an object's earlier and later stages. This article examines M. Rhemtulla and F. Xu's (2007) critique of the causal theory. The authors argue that M. Rhemtulla and F. Xu's alternative sortal proposal is not a necessary part of identity judgments, is internally inconsistent, leads to conflicts with current theories of categories, and encounters problems explaining empirical dissociations. Previous evidence also suggests that causal factors dominate spatio-temporal continuity and perceptual similarity in direct tests. The authors conclude that the causal theory provides the only existing account consistent with current evidence.

Keywords: concepts, causation, sortals, object concepts

Many objects persist over stretches of time too long to track perceptually. If you attend your thirtieth high school reunion, you are bound to run into classmates like Fred Lugbagg whom you have not encountered since graduation and whose perceptual appearance is no more similar to the 18-year-old Lugbagg than is the appearance of most of the other males at the reunion. Yet you can often be fairly confident that the 48-year-old and the 18-year-old Lugbaggs are one and the same. Much the same is true of earlier and later encounters with your neighbor's cat, your grandfather's backpack, Manny's Deli, and many other individual objects.

A variety of clues provide evidence for identification. In easy cases, one can use location or similarity, because most objects do not radically change their positions or their superficial properties over the short term. In more difficult cases, though, one has to resort to more careful methods for tracing the history of the object from its current stage to its prior ones. In an earlier article (Rips, Blok, & Newman, 2006), we described a causal continuer theory for identity judgments that was based on a previous analysis by Nozick (1981). Nozick's theory aimed to provide a metaphysical analysis of identity in terms of causal proximity. The goal of our own theory is to describe people's beliefs about the nature of identity for physical objects, beliefs that inform particular identi-

fication judgments. The central notion is that people conceive the later stages of an object as causal outgrowths of earlier stages, and they take closeness of the causal connection as the underlying factor in establishing identity (see also Gelman, 2003). Evidence favors identity if the evidence favors close causal connectedness. Rips et al. (2006) provided qualifications to this account, a mathematical version of the theory, and predictions for experiments in which participants judged whether an object was the same as an earlier one.

In their comments on Rips et al. (2006), Rhemtulla and Xu (2007) approach singular concepts from a different tradition. The central notion here is a special class of terms called *sortals*—certain count nouns in English—whose meaning includes identity conditions for the objects they describe. Exactly which count nouns are sortals is a matter of debate (see Hirsch, 1997; Wiggins, 1997; Xu, 1997); however, most sortalists would probably agree that basic-level count nouns, such as *daisy* or *dog,* qualify as sortals.

The identity conditions that these sortals furnish are necessary and sufficient relations for identifying objects at different times, and they take the following form (see Lowe, 1989):

Principle 1: If object x at time t_1 and object y at time t_2 are members of sortal category S, then $x = y$ if and only if $R_S(x,y)$.

In this formulation, R_S is an equivalence relation (i.e., it is reflexive, symmetric, and transitive). In addition, R_S must be an informative relation (one that does not merely paraphrase or presuppose identity for Ss) to avoid trivializing the analysis. Most sortal theories also assume that objects in distinct sortals cannot be identical. In other words:

Sergey V. Blok, Department of Psychology, University of Texas at Austin; George E. Newman, Department of Psychology, Yale University; Lance J. Rips, Psychology Department, Northwestern University.

Correspondence concerning this article should be addressed to Lance J. Rips, Psychology Department, Northwestern University, 2029 Sheridan Road, Evanston, IL 60208. E-mail: rips@northwestern.edu

Principle 2: If object x at time t_1 is a member of sortal category S and object y at time t_2 is not a member of S, then $x \neq y$.[1]

Rhemtulla and Xu (2007) have a theoretical argument that attempts to show that sortal concepts take cognitive precedence over a purely causal account of identity judgments. They argue, too, that previous evidence (Blok, Newman, & Rips, 2005; Liittschwager, 1995; Rips et al., 2006) that seemed to disconfirm a psychological version of the sortal view fails to do so. In the first section of this reply, we examine Rhemtulla and Xu's sortal theory because it differs in important ways from its philosophical predecessors, and we try to discern whether the theory is strong enough to support their central argument. (For critiques of sortal theories in philosophy, see Ayers, 1974, 1997; Campbell, 2006; Mackie, 2006; Merricks, 1998.) In the second section, we examine the available empirical evidence to see whether the sortal theory is able to withstand it, and in the third section we return to the causal continuer model to assess its strengths and limitations in view of Rhemtulla and Xu's remarks. Thus, we consider Rhemtulla and Xu's points in an order opposite that in their article.

Sortals and Identity Conditions in a Psychological Context

Merely knowing an object's sortal category is rarely of direct help in deciding questions of identity. If you are worried that the man at the reunion party might not be Fred Lugbagg, it will not aid you much to learn that he is a person. More helpful will be tests that trace Fred's causal path: Does the person claiming to be Fred recall the name of your chemistry teacher? Does he have a scar on his left hand? Could he really have survived the fall from a fourth floor window? Is there evidence to show he ever returned from Brazil? These questions seem like ordinary causal hypothesis testing (or detective work), with its characteristic open-endedness about the nature of relevant evidence (Fodor, 1987). They do not seem like ticking off necessary and sufficient conditions for person identity. (For a historical account of the struggle to find reliable means of re-identifying people, see Cole, 2001.)

Rhemtulla and Xu (2007), however, believe that although causal factors are important for identity judgments, sortal concepts are necessary to determine which causes pertain to which objects. "Knowledge about particular sortals is necessary for a person to decide when a causal connection is enough to justify persistence, and when it is not" (p. 1091). To evaluate this idea, though, we need to take a closer look at Rhemtulla and Xu's theory of sortals.

Are Sortal Concepts Necessary for Identity Judgments?

Experiments by Xu and Carey (1996) suggest that it is not until around 12 months of age that infants have acquired enough information about categories like trucks and cats to determine that a truck and a cat cannot be the same object. Younger infants, however, can also discriminate objects in slightly simpler paradigms by taking into account the objects' properties (Wilcox & Baillargeon, 1998; Xu & Baker, 2005). For example, Xu and Baker (2005) removed a toy of one kind (say, a truck) from a box and then replaced it while an infant watched. The infant then retrieved either the same toy (the truck) or a toy of a different kind

(e.g., a cat) from the box. The infant was then allowed to search the box again. Those who had retrieved a different toy searched longer than those who had retrieved the same toy, as if the infants inferred that the new toy signaled that two objects were in play. "Thus by 10 months of age, infants are able to use perceptual property information for object individuation in both a looking time paradigm [that of Wilcox & Baillargeon, 1998] and a manual search paradigm" (Xu, 2005, p. 70).

What is of interest about this conclusion for present purposes is that sortal concepts such as TRUCK or CAT could not be responsible for the 10-month-olds' success in the simpler paradigms because, according to Xu (2005), infants only acquire these concepts some months later. Instead, property information—that is, nonsortal concepts associated with adjectives, such as *furry* or *shiny*—must be responsible for the discrimination. Xu (1997) believes that 10-month-olds do have an abstract sortal, *physical object,* which provides identity conditions, where a physical object in this sense is something that moves as a whole and is spatially separate from its background (see Hirsch, 1997, and Wiggins, 1997, for criticisms of this claim). PHYSICAL OBJECT, however, could not be responsible for infants' correct performance in Xu and Baker's (2005) experiment, because they specifically designed the study to eliminate spatial cues that would allow infants to decide whether there were one or two physical objects in the box.

According to Xu's (2005) theory, then, neither lower level sortal concepts (e.g., TRUCK) nor higher level sortal concepts (PHYSICAL OBJECT) could individuate the objects in this experimental set-up; hence, these sortal concepts cannot be necessary for correct identity decisions. Rhemtulla and Xu (2007) appear to think that combining property information with PHYSICAL OBJECT is what allows infants to make the differentiation. But if so, it must be the property information that is doing all the work because the infants had no cues about the status of the items as distinct physical objects. In this respect, Xu's position differs sharply from that of philosophers who have championed the role of sortals in identity (e.g., Dummett, 1973; Strawson, 1959; Wiggins, 2001). The central tenet of these latter theories is that nonsortal predicates, such as *shiny* or *furry*, differ from sortal predicates precisely because nonsortals do not provide identity conditions for objects; they fail to divide the material to which they apply into distinct instances. To allow nonsortal properties to individuate objects is not to modify the theory but to abandon it.

---

[1] Sortal theories in psychology appeal almost exclusively to Principle 2, as we discuss later in this article, but it is very difficult to state this principle adequately. The main problem is that some sortal theories allow sortal categories to be nested. According to Xu (1997), for example, both *dog* and *physical object* are sortals with distinct identity conditions, R_{dog} and R_{object}. Hence, Fido can go from being a dog to a nondog as long as he is covered by the sortal *physical object.* For Principle 2 to work correctly, we need to know under what conditions an object's exiting one sortal category ends the object's existence. We are unsure whether there is a way to formulate Principle 2 that is not question begging. Principle 2 requires a means to determine when being an S at t_1 and being a non-S at t_2 imply nonidentity, without presupposing a condition on S such as the following: No x can be both an S at t_1 and a non-S at t_2. We doubt whether such a formulation is possible, but we can safely leave this problem for proponents of sortal theories.

Are Sortals Necessary to Determine Which Causes Justify Persistence?

Rhemtulla and Xu (2007) nevertheless argue that sortals are necessary to specify which of the causes associated with a particular object contribute to its identity over time: "Representations of sortals are necessary to indicate which causes to pay attention to for each kind of individual" (p. 1088). But if nonsortal concepts like FURRY and SHINY are enough to allow even infants to identify objects over time, why does causal information stand in need of sortals to do the same?

To see that sortals are not necessary, notice first that examples of sortal-relative identity conditions are in short supply. Sortal theories need these conditions to specify the R_S relations in Principle 1, but we know of no clear examples of identity conditions for typical everyday sortals such as *cats* or *trucks,* with the possible exception of (much disputed) criteria for *persons.* What are the necessary and sufficient conditions that cat x at t_1 and cat y at t_2 must possess in order for x to be identical to y, that is, $R_{cat}(x,y)$, and how do they differ from those conditions for dogs, that is, $R_{dog}(x,y)$? (See Mackie, 2006, for similar complaints.) The difficulties for sortal theories of singular concepts parallel the well-known difficulties for classical theories of category-level concepts. There are few convincing examples of necessary and sufficient properties for membership in everyday categories, and most cognitive psychologists have given up hope of uncovering them. We suspect that in the case of sortal theories, too, the shortage of plausible examples is due to the fact that people simply do not know sortal-relative conditions of identity for everyday categories.[2] If so, and if sortals are count nouns that furnish identity conditions, then there are few or no sortals.

If people do not use sortal-relative identity conditions, how do they judge identity of objects over time? According to the causal continuer theory, people take the causal–historical path of the object as the determining factor, but how do people know which causes to trace as parts of this path? Rhemtulla and Xu (2007) press this question for the causal continuer theory because they apparently believe the structure of the causes is insufficient to determine which causes are relevant. But it is unclear why they think this is so. Although it may be true that "all kinds are subject to all types of causes" (Rhemtulla & Xu, 2007, p.1088)—that many different types of causes impinge on individual objects—this does not imply that there is any indeterminacy in using causes to track identity. For one thing, not all causes are equally influential in shaping an object, so a weighting of causes by their causal strength or power may suffice for identity judgments. (Rhemtulla & Xu also suggest that the difficulty with the causal continuer theory is not that people know too many causes for a given object, but that they know too few: "even adults may not have the knowledge required to rely on understanding specific causal mechanisms in determining identity" [p. 1093]. But the theory need not assume that the people have to know the true scientific causes supporting the object's persistence or even know especially deep causes. What they need is practical information about the sorts of causal changes an object undergoes during its lifetime, e.g., in movement, growth, and decay, as in our earlier reunion example.)

As a second, related possibility, suppose that people believe that causal forces tend to cluster in ways that yield relatively stable configurations of matter over time. These equilibrium states pro-duce clustering of local properties, segregate one cluster from others, and allow inferences about the relations among the clustered properties at a time and about their likely histories. In fact, where we have good scientific models of the formation of natural objects, an equilibrating process of this sort takes center stage. The formation of a star, for example, depends on the opposing forces of gravity (which causes contraction of surrounding gases) and outward pressure (from internal nuclear reactions), which result in a stable state enabling the star's persistence over time (Kaufmann, 1994). One can expect laypeople's beliefs to be simpler than the scientific models, as we have just noted, but it does not seem unreasonable that the same type of thinking may be responsible for people's tendency to isolate and re-identify objects. What allows one to identify Fido may be a balance among forces including his birth, growth, habits, decay, physical constraints, interactions with other objects, and many other causes. What allows one to identify Rover is a distinct, but similar, mix of forces. (Rover's identity depends on the specifics of Rover's birth, not Fido's, and particular events that end Rover's career may be irrelevant to Fido's.)[3]

This equilibrium account is a purely speculative one, but we note that it has at least two advantages over sortal theories. First, it does not depend on the improbable idea that people represent identity conditions in the sense of the necessary and sufficient relation in Principle 1. The equilibrium states in question require causal forces acting in concert, but they do not entail any specific set of necessary and sufficient causes. Second, the causal equilibrium view accords well with, and may in fact explain, corresponding theories of categories. Most current theories of object categories are explicitly causal in holding that beliefs about causal relations account for people's judgments about category structure, typicality, inductive potential, and other features (e.g., Ahn, 1998; Gelman, 2003; Keil, 1989; Rehder & Hastie, 2001; Sloman, Love, & Ahn, 1998). Notice, though, that the causes in question usually act at the level of individual category members, according to these theories. If people believe, for example, that genetic components cause the category-typical properties of living things, that is presumably because they believe that these genetic elements act on

---

[2] Psychological essentialists believe that although people do not know the essential properties for a category, they nevertheless believe there are some (Gelman, 2003; Medin & Ortony, 1989). But the same tactic does not work for psychological sortalists. For individuals to identify objects over time, it is usually not enough for them to believe that a category has some identity conditions or other (i.e., to have a placeholder for these conditions); they have to know exactly what the conditions are to identify the objects via Principle 1.

[3] For a discussion of other issues having to do with possible causal ambiguity, see Rips et al. (2006, pp. 23–25). One might suggest that causal laws are defined over classes of things, not over an individual, because causal laws are (very roughly) of the following form: Things of Type C cause things of Type E. For a causal law to apply to an individual, the law must subsume the individual under a type. Thus, any causal story about individuals presupposes a type or category to which they belong. But notice that the causes we are appealing to do not relate either individual object tokens (e.g., Fido) or their types (e.g., dogs). They relate events (e.g., certain physical processes in Fido) to other events that support the persistence of the object. People may be applying causal laws (over event types) to discern these relations, but there is no reason to think that these laws presuppose sortal categories, such as dogs or pencils.

individual organisms to produce instantiations of these properties. It is Fido's dose of DNA that causes Fido to have the typical dog properties he possesses. The account we are suggesting here adds to this the idea that some of the forces acting on the object are also responsible for the object's persistence and other qualities. Fido's persistence depends on the constellation of causal forces that support his integrity over his lifetime.

The parallel between objects and their categories also highlights a striking oddity in sortal theories. The natural attitude about the relation of objects and categories is that people grasp the existence of objects before they gather these objects into categories. Sortal theories reverse this cognitive priority: "Sortals are kind concepts that come prior to individual concepts" (Rhemtulla & Xu, 2007, p. 1092). As we noted earlier, Xu's (1997) brand of sortalism softens this counterintuitive stance by supposing that the term *physical object* can function as a higher level sortal (again, contrary to sortal theories in philosophy). Thus, an individual cat can be singled out by means of the identity conditions of *physical object* before it receives the identity conditions of *cat*. But as far as we can see, the only reason for insisting that *physical object* is a sortal is a commitment to the idea that concepts of individual objects require identity conditions and that only sortals can provide them. Causal theories have an advantage here: According to such theories, objects are individuated by causes that constitute those objects. These theories can then explain category membership of the same objects in terms of overlap in the causal structure of its members, at least for natural kinds.

Finally, the same parallel brings into focus a central problem for Rhemtulla and Xu (2007). In their account, the identity of objects fundamentally depends on people's beliefs about the sortal categories to which the objects belong. But, as we have noted, most current theories of sortal categories are explicitly causal. Underlying causes are what largely produce these categories, according to these models (for reviews see Murphy, 2002; Rips, 2001). This puts Rhemtulla and Xu in a very uncomfortable position. Assuming that causal theories of categories are true, they face a version of exactly the same problem they attribute to the causal continuer theory. Which of all the causes that operate on a sortal category (e.g., dogs) determines that category? If the answer is another higher level category (e.g., animal or living thing in the case of dogs), why does the same problem not arise at the higher level as well?

Status of the Empirical Evidence

Empirical results reinforce the theoretical doubts about sortals. Blok et al. (2005, Experiments 1–2) and Rips et al. (2006, Experiment 1) present evidence for a dissociation between people's judgments that (a) an object after a transformation is the same as the object before and (b) the object after the transformation is in the same sortal category as the object before. For example, in one condition of Blok et al. (2005, Experiment 1), participants read a story about an accountant named Jim who suffered a tragic car accident and, as a result, had his brain transplanted to a new robot body while the rest of his old body was destroyed. Participants who learned that Jim's memories remained intact after the transplant tended to rate the postoperative creature as "still Jim" even though they also rated it (or him) no longer a person.

Most sortal theories predict that if an object leaves its sortal category it becomes a different object (see Principle 2). For example, if *person* is the relevant sortal for preoperative Jim, then it should be impossible for the postoperative creature to be both a nonperson and still Jim, contrary to our participants' judgments.[4] These findings, along with those of Rhemtulla (2005, Experiments 1–2), disconfirm this version of the sortal account.

Transformation Experiments

Rhemtulla and Xu (2007) are suspicious of these transformation results (and similar ones, such as Liittschwager, 1995) because they "do not seem to match experiences with understanding the same sorts of scenarios presented in fiction" (p. 1091). According to Rhemtulla and Xu, readers of fairy tales and similar stories suspend disbelief in the transformations, no matter how extreme they are, whereas participants' judgments vary across the transformational levels in the experiments. However, this comparison to fictional cases cuts both ways. We are uncertain, to begin with, whether Rhemtulla and Xu's assumption about fiction is correct.[5] But whatever the truth of the matter, differences between performing the experimental tasks and understanding fiction could be taken to show that participants were treating the transformations seriously as life-like possibilities and subjecting them to the same sort of scrutiny that they would reserve for matters of fact, rather than merely accepting them at face value. The purpose of these studies is to determine if (and how) people deploy their actual knowledge of sortal concepts in thinking about individual persistence, and within this context, the more analytic approach would seem to be an advantage.

Rhemtulla and Xu (2007) also note that the same participants judged different types of transformations in Blok et al. (2005), Liittschwager (1995), and Rips et al. (2006), and they believe that this within-participant design may have been responsible for some of the results. But although this feature may help explain differences between the shapes of the functions in Liittschwager (1995)

---

[4] Our participants explicitly denied that postoperative Jim was still a person, so it is not possible to explain these findings on the grounds that the participants believed that *person* extended loosely to cover Jim's later condition. This study and the others in Blok et al. (2005) and Rips et al. (2006) explicitly assume that basic-level terms—such as *person*, *truck*, and *cat*—are sortals. Basic-level terms are the usual examples of sortals in both the philosophy and psychology literature. Moreover, one prominent sortal theory (that of Wiggins, 2001) identifies the sortal for an object as the answer to the question "What is it?" for that object. Because Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976) found that basic-level terms are the usual answers to such questions in neutral contexts, it seems reasonable to think that basic-level terms are among, if not identical to, sortals.

[5] Kelly and Keil's (1985) study of transformational narratives shows that fiction writers produce more transformations between causally close items (e.g., between a person and an animal) than causally distant ones (e.g., between a person and a brook). Of course, a correlational study like this one is limited in its ability to rule out confounding factors. But we suspect that an experimental analysis of people's comprehension of such transformations would show that they have an easier time interpreting close than distant transformations, just as readers have an easier time understanding stories that preserve causal continuity than stories that contain causal discontinuities (e.g., Zwaan, Magliano, & Graesser, 1995).

and in Rhemtulla (2005), it leaves the crucial counterevidence intact. Dissociations between individual persistence and category persistence (as in the case of Jim outlasting his membership in the category of persons) hold for experiments using both within-participant designs (Blok et al., 2005; Rips et al., 2006) and between-participants designs (Rhemtulla, 2005, Experiments 1–2).[6]

Developmental Findings

As we have noted, the model in Rips et al. (2006) aimed to capture adults' deliberate and common-sense judgments about object identity over the long term. Sortal theory in psychology, however, was originally applied to research on infants and children by Macnamara (1986) and Xu and Carey (1996). Although this work may be relevant to adult judgments, there is no empirical evidence that directly links these two programs of research, and the lessons from the infant research are ambiguous. The developmental results on naming that Rhemtulla and Xu (2007) describe in the final section of their article illustrate this ambiguity. The main finding (from Xu, 2002) is that labeling objects with distinct nonsense count noun phrases (*Look, a toma* vs. *Look, a fendle*) leads 9-month-olds to conclude that two objects are present. Obviously, nonsense nouns do not possess identity conditions for either babies or adults (because they are meaningless). Moreover, the main thrust of Xu and Carey's (1996) earlier work is that babies do not understand the meaning of ordinary sortals, like *duck* or *ball*—in particular, babies do not understand the identity conditions of these nouns—until they are 12 months old, so they cannot be analogizing to such concepts to determine the identity conditions.

Instead of applying a sortal's identity conditions, the 9-month-olds must simply be assuming (incorrectly) that two different count nouns cannot label the same object—a type of mutual exclusivity bias (Markman, 1989), as Xu, Cote, and Baker (2005) suggest. Rhemtulla and Xu (2007) appear to believe, though, that infants classify the objects as members of different sortal categories and attribute to these new categories distinct but empty placeholders as identity criteria because the infants do not know what these identity criteria are. The infants then reason (via Principle 2) that because one object is a toma (fendle) and the other is not, the objects must be different. But it is not plausible to suppose that infants who, according to Xu and Carey (1996), have no concepts for basic-level categories, would decide in the course of a single experiment that (a) a toy duck belongs in the category toma, (b) a ball belongs to the category fendle, and (c) each category—each of the infants' very first lower level categories—is a sortal associated with a unique (but unknown) identity condition. It seems much more plausible to think the contrastive labels encourage infants to expect a correlated contrast, either quantitative (Xu, 2002) or qualitative (Dewar & Xu, in press). Thus, rather than establishing that sortals are the basic providers of identity conditions, these findings suggest the opposite. Like the results on individuation by properties (Xu & Baker, 2005), these data present further difficulties for the view that sortals are necessary for identity.

Along the same lines, we note that the developmental studies provide little or no evidence that infants ever consult a sortal's identity conditions to identify or to distinguish objects. In fact, evidence from Xu, Carey, and Quint (2004) suggests that even at 12 months, when infants are supposed to have acquired sortals and their associated identity conditions, they are unable to use them to decide that two objects within the same sortal category are distinct. For example, these infants cannot distinguish "a sippy cup with a top and two handles vs. a regular cup with one handle" (Xu et al., 2004, p. 175), despite their presumed knowledge of R_{cup}. The evidence for infants' use of basic-level sortals relies entirely on Principle 2, which has the problematic features we mentioned earlier (see Footnote 1). No evidence shows that infants use Principle 1, marshalling identity conditions to decide on identity (or difference) for members of basic-level categories.

Status of the Causal Continuer Theory

Belief in the identity of an object over time does not always depend on its sortal category. Sortals can still be useful in providing evidence for identity, but so do perceptual similarity, spatial and temporal continuity, and other cues. Can the causal continuer theory do better in explaining the basis of these judgments?

Rhemtulla and Xu (2007) are skeptical of the causal approach because they believe that causality is simply a cover term applied to many different situations, situations in which the effective variables may be similarity, spatiotemporal continuity, continuity of names, memory, and other factors. We count the fact that the causal continuer theory provides a unified account of identity as an advantage over approaches that rely on a combination of similarity, continuity, and other variables. Moreover, the experiments reported in Blok et al. (2005) and Rips et al. (2006) are systematic attempts to pit the causal theory against alternatives, deriving specific predictions from the theory and testing them both quali-

---

[6] Rhemtulla's (2005) studies, however, add a novel finding. In Rhemtulla (2005, Experiment 1), for example, participants learned that a specific pictured animal (e.g., a picture of a poodle named *Gax*) was disassembled into its individual atoms and then reassembled. In one condition, the resulting object looked identical to the original; in a second condition, it looked like a distinct member of the same subordinate category (another poodle); in a third, it looked like a distinct member of the same basic-level category (a dog that is not a poodle); and in a fourth, it looked like a distinct member of the same superordinate category (an animal that is not a dog). Participants in all conditions then rated whether the resulting creature was still the same individual ("Still Gax?") and whether it was still a member of the same basic-level category ("Still a dog?"). Ratings of whether the creature was the same individual followed a nonmonotonic function of taxonomic distance, with high ratings for the identity transformation and for the same-superordinate transformation, but lower ratings for the same-subordinate and same-basic conditions. The results for identity and same-superordinate conditions replicate Blok et al. (2005) in producing the dissociation just discussed. The remaining conditions were not included in Blok et al., and they present a puzzle to all accounts. Sortal theories predict a decreasing function, with identity possible only for members of the same basic-level category. Causal theories also predict a decrease, on the assumption that causal distance increases as a function of taxonomic distance. Rhemtulla and Xu (2007) suggest that pragmatic factors may be responsible for the failure of these predictions and that these factors cast doubt on the validity of the transformation paradigm. But their pragmatic account focuses on the role of the basic level in highlighting individual distinctiveness. Because Blok et al. did not include transformations within the basic level, this explanation does not apply to the experiment and does not affect the interpretation of the crucial dissociation.

tatively and quantitatively. Contrary to Rhemtulla and Xu, these predictions do not involve "post hoc descriptions, in each case, of how the observed results could have been due to causal reasoning" (p. 1088) but are consequences of the causal continuer theory. For example, Experiment 1 in Rips et al. contrasts causality with spatiotemporal continuity and sortal membership, and Experiment 2 contrasts causality with similarity. Although it is conceivable that a factor that fails to account for identity judgments in one experiment manages to do so in another, a theory that posits one underlying factor is more plausible and more attractive than a theory in which a changing set of variables explains the findings. The burden of proof would seem to be on those who claim that an alternating cast of (noncausal) factors collectively accounts for the results. But the only positive evidence that Rhemtulla and Xu offer—results from Rhemtulla (2005)—is, on their own account, not a reflection of identity decisions (see Footnote 6).

The strategy Rhemtulla and Xu (2007) use in defense of the cover term idea is to assume without evidence that certain task factors, mediating variables, or unobserved variables are directly responsible for participants' judgments. They then assert that causal connectedness is a cover term because it must apply broadly enough to describe these factors. For example, Rhemtulla and Xu state that participants "judged identity to be dependent on physical similarity" (p. 1089) in experiments in which the original object differed in its physical appearance from the later object (e.g., Blok et al., 2005, Experiment 2; Liittschwager, 1995; Rhemtulla, 2005). But "dependent on" is ambiguous in this context. Although the responses may have depended on physical similarity, in that participants have to recognize the physical properties of the stimuli to understand the task, there is no evidence that physical similarity is what was underlying their judgments. Consequently, there is no reason to think that causality is serving as a cover term for physical similarity.[7]

Rhemtulla and Xu (2007) use the same strategy for each of the experiments they discuss in their Section 1. In Rips et al.'s (2006) Experiment 2, for example, participants read scenarios about scientists who were studying a $3 \times 3 \times 3$ m iceberg and, after an interval away, returned to find an iceberg with different dimensions. Participants tended to rate the latter iceberg more likely to be identical to the former if it had smaller dimensions. Rhemtulla and Xu assert that "in this situation, causal knowledge acted as a cue to spatiotemporal continuity" (p. 1088). However, participants received no information about spatiotemporal continuity of the icebergs, and there is no evidence that they considered spatiotemporal continuity rather than causal connectedness in their decisions.[8]

This strategy works for any experiment. Because theoretical factors, such as causality, are not directly observable, experiments must manipulate more concrete proxy variables. You can always claim that such a variable, rather than the theoretical factor, is responsible for the effect. You can also claim that some unobserved variable is responsible (as in the case of spatiotemporal continuity, just discussed) or even that the dependent variable is responsible.[9] You can then claim that the theoretical factor is merely a cover term because the construct must be stretched to include the assumed factors. To make matters worse, the greater the variety of experiments supporting the theoretical construct, the more stretching must occur. But this strategy is unreasonable.

Without evidence that the alleged factors really are responsible for the effect, the strategy merely assumes what it wants to prove.

Concluding Comments

Successful cognitive theories produce predictions for a range of stimulus conditions, exhibit advantages over rival theoretical mechanisms, and unify new results with old ones. The causal continuer theory aims at these goals. We hope to have shown in

---

[7] There may be a lingering doubt about the causal continuer theory because of the variety of causes it could encompass. We argued earlier that the causes that are critical for object identity depend on the causal structure itself rather than on the categories to which the object belongs. But what if entirely different causes are important for objects of different kinds (e.g., the causes that support the identity of a robin are completely different from those that support the identity of a book)? Would that not show that the model uses causation as a cover term or that the model is incomplete? However, we do not see this generality as a defect of the model. The theory specifies a particular schema that determines object identity, given a set of causes, and in this respect it is exactly on par with other well-known theories. For example, decision theory predicts a particular alternative as optimal, given a decision maker's utilities and subjective probabilities, but it does not dictate which utilities or probabilities the decision maker ought to hold. Although decision theory may have its problems, generality is not one of them.

[8] Likewise, Rips et al.'s (2006) Experiment 1 described a machine that transmitted a variable proportion of copied particles from one object to another (perhaps identical) object. Thus, the experiment provided a measure of causal connectedness in terms of the proportion of transmitted particles. Rhemtulla and Xu (2007) comment that "in this experiment, causal connectedness looks like a compositional criterion because it is measured by the percentage of the original's (copied) particles" (p. 1088). But although the connectedness is measured by the percentage of particles, there is no evidence that composition rather than causal connectedness was the criterion participants were applying.

[9] This problem occurs in Rhemtulla and Xu's (2007) discussion of Experiment 3 in Rips et al. (2006). In this study, we presented participants with Y-shaped diagrams that we described as maps of streets or rivers. We labeled the vertical segment *Green Street* or *Green River*, and participants decided which of the branches (which varied in width and angle) was the likely continuation of Green Street or River. According to Rhemtulla and Xu, participants "judged street and river identity on the basis of whether or not the continuing forks were likely to have the same proper name as the original" (p. 1089). However, this statement contains the same equivocation discussed above. In one sense, all of the adult studies of identity, including Rhemtulla (2005), are "based on" whether later objects have the same proper name as the original, because in each of these studies participants had to decide whether the later object was "still Fred," "still Maf," and so on. This simply describes the experiments' dependent variable. However, decisions in these studies were not "based on" proper names in the sense that participants used proper names to guide their answers, because no proper names were given to the successor objects (the branches of the street or river). What Rhemtulla and Xu (2007) appear to mean is this: Judgments of the identity of streets and rivers, unlike that of other objects, depend on naming conventions that dictate which angles and widths continuations can have ("The identity of streets over space may be an exception: It may be that the identity of an individual street is completely determined by its proper name, such that arbitrary redesignation can sever identity" [p. 1089].) But this is entirely ad hoc. Why would the power of naming conventions for streets differ from that for other artifacts and natural kinds, for which naming conventions typically reflect deeper relations?

our earlier articles that the theory delivers correct predictions for situations in which participants have to choose which later object is identical to an earlier one, both in relatively naturalistic cases and in more extreme ones. In direct tests, the theory outperforms those based on sortal category membership (Blok et al., 2005, Experiments 1–2; Rips et al., 2006, Experiment 1) and similarity (Rips et al., 2006, Experiment 2). And the theory provides a plausible account of earlier results (e.g., Liittschwager, 1995). Generality for the theory entails that causal continuity will be manifest in different ways in different experiments. Different experimental conditions are necessary to eliminate confounds from the many factors that accompany causal closeness because it is impossible to vary all potential factors in a single experiment. However, we think that the theory provides a consistent way of explaining all of the results on adult judgments of identity, and it dovetails nicely with current theories of categories.

References

Ahn, W.-K. (1998). Why are different features central for natural kinds and artifacts? *Cognition, 69,* 135–178.

Ayers, M. (1974). Individuals without sortals. *Canadian Journal of Philosophy, 4,* 113–148.

Ayers, M. (1997). Is *physical object* a sortal? A reply to Xu. *Mind & Language, 12,* 393–405.

Blok, S., Newman, G., & Rips, L. J. (2005). Individuals and their concepts. In W.-K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 127–149). Washington, DC: American Psychological Association.

Campbell, J. (2006). Sortals and the binding problem. In F. MacBride (Ed.), *Identity and modality* (pp. 203–218). Oxford, England: Oxford University Press.

Cole, S. A. (2001). *Suspect identities: A history of fingerprinting and criminal investigation.* Cambridge, MA: Harvard University Press.

Dewar, K., & Xu, F. (in press). Do 9-month-old infants expect distinct words to refer to kinds? *Developmental Psychology.*

Dummett, M. (1973). *Frege: Philosophy of language.* Cambridge, MA: Harvard University Press.

Fodor, J. A. (1987). Modules, frames, fridgeons, sleeping dogs, and the music of the spheres. In J. L. Garfield (Ed.), *Modularity in knowledge representation and natural-language understanding* (pp. 25–36). Cambridge, MA: MIT Press.

Gelman, S. A. (2003). *The essential child.* Oxford, England: Oxford University Press.

Hirsch, E. (1997). Basic objects: A reply to Xu. *Mind & Language, 12,* 406–412.

Kaufmann, W. J., III (1994). *Universe* (4th ed.). New York: Freeman.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development.* Cambridge, MA: MIT Press.

Kelly, M. H., & Keil, F. C. (1985). The more things change: Metamorphoses and conceptual structure. *Cognitive Science, 9,* 403–416.

Liittschwager, J. C. (1995). Children's reasoning about identity across transformations. *Dissertation Abstracts International, 55* (10), 4623B. (UMI No. 9508399)

Lowe, E. J. (1989). What is a criterion of identity? *Philosophical Quarterly, 39,* 1–21.

Mackie, P. (2006). *How things might have been: Individuals, kinds, and essential properties.* Oxford, England: Oxford University Press.

Macnamara, J. (1986). *A border dispute: The place of logic in psychology.* Cambridge, MA: MIT Press.

Markman, E. M. (1989). *Categorization and naming in children.* Cambridge, MA: MIT Press.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–195). Cambridge, England: Cambridge University Press.

Merricks, T. (1998). There are no criteria of identity over time. *Noûs, 32,* 106–124.

Murphy, G. L. (2002). *The big book of concepts.* Cambridge, MA: MIT Press.

Nozick, R. (1981). *Philosophical explanations.* Cambridge, MA: Harvard University Press.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General, 130,* 323–360.

Rhemtulla, M. (2005). *Concepts of individuals do not allow identity persistence within the basic level.* Unpublished master's thesis, University of British Columbia, British Columbia, Canada.

Rhemtulla, M., & Xu, F. (2007). Sortal concepts and causal continuity: Comment on Rips, Blok, and Newman. *Psychological Review, 114,* 1087–1095.

Rips, L. J. (2001). Necessity and natural categories. *Psychological Bulletin, 127,* 827–852.

Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review, 113,* 1–30.

Rosch, E. H., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439.

Sloman, S. A., Love, B. C., & Ahn, W.-K. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22,* 189–228.

Strawson, P. F. (1959). *Individuals: An essay in descriptive metaphysics.* London: Routledge.

Wiggins, D. (1997). Sortal concepts: A reply to Xu. *Mind & Language, 12,* 413–421.

Wiggins, D. (2001). *Sameness and substance renewed.* Cambridge, England: Cambridge University Press.

Wilcox, T., & Baillargeon, R. (1998). Object individuation in infancy: The use of featural information in reasoning about occlusion events. *Cognitive Psychology, 37,* 97–155.

Xu, F. (1997). From Lot's wife to a pillar of salt: Evidence that *physical object* is a sortal concept. *Mind & Language, 12,* 365–392.

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition, 85,* 223–250.

Xu, F. (2005). Categories, kinds, and object individuation in infancy. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 63–89). Mahwah, NJ: Erlbaum.

Xu, F., & Baker, A. (2005). Object individuation in 10-month-old infants using a simplified manual search method. *Journal of Cognition and Development, 6,* 307–323.

Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology, 30,* 111–153.

Xu, F., Carey, S., & Quint, N. (2004). The emergence of kind-based object individuation in infancy. *Cognitive Psychology, 49,* 155–190.

Xu, F., Cote, M., & Baker, A. (2005). Labeling guides object individuation in 12-month-old infants. *Psychological Science, 16,* 372–377.

Zwaan, R. A., Magliano, J. P., & Graesser, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 386–397.