Mind & Language, ISSN: 0268-1064
Vol. 12. Nos. 3/4 September/December 1997, pp 365–392.

From Lot's Wife to a Pillar of Salt: Evidence that *Physical Object* is a Sortal Concept

FEI XU

Abstract: A number of philosophers of language have proposed that people do not have conceptual access to 'bare particulars', or attribute-free individuals (e.g. Wiggins, 1980). Individuals can only be picked out under some sortal, a concept which provides principles of individuation and identity. Many advocates of this view have argued that *object* is not a genuine sortal concept. I will argue in this paper that a narrow sense of 'object', namely the concept of any bounded, coherent, three-dimensional physical object that moves as a whole (Spelke, 1990) is a sortal for both infants and adults. Furthermore, *object* may be the infant's first sortal and more specific sortals such as *cup* and *dog* may be acquired later in the first year of life. I will discuss the implications for infant categorization studies, trying to draw a conceptual distinction between a perceptual category and a sortal, and I will speculate on how a child may construct sortal concepts such as *cup* and *dog*.

1. Introduction

1.1 What Is a Sortal?

Since Frege (1884) first observed that one cannot count without specifying what to count, various philosophers and psychologists of language have argued that certain concepts dubbed 'sortals', e.g. *car*, *person*, *dog*,[1] provide principles of individuation and numerical identity. These concepts tell us what to count as one instance of something and whether something is the same one as what we have seen before (Geach, 1962; Gupta, 1980; Wiggins, 1980; Hirsch, 1982; Macnamara, 1987; Lowe, 1989a; Macnamara and Reyes, 1994). Sortals are lexicalized as count nouns in natural languages that make the count/mass distinction. For example, a request to 'count the red in this room' cannot be complied with: Should a red shirt be counted as one or should the shirt, the two sleeves, and two pockets be counted separately so that we have five reds? In contrast, a request to 'count the red shirts in this

I would like to thank Susan Carey, Lila Gleitman, Eli Hirsch, John Macnamara, Gary Marcus, Sandeep Prasada, Cristina Sorrentino, Elizabeth Spelke, Robert Stalnaker and Josh Tenenbaum for discussions of the subject matter and/or comments on earlier drafts of the paper. Portions of this paper were presented at the 1995 Annual meeting of the Society for Philosophy and Psychology.

Address for correspondence: Fei Xu, E10-244, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
Email: fei@psyche.mit.edu.

[1] Throughout the paper I will adopt the convention that all sortals be italicized and all mention of words be in quotes.

room' will receive a definite answer: a red shirt (with its sleeves and pockets) should be counted as one shirt, not two. Hence the count noun 'shirt' gives us the principles for what to count as one shirt whereas the adjective 'red' does not provide principles of counting. In general, other predicates besides count nouns, e.g. verbs or adjectives, do not serve the logical function of providing principles of individuation. We cannot count 'sleeping' or 'blue' unless we mean, e.g. count the *naps* you took or the blue *birds*.

Sortals also provide principles of numerical identity. We cannot ask the question 'is this the same X' without using a sortal to specify what X is. When a dog dies, even though we can trace a spatiotemporally connected path from the dog to its body, we nevertheless decide that the dog has gone out of existence. Your dog and the neighbour's dog are two different dogs whereas a certain puppy and a certain grown dog may be the same dog. Again, adjectives and other grammatical categories do not provide such principles of identity. For instance, the question whether something is 'the same red' does not have a definite answer unless we mean 'the same *shade* of red' or 'the same red *sweater*'—count nouns such as 'shade' and 'sweater' provide the principles of identity.

1.2 Some Fundamentals of a Logic of Sortals

The fact that these logical functions are uniquely fulfilled by natural language count nouns has led some philosophers of language to develop a logic of sortals (Gupta, 1980; Wiggins, 1980; Macnamara and Reyes, 1994 among others). The fundamental tenet of a logic of sortals is *there are no bare particulars*; we cannot enumerate or trace identity without the support of a sortal. 'Bare particulars' are the alleged individuals that have no properties of their own whatsoever but still serve as entities on which to hang properties. Suppose someone is pointing at some part of the visual scene and uttering the word 'that'. The demonstrative 'that' may refer to a bare particular. It does not pick out an individual for which we can trace identity over time. We may be able to figure out that the person intends to pick out part of the visual scene with a table present, but we would not know whether the person is pointing to the table, a colour patch of the table, the millions of molecules of the table, or the table plus the dish that is sitting on it. The logic of sortals contrasts with standard model-theoretic semantics, in which the logical form for 'Some person is tall' in predicate calculus is '$\exists$x (person (x) & tall (x))'. In this formulation, x is a bare particular—it is an individual with no properties of its own but it still supports properties such as 'being a person' or 'is tall'. The champion of the logic of sortals denies that we have conceptual access to bare particulars.[2]

[2] The reviewers of an earlier version of this paper pointed out to me that standard model-theoretic semantics is an all-purpose tool and it should not be accused of having any inbuilt metaphysical bias to 'bare particulars'. However, if a semantic theory intends to be psychologically real, the formalisms of standard model-theoretic semantics do carry with them the commitment to bare particulars.

Besides the stricture against bare particulars, three other important tenets of a logic of sortals are relevant to the present discussion. First, a principled distinction should be made between count nouns and other grammatical categories; second, the principles of individuation and identity supported by sortals cannot be reduced to some basic notion of spatiotemporal continuity; third, within the class of sortals, a distinction should be made between substance and stage (or phase) sortals.

Consider first the distinction between count nouns and other grammatical categories. Count nouns denote kinds which provide principles and criteria for individuation and identity whereas other grammatical categories do not. Adjectives and other predicates are predicates of the individuals sorted by count nouns; their interpretations often depend on which count nouns they are predicated of, e.g. 'good' as in 'a good person' means something rather different from 'good' as in 'a good thief'.

Second, principles of individuation and identity provided by sortals may override our basic criteria of identity based on spatiotemporal continuity. To borrow an example from Hirsch (1982): A car consigned to a crusher follows a spatiotemporally continuous path in the crushing process and it gradually becomes a pile of metal and plastic, but nevertheless at some point, we decide that the car has gone out of existence. This is because the sortal *car* provides the criteria for what counts as a car.

Third, there is a distinction between substance and stage (or phase) sortals (Wiggins, 1980). In a nutshell, a count noun is a substance sortal if instances of the sortal it denotes cease to exist when they cease to be members of the sortal, e.g. *person*, *dog*, *tree*, *car*. In other words, substance sortals satisfy the condition that once something is no longer an X, it is also 'no longer'. For example, when a person dies, he ceases to be a member of the sortal *person* and he goes out of existence. Hence 'person' is a substance sortal. Substance sortals contrast with phase sortals such as *baby* or *tadpole*, which do not have this property—a baby does not cease to exist when she grows up even though she or he is no longer a member of the sortal *baby*. Similarly, a tadpole does not cease to exist when it becomes a frog although it is no longer a member of the sortal *tadpole*. For Wiggins (1980), only substance sortals stand for genuine kinds in a metaphysical sense.

A logic of sortals has been offered as an alternative for formalizing natural languages such as English. As mentioned above, standard model-theoretic semantics posits bare particulars, i.e. attribute-free individuals. Recently, Macnamara and Reyes (1994) have argued that in order to have a psychologically real semantics, we need a better formalism for a natural language like English, which would take into account the fact that individuals have to be typed by sortals. For example, the logical translation for 'the man is tall' is better represented as '$\exists$ (x: man) (tall (x: man))'—to be read as 'There is a man such that he is tall'. In this formalism, there is no x that stands by itself without the support of a sortal such as *man*. The grammatical distinction between 'man' and 'is tall' is also respected in that the interpretation of the predicate 'tall' now depends on which sortal it modifies. Employing a branch of mathematics,

namely category theory, Macnamara, Reyes, and their colleagues have developed a detailed theory of kinds/sortals. Their system of kinds is a system of underlying maps among kinds, instead of the standard formulation of kinds as class inclusions (see Macnamara and Reyes, 1994, for details).

1.3 Object *and* Thing*: Candidate Bare Particulars or General, Universal Sortals?*

It is crucial for sortal theorists such as Macnamara that people have no conceptual access to bare particulars. If they did, we would not need a logic of sortals to represent the semantics of natural languages: All count nouns as well as other grammatical categories could then be construed as predicates of these bare particulars.

Bare particulars may be devastating for the sortal approach to semantics, but even a general, universal sortal, as sortal theorists believe, would cause serious damage. Finding a general, universal sortal entails that all other sortals, e.g. *dog* or *car*, would be dispensable because they could be construed as predicates of this general, universal sortal. Note this would be a rough equivalent of finding a bare particular—other count nouns as well as predicates such as verbs and adjectives would be construed as predicates of this general, universal sortal. If such a general sortal existed, we would not need to draw a sharp line between count nouns and other grammatical categories, and predicate calculus could be only slightly modified in order to represent natural language semantics: We could replace the xs with this general sortal.

One might ask why the sortal theorists insist on *dog*, *tree* and *person* being the right level of specificity for sortal-hood. Implicitly or explicitly, all sortal theorists agree that a sortal should provide a satisfactory answer to the question 'What is it' (Wiggins, 1980). *Dog*, *tree* and *person* certainly fulfil this requirement. Furthermore, the psychological literature on kind concepts suggests that these 'basic-level' sortals enjoy a privileged psychological status (Rosch et al., 1976).

A first glance at a natural language such as English, however, seems to give us some candidates for either bare particulars or general, universal sortals. Count nouns such as 'object', 'thing' or 'entity' give the appearance of being either individuals that have no properties of their own but still serve as individuals on which to hang properties, or general sortals that other specific sortals (e.g. *dog*, *car*) may be predicates of.

Not surprisingly, sortal theorists have argued that 'object', 'thing' and 'entity' are not sortal concepts. For example, Wiggins (1980) says, 'For a formal concept like *entity* or *substance* has no autonomous individuative force of its own, and must be variously supplemented, wherever it appears in contexts of identification, according to the kind of the individual in question' (p. 63). '*Material object* is now ruled out from sortal status, and so are other dummy substantives, . . .' (p. 64). That is, Wiggins explicitly denies that 'object' and 'entity' are sortals.

Three types of arguments have been put forth against 'object' or 'thing' being a sortal. The first type is that natural language terms such as 'object'

or 'thing' do not provide principles of individuation (Wiggins, 1980; Macnamara, 1987, 1994; La Palme Reyes et al., 1994). Macnamara (1994) says it clearly: 'We cannot conceptually grasp an individual in a universal kind supposedly denoted by the count noun "thing" or "object"' (p. 20). The reason is that 'thing' or 'object' does not tell us what to count as one instance. If you were to count the things in the room, you might count the chair as a chair, four legs, plus one top. Hence six things altogether! Similarly, Hirsch (1982) claims that at least the broadest sense of the word 'object' is not a sortal: 'There may possibly be a completely permissive sense of the word "object" which applies in fact to any aggregate of matter, however spatially discontinuous' (p. 97). This sense of 'object', or even the sense of the 'object' that applies to any arbitrary continuous portion of matter, certainly does not give us countable individuals.

The second type of argument is that 'object' or 'thing' does not provide principles of numerical identity. The Old Testament (Genesis, chapter 19) tells the story of Lot's wife. Lot and his wife were told by the Lord 'Escape for thy life; look not behind thee. . . . Then the Lord rained upon Sodom and upon Gomorrah brimstone and fire from the Lord out of heaven. . . . But his wife looked back from behind [Lot] and she became a pillar of salt.' Wiggins (1980) claims that in this process, nothing persisted through time: Lot's wife ceased to exist and a pillar of salt came into being. According to Wiggins, there is no sortal concept, e.g. *object* or *thing*, that could answer the question of 'same what?' for Lot's wife and the pillar of salt. He argues that short of inventing a sortal concept *woman-pillar*, 'certainly there is no substantial sortal, . . . suitable to cover the identity between Lot's wife and the forty foot pillar of salt that is still to be encountered even now on the Jebel Usdum near the Dead Sea' (p. 61). The difficulty with *object* or *thing* as candidate sortal concept in Lot's story, as I understand Wiggins, arises from the fact that in this story, an entity starts off with the persistence principle of one kind (i.e. *woman*), then exchanges that principle for the persistence principle of another kind (i.e. *pillar of salt*). In addition, Wiggins believes that the story of Lot's wife is incoherent because it violates actual laws of nature.

Finally, terms such as 'object' or 'thing' may be less good candidates for answering the question 'what is it?'. If one were to ask the question 'what is it?', an answer such as 'a cup' or 'a person' would suffice. But if the answer were to be 'a thing' or 'an object', we still would not have any idea what that thing is! It is peculiar that 'object' and 'thing' are nevertheless lexicalized as count nouns even though they don't seem to behave in the same way as most other count nouns, i.e. they are not used to answer the question 'what is it?'. One explanation may be that these nouns are place-holders and not true sortals.

Although the English words 'object' and 'thing' may not be sortals, as argued above by philosophers of language, I will argue in this paper that for both adults and young infants, there is nonetheless a sortal, *physical object*, which is more general than *person*, *car* or *tree*. A physical object is defined as any three-dimensional, bounded entity that moves on a spatiotemporally continuous path

(Spelke, 1990; see also Jackendoff, 1983). The English word 'object' has multiple senses; this concept of physical object corresponds to one of the senses. In the rest of the paper, I will defend the claim that *physical object* (that is, one sense of the English word 'object') is a sortal. Section 2 provides evidence that *physical object* is the first sortal for young infants. Section 3 provides evidence that *physical object* is a sortal for adults and that the aforementioned arguments against 'object' being a sortal may be remedied. Section 4 discusses the relation between *physical object* and other more specific sortals. Section 5 discusses the advantages of building *physical object* as infants' first sortal concept and finally, in Section 6, some possible learning mechanisms which would allow infants to construct more specific sortals are discussed.

2. Evidence from Infants: Physical Object *as Their First Sortal*

The philosophical literature on numerical identity generally agrees that any explication of the concept of identity should satisfy Leibniz's Law (with the exception of relativists such as Geach): If x is identical to y, x and y should have exactly the same properties. Although Leibniz's Law is logically necessary and sufficient for numerical identity, it does not give us a psychological notion of how things persist through time and space. A psychologically real explication of the concept of identity, however, should include people's criteria for persistence over time, that is, how people decide whether something is the same one they have seen before.

Our criteria for individuation and identity of physical objects may be roughly divided into two types: spatiotemporal information and property/kind (or sortal) information. The spatiotemporal criteria include the following generalizations: (1) one object cannot be at two places at the same time; (2) two objects cannot be at the same place at the same time; (3) objects travel on spatiotemporally connected paths. No object can move from point A to point B without traversing a continuous path in between; if some object appears to have travelled on a spatiotemporally discontinuous path, people judge that there must be two objects involved. The property/kind criteria include the following generalizations: (1) upon seeing a member of a kind now (e.g. a cup) and a member of a different kind (e.g. a dog) at a later time, we infer there are two numerically distinct entities; (2) upon seeing a member of a kind now (e.g. a red block) and a member of the same kind with a different property (e.g. a blue block) at a later time, we (often) infer there are two numerically distinct entities. Note that the spatiotemporal criteria apply to all physical objects, regardless of kinds of object. The property/kind criteria, on the other hand, are kind-relative. Certain property changes signal a change in identity only within certain kinds of objects. For example, if you see a small chair in the corner now and a big chair there later, you infer that there are two numerically distinct chairs. But if you see a small plant in the corner now and a larger one there a few months later, it is not necessarily the case that there are two distinct plants. One way to see what sortals adults or infants represent is by probing what criteria they use in individuation and identity judgments.

2.1 Evidence for Object *as a Sortal in Infancy*

Psychological investigations have focused on what criteria are employed for individuating objects and deciding whether something is the same one as seen before. Piaget and T.G.R. Bower first raised these questions with young infants. Bower (1974/1982) presented some experimental evidence that very young infants use spatiotemporal rules to trace identity, but that before 5 months of age they do not use property differences to infer change of identity. In addition, he claimed that infants were guided by different spatiotemporal rules from adults: for infants, stationary and moving objects are numerically distinct. Bower employed a visual tracking paradigm in these studies. The rationale for the method was that if the infants detected something unexpected or novel, their tracking behaviour would be interrupted. In one experiment, infants learned to track moving objects by turning their heads. In the test trials, the moving objects would suddenly stop. Bower found that infants kept turning their heads when the object stopped. He interpreted these results as showing that the infants thought that the moving object would continue to move along its path and the stationary object (which was moving until it stopped suddenly) was a different object. In the second experiment, Bower tested whether infants used property differences to establish object identity. A moving object, say a bunny, would disappear behind a screen, then another object, say a ball, would appear from behind the screen. Bower found that young infants kept tracking the new object and did not look back to the screen. He interpreted these results as showing that infants did not make the inference that the ball was a different object from the bunny. By 5 months of age, however, infants looked back at the screen. Bower interpreted this as indicating that infants realized that the ball and the bunny were two different objects and the bunny must have stayed behind the screen.

There are a number of serious methodological flaws in these experiments: the failure to stop head turning in the first experiment may be due to the fact that young infants do not have very good neck control at this age and hence cannot inhibit the head-turning motion when the object stops; the failure to look back at the screen in the second experiment could be due to a similar factor, or alternatively, it could be due to the fact that the infants were well aware that the ball was a different object from the bunny, but kept tracking because of its novelty. Furthermore, a number of researchers either failed to replicate Bower's findings or found that the looking back behaviour was a function of the speed at which the objects were moving but not a function of the property differences; that is, the looking back behaviour occurred equally often when the same object as opposed to a different object emerged from behind the screen (Muller and Aslin, 1978; Meicler and Gratch, 1980; Gratch, 1982).

In sum, Bower's experiments do not give us any clear evidence about infants' criteria for object identity. However, his conjecture that infants use spatiotemporal information before using property information to trace identity may be on the right track. There is now considerable evidence in the

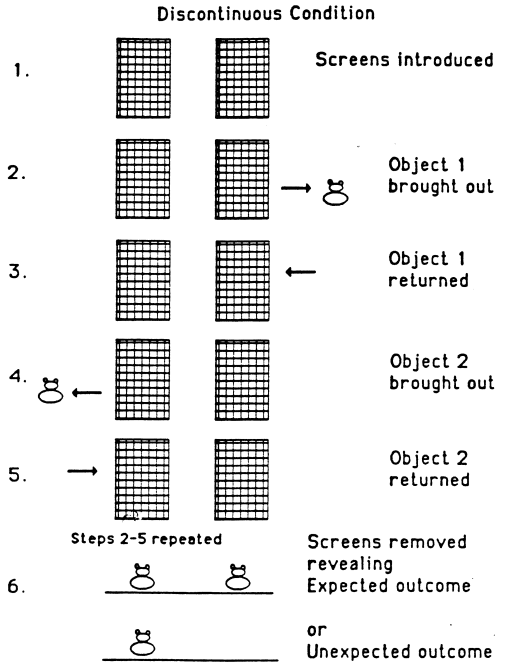

Figure 1 Schematic representation of the Spelke et al. (1995) experiment

infant cognition literature suggesting that infants as young as 4 or 5 months use all three spatiotemporal generalizations to individuate objects and trace identity, just like adults. But unlike adults, they do not use property/kind information until about 10 to 12 months of age.

In the last decade or so, a methodology has been developed to study the cognitive capacities of pre-verbal infants (Spelke, 1985). In this method, infants are shown the same event repeatedly and their looking times recorded. With each repetition their looking times decline, that is, infants 'habituate'. When infants reach a pre-set habituation criterion, they are shown two displays alternately, one consistent with adults' understanding of the event and the other inconsistent. If the infants have the same understanding of the habituation event as adults, they should look longer at the inconsistent display as opposed to the consistent one. In a seminal study, Spelke et al. (1995) showed that 4-month-old infants take evidence of spatiotemporal discontinuity as evidence for numerically distinct objects. In this experiment, two screens were lowered onto the stage with some space in between them (see Figure 1 for a schematic representation of the event using different objects). The infant saw that a rod appeared from behind one screen, say the left one, moved to the left end of the stage, then returned behind the left screen. No object appeared between the two screens. After a

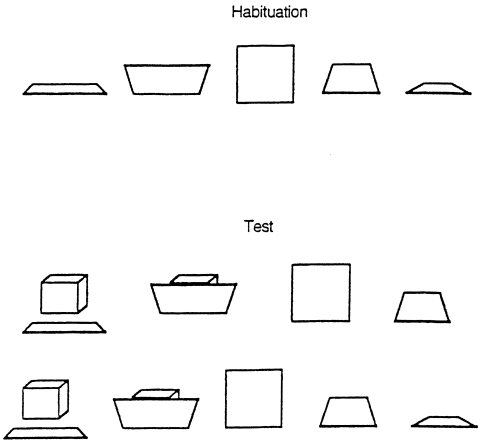

Figure 2 Schematic representation of the Baillargeon et al. (1985) experiment

short pause, a physically identical rod appeared from behind the right screen, moved to the right end of the stage, then returned behind the right screen. This event was repeated until the infant reached a habituation criterion, which was defined as the average looking time of the last three habituation trials being half of the first three trials or less. The screens were then removed to reveal one of two outcomes: the expected outcome of two identical rods or the unexpected outcome of just a single rod. The infants looked longer at the one-rod outcome, suggesting that they, like adults, had expected two rods and were surprised to see just one. When the rod did appear in the space between the two screens, on the other hand, the infants looked about equally at the one-rod and two-rod outcomes, as if undecided as to how many rods were behind the screens. Xu and Carey (1996) replicated the above finding with 10-month-olds using a rather different set of objects, e.g. toy ducks, balls, elephants and trucks.

Baillargeon et al. (1985) presented evidence that 5-month-old infants understand that two objects cannot be at the same place at the same time. In these experiments a rotating screen facing the infants was introduced. After the infant had habituated to the rotating screen going 180° towards then away from her in a drawbridge fashion, a box was introduced and placed behind the screen. The rotation resumed and sometimes the screen stopped short of 180° (the expected event because the box was behind it) and sometimes the screen rotated all the way back to 180° (the unexpected event if the infant understood that two objects cannot occupy the same space at the same time; see Figure 2). Infants looked reliably longer at the unexpected event, suggesting that they expected the box to remain behind the screen and that the two objects, the screen and the box cannot occupy the same space at the same time.

In addition, Baillargeon and Graber (1987) showed that 5-month-old

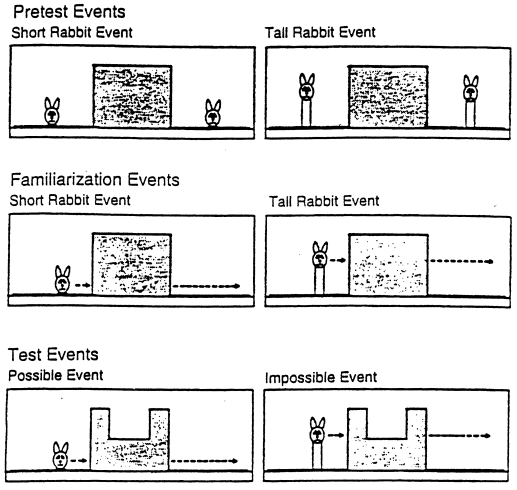

Figure 3 Schematic representation of the Baillargeon and Graber (1987) experiment

infants understand that one object cannot be at two places at the same time. In one of the experiments, infants were habituated to a tall rabbit going behind a screen and appearing on the other side (see Figure 3). Then the middle section of the top half of the screen was removed so that the tall rabbit should appear in this 'window'. If the rabbit did not appear in the window, the infants looked longer than if the rabbit did appear in the window. But if the infants were shown two identical tall rabbits simultaneously, one on each side of the screen, they did not look longer when no rabbit appeared in the window. Infants could only succeed if they interpreted the two identical-looking rabbits as two distinct rabbits using the location information. In other words, if shown two objects simultaneously, the infants set up representations of two numerically distinct objects that allowed them to resolve an apparent violation of spatiotemporal continuity.

Researchers have argued that in these experiments, infants are able to keep track of discrete physical objects. However, it is possible that the infants kept track of the amount of stuff in some of the experiments. That is, these results may not speak to the infants' representations of physical objects (which requires quantification of discrete individuals) but rather their ability to keep track of amounts of stuff (which only requires continuous quantification). However, experiments by Huntley-Fenner and Carey (1995) suggest that infants in these experiments apparently quantified over objects as discrete entities and did not simply keep track of the amount of stuff present in an array. In several experiments with 8-month-old infants, they showed that if sand is poured onto a stage behind a screen, infants fail to count the number

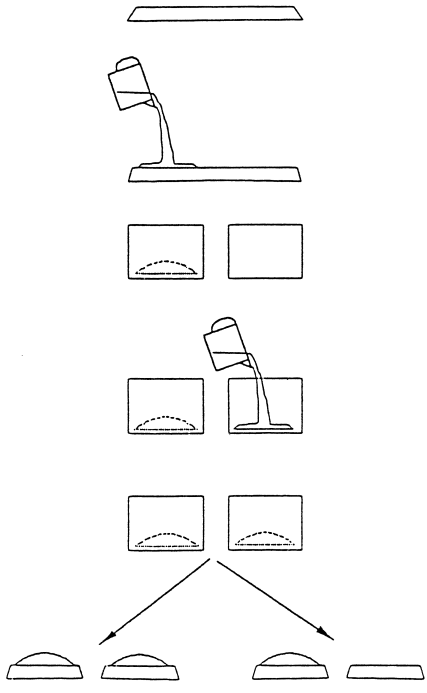

Figure 4 Schematic representation of the Huntley-Fenner and Carey (1995) experiment

of piles while succeeding in a similar experiment with objects (see Figure 4). In other words, infants do not seem to track the amount of stuff present in an event, which suggests that in the experiments involving solid objects, they quantified over individuals.

There are at least two interpretations of which sortals underlie the capacity revealed by the above experiments. One is that infants represent specific sortals such as *rod*, *box*, *rabbit*, and others. After all, all sortals (which refer to physical entities in the world) are subject to the spatiotemporal constraints on individuation and identity. However, a more conservative interpretation is that the sortal concept underlying these successes is *physical object*, since the spatiotemporal criteria apply to all physical objects regardless of whether the objects are ducks or cups. None of the above experiments *required* any representation of individuation and identity criteria for specific sortals. *Physical object* is a sortal because it provides criteria for individuation and numerical identity; these criteria are spatiotemporal in nature. Further evidence is needed to decide which interpretation of these results is correct.

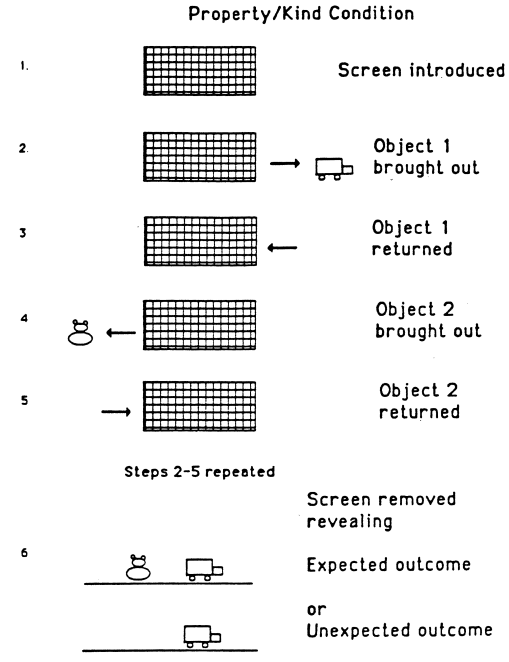

Figure 5 Schematic representation of the property/kind condition of Xu and Carey (1996)

2.2 Evidence for Object as the First *Sortal in Infancy*

Do young infants also represent more specific sortals such as *ball* or *bottle* or do they only represent the sortal *physical object*? Xu and Carey (1996) devised further experiments to address this question. For 'ball' or 'bottle' to be a sortal, minimally the infant should be able to use the differences between a bottle and a ball to set up representations of two numerically distinct individuals. In these experiments, a single screen was lowered on the stage. Ten-month-old infants saw that a ball appeared from behind the screen, moved to the left end of the stage and stopped. The infant's looking time was monitored, then the ball returned behind the screen. After a short pause, a bottle appeared from behind the screen, moved to the right end of the stage and stopped. Again the infant looked at it until she turned away, then the bottle returned behind the screen. After habituation, the screen was removed, revealing either two objects, a ball and a bottle (expected outcome), or only one of the two objects (unexpected outcome); see Figure 5. If the infant is able to use the property/kind (or sortal) difference between the ball and the bottle to infer two distinct objects, she should look longer at the one object outcome. Surprisingly, these 10-month-old infants failed to look longer at the unexpected one-object outcome. They simply exhibited a base-

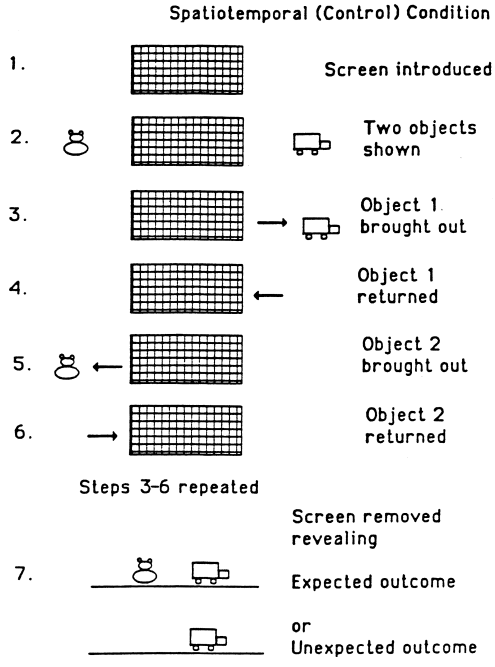

Figure 6 Schematic representation of the spatiotemporal condition of Xu and Carey (1996)

line preference for looking longer at two objects. Success in this task may or may not show that the infants represent sortals such as *ball* or *bottle*, because as adults, we can also use the property difference within a single sortal, say, a red ball and a blue ball, to infer two balls, but the failure certainly suggests that they do not represent sortals *ball* or *bottle*. To ensure that the result in this experiment was not an artifact of the method we employed, a control version was carried out. The infants were simply shown the two objects simultaneously at the beginning of the experiment; that is, they were given spatiotemporal evidence that there were two objects (see Figure 6). The experiment then unfolded as before and the outcomes were the same as before. The infants in this control group looked longer at the unexpected outcome of one object, suggesting that the method was sensitive and reliable, and supporting the conclusion that spatiotemporal criteria are used before property/kind criteria in object individuation.

One may wonder if the infants even coded the properties of the objects. If the infants did not encode the properties, one would not expect them to be able to succeed at this task. To block this alternative interpretation, Xu and Carey devised a variant of the above experiment in which infants' looking times to the habituation sequence of ball, cup, ball, cup or ball, ball,

ball, ball were recorded. The prediction was that if the infants encoded the properties of these objects, it should take them longer to habituate to the sequence ball, cup, ball, cup than to the sequence ball, ball, ball, ball. This is exactly what they found. After the habituation sequence, the screen was again removed to reveal one or two objects. Replicating the result of the earlier experiment, Xu and Carey found that infants did not look longer at the unexpected outcome of one object after seeing a ball and a cup one at a time. Xu and Carey concluded that even though the infants had encoded the properties of the objects, they did not use these differences to infer that there were two distinct objects.[3]

Infants' representations of these events were clearly different from those of adults. An important question is how exactly the 10-month-old infants represented the event in these experiments. It is likely that the infants represented the event as an object (with cup properties), an object (with ball properties), . . . and they did not commit themselves to whether there were one or two objects behind the screen. The longer looking at the two-object outcome simply reflected their intrinsic preference for two objects and may not be a positive expectation of one object. In other words, if the only sortal infants represent is *physical object*, then they can only use spatiotemporal criteria to individuate objects; in the Xu and Carey experiments, the spatiotemporal information was ambiguous, and therefore the infants were agnostic as to how many objects were behind the screen.

Further experiments in Xu and Carey (1996) showed that 12-month-old infants succeed at these tasks. This is at least suggestive that the older infants may have sortal concepts *ball* and *bottle*, as it seems likely that they represented the event as a bottle emerging from behind the screen followed by a ball emerging from behind the screen to the other side. As the two exemplars belong to two different sortals/kinds, they must be two distinct objects. However, an alternative representation not involving the sortals *bottle* and *ball* could underlie this success. Infants may have learned that a round thing with pink and green stripes does not change into a cylindrical thing with a lid, that is, they may be using the property differences between a ball and a bottle to infer change of identity. However, the following finding from this series of studies suggests that this task may be directly relevant to the infants' representation of sortals. In two versions of the experiment where 10-month-olds failed as a group (using the same paradigm as above), we found a correlation between word comprehension score and performance on the task. Within the group of 10-month-olds, the infants who were judged by their parents to

[3] The failure in using property/kind information to establish numerosity is sometimes taken as conflicting with Wynn, 1992, where 5-month-olds were shown to succeed in enumerating objects. However, in Wynn, 1992, infants were given clear spatiotemporal information about how many objects there were behind the screen, i.e. the second object clearly comes from a different place from the first. Furthermore, Wynn's objects were identical so that the infants could only rely on spatiotemporal information to establish numerosity and not property/kind information, as in the present studies.

understand at least 2 of the 4 nouns that name the objects in the task (i.e. ball, bottle, book and cup) succeeded at the task, whereas the infants who understood only one or none of these nouns failed. This rather tight relation between word comprehension and performance suggests that the infants who succeeded may have indeed used the kind difference as opposed to just property differences to make the inference. The infants' representation of the event may have been that 'a ball' (as opposed to 'an object with ball properties') appeared then disappeared, followed by 'a bottle' (as opposed to 'an object with bottle properties'), then a ball again, a bottle again, etc. If the two objects belong to different kinds, they must be two distinct objects. However, this representation was only possible for infants who already knew the nouns, i.e. a ball and a bottle. For the infants who did not know the nouns, detection of the properties such as colour, shape and texture did not help them infer that there were two objects behind the screen.[4]

The experimental evidence reviewed above suggests that *physical object* may be the first sortal concept infants represent and that it is not until 10 to 12 months of age that they represent more specific sortals such as *ball* or *bottle*.[5] Results from two series of additional experiments provide further support for this claim (Xu et al., 1995; Xu et al., 1996; Xu and Quint, 1997). The first series of studies presented 10- and 12-month-old infants with a static display which consisted of a toy duck perching on top of a toy car. After habituation or familiarization to this display, a hand grasped the head of the duck and lifted it. In the expected outcome, the duck was lifted but the car stayed on the stage floor; in the unexpected outcome, the duck and the car were lifted together as if they were parts of the same object. The idea was that if infants represented the sortals *duck* and *car*, they should parse the array into two distinct objects. They should look longer at the unexpected outcome. Xu et al. (1995) found that the 10-month-olds did not look longer at the unexpected outcome whereas the 12-month-olds did, providing converging evidence for the Xu and Carey (1996) claim. The second series of studies used number of reaches as the dependent measure instead of looking time. Twelve-month-old infants were trained to reach into a small hole on top of a box to retrieve objects. Sometimes two different objects, say a duck and a ball, were pulled out of the box, *one at a time so they never saw the two objects together*; sometimes the same object, say a duck, was pulled out of the box twice. If the infant can use the property differences between a duck and

---

[4] It is possible that the two processes, noun comprehension and success at the individuation task, are merely correlated but not causally related. I cannot rule out this possibility at the moment. Future experiments will address this issue.

[5] Xu and Carey make the claim that *object* may be the only sortal young infants represent, but of course they have not exhausted all possible kind contrasts. However, they have surveyed a fairly representative repertoire. Some of the stimuli (e.g. toy duck vs. ball, and toy truck vs. toy elephant) span categories of animate vs. artifact and vehicle vs. animate, and the other stimuli were highly familiar to 10-month-old infants (e.g. bottle, cup, ball, book). Current studies are underway to investigate changes that cut across certain ontological boundaries, e.g. a live gerbil vs. a chair, or a person vs. a large box.

a ball to infer two distinct objects, they should reach twice. That was indeed what Xu et al. (1995) found: the infants reached roughly twice when shown two different objects and once when shown the same object twice. In the third series of studies, Xu and Quint (1997) found that shape appears to have a privileged status in object individuation at 12 months. For example, if two objects differed only in shape, say a red ball and a red cup, infants inferred two objects; if two objects differed only in colour, say a red ball and a green ball, infants did not infer two objects. Shape may be privileged because shape changes are well-correlated with sortal changes.

2.4 Distinguishing a Perceptual Category and a Sortal

The results of Xu and Carey (1996) appear to conflict with other data that have been taken to show that 3- and 4-month-old infants represent basic-level *kinds/sortals*. Several studies have shown that young infants can form *perceptual categories* such as table, cat, horse or animal (e.g. Cohen and Caputo, 1978; Cohen and Younger, 1983; Quinn and Eimas, 1993; Eimas and Quinn, 1994). In these experiments, infants were habituated to exemplars from a given category, say, cat, then they were shown a new exemplar from the same category, another cat, or an exemplar from a different category, say, a cup. Infants dishabituated to the exemplar from the new category (the cup) but remained uninterested in the new exemplar from the familiar category (the cat). Furthermore, the infants were able to discriminate among the exemplars from the familiar category. That is, they can habituate to the same duck then dishabituate to another exemplar of a duck. These findings have been used to argue that infants represent basic level kinds / sortals, and that this representation supports the learning of natural language count nouns (Roberts and Horowitz, 1986; Macnamara, 1987; Roberts, 1988; Mandler et al., 1991).

I want to suggest, however, that these studies do not address the issue of sortal representation. For one thing, three possible underlying representations are consistent with the habituation results and only one of them supports the representation of sortals; these three representations are at the level of properties, persisting objects and specific sortals respectively.

Possibility One: The infants could be applying some similarity metric in a Quinean quality space which distinguishes cat-shape from cup-shape (or cat-properties from cup-properties), with no commitment to these exemplars being persisting objects. The exemplars are represented as mere bundles of features or properties. Just as the infants can habituate to red-ness then dishabituate to green-ness, they can also habituate to cat-ness then dishabituate to cup-ness. The infant's conceptual system, if there is one at all, bears almost no resemblance to the adult's, for it has no individuals that persist through time and space and the infant has no criteria for individuation and numerical identity, not even at the level of *object*.

Possibility Two: The infants could represent these exemplars as 'an *object* with cat-properties, an *object* with somewhat different cat-properties, . . . an

object with cup-properties', with no commitment as to whether the first cat, the second cat, and the cup were numerically distinct objects. But these objects, *qua* objects, persist through time and space, and spatiotemporal criteria are used to decide whether there is one or more of them. When the infant is shown a cat, another cat, etc. they may represent these stimuli as one object with changing properties, or perhaps because there is no clear spatiotemporal information the infant remains agnostic as to how many objects there are.

Possibility Three: The infants may have represented these events as 'a cat, another cat which is numerically distinct from the first cat, a third cat that is numerically distinct from the first two cats, . . ., a cup which is numerically distinct from the cats', in which sortals/kinds such as *cat* and *cup* play a role. Evidence for the third representation is what is needed if one were to claim that infants represent kinds/sortals that underpin count nouns.

Most of the infant habituation/categorization studies conclude that the data showed that infants have the concepts needed for learning count nouns such as 'cat' or 'cup', that is, the infant represents sortals/kinds. However, as the above discussion shows, the experimental paradigm used in these habituation studies does not tease apart which of the three possible representations underlies the habituation–dishabituation looking time pattern. Spelke's (1990) and Xu and Carey's (1996) findings suggest that Possibility Two is correct: Infants represent persisting objects using the spatiotemporal criteria but they do not represent specific sortals using kind criteria. This suggests that the young infants' representation was 'an object with cat-properties, an object with somewhat different cat-properties, . . . an object with cup-properties', where the indefinite determiner 'an' signals the agnosticism as to whether the objects were numerically distinct from each other. In the absence of clear spatiotemporal information, the infants had no basis for deciding whether the cat and the cup were numerically distinct.

Furthermore, these earlier categorization/habituation studies were inadequate in principle for addressing the question of representations of sortals for two reasons. One is that these studies confused two distinct psychological processes: our capacity to discriminate and correlate properties (perhaps a purely perceptual phenomenon) and our capacity to use discriminable properties to infer distinct individuals. The other reason is that different grammatical categories encode different types of meaning. A task aimed at probing the underlying representations of count nouns should demonstrate the conceptual representations that are unique to count nouns. Let me spell out these points.

First, Xu and Carey (1996) suggest that our capacity to discriminate properties and our capacity to use discriminable properties to infer distinct individuals are two psychologically distinct processes. They showed that these processes are dissociable; in one of the experiments, infants were able to discriminate between a duck and a ball, as shown by the slower habituation rate with the sequence duck, ball, duck, ball compared to the sequence

duck, duck, duck, duck, yet they failed to use these discriminable differences to infer two numerically distinct individuals.

This idea may be made intuitive with some everyday examples. For instance, a person may wear glasses one day and no glasses on another day; the two are clearly perceptually different, yet one and the same person can have these two different appearances. A human hand can take the form of a clenched fist or with the palm showing—again the two are perceptually different but the same hand (that is, the same individual) may take either form. We may also be agnostic about how many individuals are involved in a given scenario even though we have noted the perceptual differences. For instance, we might spot a beautiful green leaf on a maple tree on a fall day. A few days later we come by the same place and there is a leaf on the ground that is orange. We may remain agnostic as to whether it was the same leaf or not since we do not have sufficient information, although we certainly perceive the differences between them.

Once we have noted the differences in appearance, there is a further step to be taken to decide whether there are two distinct individuals. Suppose that we see a person with long black hair, wearing a red jacket and a black skirt. She then walks out of the room. A few minutes later a woman with a pony tail and glasses comes in, wearing a blue jacket and a blue skirt. We first note the differences in appearance; our visual system tells us that we are registering in one case human shape, blackness, redness, etc. and in the other case human shape, glass-ness, blueness, etc. Is she the same person or not the same person? These two logical possibilities are still open after noting these perceptual differences. She may be the same person who just changed clothes and hairdo, or she could be a different person. The point is that detecting perceptually discriminable properties does not warrant the inference of numerically distinct individuals. The earlier habituation/categorization studies have only shown that young infants are able to discriminate property differences; they do not speak to the question of whether the infants represented these differences as signalling numerically distinct objects.

Second, most researchers in lexical semantics (e.g. Grimshaw, 1981; Jackendoff, 1983; Pinker, 1984; Gleitman, 1990) and some philosophers of language (e.g. Gupta, 1980; Wiggins, 1980; Macnamara and Reyes, 1994) are committed to the idea that different grammatical categories, e.g. noun, verb and adjective, encode different types of meaning. Roughly speaking, there exist language universals where certain conceptual categories correspond to certain grammatical categories: object names are usually lexicalized as nouns, actions are usually lexicalized as verbs, and properties are usually lexicalized as adjectives. As discussed in the first section of this paper, count nouns fulfil two logical functions that distinguish them from other grammatical categories, i.e. they provide principles and criteria for individuation and numerical identity. Therefore, empirical studies that attempt to show that infants have the conceptual representations required for learning count nouns should satisfy the following condition: the experiment should show

that infants can individuate objects and trace their identity over time; it should uniquely probe these logical functions of count nouns and should not be equally compatible with conceptual representations for other grammatical categories. The infant categorization / habituation studies fail to meet this criterion. On the one hand, they do not address the issue of individuation and identity; on the other hand, these experiments do not uniquely probe the type of conceptual representations that underlie count nouns. For example, using these types of procedures, we can also show that infants discriminate red things from green things. That is, if the infants are shown red square, red triangle, red circle, etc. they would then dishabituate to a green circle, but not a red rectangle. This method would show equally well that infants are sensitive to colour contrasts such as red vs. green, but most of us would not claim that the infant represents *sortals* such as red and green, but rather red-ness vs. green-ness. In other words, these infant categorization studies do not address the question of whether infants have the conceptual representations that underpin count nouns and only count nouns.

As we have seen, a conceptual distinction needs to be drawn between discriminating properties and being able to use these property differences to infer distinct individuals. Furthermore, different grammatical categories require rather different kinds of conceptual representations. The task of a cognitive scientist interested in how conceptual representations are related to language acquisition is to devise empirical tests which probe the different logical functions of various grammatical categories.

3. Arguments for Physical Object *as a Sortal in the Adult's Conceptual System*

The infant studies discussed above show that *physical object* is a sortal concept, and it may be the *first* sortal infants represent. The studies suggest that the sortal *physical object* provides a set of criteria for individuation and numerical identity independent of the criteria provided by the specific sortals such as *cup* and *ball*. However, it is possible that *object* is not a sortal in the adults' conceptual system. That is, in the course of development, infants may come to construct more specific sortals such as *bottle*, *ball* and *dog*, then abandon *physical object* completely. Below I will present several arguments that *physical object* functions as a genuine sortal even for adults. As I have claimed that *physical object* is one sense of the English word 'object', I will suggest that if we apply this narrower sense, the aforementioned arguments against 'object' being a sortal become unconvincing.

3.1 Does Object *Provide Principles and Criteria for Individuation?*

First consider the question of criteria for individuation. Imagine that you are asked to count the objects in a room that contains tables, chairs, mirrors, books, and other household items. A moment of introspection will tell you

that each chair or each table counts as one object and nobody will count a table as five objects because the table has a top and four legs. To test this empirically, my colleagues and I have recently carried out an experiment where adults were asked how many objects there were in a display which consisted of a toy duck perching on top of a toy car (Xu et al., 1996). All subjects reported that they saw two objects; none counted the eyes, the head and the wings of the duck, or the wheels, the doors, and the windows of the car. A consequence of perceiving the display as two objects was that subjects predicted that the objects would move independently of each other. Of course when asked to carry out the same task in a normal room with hundreds of items we may find that subjects vary somewhat on what exactly counts as an object, but I suspect that pages of a book and the letters written on a blackboard would not be counted as objects. In other words, *letter* and *page* are perfectly good sortals, yet when one is asked to count objects, separate movability becomes a very important criterion. What makes people count tables, chairs, and people but not the legs of tables and the tops of chairs if 'object' is not a sortal that provides criteria for individuation? When asked to 'count objects', people readily apply the word 'object' in accord with the definition of a physical object that I provided above: to constitute a countable object, an aggregate of matter has to be coherent and retain its boundaries as it moves in space and time. The sortal *physical object* (which is one sense of the English word 'object') provides principles of individuation, just like any other sortal. This narrow sense of object, as defined above, seems the default interpretation of the natural language term 'object'.

Now consider the second argument (Hirsch, 1982). If there is a sense of the word 'object' that means 'any aggregate of matter', it is true that this meaning is not a sortal. However, my intuition is that there is no sense of the word 'object' that is defined by 'any aggregate of matter'. For example, even if one is instructed to count the objects in a room using the broadest sense of 'object', nobody will count the table plus the wall, or the table plus the chair, as one single object, although according to Hirsch 'any aggregate of matter' should suffice for being an object. But even Hirsch's narrower definition of object—spatially continuous aggregates of matter—is not used in everyday life. Imagine that I put a chair on top of a table such that there is direct contact between the two. Together they constitute a spatially continuous aggregate of matter. We would still not count the table plus the chair as one object. Or suppose that I poured a jar of sand on the cover of a book, nobody would count that the book plus the pile of sand as one object. In general, being spatially continuous is far from being a sufficient condition for being an object in our conceptual system.

3.2 Does Object *Provide Criteria for Numerical Identity?*

Next consider the story of Lot's wife, which seems to pose a special challenge for all sortal theorists. Indeed, given our knowledge about the physical world, the story violates the laws of nature. However, as we do not have any trouble

conceptualizing the event and understanding the story, it would be hard to argue that the story is 'incoherent'. If it were incoherent, our conceptual system should simply break down, as in the case of 'a round square'.

The example of Lot's wife is one of many metamorphoses ubiquitous in fairy tales. Adults and children appear to have little trouble tracing the identity of a fictional character who at one moment is the prince and the next moment turns into a frog. La Palme Reyes (1994) proposed that the prince and the frog are the same *living being*, which persists through the dramatic appearance changes. She generalizes this case to all cases of metamorphoses, arguing that we sometimes invoke rather abstract sortals such as *soul* as our covering concept (which is the concept that specifies what persists through the changes). The key here is that a sortal concept is needed in order to understand these transformations that seem to express identity across kinds. In some of the cases, we need to invoke the sortal *physical object*. In the movie *Cinderella*, for example, a pumpkin is turned into a chariot. In this case, *animal* or *soul* would not be the appropriate sortal concept, yet *physical object* will do the job. Different parts of the pumpkin have been turned into parts of the chariot by some imaginary and magical mechanism, but as a whole it is still the same physical object. This line of argument can be applied to the story of Lot's wife as well since it is a case of metamorphosis. I suggest that *physical object* is the sortal that covers both the phase in which Lot's wife was a person and the phase in which Lot's wife was a pillar of salt.

Hirsch's car-crusher makes an interesting example. As the car is consigned to the crusher, the sortal *car* allows us to override spatiotemporal continuity; we decide that at some point the car goes out of existence and has been replaced by a pile of metal and plastic. However, intuitively we think something has persisted through this process because we can point to the pile of metal and plastic and say, 'That used to be a car.' The demonstrative 'that' requires a reference that persisted through the changes; *physical object* appears to be a good candidate. *Car* as a substance sortal satisfies the condition that when a car is no longer a car, it has also gone out of existence. Although no one would deny that *car* is a substance sortal and not just physical object plus some accidental properties, under these circumstances we behave as if we temporally suspend the status of *car* as a substance sortal, and think of it as a phase/stage sortal so that the life of the physical object has been divided into a car phase and a pile of metal and plastic phase.

Finally, an example from daily life also supports this view. Take 'gobots' or 'transformers'—something one can buy in a toy store which can take the form of a robot as well as, say, a space ship. A gobot is neither a robot nor a space ship. However, we judge that it is the same object when the robot is transformed into a space ship.

Interestingly, in virtually all metamorphoses, the transformations do not violate the criteria for object-hood such as spatiotemporal continuity or the constraint that one object cannot be at two places at the same time. For example, it seems incomprehensible to us that the frog could vanish on Tuesday then the prince (supposedly transformed from the frog) would appear

somewhere else on Thursday and *nothing* exist in between. It seems a conceptual necessity that continuity be preserved in order for us to understand these fictional cases; this is consistent with the sortal *object* being the underlying covering concept. Wiggins and others have no explanation for this fact.[6] After all, Lois Lane finally figured out that Superman and Clark Kent were one and the same person because they never appeared in the same room at the same time!

3.3 'What Is It?'

The last argument against 'object' being a sortal is that it does not answer the question 'what is it?'. It is true that once we know the names *table*, *computer*, *tree* or *dog*, a question such as 'what is it?' demands a specific sortal as an answer. However, when we encounter novel objects that we have never seen before, the answer 'it is an object' is acceptable. This answer is not as vacuous as it initially appears, because we would only apply this label 'an object' if we believe that this novel individual is three-dimensional, bounded, coherent and it can be quantified discretely. A puddle of some unknown liquid will not be called 'an object'!

Spelke (personal communication) pointed out to me that it may simply be pragmatically odd to answer 'an object' when asked 'what is it?'. The 'it' has already picked out some object in the array and the question 'what is it?' asks for further information. Of course the 'it' might simply pick out an individuated entity, e.g. a puddle, but it seems that 'object' is the default interpretation of 'it'. In discussing the primacy of the notion of a physical object, Kahneman et al. (1992) use the following example: as something is approaching from a distance, one might hear the sentence 'It is a bird, it is an airplane, it is Superman.' Here the reference of the pronoun 'it' does not change; it refers to the same physical object.

In sum, the arguments against 'object' being a sortal can be countered by applying the narrow sense of the word 'object', defined as *physical object*. One remaining question is addressed in the next section, namely why *physical object* is not equivalent to some basic notion of spatiotemporal continuity.

3.4 Spatiotemporal Continuity or the Sortal Object*?*

One might suggest that perhaps all we need in addition to sortals such as *car*, *duck* and *person* is some general notion of spatiotemporal continuity. Hirsch (1982) has argued, convincingly in my view, that mere spatiotemporal continuity is not sufficient. An example involving novel objects makes this clear.

Most objects that we encounter in daily life, e.g. cars, trees, rocks and dogs, are relatively familiar. Our conceptual system assigns objects to kinds; the inferences we make are based on what kind of object we are dealing with. We individuate and trace the identity of objects by taking into account

[6] I am indebted to Elizabeth Spelke for this observation.

what kind of object they are. Now imagine encountering something that we have never seen before, something that cannot be assimilated into one of the kinds we know about; can we still trace its identity over time? Consider Hirsch's (1982) example of a child who grows up on a farm and who has never seen a car before. Upon seeing a car moving in open field, can the child make any judgement about the identity of the car over time? As Hirsch argues, the following would never happen: after a while the child says, 'The object I was looking at was first rather wide and very oddly shaped, but now it is much more square.' Meanwhile he gestures with his hands to show that he was in fact piecing together the whole car and then just the middle section of the car without its front hood and back trunk even though the car has not undergone any qualitative or quantitative changes. Note that the child's hypothetical sentence preserves spatiotemporal continuity of the object, that is, the imaginary process of shrinking allows one to trace a spatiotemporally continuous path between all time slices of the car.

What prevents the child, or any adult, from making such a judgement? Hirsch (1982) argues for the existence of a basic rule, 'trace an object's career by following a spatiotemporally and qualitatively continuous path which minimizes changes as far as possible'. However, it is difficult to make precise the notion of minimizing changes because one can minimize change along different dimensions, e.g. minimizing a change of size may not be minimizing change of colour. I suggest that it is the concept of a *physical object* that guides the child's judgement in this case. The child applies the criteria for being the same *object*. That is, he picks out the object (which happens to be a car) and traces its identity over time. The imaginary shrinking process is prohibited because part of an object is not an object so long as that part does not fall off and starts to move independently on its own. In other words, object-hood requires an additional 'boundedness' condition that mere spatiotemporal continuity lacks.

In sum, several arguments show that adults utilize the sortal *physical object* in their construal of the world. Although we are not often forced to conceptualize the world in terms of physical objects instead of cars, trees, and tables, we resort to the sortal *physical object* under many somewhat out of the ordinary circumstances.

4. The Relation Between Physical Object *and Other More Specific Sortals*

Macnamara et al. (1994) proposed that the relations among kinds are best characterized by underlying maps. My proposal for the relation between *physical object* and more specific sortals will be an extension of their theory.

The relation between a phase sortal such as *passenger* and a substance sortal such as *person* is one of underlying map as opposed to class inclusion. An underlying map is simply a function that maps members of a set A to members of a set B. Note that unlike class inclusion relations, it is possible to map more than one member of A onto a single member of B. For example, a passenger from set A is mapped onto a person from set B—u:

passenger → *person*. In this process, a member of the kind *passenger*, p, is identified with u(p), a member of the kind *person*, a. Thus u(p) is the person underlying the passenger p. It is also possible to have more than one passenger mapped onto one person. If a passenger, p, is identified with a person, u(p), and another passenger, p′, is identified with person, u(p′), then if u(p) = u(p′), the two passengers would map onto one and the same person.

Similarly, the relation between *cow* and *animal* is also one of underlying map even though *cow* is not a phase/stage sortal. The mechanism is the same in the two cases (see Macnamara et al., 1994, for a detailed discussion). This proposal may be extended to the relation between a kind such as *car* and the kind *physical object*. There is an underlying map between members of the kind *car* and members of the kind *physical object*. A member of the kind *car* is identified with u(c), a member of the kind *physical object*.

Two issues should be clarified at this point. First, a car is *not* identical to a physical object. This is important because if a car were identical to a physical object, we would run into the following problem: Suppose we consign the car to a crusher, at a certain point the car goes out of existence. But as I argued in the last section, the physical object has survived the crushing process. If the car were identical to the physical object, it would be absurd to say that the car has gone out of existence but the physical object has not.[7] Second, *physical object* is not a universal kind that all other kinds are predicates of as the relation between *physical objects* and other more specific sortals is not construed as one of class inclusion, as suggested above. More importantly, *physical object* is certainly not a bare particular in the sense of an attribute-free support for attributes because *physical object* has its own criteria of individuation and identity. It is, after all, a sortal.

5. Implications for Cognitive Development

Adults generally conceptualize the world in terms of cars, people, trees and telephones and they do not see them as types of physical objects. However, in a number of situations (discussed above), adults resort to the sortal *physical object* and temporarily suspend their beliefs that *car* and *tree* are substance sortals. Young infants, on the other hand, may begin with an innate mechanism to give them individuated entities, i.e. physical objects, and acquire more specific sortals later (see Lowe, 1989, for a similar proposal from a philosophical standpoint). This way of building a baby has three advantages.

First, it gives the baby an initial conceptual state that can develop into an adult-like system because the initial state has *some* criteria for individuation and identity. The traditional view holds that the infant's world does not contain any objects persisting through time and space, hence there is no criterion for individuation and identity (Piaget, 1954; Quine, 1960). An important question to ask is: if this 'blank slate' characterization of the initial

[7] I am indebted to Robert Stalnaker for this point.

state were correct, what kind of learning mechanisms would enable the child to later develop a conceptual system that is based on adult's ontology of individuated entities? There are two proposals in the literature.

The first, suggested by Quine, is that by acquiring certain quantificational properties of a natural language the infant comes to reconstruct her ontology. However, Soja et al. (1991) have shown that children has fully distinguished individuated and non-individuated entities long before they acquire the quantificational properties reflected in the syntax of English. In particular, well before children have mastered the count/mass distinction (as determined by their use of determiners), they distinguish individuated and non-individuated entities; this conceptual distinction guides the child's inferences in acquiring the meanings of nouns.

A second proposal, entertained but rejected by Kohler (1947) and Hirsch (1982), goes roughly as follows. Suppose a child starts off with a conceptualization of the world that does not have any principles of individuation and identity but only a quality space of some sort, the child can then observe that sets of sensations which have the same texture and colour but differ in these respects from their environment tend to move together, i.e. behaving like units. The child then comes to learn that there are individuated units in the world. However, this proposal is a non-starter. If the child experiences the world in terms of fleeting pieces from moment to moment, what would ever enable the child to notice that certain sensations move together? In order to observe that certain sensations tend to behave as units and notice the fact that these sensations tend to appear and disappear together, one has to already have some criteria for individuation in order to pick out a set of sensations—a good candidate seems to be the spatiotemporal criteria. But how could one ever learn those criteria for individuation? The answer seems to be that one couldn't. They would need to be 'hard wired'. Such accounts of learning simply cannot be formulated without presupposing some criteria for individuation and identity. In other words, it seems an impossible task for a child who starts with a conceptual system void of criteria of individuation and identity to ever acquire the adult conceptual system.

The second advantage of building a baby this way is an argument from learnability. From the learner's point of view, the spatiotemporal criteria that I discussed in this paper are the most reliable, whereas the property/kind criteria are less so. The spatiotemporal criteria apply across the board to all physical objects whereas the property/kind criteria vary from kind to kind. Presumably these kind-relative criteria have to be learned through interacting with the world.[8] Therefore the baby may be better off to start with

[8] One could imagine a maturational process that would enable the child to represent more specific sortals. For instance, it is possible that the 10-month-old infants have distinct visual pathways representing the 'where' and the 'what' of objects. Integration of these two sources of information may be required for representing specific sortals and the infants' failure is due to their inability to integrate these two types of information (see Xu, 1997, for a more detailed discussion).

the spatiotemporal criteria—having initially individuated the objects, they are in a good position to gradually learn about which of their properties stay stable over time.

A third advantage is an evolutionary consideration suggested by Kellman (1993). He argues that because infants have certain limitations in attention span and motor control, it is difficult for an infant to correct a mistake in perception. So it is desirable that the initial perceptual capabilities of the child be the ones that have the highest ecological validity and that are the least misleading. Spatiotemporal information is, in general, a more valid indicator of the environment than information from static arrays. It is possible that the evolution of human perceptual system may have exploited such a relationship between spatiotemporal information and the environment. An infant who initially depends heavily on spatiotemporal information in perceiving the world may be getting the most accurate picture of the environment and may therefore survive more readily.

6. How Do Infants Acquire Sortals such as Cup *and* Dog*?*

Finally, a few words on how a child might acquire the more specific sortals such as *car* or *dog*. A maturational process may allow the infant to integrate information about object properties with information about object locations. Moreover, infants may be built to expect that there are more specific substance sortals and they may have an innate syntax–semantics mapping between count nouns and sortals. Thus count nouns of a language may provide pointers for the child as to what sortals there are (see Xu, 1997, for a more detailed discussion and suggestive evidence). The process of learning a count noun may also be the process of identifying a more specific sortal, provided that it is transparent to the child that these nouns pick out objects in the world.

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology

References

Baillargeon, R. and Graber, M. 1987: Where's the Rabbit?: 5.5-month-old Infants' Representation of the Height of a Hidden Object. *Cognitive Development*, 2, 375–92.

Baillargeon, R., Spelke, E.S. and Wasserman, S. 1985: Object Permanence in 5-month-old Infants. *Cognition*, 20, 191–208.

Bower, T.G.R. 1974/1982: *Development in Infancy*. San Francisco: W.H. Freeman.

Cohen, L.B. and Caputo, N.F. 1978: Instructing Infants to Respond to Perceptual Categories. Paper presented at the Midwestern Psychological Association Convention, Chicago, IL.

Cohen, L.B. and Younger, B.A. 1983: Perceptual Categorization in the Infant. In E.K. Scholnick (ed.), *New Trends in Conceptual Representation*. Hillside, NJ: Lawrence Erlbaum Associates, pp. 197–200.

Eimas, P. and Quinn, P. 1994: Studies on the Formation of Perceptually-based Basic-level Categories in Young Infants. *Child Development*, 65, 903–17.

Fodor, J.A. 1983: *The Modularity of Mind*. Cambridge, MA.: MIT Press.

Frege, G. 1884/1959: *The Foundations of Arithmetic*. Oxford: Blackwell.

Geach, P. 1962: *Reference and Generality*. Cornell University Press.

Gleitman, L.R. 1990: The Structural Sources of Verb Meanings. *Language Acquisition*, 1:1, 3–55.

Gratch, G. 1982: Responses to Hidden Person and Things by 5-, 9-, and 16-month-old Infants in a Visual Tracking Situation. *Developmental Psychology*, 18:2, 232–7.

Grimshaw, J. 1981: Form, Function, and the Language Acquisition Device. In Baker and McCarthy (eds), *The Logical Problem of Language Acquisition*. Cambridge, MA.: MIT Press.

Gupta, A.K. 1980: *The Logic of Common Nouns*. New Haven: Yale University Press.

Hirsch, E. 1982: *The Concept of Identity*. Oxford University Press.

Huntley-Fenner, G.N. and Carey, S. 1995: Individuation of Objects and Non-solid Substances: A Pattern of Success (Objects) and Failure (Non-solid Substances). Invited poster presented at the Biennial Meeting of the Society for Research in Child Development, Indianapolis, IN.

Jackendoff, R. 1983: *Semantics and Cognition*. Cambridge, MA.: MIT Press.

Kahneman, D., Treisman, A. and Gibbs, B. 1992: The Reviewing of Object Files: Object-specific Integration of Information. *Cognitive Psychology*, 24, 175–219.

Kellman, P. 1993: Kinematic Foundations of Infant Visual Perception. In C. Granrud (ed.), *Visual Perception and Cognition in Infancy*. Hillsdale, NJ.: Lawrence Erlbaum Associates, pp. 121–73.

Kohler, W. 1947: *Gestalt Psychology*. New York: Liveright Publishing Corporation.

La Palme Reyes, M., Macnamara, J., Reyes, G.E. and Zolfaghari, H. 1994: Proper Names and How They Are Learned. *Memory*, 1:4, 433–55.

La Palme Reyes, M. 1994: Referential Structure of Fictional Texts. In Macnamara and Reyes (eds), *The Logical Foundations of Cognition*. Oxford University Press.

Lowe, E.J. 1989a: *Kinds of Being: A Study of Individuation, Identity, and the Logic of Sortal Terms*. Oxford: Blackwell.

Lowe, E.J. 1989b: What Is A Criterion of Identity? *The Philosophical Quarterly*, 39:154, 1–21.

Macnamara, J. 1987: *A Border Dispute: The Place of Logic in Psychology*. Cambridge, MA.: MIT Press.

Macnamara, J. 1994: Logic and Cognition. In Macnamara and Reyes (eds), *The Logical Foundations of Cognition*. Oxford University Press.

Macnamara, J. and Reyes, G.E. 1994: *The Logical Foundations of Cognition*. Oxford University Press.

Macnamara, J., Reyes, G. and La Palme Reyes, M. 1994: Reference, Kinds, & Predicates. In Macnamara and Reyes (eds), *The Logical Foundations of Cognition*. Oxford University Press.

Mandler, J., Bauer, P.J. and McDonough, L. 1991: Separating the Sheep from the Goats: Differentiating Global Categories. *Cognitive Psychology*, 23, 263–98.

Meicler, M. and Gratch, G. 1980: Do 5-month-olds Show Object Conception in Piaget's Sense? *Infant Behavior and Development*, 3, 265–82.

Muller, A. and Aslin, R. 1978: Visual Tracking as an Index of the Object Concept. *Infant Behavior and Development*, 1, 309–19.
Piaget, J. 1954: *The Construction of Reality in the Child.* New York: Basic Books.
Pinker, S. 1984: *Language Learnability and Language Development*. Cambridge, MA.: Harvard University Press.
Quine, W.V. 1960: *Word and Object*. Cambridge, MA.: MIT Press.
Quinn, P.C. and Eimas, P.D. 1993: Evidence for Representations of Perceptually Similar Natural Categories by 3- and 4-month-old Infants. *Perception*, 22:4, 463–75.
Roberts, K. and Horowitz, F.D. 1986: Basic Level Categorization in Seven- and Nine-month-old Infants. *Journal of Child Language*, 13, 191–208.
Roberts, K. 1988: Retrieval of a Basic-level Category in Prelinguistic Infants. *Developmental Psychology*, 24, 21–7.
Rosch, E., Mervis, C.B., Gray, W.D., Johnson, M.D. and Boyes-Braem, P. 1976: Basic Objects in Natural Categories. *Cognitive Psychology*, 7, 573–605.
Soja, N., Carey, S. and Spelke, E. 1991: Ontological Categories Guide Inductions of Word Meaning: Object Terms and Substance Terms. *Cognition*, 38, 179–211.
Spelke, E.S. 1985: Preferential Looking Methods as Tools for the Study of Cognition in Infancy. In Gottlieb and Krasnegor (eds), *Measurement of Audition and Vision in the First Year of Postnatal Life*. Norwood, NJ: Ablex.
Spelke, E.S. 1990: Principles of Object Perception. *Cognitive Science*, 14, 29–56.
Spelke, E.S., Kestenbaum, R., Simons, D. and Wein, D. 1995: Spatiotemporal Continuity, Smoothness of Motion and Object Identity in Infancy. *British Journal of Developmental Psychology*, 13, 113–42.
Wiggins, D. 1980: *Sameness and Substance*. Oxford: Blackwell.
Wynn, K. 1992: Addition and Subtraction by Human Infants. *Nature*, 358, 749–50.
Xu, F. 1997: What Is the Mechanism Underlying Infants' Construction of Sortal Concepts Between 10 and 12 Months of Age? *Proceedings of the 21st Annual Boston University Conference on Language Development*.
Xu, F. and Carey, S. 1996: Infants' Metaphysics: The Case of Numerical Identity. *Cognitive Psychology*, 30, 111–53.
Xu, F., Carey, S. and Welch, J. 1996: Infants' Ability To Use Object Kind Information in Object Individuation. Manuscript under revision.
Xu, F., Carey, S., Raphaelidis, K. and Ginzbursky, A. 1995: Twelve-month-olds Have the Conceptual Resources to Support the Acquisition of Count Nouns. *Proceedings of the 26th Stanford Child Language Research Forum*.
Xu, F. and Quint, N. 1997: Object Individuation at 12 months: Shape and Other Properties. Poster presented at the Society for Research in Child Development, Washington, DC.