Lingua 92 (1994) 143–167. North-Holland

Does learning a language require the child to reconceptualize the world?

Susan Carey

Department of Brain and Cognitive Science, Massachusetts Institute of Technology, 77 Massachusetts Ave., E25-406, Cambridge, MA 02139, USA

Sortal concepts, lexicalized as count nouns in languages with a count/mass distinction, provide criteria for individuation and numerical identity. This paper examines Quine's and Piaget's claims that babies and young children lack the logical resources to represent sortal concepts.

Evidence is marshalled against the Quine/Piaget position, in favor of a view that even young infants represent at least one sortal concept, physical object, which provides spatiotemporal criteria for individuation and identity. Evidence is also provided that babies below 11 months of age may not represent more specific sortals such as cup, animal, bottle, or book. Rather, they may conceptualize these entities in a way closely related to Quine's hypothesis.

1. Introduction

Many students of language acquisition and cognitive development argue that the *continuity hypothesis* should be the default, to be defeated only in the face of extraordinary evidence (e.g., Pinker 1984, Macnamara 1982). The continuity hypothesis is that representational format is constant throughout development; that the child has innately the logical and conceptual resources to represent his or her world as do adults. The continuity hypothesis denies stage changes of the sort envisioned by Piaget, denies changes in the child's linguistic representations such as the putative 'semantic category/syntactic category' shift posited some years ago. According to the continuity hypothesis, language learning is a very complex mapping process; the child must learn which syntactic devices his/her language employs, and which of a universal set of semantic distinctions are expressed in the syntax of his/her language. What the child need not do, on the continuity hypothesis, is construct genuinely new representational resources.

Of course, whether the continuity hypothesis is true or not is an empirical question, and to examine it, one must entertain possibilities as to what types

SSDI 0024-3841(93)E0058-F

of discontinuities *could* possibly obtain in the course of development. If evidence for discontinuities is found, several further questions are then licensed, including: (1) by what mechanism is the change effected (e.g., maturational, learning by some other process than currently understood parametersetting or hypothesis testing methods). (2) What is the relation between the discontinuity and language learning? Is some change in representational resources required as a prerequisite to some aspect of language learning? Alternatively, does language learning play a role in causing the change?

Here I examine an important *discontinuity* proposal of Quine's, versions of which are endorsed by thinkers as diverse as the British empiricists and Piaget. Quine, Piaget, and others maintain that early representations of the world are formulated over a perceptual quality space (Quine, the empiricists) or sensori-motor representational system (Piaget). On both Quine's and Piaget's views, the baby is not capable of formulating any representations with the properties of adult concepts such as *object, dog, table.*

Quine's proposal is that the ontology that underlies language is a cultural construction. 'Our conceptual firsts are middle-sized, middle distanced objects, and our introduction to them and to everything comes midway in the cultural evolution of the race' (Quine 1960: 5). Before the child has mastered this cultural construction, the child's conceptual universe consists of representations of histories of sporadic encounters, a scattered portion of what goes on. Quine speculates as to the representations underlying the toddler's uses of the words 'water', 'red', and 'Mama'. 'His first learning of the three words is uniformly a matter of learning how much of what goes on about him counts as the mother, or as red, or as water. It is not for the child to say in the first case, 'Hello, Mama again', in the second case 'Hello, another red thing', and in the third case, 'Hello, more water'. They are all on a par: Hello, more Mama, more red, more water' (Quine 1960: 92). The child masters the notion of an object, and of particular kinds of objects, in the course of getting the hang of what Quine calls 'divided reference', and this through the process of mastering quantifiers and words like 'same'. 'The contextual learning of these various particles goes on simultaneously, we may suppose, so that they are gradually adjusted to one another and a coherent pattern of usage is evolved matching that of one's elders. This is a major step in acquiring the conceptual scheme that we all know so well. For it is on achieving this step, and only then, that there can be any general talk of objects as such' (Quine 1969: 9–10). And in another place he finishes the same idea with a bootstrapping metaphor, underlining the degree of conceptual change he thinks is occurring:

'The child scrambles up an intellectual chimney, supporting himself against each side by pressure against the others' (Quine 1960: 93). Quine also states that once the child has mastered the notion of an object, and got the trick of divided reference, he goes back and reanalyzes 'Mama', so that it is now the name of a unique enduring person.

Quine's view can be schematized as follows. Imagine a portion of bottle experience that we adults would conceptualize as a single bottle. Babies respond to bottleness or bottlehood also, and can learn many things about bottlehood; for instance, they can come to associate bottlehood with milk, or with the word 'bottle'. Now imagine a portion of bottle experience that we would conceptualize as three bottles. The infant would also expect to obtain milk (indeed, more milk) from this bottleness and could also refer to it with the word 'bottle'. Note that shape is important to the identification of bottlehood, just as the shape of the individual grains is important for distinguishing rice from spaghetti from macaroni. Similary, even if Mama is a scattered portion of what goes on, shape is important for distinguishing Mama from Rover or from Papa. That shape is important for distinguishing what scattered portion of experience constitutes bottlehood does not mean that the baby is capable of representing 'a bottle', 'two bottles', or 'the same bottle I had yesterday'. Thus, demonstrations that toddlers are sensitive to shape in inductions of word meanings when new words are ostensively defined over objects (e.g., Landau, this volume) do not bear on Quine's proposal.

In this discussion I will not make contact with Quine's radical philosophical views such as the indeterminacy of translation. I assume that we can characterize the adult's ontological commitments, that these include middle-sized physical objects, and that words such as 'table', 'dog' and 'person', function as sortals in the adult lexicon, in Wiggins' (1980) sense. Sortals refer to kinds of individuals (i.e., divide reference), providing conditions for individuation (establishing the boundaries of entities) and for numerical identity (establishing when an individuated entity is the *same one* as one experienced at some other time, or in some counterfactual world). One way of stating Quine's hypothesis, as I construe it, is that babies and toddlers represent no sortal concepts, no concepts that provide conditions of individuation and numerical identity, no concepts that divide reference.

Two reviewers of this paper raised the objection that representations of shapes presuppose representations of individuals that have those shapes, claiming therefore that Quine's proposal (at least as construed above) is incoherent. This is not so. Please dwell on the spaghetti, macaroni case. It's

true that if the contrast between the two types of stuff is based on the shape differences of individual pieces, then some representation of those individual pieces must enter into the representation of *shape*. But our concepts of spaghetti and macaroni (and the words *spaghetti, macaroni*) do not quantify over those individuals. Similarly, we can represent the shape of a scattered portion of sand, arranged, for example, into an S, and when we refer to it as 'a portion' or 'an S' we are quantifying over that individual. But when we think of it as sand, we are not. Quine's proposal is that the child's *conceptual/linguistic* system has only the capacity to represent the world in terms of concepts like *furniture, sand, bottlehood.* Of course the child's perceptual system must pick out individuals in order to represent shape, to determine what to grasp, and so on. This is part of what Quine meant when he claimed that the child is inherently 'body minded' (Quine 1974).

Piaget, like Quine, believed that that baby must construct the concept of enduring objects, although he differed from Quine as to the mechanisms he envisioned underlying this construction. Quine saw the child's mastery of the linguistic devices of noun quantification, the machinery by which natural languages such as English manage divided reference, as the process through which the child's ontology comes to match his or her elders'. Piaget held that the baby constructs the concept *object* during the course of sensori-motor development by the age of 18-months or so, and that this construction is the basis for the child's mastery of natural language. Since Piaget did not frame his discussion in terms of an analysis of the logic of sortals, it is not clear when he would attribute full sortals to the child.[1]

The Quine/Piaget conjecture about the baby's representational resources is a serious empirical claim, and as I will show, it is difficult to bring data to bear on it. In what follows, I first consider Quine's views, contrasting his hypothesis that children come to represent sortals only upon learning the linguistic devices of noun quantification with what I will call the 'Sortal First' hypothesis. The Sortal First hypothesis is that babies represent sortal concepts, that toddler lexicons include words that express sortals, and that these representations underly the capacity for learning quantifiers rather than resulting from learning them. I then turn to early infancy, and explore the contrast between the Quine/Piaget hypothesis and the Sortal First hypothesis as regards the earliest phases of word learning. A preview of my conclusions: whereas the Sortal First hypothesis is ultimately favored, evi-

[1] For example, Piaget thought that the logical prerequisites for representing the adult concepts *all* and *some* are not acquired until after age 5.

dence is presented for a decidedly Quinian discontinuity in infant conceptual development.

2. The toddler's mastery of count-mass syntax

Quine's hypothesis is that the child masters the logic of sortals through a process of adjusting the meanings of nouns and of natural language quantifiers to each other (scrambling up an intellectual chimney, the walls of which are the child's currently worked out representations of the quantifiers he/she knows). To address Quine's conjecture experimentally, we must first know when in the child's life the putative scrambling is going on. Even by age 3 the child is not producing all the quantifiers that constitute the sides of Quine's chimney. The very beginnings of the English count/mass distinction are mastered in the months leading up to age $2\frac{1}{2}$. Many children age 2:0 produce nouns with no determiners or plurals, but some have begun to produce plurals and a few determiners and quantifiers (usually possessives such as 'my', plus 'a' and 'the'). Many 2-year-olds beginning to use determiners do not distribute them differently according to the noun's count/mass status in the adult lexicon. They still omit many determiners, and use others like 'the' and 'my' that do not differentiate count nouns and mass nouns. By $2\frac{1}{2}$, virtually all children distinguish in some ways the syntactic contexts in which words like 'table' and 'dog' appear from those in which words like 'water' and 'playdoh' appear (Gordon 1985, Soja et al. 1991). Gordon (1982) showed that between $2\frac{1}{2}$ and 3 years of age the distinction becomes marked in syntax, as the child's speech abruptly comes to reflect the arbitrary rule that determiners are obligatory for singular count nouns, but not for mass nouns (that is, one can say 'I like juice', but not *'I like dog').

The developmental facts summarized above determine the relevant ages for an empirical test of Quine's speculations. Data bearing against Quine's claims could be of several types: e.g., data showing children age 2 or under take proper nouns to refer to individuals of a kind or that they take count nouns to refer to kinds of individuals. But, as already mentioned, the trick is figuring out how we can know whether toddlers' 'Mama' refers to entities they conceptualize as individuals or whether their 'bottle' divides reference, referring to each individual of a certain kind, as opposed to *bottlehood*.

Another type of evidence could be relevant. If it can be shown that upon first learning 'a' or the plural '-s', toddlers interpret them correctly, as signalling an individuated entity of a kind or a plurality of individuals of a

kind, respectively, this would tell against Quine. This is because these are the first relevant quantifiers the child learns. If he or she interprets them correctly from the beginning, the interpretation could not have been acquired through an adjustment process involving the entire set of quantificational devices of noun syntax. This last point is important. In the beginnings of language learning, on Quine's view, children will not interpret those few quantifiers in their lexicons as adults do. The scramble will have just begun. Data showing that children use 'a' and plurals will not be itself relevant to Quine's hypothesis; it must be shown that such quantificational devices are doing the same work as they do in the adult language.

3. The composition of the toddler lexicon

A large proportion of the baby's first words are words for middle-sized physical objects, such as 'bottle', 'book', 'dog', 'cup', and 'banana'. But that babies have words in their lexicons that refer to object kinds in the adult lexicon tells us nothing of what these words mean to the babies. Many have argued that the earliest words are often complexive (e.g., Bowerman 1978, Dromi 1987, Vygotsky 1962). That is, children appear to extend words to new referents on the basis of any of the salient perceptual properties of the original experiences in which the word was heard. These complexive uses often cut across what are for adults distinct ontological categories, as when 'paper' apparently refers to the act of cutting, the act of drawing, to pens and pencils and to paper (Dromi 1987). If such complexive uses reflect unconstrained (from the point of view of adult lexical categories) projection of word meanings, Quine's views receive support. But it is important to see that such complexive uses are not *necessary* for Quine's conjecture to be correct.

Indeed, others deny that toddlers construct complexive meanings; Huttenlocher and Smiley, 1987, for example, present evidence that from the beginning babies use each word for middle-sized objects appropriately: 'bottle' to refer to bottles, 'book' to books, and so on. But even if Huttenlocher and Smiley are right, this fact does not disconfirm Quine's conjecture. In fact, Quine presupposes that the baby uses the words in contexts adults would. His point is that, even so, the baby might not be individuating the words' referents as we do. The baby could refer only to what we conceptualize as bottles when she uses 'bottle', but could be referring to bottlehoods. She could be using the word to refer to a scattered portion of what goes on, determined by perceptual similarity to the portions of her experience when adults use 'bottle'.

4. Toddler sensitivity to noun syntax

Children as young as 17 months (at least girls that young) are sensitive to the syntactic context in which a new noun is heard in their projection of noun meaning (Katz et al. 1974, Macnamara 1982). Specifically, if ostensively taught a new word in a count noun context, referring to an unfamiliar doll ('See this. This is a dax. Can you touch the dax? Can you put the dax on your head ...'), they assume that other dolls of the same type are also daxes. But if taught in a proper noun context ('See this. This is Dax. Can you touch Dax. Can you put Dax on your head ...'), they assume that other dolls of the same type are *not* Dax, reserving 'Dax' for the original doll only.

Do these data establish that young children distinguish kinds from individuals, and use count nouns to achieve divided reference? Certainly not. They *do* establish that toddlers are sensitive to the syntactic distinction between nouns following determiners and those not following determiners, but this distinction could be signalling a different semantic distinction than that between individuals and kinds. For a sample Quinian interpretation: babies could take nouns without determiners such as 'Dax', 'Rover', and 'Joan', to refer to portions of experience defined by a stricter similarity metric than that referred to by nouns with determiners. Suppose a Quinian baby, Alice, has a brother whom she hears called both 'Rupert' and 'a boy'. Suppose also that she relies on shape to determine Rupertness and boyness. She could have learned from others' usage of the words that to be called 'Rupert', a given portion of experience must be very similar in shape to the original portions of experience to which the term was heard to refer, whereas to be called 'a boy', 'the boy', something need look only somewhat like the original referent. A generalization of this pattern of distinction, across 'Alice' and 'a baby', 'Rover', and 'a dog', and so on, could underly the patterns of projection found by Katz et al., 1974, and subsequent replications.

This interpretation of the Katz et al. data attributes to the baby a different meaning for 'a' from the adult as well as different meanings for 'bottle', 'boy', 'Rupert'. This is, of course, Quine's position. On his view, it is only in the course of learning other quantifiers, plural markers, and so on, and adjusting to all the contrasts in usage they mark (the process of scrambling up the intellectual chimney cited above) that the baby works out the meaning of 'a', 'the', 'another', 'some', 'more', 'all', 'many', 'same',[2] etc.

[2] Bloom, personal communication, provides the following argument against the Quinian interpretation of the Katz et al. data. Among children's very first words are some pronouns (e.g.,

5. Words for novel objects and words for non-solid substances

In several studies, my colleagues and I have attempted to address Quine's proposal by comparing children's representations of solid physical objects, such as cups, with their representations of non-solid substances, such as sand or gels or creams. Our idea is that since adults conceptualize the former as kinds of individuals (i.e., in terms of sortals that divide reference), but do not conceptualize the latter in this way, we might be able to find evidence that infants and toddlers respect the quantificational distinction between the two as well.

In the first studies, Soja et al. (1991) compared 2-year-olds' projection of newly learned words ostensively defined by reference either to novel solid physical objects (e.g., a brass plumbing T) or novel non-solid substances (e.g., a hair-setting gel with grapenuts embedded in it). The objects were made of unfamiliar materials and the non-solid substances were presented formed into distinctive novel shapes. The child was introduced to the novel entity and provided a word for it (e.g. 'blicket' for a novel object; 'stad' for a novel non-solid substance). The child was then presented two new sets of stimuli and asked to give the experimenter the blicket or the stad. For each object trial, the choices consisted of another object of the same shape made of a different material (e.g., a plastic plumbing T) or three small pieces of the original material (brass). For each substance trial, the choices consisted of a new substance formed into the original shape, or three small pieces of the original substance. Figure 1 shows the design for one trial of each type. There were four object trials and four non-solid substance trials. Of course, which words were assigned to which entities varied across subjects, but for expository clarity I will use 'blicket' as my sample object name and 'stad' as my sample non-solid substance name.

'he', 'it') and these are treated from the onset as belonging to the same category as 'Rupert' and not as 'dog'. Children do not use them with determiners or modifiers (Bloom 1990). The adult-like analysis works well in accounting for this finding; pronouns, like proper names, denote individuals. But the Quinian analysis fails, since a far broader range of referents are called 'it' or 'he' than are called 'dog'. I agree with this argument, as I reject the Quinian proposal in favor of the Sortal First hypothesis. However, Quine could reply that the child has two sets of words: words like 'dog' and words like 'Rupert', and a few singularities like 'he', each tagged with some of the syntactic contexts in which they appear and each tagged with prototypical referents and a similarity metric that determines usage. Coming to recognize that the syntactic contexts in which 'he' appears are the same as those for 'Rupert' might be part of the scrambling process.

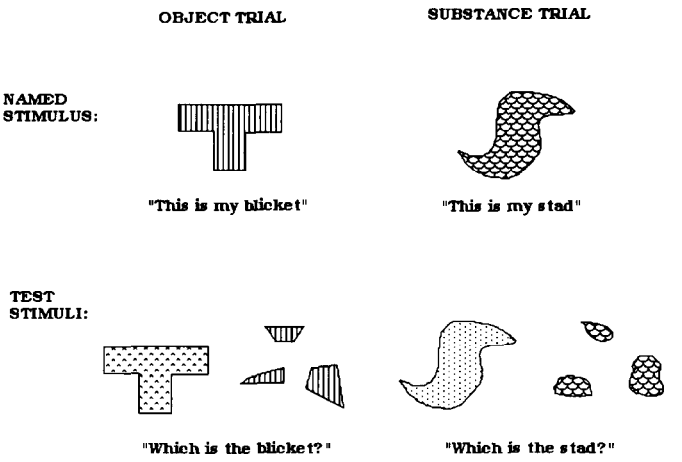

Fig. 1. Object trials and substance trials in Soja et al. (1991).

Soja et al. carried out two analyses to assess whether children's representations of the referents of the words were influenced by the status of their knowledge of count/mass syntax. First, they collected production data and assigned each child a value corresponding to the degree to which count nouns and mass nouns appeared in selective syntactic frames (e.g., 'a NOUN', 'NOUNs' 'too much NOUN'). Scores ran from 0 to near 1.0. Second, they introduced the new words in two different ways. In the neutral syntax condition, no syntactic information as to the count/mass status of the word was provided; the words were introduced as 'my blicket, my stad' and subsequently appeared in the context 'the blicket, the stad'. In the informative syntax condition, the words were introduced as 'a blicket, some stad', and further differentiated syntactically, e.g. 'another blicket, some more stad'.

As figure 2 shows, children at age 2:0 and 2:6 used different bases for their projection or words for the two different types of entities. They projected 'blicket' to the other whole object the same shape as the original referent and they projected 'stad' to the scattered portion of substance the same texture and color as the original referent. For object trials, children were sensitive to matches in shape and number; for non-solid substance trials, children ignored matches in shape and number. Performance was more adult-like on the object trials, but performance on both types of trials was better than chance at both ages. Also apparent on figure 2, the syntactic context made no difference. The children were no more likely to interpret 'blicket' as the word for a kind of individual when it was heard in a count noun context. Similarly, hearing

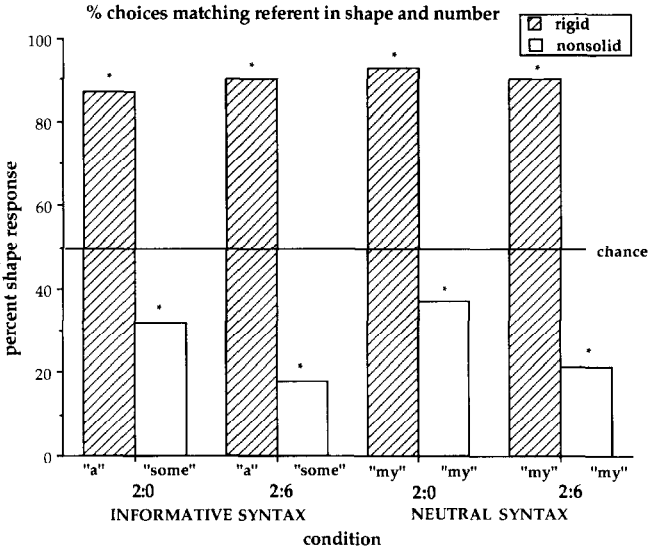

Fig. 2. % trials in which test stimulus chosen matched original referent in shape and number (Soja et al. 1991).

'stad' in a mass noun context made them no more likely to conceptualize stad as a substance that can appear either in scattered or singly bounded portions. Further, the child's productive control of count/mass syntax did not influence the pattern of projection: children with differentiation scores of 0 showed the same pattern as those with differentiation scores close to 1.

We can conclude from these results that an entity's status as a solid physical object (or not) influences which of its properties are salient in determining what other entities are referred to by the same word. We can also conclude that this distinction between objects and non-solid substances predates mastery of count/mass syntax. These data are consistent with the Sortal First hypothesis, for they are consistent with the child's taking 'blicket' to refer to each individual whole object of a kind, and 'stad' to refer to a kind of substance, conceptualized as a non-individuated entity. But the data are also consistent with the following more Quinian interpretation of the child's representations of the blicket and the stad.

Babies, being 'body-minded' (Quine 1974) could be sensitive to the perceptual experiences that determine objecthood: boundedness, rigidity, coherence through motion. Whenever these are detected, they could heavily weight such features as shape in their representation of these experiences. Shape would

thus be a salient feature of the blicket, but not of the stad, for non-solid substances do not maintain their shapes when manipulated. For non-solid substances, properties such as texture and color might be salient, for these stay constant over experiences with substances. In other words, the two-year-old could be using 'blicket' to refer to blicketness, and recognize blicketness by shape. The differential patterns of projection do not establish that the toddler is using 'blicket' to refer to any individual whole object of a certain kind, that the toddler divides the reference of 'blicket'.

One detail of the data from figure 2 favors the Sortal First over the Quinian interpretation, and that is that toddlers performed more like adults on the object trials than on the substance trials. Quine's interpretation of this would have to be ad hoc, perhaps that the baby has had more object experience than substance experience. But the Sortal First hypothesis *predicts* this asymmetry. To see this, suppose the Sortal First hypothesis is true, and suppose that upon first hearing the word 'blicket' the child assumes that it refers to each individual object of a certain kind. The choices for testing how the child projects 'blicket' included another single object, and 3 small objects. Even if the child isn't exactly sure of which features of the blicket establish its kind, the child can rule out that the 3 small objects are a blicket, for under no interpretation can they be an individual object of the same kind as the original referent. Children should then be at ceiling on the object trials, which they are. The substance trials are another story. If upon first hearing 'stad', the child takes it to refer to the kind of substance of the original referent, then scattered portions have no different status from unitary portions. There is no clue from number of piles which of the choices on the test trials is the stad. If children are not certain what properties of the original sample of stad determine the kind *stad*, they might do worse on the stad trials. And indeed, they do.

The key issue here is the role of number in determining how to project 'blicket'. If the Quinian interpretation of the data is correct, the baby should project 'blicket' on the basis of shape similarity, no matter whether the choice that does not match in shape consists of one object or three objects. That is, the baby should succeed on an object trial as on figure 3 as well as on an object trial as in figure 1. The Sortal First interpretation predicts that performance on the object trials will fall to the level of performance on the substance trials if the cue from number is removed (figure 3). In an object trial such as that on figure 3, 'blicket' is ostensively defined as before, but the choices for projection are changed: another blicket of a different material (as before) and another whole object of a different kind made of the same

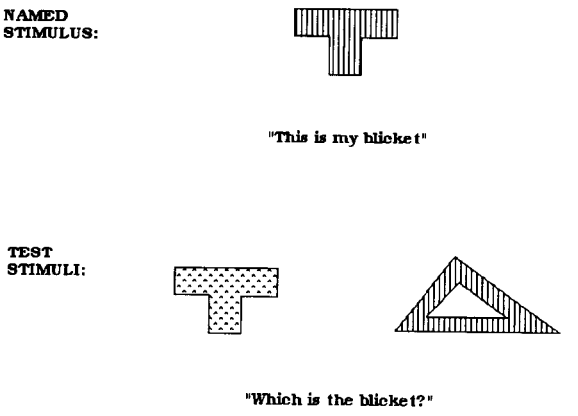

Fig. 3. Object trial in Soja (1987).

material as the original referent (instead of the three small objects). Now the child has no clues from number of objects as to which is the correct choice. Performance should fall to the level of the substance trials, and indeed, this is what happens (Soja 1987).

Apparently, the child uses the information provided by number on the object trials, but not on the substance trials. We take this as evidence that the child conceptualizes some entities as individuals (such as kinds of objects) and conceptualizes other entities as non-individuated (such as kinds of substances). These distinct ways of conceptualizing objects and substances predates mastery of count/mass syntax. Toddlers do not merely project 'blicketness' on the basis of shape of individual pieces of blicketness, as we determine whether some pasta is spaghetti on the basis of the shape of individual pieces. Instead, the pattern of projection suggests toddlers divide reference of 'blicket', and take it to refer to any individual of a certain kind.

6. Toddlers' understanding of 'a', 'some NOUN_'

I take the data reviewed in the previous section to show that by age 2:0 children take 'blicket' to refer to individual objects of a certain kind and 'stad' to refer to non-solid substances of a kind, and that the toddlers' representations of blickets and stads have the same quantificational structure as would adults'. 'Blicket' is a sortal term. These data disconfirm Quine only on the assumption that the baby did not acquire these representations from

learning English noun quantifiers. This assumption seems warranted, given that as a whole toddlers at 2:0 do not produce quantifiers, and given that the pattern of projection was independent of whether the individual subjects produced any noun quantifiers selective for count nouns. A worry, though, is that babies may have better comprehension than production of the quantifiers.

We attempted to address that possibility by manipulating the syntactic context in which the word appeared. As mentioned above, the syntactic environment in which the new word appeared had no effect in Soja et al.'s experiments, even at age $2\frac{1}{2}$ when many children did produce quantifiers differentially for what are count and mass nouns in the adult lexicon. The Quinian interpretation of this fact is that quantifiers like 'a', 'another', 'some NOUN_', 'some more NOUN_' do not yet signal the distinction between individuated and nonindividuated entities, just as the child is not projecting 'blicket' and 'stad' on the basis of that distinction. The Sortal First interpretation: objects are naturally construed as individuals of a kind and non-solid substances are naturally construed as non-individuated entities, even by toddlers, as shown by performance in the neutral syntax condition. Informative syntax merely reinforces the child's natural construal of the two types of entities.

A study by Soja (1992) decided between these two interpretations, and also established that our production data did not underestimate toddlers' interpretation of the quantifiers. Soja taught toddlers words for the objects and substances in a new condition: contrastive syntax. 'Blicket' was introduced in a mass noun context; 'stad' in count noun context. That is, when shown a novel solid object, the child was told, 'Here's some blicket ... Would you like to see some more blicket?' And when shown a non-solid substance fashioned into a distinctive shape, the child was told, 'Here's a stad ... Would you like to see another stad?' As can be seen from figure 4, at both ages 2 and $2\frac{1}{2}$, the pattern of projection was markedly different from that seen in the neutral and informative syntax conditions (figure 2). At both ages, the syntactic context 'some NOUN_', 'some more NOUN_' made children slightly less likely to construe 'blicket' as referring to an individual whole object of a kind. There was a slight tendency towards interpreting it to mean something like *brass*. The syntactic context 'a stad' made children significantly less likely to construe the non-solid substance as a non-idividuated entity. Rather, they interpreted the word as meaning something like *s-shaped pile*.

Wait, you might say, doesn't this show that children at these ages *do* know the force of 'a', 'another', and so might have learned to represent sortal

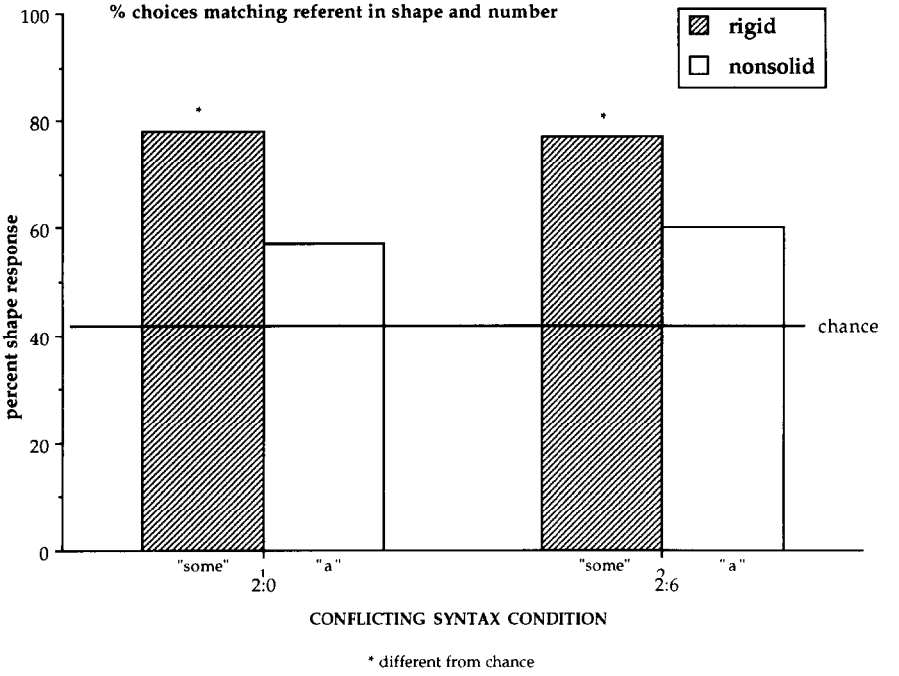

Fig. 4. % trials in which test stimulus chosen matched original referent in shape and number in Soje (1992).

concepts in conjunction with bootstrapping the meaning of the quantifiers? No, because of one further aspect of the data. At the younger age, sensitivity to conflicting syntax was shown solely by those children who had differentiated count and mass nouns in their production. Those whose differentiation scores were low performed just as did toddlers in the informative and neutral conditions, projecting 'blicket' to the other object of the same kind as the original referent and 'stad' to the other substance of the same kind as the original referent. This shows that the interpretation of 'blicket' as a sortal predates learning the meaning of 'a', 'another', and presumably underlies the latter achievement, as predicted by the Sortal First hypothesis.

These data tap the very moment children first learn the meaning of 'a'. They have only begun the scramble up Quine's chimney, and have not had time to adjust their interpretation of 'a' to many other quantifiers. Yet, 'a' signals an individuated entity of some kind. Together these data provide converging support for the Sortal First hypothesis. The child naturally

construes physical objects as individuals in distinct kinds, and naturally construes non-solid substances in terms of kind of non-individuated entities. These natural construals support adult-like projection of word meaning (figure 2), and support adult-like interpretation of newly learned quantifiers like 'a', 'some' and plurals.

7. Younger infants

Altogether the data support the Sortal First hypothesis over Quine's conjecture, but they do not establish when the child first begins to represent sortal concepts. As noted earlier, it is not clear when Piaget would attribute sortal concepts to children, but it is certain that he would deny them to young infants. The argument I have developed so far does not bear on Piaget's claims about the representational capacities of infants, as it concerns children age 24 months and older. Of course, a demonstration that young infants represent sortal concepts would defeat Quine's conjecture as well as Piaget's characterization of the infants' conceptual resources.

Studies by Cohen and his colleagues (e.g., Cohen and Younger 1983) show that quite young babies will habituate when shown, for example, a series of distinct stuffed dogs, and that they generalize habituation to a new stuffed dog and will dishabituate when shown a stuffed elephant. Similarly, when shown a series of distinct stuffed animals, babies of 8 or 9 months habituate, generalize habituation to a new stuffed animal, but dishabituate to a toy truck. Do these data not show that babies of that age represent concepts such as 'dog' and 'animal?'

Certainly not. Babies may be sensitive to dog shapes or animal shapes; babies may be habituating to doghood or animalhood. To credit the baby with sortals such as 'dog', or 'animal', we must show that such concepts provide the baby with criteria for individuation and identity.

My discussion of this question has two steps. First, I argue that babies represent at least one sortal, *object*. Second, I present some recent data from my lab that suggest that as late as 10 months of age, the baby may have no more specific sortal concepts – not *cup*, *bottle*, *truck*, *dog*, *animal* .... Thus, a Quinian interpretation of the above habituation data may well be correct.

8. Principles of individuation: Younger infants

Piaget's characterization of infants' cognitive capacities was based on tasks in which the baby must solve some problem, often involving means–end analysis, and often involving planning some action. For example, Piaget's conclusions that babies do not represent objects as continuing to exist when out of view were based on the robust finding that babies under 8 or 9 months cannot remove a cover to get a hidden object. The babies' failure might be due to their failure to realize the object still exists, as Piaget thought, or equally might be due to their inability to carry out one action (remove a cover) to achieve some other goal (obtain the object). What is needed is some reflection of the baby's conceptualization of the world that relies on behaviors well within the repertoires even of neonates. Over the past 15 years or so, such a method has been developed and is now very widely used. It relies on babies' ability to control what they attend to.

The basic idea is simple. Under most circumstances babies will look longer at what is unfamiliar or unexpected compared to what is familiar or expected. Researchers use this fact to diagnose how the baby represents some situation, especially what the baby considers surprising given his or her current state of physical knowledge. The selective looking paradigm has been used extensively to probe babies' representations of objects, and the data from a subset of these studies can be recruited to bear on the question at hand. They establish that by four months of age the baby represents at least one sortal concept – the concept of a physical object. The baby has criteria for individuation and for numerical identity of objects.

Spelke and her colleagues have shown that babies establish representations of objects on the basis of criteria which individuate them – an object is a coherent, bounded, entity that maintains its coherence and boundaries as it moves through space (see Spelke, 1990, for a review). The baby predicts the motion of objects according to principles such as that one object cannot pass through the space occupied by another (Spelke et al. 1992, Baillargeon 1990). Most relevant to the present discussion are studies showing that babies *count* objects.

These are of two types. In the first, babies are simply presented with arrays containing a fixed number of objects, say 2 of them, one after another. For example, two cups, followed by two shoes, two bottles, two hats, two pens, and so on. The pairs of objects are never repeated, so the arrays have nothing in common but twoness. The baby's looking is monitored, and after a while, the baby's attention to each new array decreases, relative to his or her

original looking time. The baby is getting bored. After looking time has decreased to $\frac{1}{2}$ its original level, the baby is presented with an array containing one object, or three objects. In both cases, looking time recovers to its original level. The baby notices the difference between two objects, on the one hand, and a single object or three objects, on the other. This result, or one very like it, has been obtained with neonates (Antell and Keating 1983).

In fact, the baby's capacity to detect similarity in number across distinct arrays serves a methodological wedge into the problem of how babies individuate objects. The baby can be habituated as described above, to two objects, and then presented with an array as in figure 5, consisting of two distinct objects sharing a common boundary. Babies dishabituate to this array, showing that they perceive it as one object, rather than two. These data support the conclusion, derived from other types of data as well, that babies are not sensitive to shape or texture regularity in individuating objects; they need positive evidence of distinct boundaries, such as one object moving with respect to the other, or the objects' being separated in space.

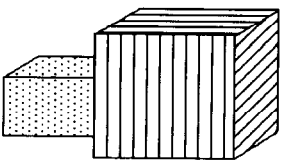

Fig. 5. Test stimulus of two adjacent blocks of different size, texture and color.

A second source of evidence that babies count objects derives from data showing that babies can add and subtract. Wynn (1992) showed four-month-olds events in which a second object was added to an array already containing one object. An object was placed on an empty stage while the baby watched and then a screen was raised that covered the object. A hand carrying a second object was shown going behind the screen and returning empty. The screen was then lowered, revealing either one object (unexpected outcome, even though that was what the baby had last seen) or two objects (expected outcome, if the baby knows $1+1=2$). Babies looked longer at the unexpected outcome.

A further experiment showed that babies expected exactly two objects, rather than simply more than one object. In this study, the expected outcome was two objects, as before, but the unexpected outcome was three objects. Again, babies were bored at seeing two there, and looked longer at the

unexpected outcome of three objects. Experiments of the same sort demonstrated that babies expected 3 − 1 to be 2, and 2 − 1 to be 1.

Whereas these studies were performed to explore the baby's concept of number, they bear on our question as well. Babies, like anybody, cannot count unless they have criteria that establish individuals to count. Babies clearly have criteria that establish small physical objects as countable individuals.

9. Principles of numerical identity: Younger infants

That babies individuate and count objects does not show that they trace identity of objects through time, that they have the representational capacity to distinguish one object seen on different occasions from two numerically distinct but physically similar objects. However, there are now two demonstrations of this capacity in infants age 4 months or younger. Spelke (1988) showed babies objects moving behind and reemerging from two separated screens, screen A to the left of screen B (figure 6). An object emerged to the left of screen A and returned behind it, and then an object emerged to the right of screen B and returned behind it. At any given time, at most one object was visible, and no object ever appeared in the space between screens A and B. Under these conditions, 4-month-olds inferred there must be two objects, as shown by the fact that when the screens were removed, revealing two objects (expected outcome), they looked less than when the screens were removed revealing one object (unexpected outcome). Baillargeon (1990) showed infants two objects at once, one on either side of a screen. The babies then used the existence of two numerically distinct objects to make sense of what would be an impossible event if only one object were involved. Together these studies show that babies use two spatiotemporal principles to individuate and trace identity of objects: one object cannot be in two places at the same time, and one object cannot go from one place to another without tracing a spatiotemporally continuous path.

In sum, infants have a concept *physical object* that functions as a sortal; they have at least one concept that divides reference, that provides criteria for individuation and numerical identity. These criteria are spatiotemporal.

10. A major conceptual difference between young infants and adults

Adults also look longer at the unexpected events in the experiments described above. Further, they ask how the magic tricks are done. This is

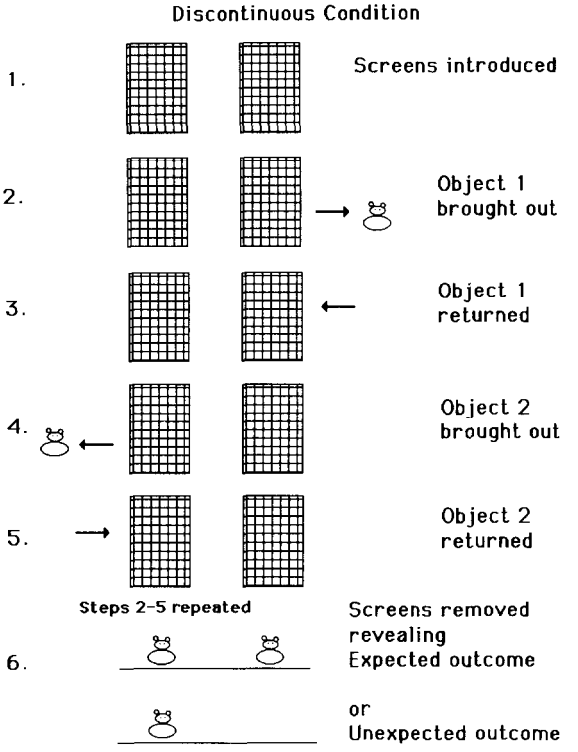

Fig. 6. Numerical identity trials, after Spelke (1988).

because adults use spatiotemporal information in just the same way as do the infants. But adults use other types of information in establishing individuals and tracing their identity through time: property information and membership in kinds more specific than *physical object*. We use property information – if we see a large red cup on a window still, and later a small green cup there, we infer that two numerically distinct cups are involved, even though we have no spatiotemporal evidence to that effect. And, as philosophers point out, our identity judgements are relative to sortals more specific than *object* (Wiggins 1980, Hirsch 1982, Macnamara 1986). Imagine a junk car, consigned to the crusher. The process of crushing is a spatiotemporally continuous, gradual process. Any changes in the car's properties are also continuous; it changes shape continuously, for example. Yet we say that at a certain point the car goes out of existence, and is replaced by a lump of metal and plastic. We trace identity relative to kinds more specific than *object*, kinds

such as *car, person, table.* Such concepts (sortals), typically lexicalized as count nouns in languages that have a count/mass distinction, provide additional criteria for individuation and identity to the spatiotemporal criteria that apply to bounded physical objects in general, and to the general assumption that an object's properties stay stable over time, or change continuously. When a person, Joe Shmoe, dies, Joe ceases to exist, even though Joe's body still exists. The sortal *person* provides the criteria for identity of the entity referred to by the name 'Joe Shmoe'.

In collaboration with Fei Xu, I have been exploring the question of whether babies represent any sortals more specific than *object*, or whether babies can use property/kind information to individuate and trace identity of objects (Xu and Carey 1993).

Consider the events depicted in figure 7. An adult witnessing a truck emerge from behind and then reenter a screen and then witnessing an elephant emerge from behind and then reenter the screen would infer that there are at least two objects behind the screen: a truck and an elephant. The adult would make this inference in the absence of any spatiotempotal evidence for two distinct objects, not having seen two at once nor any suggestion of a discontinuous path through space and time. Adults trace identity relative to sortals such as 'truck' and 'elephant' and know that trucks do not turn into elephants.

Xu and Carey (1993) have carried out four experiments based on this design. Ten-month-old babies were shown screens from which two objects of different kinds (e.g., a cup and a toy elephant, a ball and a truck) emerged from opposite sides, one at a time. Each object was shown a total of four times. After this familiarization, the screen was removed, revealing either two objects (expected outcome) or one object (unexpected outcome). In all four studies, babies looked longer at the *expected* outcome. They could not use the difference between a cup and an elephant to infer that there must be two objects behind the screen.

Another group of 10-month-olds was run in a parallel version of this study based on Spelke's design (figure 6). That is, babies were shown two identical objects emerging from the two screens a total of four times each, and the timing of the events was the same in the one screen/two kinds studies. Babies succeeded, looking longer at the unexpected outcome of one object. Apparently, babies can use spatiotemporal information to individuate objects before they can use kind information.

We have ruled out several uninteresting interpretations of the failure in the property/kind conditions of these studies. For example, it is not that babies

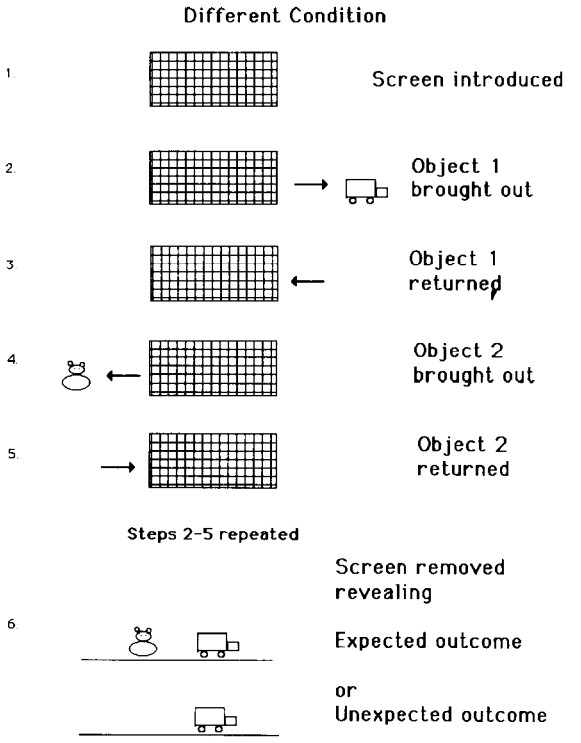

Fig. 7. Numerical identity trials (Xu and Carey 1992).

do not notice the difference between the two objects. In one version of the study, babies were allowed to handle each object (one at a time of course, for we didn't want to provide spatial information that there were two) before beginning the events. This made no difference to the results. In another, we compared looking time to the familiarization events when the objects are of different kinds (e.g., a cup and an elephant) to looking times during familiarization in a condition where the objects emerging from each side of the screen are of the same kind (e.g., two elephants). Babies habituated much faster in the latter condition. That is, they noticed that the elephant and the cup are different from each other. After habituation, we removed the screen, revealing either one object or two objects. Babies in both conditions (cup/elephant; elephant/elephant) looked longer at the outcomes of two objects (unexpected in the elephant/elephant condition; expected in the elephant/cup condition). The preference for two objects was equal in the two conditions. Thus, although babies notice the difference between the elephant and the cup, they

simply do not use this information to drive the inference that there must be two numerically distinct objects behind the screen.

In appears, then, that in one sense Quine was right. Very young infants have not yet constructed concepts that serve as adult-like meanings of works like 'bottle', 'ball', and 'dog'. How *are* the babies representing these events? We can think of two possibilities. First, the babies may actually establish a representation of a single individual object (OBJECTi) moving back and forth behind the screen, attributing to this object the properties of being yellow and duck-shaped at some times and white and spherical at other times. The basis for such a representation could be spatiotemporal: the infants may take the oscillating motion as a single, continuous, path.

A second possibility is that the baby is making no commitment at all concerning whether the objects emerging to the left and right of the screen are the same or different. That is, the baby is representing the event as OBJECT emerging from the left of the screen, followed by OBJECT emerging from the right of the screen, and represents these neither as a single object (OBJECTi) nor as distinct objects (OBJECTi, OBJECTj). Suppose you see a leaf on the sidewalk as you walk to class, and you see a leaf on roughly the same place on the sidewalk as you return from class. That may be the same leaf or it may not; your conceptual system is capable of drawing that distinction, but you leave the question open. If the infant is leaving the issue open in this case, then why does he/she appear surprised when the screens are removed and two objects are revealed? On this hypothesis, the longer looking time at two objects is a familiarity effect; the infant has been familiarized with instances of single objects, and thus seeing two objects is different. After all, babies can be habituated to 'oneness' by being shown a series of objects, one at a time. Even if you were not sure whether that leaf was the same as the one you had seen earlier, if you returned to the classroom later in the day and encountered two leaves on the sidewalk, you would see this state of affairs as different from ones in which you encountered cases of single leaves on the sidewalk.

We do not know which possibility is correct. The baby actually may be representing the events as if a duck-shaped object is turning into a ball-shaped object (possibility one) or simply may be failing to establish representations of two distinct objects (possibility two). The take-home message is the same whichever possibility is correct; 10-month-old infants do not use the property/kind differences between a red metal truck and a gray rubber elephant to infer that there must be two numerically distinct objects involved in the event.

At 11 months, about half of the babies we test succeed at our task. When babies do succeed, are they doing so on the basis of kind information or property information? That is, are they representing the events as do adults, as involving a duck and a ball, or are they individuating the objects on the basis of property differences? Further experiments could bear on this question. For example, habituation studies show babies to be sensitive to color changes and size changes, but color and size are not the types of properties that signal kind differences, at least in the adult conceptual system. Would babies of the age succeeding at this task be as likely to infer two objects when shown a blue and red cup, or a big and small cup, emerging from either side of the screen, as when shown a blue cup and a blue elephant of equal sizes emerging from either side of the screen? A difference in success rate favoring the latter pair would be suggestive that babies, just like adults, come to represent kinds of objects, and individuate objects relative to kinds. These experiments together would provide information about the developmental course of this representational capacity.

It is significant that babies begin to comprehend and produce object names at about 10 to 12 months of age, the age at which they begin to use the differences between cups and elephants to individuate objects. Again, this pattern of results is consistent with the Sortal First hypothesis. That is, babies do not seem to learn words for bottlehood; they begin to learn words such as 'bottle' just when they show evidence for sortal concepts such as *bottle* which provide conditions for individuation and numerical identity. Current studies in our lab are exploring the relations between specific words understood and success at individuation based on the kinds expressed by those words.

It is not surprising that babies use spatiotemporal information before kind information to individuate and trace the identity of objects. All physical objects trace spatiotemporally continuous paths; no physical object can be in two places at the same time. However, what property changes are possible in a persisting object depends upon the kind. An apparent change of relative location of the handle to the body of a ceramic cup signifies a different cup; an apparent change of relative location of a hand to the body of a person does not signify a different person.

In sum, these data suggest that babies have at least one sortal concept innately – *physical object*. Their object concept provides spatiotemporal conditions for individuation and numerical identity. They can use spatiotemporal information to identify individuals in their environment, and can then learn more specific sortals for kinds of these objects. Exactly how this is

accomplished is the big question, of course. The present data suggests that they spend most of their first year of life on this accomplishment.

11. A few concluding remarks

Where does this leave us vis-à-vis the continuity assumption? The major discontinuity posited by Quine and Piaget does not receive support; there is no reason to think that babies lack the logical resources to represent sortals, and indeed, *object* functions as a sortal at least from 4 months on. But if the interpretation of the Xu and Carey data suggested above is correct, then an important Quinian discontinuity is supported. Babies may be setting up a representation of an object which sometimes is round, white, and styrofoam and at other times red, metal, and truck-shaped. This is a representational system very different from yours and mine.

My story is not complete. I do not know if Xu and I are interpreting our data correctly. Also, we have as yet no account of the mechanism by which babies might begin to acquire specific sortals at around 11 months. But as I am convinced important conceptual changes occur later in life (cf. Carey 1991, Carey and Spelke, in press), I would not be shocked to find interesting discontinuities in the conceptual histories of infants, even in arenas so closely implicated in language as the conceptual underpinnings of count nouns.

References

Antell, S. and D.P. Keating, 1983. Perception of numerical invariance in neonates. Child Development 54, 695–701.

Baillargeon, R., 1990. Young infants' physical knowledge. Paper presented at the American Psychological Association Convention, Boston.

Bloom, P., 1990. Syntactic distinctions. Child Language 17, 343–355.

Bowerman, M., 1978. The acquisition of word meaning: An investigation into some current conflicts. In: N. Waterson, C. Snow (eds.), Development of communication. New York: Wiley.

Carey, S., 1991. Knowledge acquisition: Enrichment or conceptual change? In: S. Carey, R. Gelman (eds.), The epigenesis of mind: Essays in biology and cognition. Hillsdale, NJ: Erlbaum.

Carey, S. and E. Spelke, in press. Domain specific knowledge and conceptual change. In: L. Hirschfeld, S. Gelman (eds.), Cultural knowledge and domain specificity. Cambridge, UK: Cambridge University Press.

Cohen, L.B. and B.A. Younger, 1983. Perceptual categorization in the infant. In: E.K. Scholnick (ed.), New trends in conceptual representation, 197–200. Hillsdale, NJ: Erlbaum.

Dromi, E., 1987. Early lexical development. London: Cambridge University Press.
Gordon, P., 1982. The acquisition of syntactic categories: The case of the count/mass distinction. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
Gordon, P., 1985. Evaluating the semantic categories hypothesis: The case of the count/mass distinction. Cognition 20, 209–242.
Hirsch, E., 1982. The concept of identity. New Haven, CT: Yale University Press.
Huttenlocher, J. and P. Smiley, 1987. Early word meanings: The case of object names. Cognitive Psychology 19, 63–89.
Katz, N., E. Baker and J. Macnamara, 1974. What's in a name? A study of how children learn common and proper names. Child Development 45, 469–473.
Macnamara, J., 1982. Names for things: A study of human learning. Cambridge, MA: MIT Press.
Macnamara, J., 1986. A border dispute. Cambridge, MA: MIT Press.
Pinker, S., 1984. Language learnability and language development. Cambridge, MA: Harvard University Press.
Quine, W.V.O., 1960. Word and object. Cambridge, MA: MIT Press.
Quine, W.V.O., 1969. Ontological relativity and other essays. New York: Columbia University Press.
Quine, W.V.O., 1974. The roots of reference. New York: Columbia University Press.
Soja, N.N., 1987. Ontological constraints on 2-year-olds' induction of world meanings. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
Soja, N.N., 1992. Inferences about the meanings of nouns: The relationship between perception and syntax. Cognitive Development 29–45.
Soja, N.N., S. Carey and E.S. Spelke, 1991. Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. Cognition 38, 179–211.
Spelke, E.S., 1988. The origins of physical knowledge. In: L. Weiskranz (ed.), Thought without knowledge, 168–184. Oxford, UK: Oxford University Press.
Spelke, E.S., 1990. Principles of object perception. Cognitive Science 14, 29–56.
Spelke, E.S., K. Breinlinger, J. Macomber and K. Jacobson, 1992. Origins of knowledge. Psychological Review 99, 605–632.
Vygotsky, L.S., 1962. Thought and language. Cambridge, MA: MIT Press.
Wiggins, D., 1980. Sameness and substance. Cambridge, MA: Harvard University Press.
Wynn, K., 1992. Addition and subtraction by human infants. Nature 358, 749–750.
Xu, F. and S. Carey, 1993. Infant metaphysics: The case of numerical identity. MIT Center for Cognitive Science Occasional Paper.