12

The Essence of Artifacts: Developing the Design Stance

Deborah Kelemen and Susan Carey

1. THE THEORY-THEORY OF CONCEPTS

First some terminology: by 'concept', we simply mean a mental representation. We endorse a 'two factor' theory of conceptual content: concepts are individuated both by causal processes that relate mental symbols to their referents and also by internal inferential role. Philosophers sometimes treat these two factors as determining distinct kinds of content (wide and narrow), but we will not take a stance on this issue here. We assume that the meanings of terms in natural language are fixed in the same way, and thus we will sometimes speak of the meaning of a word such as 'accordion' and sometimes speak of the concept *accordion*.

According to the 'theory-theory' of concepts, concepts are analogous to theoretical terms in the following straightforward sense: whatever determines the content of a theoretical term such as 'gene' also determines the content of at least some ordinary concepts, such as *dog* or *think*. Of course, so described, the theory-theory places almost no constraints on conceptual representation. The theory-theory merely posits a continuity in the content determining mechanisms in the two cases—whether these turn out to be classical definitions, patterns of use, conceptual/inferential role, socially based causal connections, or something else.

Most psychologists, including ourselves, who endorse the theory-theory are committed to some particular aspects of conceptual role being central to concept individuation. They hold that the causal-explanatory principles embodied in intuitive and scientific theories provide the most important inferential machinery for determining the content of both scientific theoretical terms and of concepts that articulate intuitive theories(e.g. Carey 1991; Gelman, 2003; Keil 1989; Murphy and Medin 1985). Work in this tradition has uncovered the phenomenon of 'psychological essentialism' (e.g. Gelman 2003; Gelman, Coley, and Gottfried 1994; Keil 1989; Medin and Ortony 1989), and has demonstrated that causal-explanatory features of objects are indeed the most heavily weighted in category membership decisions (e.g. Ahn 1998; Ahn and Kim 2001). Work in this tradition has also analyzed conceptual change in childhood on analogy with conceptual change in the history of science (Carey 1985, 1991; Chi, Glaser, and Rees 1982; Thagard 1992).

As indicated above, we endorse a dual theory of conceptual content, believing in both a wide determining factor (Burge 1979; Fodor 1998; Kripke 1972/1980; Putnam 1975) and narrow determining factor (see Block's 1986 argument for a dual theory). One way of reconciling arguments for both wide and narrow content is to assume that internal conceptual role will turn out to be part of the causal link between entities in the world and mental representations (see Harman 1987; also Margolis 1998, on 'sustaining conditions'). Like others, we have no psychologically adequate analysis of the causal connections between entities in the world and the symbols that refer to them, but we do believe that a full theory of conceptual content will detail these causal connections. We focus here on internal conceptual role, for this is where psychological methods can shed light.

2. ARE ARTIFACT CONCEPTS IN THE DOMAIN OF THE THEORY-THEORY?

What is the domain of the theory-theory? Nobody would expect the theory-theory to provide an analysis of the concepts *over* or *of* or *seven.* These concepts are not embedded in intuitive theories, and do not engender the assumptions of psychological essentialism. Convergent evidence from many sources suggests that natural kind concepts such as *tiger*, *gold*, and *star* fall under the theory-theory. That is, adults adopt an essentialist stance when reasoning about natural kinds (Gelman, Coley, and Gottfried 1994; Keil 1989; Medin and Ortony 1989), assuming they have causally deep, hidden properties (i.e. their essence) which explain the existence of individual members of the kind, determine their surface properties and their behavior in causal interactions with other entities in the world. Since essences determine kind membership for natural kinds (a metaphysical assumption), the representations of essences or of essence placeholders are also at the core of the meaning of natural kind terms (a psychological fact).

Lexical categorization practices have provided one source of data in support of psychological essentialism. For example, Keil (1989) showed that adults are sensitive to the origin of surface properties in deciding animal kind: an animal that looks identical to a skunk and acts like a skunk, spraying smelly stuff at enemies, is not judged a 'skunk' if these features are the result of plastic surgery or a mistaken injection of some mystery chemical during the life of the animal.

Moreover, if this animal's parents and babies are not skunks, it is not judged to be a 'skunk'.

But while theory-theory adherents are in broad agreement that natural kinds fall within the purview of theory-theory, many have been at pains to show that it does not apply to artifacts (Gelman 1988; Keil 1989; Keil, Greif, and Kerner, this volume; Wellman and Gelman 1992; see also Schwartz 1979, for an argument that the causal theory of reference applies only to natural kinds and that artifacts are not natural kinds). For instance, a point often made is that while theoretical developments (e.g. the discovery of genes) are highly relevant to understanding the true nature of natural kinds such as tigers, such developments are irrelevant to the understanding of the true nature of artifacts such as baseball bats (although theoretical developments may allow the successful design of a more ergonomic bat). This absence of a science illuminating the underlying nature of artifacts along with the relative irrelevance of an artifact's underlying material constitution to its identity, has thus led many writers to treat natural kinds alone as falling under the assumptions of psychological essentialism (e.g. Gelman 1988; Keil 1989; Schwartz 1979).

In spite of these considerations, in this chapter we argue that artifact concepts can be readily analyzed within the framework of the theory-theory and, contrary to the emphases on their distinctiveness, that artifact concepts function in everyday life very much as do natural kind concepts. A causal-explanatory structure, with some similarity to what Dennett (1987, 1990) calls 'the design stance', underlies adult concepts of artifacts much as the causal structure of a theory of mind ('the intentional stance') underlies concepts of belief, knowledge, and perception, and much as an intuitive vitalist biology underlies concepts of living things. (Although we are using Dennett's terminology, we do not go along with his view that there is no fact of the matter as to whether an intentional system has intentionality—we do not think that attributions of intentionality are a mere 'stance'.)

3. THE DESIGN STANCE AND EVIDENCE THAT IT STRUCTURES ADULT ARTIFACT CONCEPTS

According to our version of the design stance, an artifact is intentionally created by a designer to fulfill some function. The intended function is the factor which determines the artifact's surface properties, the actual uses it can serve (the intended function as well as others), and its kind. In that sense, the original intended function is the artifact's essence (Bloom 1996, 1998, 2000; Keil 1989; Putnam 1973). Thus, a coffee mug is capable of containing liquids because that is what its designer intended. This intended function in turn constrains its form (it must be closed at the bottom, open at the top, graspable when filled with hot liquids, and so on) and also constrains the material from which it can be made (e.g. not ice). Note that the properties which make it function as a coffee mug also allow it to be used as a pencil holder. Nevertheless, the ability to hold pencils is not the reason the mug came into existence. The cause of its coming into existence is the intention of its designer that it function as a coffee mug.

If it is correct that adults reason about artifacts in terms of the design stance, then just as essential properties are weighted over more superficial properties in judgments of natural kind categories, information about original intended function should be weighted more heavily than superficial properties in judgments of an artifact's kind or purpose. And, indeed, research indicates that adults do just that.Rips (1989) showed that adults weight the original function of an artifact over its form in kind judgments. For example, adults judged an object that had the features of an umbrella but whose creator had intended it to be a lampshade to be a lampshade. Richards *et al.* (1989) found that adults exclude objects from familiar artifact categories if, despite appropriate overall form, a central feature suggests an alternative intended function. Thus, adults judge that an object that looks like a shower cap but is made of paper is not a shower cap (also Kemler Nelson, Herron, and Morris 2002). Kelemen (1999*a*) and German and Johnson (2002) have both shown that adults weigh intended function over current function in judging a novel artifact's purpose. For example, adults judge that are object that was made for one activity (e.g. exercising backs) but used everyday to perform another (e.g. stretching clothes) is 'for' the design function (i.e. back exercising). This same weighting has also been observed by Hall (1995) and Matan and Carey (2001) in adults' kind judgments: an object used for watering flowers but made for making tea is judged a 'teapot'. In short, when making kind and purpose decisions, adults weight the intended function of an artifact over both a current function (Hall 1995; Kelemen 1999*a*; Matan and Carey 2001; German and Johnson 2002) and other properties such as its form (Rips 1989; Richards *et al.* 1989).

Despite this body of research, some researchers have failed to find the expected salience of causally deeper features over more superficial ones in kind determination. For example, Malt and Johnson (1992) argue that both the intended function and the physical properties of the artifact are important features that influence kind decisions, but that neither absolutely pre-empts the other. Consistent with this, they found that a 'thing manufactured and sold to carry one or more people over a body of water for the purposes of work or recreation' (the function associated with boats) but which is 'spherical and made of rubber, is hitched to a team of dolphins, and has a large suction cup that can keep it in one place' (physical features not typically associated with boats) was not judged to be a boat, despite the clearly stated intended function.

Our response to this is that reasoning in terms of the design-stance schema, like all causal reasoning, is a form of inference to the best explanation. People infer function from form, and intended function from possible function, and draw inferences in the other direction too. According to the theory-theory of concepts, intended function is not the most heavily weighted feature because it provides a definition of artifact kind, or because it is the most reliable feature in a prototype structure, but rather because people try to rationalize all they know about an artifact, and knowledge of the intended function constrains this process. It is likely that in the Malt and Johnson (1992) boat example participants did not accept that somebody would design something to carry people over water in such a manner, given that they know that better boat designs are available (see also Bloom 1996, 1998).

In sum, there is considerable evidence that adults reason about artifacts in terms of the design stance, and that intended function plays the same role in reasoning about artifact kinds as representations of essences play in reasoning about natural kinds. This is not to deny important differences between artifact kinds and natural kinds. For example, natural kinds have true essences that are the object of scientific endeavors that do not apply to artifacts: there is no need to study what causes a telephone's surface properties in the same way as one might study what causes an animal's surface features. Furthermore, despite people's commitment to the existence of hidden causes to the surface properties of entities such as gold and tigers, adults' representations of natural kind essences are often under-determined in ways that representations for artifacts are not. That is, people allow that they lack accurate or any ideas concerning what causes a tiger or gold to look as it does—indeed, until relatively recent history they could not know that atomic and genetic structure held the key to each respectively. By contrast placeholders do not play much role in adults' artifact concepts since adults have explicit representations of the causal-explanatory structure underlying chairs and cups. They understand intended function, hence, the design stance.

This noted, our contention is that these differences are less important than they seem and that one need only look to development for evidence that this is the case. Developmental parallels exist to indicate that just as children have to construct a vitalist understanding of living things (Hatano and Inagaki 1999), along with an understanding of species based on reproductive transmission (Solomon *et al.* 1996), so too children must construct the design stance—the intentional-historical scheme that makes full sense of artifact kinds in terms of their intended function. In other words, full insight into artifact kinds is not a given. Early in childhood, all essences are placeholder essences, including those for artifacts.

4. DEVELOPMENT OF A DESIGN STANCE ABOUT ARTIFACTS: CONSTRUCTED OR CORE KNOWLEDGE?

When in development does the design stance become available to organize children's understanding of artifacts and to provide the core of the meaning of artifact terms? There is a natural alternative to the proposal that the intentional-historical understanding of artifact kinds is not constructed until well into childhood—namely that it is one of the domains of innate core knowledge. It may be available early in development, perhaps even in late infancy, as are the physical stance toward inanimate objects and the intentional stance towards agents (Baillargeon 1993; Baron-Cohen 1995; Gergely, *et al.* 1995; Leslie 1994; Spelke 1991).

What is core knowledge? Baillargeon, Carey, Leslie, Spelke, and colleagues have made an empirical claim that there are systems of core knowledge with the following properties: core knowledge is articulated in terms of conceptual representations, some of which are innate. The identification of entities in the world that fall in a domain of core knowledge is supported by innate perceptual input analyzers. Core knowledge systems are learning mechanisms, they support learning about the entities in their domain. Finally, core knowledge continues to articulate our representations of the perceived world throughout development. It is never overturned (e.g. see Carey and Spelke 1994, and Hauser and Santos, this volume, for further characterization of core knowledge, and see Mandler, this volume, for a critical perspective on the existence of systems of knowledge that meet the specification of core knowledge).

Although our concern in this chapter is mainly the narrow determiner of content, it is worth noting that core knowledge systems provide a partial account of the causal connections between the entities in their domain and the symbols for them. There are dedicated perceptual input analyzers for the entities in the extension of core knowledge which take specific kinds of spatio-temporal information as input and yield specific representations as output (e.g. representations of objects (Baillargeon 1993; Carey and Spelke 1994; Spelke 1991), goal-directed action (Csibra and Gergely 1998; Gergely *et al.* 1995; Johnson 2000; Watson 1979), contact causality (Leslie and Keeble 1987; Oakes and Cohen 1990)). We look to evolution to explain how these perceptual analyzers came to be, but their operation satisfies Fodor's explication of asymmetric dependency (Fodor 1998). That is, these analyzers may be fooled into outputting a representation *object* when there is no object present (perhaps just a pattern on a computer monitor), but this mistake depends upon the relations between the spatio-temporal specification of ordinary 3D objects and the processes that build representations of them.

Consider Michotte launching causality as a specific example of a piece of core knowledge. Michotte (1963) carried out elegant psychophysical studies that specified the spatio-temporal properties of events that led to the perception of causality as one object hits a stationary one, leading the latter to go into motion. These events can be described entirely in spatio-temporal terms, but the mind provides a causal interpretation. In addition, elegant studies by Leslie and Keeble (1987), and by Oakes and Cohen (1990), suggest that young infants make this causal attribution, which is supported by the same perceptual analyzers of the spatio-temporal relations among events throughout life.

Or consider representations of intentional agents. Infants analyze patterns of motion (again specifiable in spatio-temporal terms), especially contingency among moving entities and contingency between moving entities and stationary objects in their environments, and attribute goals and attentional/perceptual states to agents or agentive action from this information (e.g. Gergely *et al*. 1995; Johnson 2000; Johnson, Booth, and O'Hearn 2001; Spelke, Phillips, and Woodward 1995; Watson 1979, 1985). As Heider and Simmel (1944) and Durgin and Gelman (see Gelman, Durgin, and Kaufman 1995) elegantly showed, adults also create representations of intentional agents from such displays.

Evolutionary considerations would justify the hypothesis that an intentional-historical understanding of artifact kinds might be part of core knowledge. Human beings, alone among animals, are prolific tool-makers and users. Just as natural selection endowed us with an innate language-acquisition device, and with an innate intentional stance, so too she may have endowed us with core knowledge of artifacts. Such a system of knowledge would enable infants to identify artifacts, guide them in inferring their functions from the uses to which adults intentionally put them, guide them in explaining their properties in terms of those functions, and explain all of this in terms of intentional design. It might be so, but is it?

5. DEVELOPMENTAL DATA FROM STUDIES OF CHILDREN'S ARTIFACT UNDERSTANDING

Developmental research certainly suggests that several prerequisites to a design stance are present from early in development. That is, even if the design stance itself is not innate, components of it probably are, and may form parts of other systems of core knowledge, for example the physical stance (e.g. sensitivity to the structural properties of physical objects) and the intentional stance (e.g. the ability to attribute goals).

With respect to the physical stance, toddlers (Brown 1990; DiYanni and Kelemen 2006; also McCarell and Callanan 1995) and even infants (Caron, Caron, and Antell 1988; Hespos and Baillargeon 2001; Mandler and McDonough 1998*b*; Mandler, this volume) can analyze the functional affordances of objects, recognizing those structural properties that make objects appropriate means to ends. For example, 8-month-olds consider physical width and contact relations when reasoning about containment and pulling events (Aguiar and Baillargeon 1998; Willatts 1999) and 1–3 year-olds are sensitive to shape, rigidity, and length requirements when selecting tools for pushing, pulling, and crushing (Brown 1990; Casler and Kclemen 2005, 2006; DiYanni and Kelemen 2006). With respect to the intentional stance, research focused on early theory of mind suggests that children between 1 and 2 years of age recognize the relevance of monitoring intentional cues from others as the basis for figuring out how to make an artifact work or what an artifact does (e.g. Carpenter, Nagell, and Tomasello 1998; Carpenter, Call, and Tomasello 2002; Gergely, Bekkering, and Király 2002; Hanna and Meltzoff 1993; Meltzoff 1995 for work focused on theory of mind; see Casler and Kelemen 2005; DiYanni and Kelemen 2004; Tomasello 1999; for work focused on artifacts). Finally, by at least 2 years children have built on their intentional stance such that their behavior towards objects reflects a functional construal in which children presume that novel artifacts are 'for' a single, privileged purpose. For example, after only one exposure to an adult intentionally using a novel tool, children will fast map this goal-directed action as the tool's enduring function, consistently returning to the same kind of artifact as the 'tool' for the task and avoiding use of it for any other activity (Casler and Kelemen 2005, 2006; also Markson 2001).

However, while this initial functional construal provides a substantial basis for explanation and inference about objects (e.g. Kelemen *et al.* 2003; Kelemen 2006), it is still not equivalent to the causally rich explanatory structure represented by a fully fledged, intentional-historical design stance based on intended function. Indeed, for a long time it seemed unlikely that any evidence of a design sensitivity would be found until late childhood. This was because of the much-replicated finding that, until at least 6 or 7 years of age, children do not attend to shared function but rather shared shape when lexically categorizing artifacts. Thus, it was repeatedly found that on being shown an exemplar artifact and told its category name, children would extend that name to other artifacts that looked alike but did not serve the same function, eschewing dissimilar objects that could actually do the same thing (e.g. Gentner 1978; Graham, Williams, and Huber, 1999; Landau, Smith, and Jones 1998; Merriman, Scott, and Marazita 1993; Smith, Jones, and Landau 1996; Tomikawa and Dodd 1980). Although adult subjects in these studies sometimes demonstrated the same pattern (e.g. Gentner 1978), children's apparent indifference to what artifacts did in these tasks seemed to render it unlikely that the deeper principle of intended function could play much of a role, let alone a core role, in their artifact concepts.

Recent work demonstrates, however, that these experiments underestimated the weighting children give to functional considerations when classifying artifacts into kinds. This was because, in many such studies, form was, unnaturally, treated as dissociable from function, and this frequently led to comparison stimuli whose functions were arbitrary with respect to shape and were, instead, tied to material or other properties—an approach giving rise to somewhat uncompelling 'functions' (e.g. the capacity to rattle, roll, absorb, and be stickable by pins). In contrast, more recent research finds that when studies use comparison artifacts whose structural properties clearly relate to their functions, children can generalize labels on the basis of function rather than shape similarity as early as 2 years of age (e.g. Kemler Nelson 1999; Kemler Nelson, Russell, *et al.* 2000; Kemler Nelson, Frankenfield, *et al.* 2000). For example, when they are allowed to briefly explore a toy-like artifact called a 'gidget' (e.g. a rectangular object with a hinged flap that slots into the base like a puzzle piece) and are then asked to find another gidget out of a pair of objects, 2-year-olds will select a less similar but functional object (e.g. differently shaped base and flap) rather than a similar object with a dissimilar function (e.g. same shape as exemplar but with a flap that slides into a base like a drawer) (Kemler Nelson,Russell, *et al.* 2000). Three-year-olds will perform in this manner even when the object functions are never directly observed and therefore have to be inferred (also Kemler Nelson and students 1995).

Young children can, then, categorize artifacts on the basis of functional properties from quite early on, as long as the perceptual information is clearly consistent with a specific function. Moreover, recent work suggests that other precursory components to an adult-like design stance are also available quite early. Several studies now reveal that in addition to knowing that artifacts—rather than natural kinds—are made by people (Gelman 1988; Gelman and Kremer 1991; Keil 1989; Petrovich 1997; Kelemen and DiYanni 2005), pre-schoolers recognize the special role that a designer's intention plays in designating an artifact's category and name. Some indication of this first came from studies exploring children's naming of representational artifacts like drawings. Bloom and Markson (1998) found that 3- and 4-year-olds named pictures whose referent was objectively ambiguous (potentially lollipop or balloon) based on what the creator (themselves) intended the picture to depict. Similarly,Gelman and Ebeling (1998) found that 2- and 3-year-olds were only likely to label familiarly shaped drawings with familiar object names when they were told that the pictures were the products of intentional creation as opposed to accidental action (e.g. someone spilling paint). Importantly, this finding has now also been extended to non-representational artifacts with Gelman and Bloom's (2000) finding that, from 3 years of age, children are more likely to generate an artifact term (e.g. 'belt', 'hat') for familiarly shaped objects described as purposefully created, but more likely to generate a material composition term (e.g. 'paint', 'clay') for the same objects if told that they were accidentally originated. However, while these results reveal significant competence relevant to the design stance, they provide direct evidence only of young children's knowledge that (i) artifacts are created by people, (ii) it is appropriate to extend a familiar artifact category label to an intentionally created object, (iii) the intentional designer has the right to designate the name and category membership of their creation, in other words, the designer possesses 'baptism rights' (see German and Johnson 2002, for discussion). This still stops short of evidence for the full adult design stance, which requires that the above elements are drawn together and given cohesion by a notion that subsumes them—the idea that the designer creates an artifact category member with the intention that it perform a particular function. Children who have this insight should consequently weight the design function over any non-designed, salient current use in their artifact reasoning. So, when exactly do children demonstrate this level of understanding?

As with most developmental research, a precise answer to this kind of question is up for debate, because different methods have produced slightly different findings for different researchers. Nonetheless, current evidence converges to suggest that a full understanding of the design stance does not emerge until late in the pre-school years, somewhere in the 4- to 6-year age range. For example, on the older side of this range, Matan and Carey (2001) and German and colleagues (e.g. German and Johnson 2002; Defeyter and German 2003) both find evidence that the design stance is not constructed until close to 6 years of age. Matan and Carey (2001) presented children with scenarios in which one character made an artifact for a particular purpose (e.g. to eat dinner on), but before it was ever used for that purpose, another character used it for another purpose (e.g. to play a game throwing it to other people in the park). When asked what the object was—a plate or a frisbee—the 4- year-olds were at chance in one study and just above chance in selecting the intended function in another. Few 4-year-olds appealed to intended function when justifying their judgments and, while they were not tested unless they did, many had difficulty remembering who made the object and for what purpose, as if this information was not naturally relevant to organizing their representations. Indeed, Matan and Carey suspected that the apparent design-based responding of some of their 4-year-old children may have only reflected shallower knowledge of a creator's 'baptism rights' rather than any understanding of the design stance. Six-year-olds, in contrast, differed from the 4-year-olds in each of these respects, resembling adults in being able to remember the information upon first being told it, in categorizing the artifact on the basis of original intent of the designer, and justifying their responses in terms of the design stance.

The question of whether those of Matan and Carey's younger children who did categorize consistent with intended function were solely responding on the basis of a creator's right to name their creations rather than insight into the significance of intended function remains open, in part because, beyond the function information presented in the stories, half of Matan and Carey's trials involved items with function-based names (e.g. shopping cart vs. stroller, baseball bat vs. rolling pin) (see Kelemen 2004 for additional discussion). Nevertheless, findings by German and Johnson (2002) accord with Matan and Carey's conclusion that a design stance based on intended function is not present until quite late. In one condition of German and Johnson's studies, 5-year-old children were told stories about an object that was made for one purpose, given away, and then intentionally used by a new owner for something else. German and Johnson found that while 5-year-olds showed a sensitivity to baptism rights by weighing a designer's label over another agent's label when deciding what the novel artifact's category name was, they did not reliably use designer intent over current intentional use when judging what a novel object was 'really for'.

German and Defeyter (German and Defeyter 2000; Defeyter and German 2003; Defeyter 2003) have gone on to argue that further evidence of pre-school children's lack of a design-based construal is also provided both by 5-year-olds' performance on tests of functional fixedness and by their approach to functional fluency tasks. For example, in context of functional fixedness, German and Defeyter (2000) have found that while 6- and 7-year-olds have difficulty disregarding an artifact's design function when asked to creatively problem-solve with it (e.g. figure out how to use a box to reach a shelf), 5-year-olds have far less problem-solving difficulty, more readily seeing an artifact as 'for' whatever someone wants it to be 'for' (i.e. seeing the box as a mounting block, not as a container). They suggest that this immunity to functional fixedness occurs because design function is not yet core to pre-schooler's artifact conception. Similarly, in functional fluency tasks in which children are asked to generate possible uses for familiar artifacts (e.g. bricks), Defeyter (2003) has found that while both 5- and 7-year-old children are uniformly relatively poor at the task, 7-year-olds are more likely to remain fixed on conventional 'design' functions than 5-year-olds, who are more likely to generate entirely novel uses—a tendency that, again, seems to indirectly suggest that intended function is less central to their artifact thinking. German and Johnson (2002) point out that results suggesting that 5-year-olds do not have a design stance are potentially unsurprising if the computations involved in reasoning about design intentions are actually considered. Specifically, they argue that design attributions involve recursive reasoning about second-order mental states (e.g. 'the maker intends (that the user intends) that X will perform Y'), something regarded as difficult for young children (e.g. Perner and Wimmer 1985; but see Sullivan, Zaitchik, and Tager-Flusberg 1994).

However, the involvement of second-order mental state reasoning in design attributions is challengeable. Computationally, design intentions may reduce to 'the maker intends that the user does X with Y' or 'the maker intends that X does Y', and children's ability to manipulate embedded mental state content of a more opaque, complex form than the goal state content of design intentions has been documented as early as 3 and 4 years of age (e.g. Chandler, Fritz, and Hala 1989; Siegal and Beattie 1991) and even, perhaps, infancy (Onishi and Baillargeon 2005). In principle then, it is not clear that there is any computational barrier to children representing and reasoning about intended function earlier than 6 or 7 years of age. Additionally, with regard to German and Defeyter's functional fixedness and functional fluency results, there are reasons to suspect that factors independent of children's artifact concepts underlie the findings. For example, other studies by Kelemen and colleagues (Kelemen 2001, 2006*b*) have directly explored whether there is any correlation between 3- to 5-year-olds children's susceptibility to functional fixedness during problem-solving tasks and their tendency to construe artifacts in terms of original design (as assessed by a task described below). These studies find none of the expected negative correlations between the two tendencies. Furthermore, follow-up work by Defeyter (2003) on functional fluency has revealed that when children are given instructions highlighting the acceptability of generating entirely novel functions, developmental differences between 5- and 7-year-olds are eliminated. Both these patterns of results therefore strongly suggest that something other than children's immature artifact concepts (e.g. age- or education-related increases in conventionality) might account for 5-year-olds' relative advantage in both German and Defeyter's functional fixedness and functional fluency findings. Finally, questions also occur regarding German and Johnson's (2002) function judgment results, for two reasons. First, aside from children's performance, even adults' tendency to judge that the novel artifacts were 'really for' the designed function rather than the intentional use was low—initially more than half of the adults made design-based judgments on 50 per cent or less of the trials suggesting possible issues with stimulus items (see Kelemen 2004). Second, using the same kind of function-judgment method, earlier work by Kelemen (1999) had already found evidence of a design stance understanding in children as young as 4- to 5-years of age.

Specifically, in Kelemen's (1999) study, adults and a mixed group of 4- and 5-year-old children were told about depicted novel artifacts that were successfully intentionally designed for one purpose (e.g. squeezing lemons) then immediately given to someone else who, depending on the experimental condition, accidentally or intentionally used the artifact for another activity on either one or many occasions (e.g. picking up snails once or repeatedly). Children were reminded about both functional activities after hearing the story and all alternative uses, whether accidental or intentional, were explicitly described as positive outcomes. Nevertheless, when asked to judge what the objects were 'for', children and adults had no overall differences from each other and showed a significant tendency to say the artifacts were 'for' their intended function in each condition. A subsequent study then replicated this effect using actual, manipulable novel artifacts (Kelemen 2001, 2006*b*). In contrast to 3-year-olds, who were at chance, separate groups of 4- and 5-year-olds not only judged the objects as 'for' their design function rather than their everyday intentional use, but also favored design function when judging what kinds of other items the novel artifacts would belong with (i.e. in a house). The tendency to favor intended function in this latter study was most marked in the 5-year-old group who averaged doing so 72 per cent of the time.

Kelemen's (1999) finding of a sensitivity to design in younger pre-schoolers does not stand in isolation. Kemler Nelson,Herron and Morris (2002) recently found that 4-year-olds are more likely to extend familiar category names to unfamiliar non-functional artifacts by making inferences about intended function rather than by attending to the objects' superficial appearance (see also Richards *et al.* 1989). Indeed, newer evidence indicates that, under certain conditions, even 3-year-olds know to weigh intended function when deciding how to categorize an artifact. For instance, Jaswal (2005) found that 3- and 4-year-old children are more likely to assign 'label-consistent' functions to hybrid artifacts resembling members of one familiar artifact category but labeled as another (e.g. a hat-like object labeled 'cup') if told that the labeler 'made' the object rather than 'found' it. Furthermore, DiYanni and Kelemen (2006; see also DiYanni 2006) found that 3-year-olds demonstrate a tool preference across contexts and users if shown two equally affordant tools successfully performing a task (e.g. ringing a bell in a cage) but hear one tool described as 'made for' the purpose and the other's intended function described as unknown. Finally, DiYanni and Kelemen (2006; see also DiYanni 2006) also found that unlike 2-year-olds, 3-year-old children can be induced to select a physically inappropriate novel tool for a task (e.g. a fluffy object for cookie-crushing) over a highly appropriate novel tool (e.g., a pestle) but only if (misleadingly) told that the inappropriate tool is 'made for' for the purpose (for related research see Diesendrunk, Marleson, and Bloom 2003, but also see Truxaw *et al.* 2006).

In summary, minor age differences aside, the body of contemporary research converges on 4 to 6 years as the age range when an explicit understanding of design becomes progressively more evident, with studies also beginning to find that, under certain conditions, design-stance insights can be elicited from children during the early pre-school years. Nevertheless, despite these findings of earlier competence, even the most sympathetic interpretation of the current evidence indicates that the answer to the question of whether humans are innately endowed with the design stance as an aspect of core knowledge is 'no'. Instead, using the innate building blocks that core knowledge provides them, children's design stance seems to be gradually constructed. The developmental progression can be crudely characterized (Casler and Kelemen 2005, 2006; Kelemen 2004*a*, 2006*b*) as one in which children move from understanding an artifact as a means to an intentional end (thus 'for' a user's current goal), to viewing it as the embodiment of a goal (thus 'for' a privileged, intrinsic, enduring, function), to finally understanding it in terms of a full-blown design stance—an explanatory structure that is anchored by an understanding of intended function and supports rich inferences about the artifact's raison d'être, kind, properties, and future activity.

6. CAUSES AND IMPLICATIONS OF THE DEVELOPMENTAL CHANGE

The evidence surveyed above suggests that young toddlers attend to artifact function and conceptualize artifacts in terms of an explanatory structure that is derived from their understanding of goal-directed action and supports function-based reasoning and inference about objects. Nonetheless, it is not until several years later in development that children draw together various elements of their understanding to construct a fully elaborated design stance.

This brief summary of the current state of the art regarding children's conceptions of artifacts raises several questions. First, how should we understand the transition from artifact concepts not rooted in the design stance to artifact concepts that are so rooted? In descriptive terms, does this transition involve a conceptual change in artifact concepts? Second, how does the child construct the design stance? What contributes to such a transition? Third, how does this transition relate to developmental transitions within representations of essentialized natural kinds? Are there parallels, or perhaps even direct influences of one on the other?

6.1. Conceptual Change in Artifact Concepts?

Whether the transition we have argued for constitutes a conceptual change depends, of course, on one's analysis of concepts and of conceptual change. As we use the term, concepts are representations, and representations persist through time (one thinks thoughts about the same entities on different occasions). Conceptual change, then, occurs when whatever determines the content of a given representation changes. On many analyses of concepts, as different from one another as the classical empiricist view and Fodor's atomistic view (e.g. Fodor 1990, 1991), the very notion of conceptual change is incoherent. On the classical view, in which concepts are individuated by definitions that provide necessary and sufficient conditions for category membership, it is more natural to think of concept replacement than conceptual change (i.e. a change in a definition results in a new concept; see Katz 1972). Similarly, on the atomistic view, if the causal laws that relate a symbol to entities in the world change such that a new set of entities is picked out by a given symbol, it may seem more natural to think of concept replacement rather than conceptual change. It should be noted, however, that this assumption of conceptual replacement depends upon the way the concept-to-world causal laws are characterized. If the extensions of successive concepts overlap and if the processes leading to the change of extension involve some transformation of those very causal connections, then it makes sense to talk of conceptual change.

At any rate, the current work is placed in a theory-theory framework that endorses a place for conceptual role in determining conceptual content. Indisputably, theories change, and the most deeply entrenched causal schemata that structure theories change. Indeed, it is this fact that has led students of conceptual development from Piaget on to look to historical theory change for insights into the process of conceptual development in childhood, and parallels between historical theory change and conceptual development motivate the theory-theory of conceptual development, as well as the theory-theory of concepts. The theory-theory of concepts speaks of conceptual change rather than replacement because much of the inferential role that partly determines conceptual content remains constant over the change; as Kuhn (1983) puts it, incommensurability is always local.

Returning to the specific case of artifact kinds, according to the present analysis, the identification of something as a member of an artifact category (e.g. a chair) is not a process of applying a classical definition but one rooted in inferences to best explanation based on the deepest causally relevant features known. Conceptual change can be said to take place to the extent that that these causally deepest structures—those used to explain an entity's existence, properties, and activities—undergo modification. This is because, on the theory-theory view (with its emphasis on internal conceptual role), such structures are central to the concept-to-world sustaining mechanisms determining reference and, as such, re-analyses may cause changes in the extension of the concept (see Kitcher 1988, on mismatch of referential potential across episodes of conceptual change).

Pre-school developments in the artifact concepts between 3 to 6 years reviewed earlier satisfy this analysis of conceptual change. For sure, the explanatorily deepest features change: artifacts move from being explained in terms of an intentional stance (how a person might use an object), to a deeper functional stance (what the object itself is for), to an even deeper design stance (what it is made for). A consequence of this change is that the referential potential of the term changes throughout these years. For a 2-year-old, an object that someone drinks out of, but which was made as a flower vase, might fall in the extension of the concept *glass*. Of course, the child also defers to experts, and a correction that it is a vase not a glass might be part of the input that leads to a change in the core of the concept. These changes would be expected to reflect in children's categorization decisions, and such changes are the primary evidence of the transition under discussion

This is not to deny that theories change in many ways, and conceptual change can be a matter of degree—there will be a continuum of changes between mere changes in beliefs about the entities in some domain and changes in the very concepts of those entities. Thus while we argue that the construction of the design stance does entail conceptual change within artifact concepts, we also importantly note that it provides a relatively weak case. It does not involve the multiple interdependent differentiations and coalescences that constitute conceptual change accompanying radical theory changes (e.g. Carey 1985, 1988, 1991; Kitcher 1988; Kuhn 1962).

6.2. Where does the Design Stance Come From?

Let us turn now to the second question. Whether or not the creation of the design stance contributes to conceptual change within artifact concepts, how does the child manage it? Our answer appeals to two very different types of influences. First, we must account for the origin of conceptual components of the design stance, the conceptual stuff from which it is constructed. And second, we appeal to domain general factors that enable children to construct kind representations overall, both natural kinds as well as artifact kinds.

As we indicated in the above review, we trace the ultimate origin of the components of the design stance to two innate systems of core knowledge—the system that provides representations of intentionality and the system that provides representations of objects and their causal potentials. It is because humans are innately endowed with the capacity to analyze their own and others' actions in terms of goals, as well as with the capacity to analyze events in terms of the causal relations among objects, that they can also analyze the role and properties of external objects in terms of human goals. It is these abilities that get artifact representations off on the right foot, so to speak.

Several domain general tools that support concept acquisition also play a role in the process. First and foremost, children have the capacity to distinguish kinds from other types of categories (e.g. those united by properties), and they are sensitive to several types of information in establishing whether a given term refers to a kind or not. One type of information is linguistic—kinds are lexicalized as nouns and properties typically by adjectives, and children as young as 13 months are sensitive to this contrast (Waxman and Markow 1995; we speak of 'spoons', not 'spoonish things'). Also, kinds are referred to by generics (e.g. 'cars need gasoline, the radio is a wonderful invention'), and children as young as age 2 take generics to refer to kinds (Gelman 2003). Other information is conceptual—kinds have more inductive potential and are more causally potent, on average, than are properties (Gelman,Collman, and Maccoby 1986). Evidence of this sort helps the child establish that *cup* is a kind concept and *red* is not. These assumptions lead children to weight causal explanatory features most heavily in their representations of concepts such as *cup*, which is why functional features are weighted more heavily than purely perceptual features by children as young as 2.

Once the child has evidence that a given concept is a kind concept, the child's first assumption is that it is a basic level substance sortal (e.g. Carey 1994; Hall 1993; Macnamara 1986; Xu and Carey 1996). Substance sortals are contrasted with phase sortals (e.g. passeng*er*) or stage sortals (e.g. *puppy*) because, unlike these other sortal types, substance sortals trace identity throughout an entity's entire existence. As a result of this identity-tracking property, children's assumption that the kind concept is a substance sortal may lead them to focus their attention on origin, which may be why the child begins to attribute the maker with baptism rights. That is, if a cup is a cup throughout its whole existence (if *cup* is a substance sortal), any explanation of how it comes into existence that coheres with other, already analyzed, explanatory features will become entrenched. Certainly, the child must learn about manufacture, but he or she has ample opportunity to do so from very early in development—the child creates drawings and participates in making meals, creating towers with sticks, blocks, and so forth. Children's participation in the kinds of easily observable 'manufacturing' activities that pervade all human cultures (e.g. cooking, building) readily provide information about the relevant aspects of origin tied to intention. Indeed, a prediction can be made that, among children who are equivalent with respect to meta-cognitive skill (thus equivalent in their abilities to reflect on their own and others' creative actions), children from more 'self-reliant' (do-it-yourself) rather than 'consumerist' (just buy it) social/cultural backgrounds may show precocious development of a design stance.

In sum, we offer two types of answers to the question of the origin of the design stance: we appeal to the systems of core knowledge that provide part of the material from which it is constructed, and we appeal to domain general theory-building processes that guide the child toward essentializing and theorizing about artifacts in terms of their origins.

6.3. Artifact Concepts and Natural Kind Concepts: Relations in Development

We turn now to the third question raised by our brief review, namely, how the construction of the design stance relates to developments within natural kind categories. Structurally, there is a close analogy between the changes within artifact concepts sketched here and some of the conceptual changes within concepts of animal kinds that have been described in the literature. Studies of switched-at-birth animal adoption show that by age 4 children know that cats give birth to cats, and that even if a baby born to a cat mother is raised by dogs, it will grow up to be a cat (Wellman and Gelman 1992; Johnson and Solomon 1997; Lopez, Atran, and Coley 1997). However, this origins knowledge is only gradually elaborated into a causal schema of inherited essences, as the mechanisms of biological transmission of traits are differentiated from social transmission, and the traits that fall under each type are differentiated from each other. This process takes place, in Western culture, over the years of 4 to 7, as shown by Keil's transformation task and discovery tasks (Keil 1989) and in inheritance tasks (Johnson and Solomon 1997; see also Solomon *et al.* 1996; and Springer and Keil 1989).

Insofar as this process reflects a gradually deepening understanding of the relevance of origin to explaining object properties and kind, this is the same kind of conceptual change as that described for artifact terms, and it may not be coincidental that it is taking place at roughly the same time. Analogical transfer of knowledge derived from developments in the domain of intuitive biology might contribute to developments in the artifact domain—increasing weighting of details of origin in determining species kind may reinforce increasing weighting of details of origin in determining artifact kinds. For example, as 5- to 8-year-old children increasingly reorganize their understanding of the identity of living things in terms of reproduction and birth (e.g. Solomon, *et al.* 1996; Johnson and Solomon 1997), their attention to origins in the biological realm may inform their attention to origins in the artifact realm (Matan and Carey 2001).

Alternatively, the direction of analogical transfer might be the reverse. Children have an early toehold into understanding the artifact domain, given their precocious abilities in relation to the intentional stance. Perhaps it is unsurprising then that Mandler (this volume) finds that 19-month-olds have greater expertise identifying the specific properties of artifact categories versus biological kinds, and as we reviewed above, there is some evidence for the beginnings of the construction of the design stance as early as 3 years. These data suggest that insights into artifacts and artifact origins are developmentally antecedent to insights into the domains of biology and other natural phenomena (see Keil Grief, and Kerner, this volume for a proposal why this should not be the case). Indeed, artifact knowledge may not only help to facilitate children's understanding of the biological and natural world but also potentially obfuscate it. Specifically, it may be via the influence of their privileged sensitivity to intentionality and hence their deepening artifact knowledge that children become prone to develop a 'promiscuous teleology'—the tendency to treat natural objects of all kinds as occurring for a purpose; a cross-culturally documented bias that impacts the ease with which scientific ideas, such as those inherent to evolutionary theory, are ever truly acquired (Kelemen 1999*b*, *c*, *d*, 2003, 2004; see also Evans 2000, 2001; but see Greif *et al.* 2006; and Keil 1992).

Evidence of the influence of artifact knowledge on reasoning about natural kinds is further provided by findings that children both endorse and spontaneously generate artifact-like, other-serving functional explanations for natural objects and their properties (e.g. animals have wide backs so that they can be physically sturdy *and* so that other animals can ride around on them) (Kelemen 1999*b*, 2003, Kelemen and DiYanni 2005; but see Kerner and Keil cited by Keil Greif, and Kerner, this volume, for possible conflicting evidence) and regard natural entities that cannot perform other-serving activities (e.g. a mountain that can no longer be climbed) as 'broken' and in need of being fixed or replaced (DiYanni and Kelemen 2005). Finally, elementary-school children's tendency to ascribe purpose to nature is significantly correlated with their tendency to view natural phenomena as being 'made by someone' (Kelemen and DiYanni 2005). In short, at developmental points when children's design stance on artifacts seems well established, there are results suggesting that children are using artifact knowledge to make sense of domains where they have less expertise.

Of course, ultimately the question of whether biological knowledge influences artifact knowledge, or whether the reverse is true, may never be decided, since both possibilities could be accurate—analogies of this sort may serve to reinforce explanatory schemata back and forth across both domains.

In considering these similarities between the biological and artifact domains, it is important to note how the kinds of conceptual changes involved in each of these domains are also different in degree, if not in kind. We have argued that changes in the core of a specific artifact concept (e.g. *broom*, *cup*) parallel those that take place in the core of a living thing concept (e.g. *baby*, *dog*), insofar as, in both cases, the child constructs an explanatory schema that privileges origin at the core of the kinds in the domain. However, it is also true that the biological case involves much deeper and far-reaching conceptual change, such that in some cultural contexts it is not complete until adolescence or even adulthood (Astuti 2001; Astuti, Solomon, and Carey 2004; Bloch, Solomon, and Carey 2001). On Carey's analysis (1985, 1995), this is due to the fact that a vitalist biology, as well as biological understanding of naturalized kind essences, has no direct precursors in core knowledge. In contrast, as we have argued here, the infant's understanding of intentionality places artifact concepts in an inferential structure that maintains some fundamental continuities throughout development, and this has implications for the degree of incommensurability between child and adult artifact representations and talk. In sum, while acknowledging their differences, an analysis of artifact concepts in terms of the design stance has many deep parallels with the analysis of natural kind concepts in terms of psychological essentialism and the theory-theory of concepts. These analyses place an explanatory schema at the core of each type of concept, and pose parallel questions for development. From the point of view of development, they yield many parallel answers.