PHILOSOPHICAL PSYCHOLOGY, VOL. 13, NO. 3, 2000

Natural epistemology or evolved metaphysics? Developmental evidence for early-developed, intuitive, category-specific, incomplete, and stubborn metaphysical presumptions

PASCAL BOYER

ABSTRACT *Cognitive developmental evidence is sometimes conscripted to support "naturalized epistemology" arguments to the effect that a general epistemic stance leads children to build theory-like accounts of underlying properties of kinds. A review of the evidence suggests that what prompts conceptual acquisition is not a general epistemic stance but a series of category-specific intuitive principles that constitute an evolved "natural metaphysics". This consists in a system of categories and category-specific inferential processes founded on definite biases in prototype formation. Evidence for this system provides a better understanding of the limited "plasticity" of ontological commitments as well as a computationally plausible account of their initial state, avoiding ambiguities about innateness. This may provide a starting point for a "naturalized epistemology" that takes into account evolved properties of human conceptual structures.*

An ontology specifies kinds of stuff in the world. There is now considerable psychological evidence that mentally represented conceptual knowledge includes ontological commitments. The latter not only specify large categories of objects (e.g. "artefact", "animal", "person", "event" or "number") but also produce spontaneous inferences from perceived states and events as well as predictions of future possible states and events. These ontological commitments, expectations and inference engines are crucial in the acquisition of new basic-level kinds (e.g. "telephone", "giraffe", "child", "dozen") and they are now the object of fine-grained developmental accounts (see Hirschfeld & Gelman, 1994, for illustrations). How these broad ontological concepts and their associated inferential principles are acquired and mentally represented is a crucial question for epistemology. Indeed, some recent

Pascal Boyer, Centre National de la Recherche Scientifique, France; Institut des Sciences de l'Homme, 14 Avenue Berthelot, 69363 Lyon, France, tel.: + 33 47 272 64 60, fax: + 33 47 272 65 90, email: pboyer@ish-lyon.cnrs.fr; until 31 July 2000: ISBER, 2201 North Hall, University of California, Santa Barbara, CA 93106, USA, tel.: + 1–805–893–7749, fax: + 1-805-893-7995, email: pboyer@isber.ucsb.edu

ISSN 0951-5089 (print) ISSN 1465-394X (online)/00/030277–21 © 2000 Taylor & Francis Ltd

arguments in "naturalized epistemology" take their illustrations and derive some of their arguments from the experimental literature on conceptual development. However, this use of the evidence is often misleading. I will argue here that a more precise description of the developmental evidence suggests a more specific picture of early ontology (and of its consequences for adult conceptual structure). This may help construct a natural epistemology supported by computationally tractable, empirically based accounts of our cognitive dispositions rather than by *a priori* assumptions about what they should be. It may also illuminate traditional debates on processes of concept acquisition, on the conceptual continuity between infants and adults and on the putative "innateness" of particular ontological commitments.

Epistemic stance vs natural metaphysics

That our conceptual sensitivities are dramatically off target can be in principle excluded on evolutionary grounds (Quine, 1969). However, this leaves open the possibility of a vast range of sub-optimal designs whereby organisms survive on the basis of false but innocuous beliefs. A "sound" evolved cognition, in this view, would have two main characteristics. First, it would be able to recognize relevant features of objects and identify their correlations. Second, it would go beyond surface correlations and identify underlying structures that cause them.

Developmental psychological evidence seems to bear good news for this version of the naturalized epistemology project. Studies of pre-schoolers' induction and categorization behaviour seem to show that children are certainly not driven by a pure sensitivity to correlations of external properties in objects. Summing up developmental research by Gelman and Wellman and others, Kornblith argues that pre-schoolers show "an understanding of the relevance of underlying properties of natural objects to their kind and function, even in cases where they do not know what those underlying properties are" (Kornblith, 1993, p. 70). This results from an epistemic stance that appears very early: "even extremely young children look for regularities beyond superficial characteristics in their classificatory schemes" (Kornblith, 1993, p. 78). According to Kornblith, developing minds are essentially on the right epistemic track, given the kinds to be discovered in nature.

This kind of developmental evidence, however, is compatible with two rather different interpretations, which I will call the epistemic stance and natural metaphysics views, respectively.

The *epistemic stance* view is that children attend to underlying properties of objects as a consequence of a general, principled cognitive strategy. They have a general disposition to attend to underlying properties, causally connected properties in particular, as well as correlated surface properties of objects.

The alternative, *natural metaphysics* view, is that children are led to attend to underlying properties in specific ways for each ontological domain. In this view, they, for instance, stipulate common internal, functional processes in members of an animal species, not because of a general essentialistic stance, but because of particular, domain-specific presumptions about objects in the ANIMAL domain.

These presumptions, in this view, would not be activated when dealing with objects in another ontological domain, such as artefacts, for instance.

Kornblith himself does not choose between these two interpretations. By contrast, R. Millikan presents a view of early concept development that strongly suggests an "epistemic stance" interpretation. Millikan first argues that a single "substance" notion underpins the development of concepts traditionally viewed as different: concepts of kinds (TIGER, TELEPHONE), of "stuffs" (MILK) and of persons (BILL, HILLARY). We should view concepts as cognitive "practical skills" that allow us to recognize substances rather than descriptions of these substances. Now many properties of objects, though not necessarily hidden, are in fact undetectable in each particular context. So a natural cognitive system needs a general "conceptual tracking" capacity. This conceptual analogue of perceptual tracking "makes it possible for the organism to store away knowledge about the thing collected on earlier encounters for use on later occasions" (Millikan, 1998, p. 62). Conceptual tracking includes the disposition to go beyond observable encounters and to *assume* non-obvious stable properties in objects of a same substance. This is construed as a *general* strategy that applies across domains, indeed allows the child gradually to acquire representations of the differences between domains. For Millikan, the concept of a "substance" is more general than that of KIND or STUFF or INDIVIDUAL, and probably comes first in conceptual development. "Putting it Quine's way, the child's first recognitions … *must* be merely of more Mama, more milk, and more mouse" (Millikan, 1998, p. 58, my emphasis). This then develops into domain-specific knowledge: "Tracking in this property-blind way would make it possible to observe, for various kinds of objects, what sorts of things tend to remain the same and what may change within a short period" (Millikan, 1998, p. 63). All this seems to suggest that early cognitions are indeed the outcome of a single, cross-domain epistemic stance. In the alternative, "natural metaphysics" view the child would be described as developing first a set of domain-specific understandings of underlying properties.

On the whole, naturalizing epistemology arguments are compatible with either Millikan's general "epistemic stance" interpretation or the "natural metaphysics" interpretation. However, choosing between these two—that is, evaluating which of them is better supported by the developmental evidence—does make a difference as regards the initial form and the plasticity of ontological commitments.

If we construe early development as driven by one general epistemic attitude, we should expect that the infant's conceptual baggage consists not in specific presumptions about what kinds of objects there are in the world but of a set of good recipes about how to get sound evidence beyond misleading surface features. This in turn implies that ontological commitments should be rather flexible, depending on what kinds of evidence children are presented with. Some aspects of conceptual development seem to support this view. For instance, evidence for early "essentialism" shows that the child's ontological understandings are intrinsically *theoretical*, underdetermined by available evidence (Gopnik & Wellmann, 1994; Wellmann & Gelman, 1992). Now if cognitive processes are intrinsically theoretical, they are also *defeasible*, as Churchland points out (Churchland, 1989, p. 188; see also Churchland, 1979, p. 63ff).

If, on the other hand, we interpret conceptual development as the consequence of natural metaphysics, the infant is described as starting with a set of definite presumptions about broad ontological categories. This in turn suggests that developing ontological commitments should be less flexible, as they are partly informed by these prior ontological presumptions.

A more precise consideration of the evidence does not just help choose between these interpretations of early development. It also suggests a precise, computationally tractable description of what "prior ontological presumptions" may consist of and of their connection to evolved computational properties of human brains.

Psychological evidence: categories and inferences

That ontological categories are real psychological structures is not really in doubt. Concepts such as ANIMAL, PERSON or ARTEFACT are activated by specific cues in the environment and allow specific inferences from these cues. Presentation of new instances of these different categories triggers different neural activation (see e.g. Martin *et al.*, 1996) and some neural pathologies sometimes result in selective impairment of one of the categories (see e.g. Sartori & Job, 1988; Damasio, 1990). Developmental evidence demonstrates the early emergence of a system of fast, spontaneous identifications, categorizations and inferences at the level of broad ontological categories. This occurs at an age where these processes cannot be much affected by conscious deliberation or by explicitly available knowledge.

There is evidence for the presence and salience of categories like HUMAN, ARTEFACT, ANIMAL, from infancy, especially from spontaneous sorting behaviour (Younger, 1985) and serial touching (Starkey, 1981). That infants possess the capacities underlying complex categorization activities is not really controversial (Nelson, 1973; Younger, 1990). Infants use those general capacities in the spontaneous representation of their environment. Obviously, human infants can, like most animals, distinguish con-specifics from other types of beings on the basis of complex perceptual cues (Morton & Johnson, 1991) that trigger particular forms of interaction (see Meltzoff, 1994, for instance). But less obvious ontological distinctions appear very early, like a distinction between living and non-living things (Bullock, 1985; Gelman *et al.*, 1983; Richards & Siegler, 1986) that is grounded in an early sensitivity to the perceptual difference between self- and non-self-generated movement in physical objects (Massey & Gelman, 1988; Premack, 1990; Premack & James-Premack, 1995). The salience of the category ANIMAL is also demonstrated in habituation tests (Eimas, 1994; Faulkender, 1974; Ross, 1980). Mandler and her colleagues have demonstrated the existence of a variety of ontological categories such as ANIMAL and ARTEFACT in 18-month-old children (Mandler *et al.*, 1991).

To say that ontological categories include perceptual cues does not entail that they are the direct outcome of simple feature distinctions. Indeed, ontological identification often overrides superficial similarity. Although young children focus on shape similarity when trying to constrain the reference of a novel noun (Landau *et al.*, 1988), the salience of shape is dependent on prior ontological identification [1]. Ontological categories allow children to make a clear intuitive distinction between superficial similarity and same kind membership.

Ontological categories are intimately associated with domain-specific *inference engines*, also called "modes of construal" (Keil, 1994) or "skeletal principles" (Gelman, 1990) in the literature. Inference engines focus on particular aspects of the objects encountered. They are activated by the identification of an object as a member of one of the ontological categories. They deliver inferences about various non-obvious aspects of these objects. They can be concurrently activated in the representation of a single object. To take a simple example, seeing someone walk towards you and anticipating what will happen next simultaneously involves (1) physical predictions, e.g. how fast this person's current motion will get her closer to you, (2) biological expectations, e.g. that this person is propelled by some internal source of energy, and (3) psychological inferences, e.g. that this person may want to talk to you. There is developmental evidence for a variety of such aspect-based principles, of which I will only mention a few (for a general survey, see Atran, 1989; Gelman, 1990; and the essays collected in Hirschfeld & Gelman, 1994).

Number

The domain of *number* provides a good illustration of specificity of informational cues and inference principles. Experiments by Wynn demonstrated that infants are sensitive to the pure "numerosity" of displays (Wynn, 1992). "Impossible" changes in the number of items displayed (e.g. $1 + 1 = 1$, $2–1 = 2$) surprise infants in dishabituation experiments (see also Antell & Keating, 1983; Starkey *et al.*, 1990). Early sensitivity to numerical aspects of displays underlies the subsequent ease with which children acquire counting systems (Gallistel & Gelman, 1992).

Physics of solid objects

Many of the principles underlying adult intuitive expectations in the domain of "intuitive physics" appear early in infancy, e.g. "continuity" (objects move in continuous paths), "solidity" (objects do not coincide in space), "support" (unsupported objects fall downwards) (Baillargeon, 1987; Baillargeon & Hanko-Summers, 1990; Spelke, 1990), or a rule of "no action at a distance" that excludes non-contiguous physical causation (Leslie, 1979, 1988).

Features of animacy

Identification of an object as an ANIMAL or HUMAN triggers inferences about agency. For instance, pre-schoolers assume that movement is externally generated for artefacts and internally generated for animals. This, even in a case where both kinds of items are unfamiliar and even in a case where the child can actually see a human hand moving the animal (Gelman *et al.*, 1994). Children are reluctant to judge artefacts capable of self-propelled motion and resist extending animate properties to inanimate objects such as dolls and stuffed animals, although they judge them very similar to real people and animals (Gelman *et al.*, 1983).

Structure from function inferences

These are based on the "teleological" principle that the typical function of an object (or part of an object) is an explanation for its structure. Young children seem to adopt this teleological stance both for artefacts and for some aspects of the structure of living organisms (Keil, 1994). However, in both domains they are very selective and often wrong in their choice of features. For instance, they may identify artefacts in terms of function yet fail to identify crucial functional affordances (Tversky & Hemenway, 1984).

Biological processes

Specific principles include the assumption that members of a living kind have some internal similarity beyond their common external features and sometimes in conflict with observable features (Becker & Ward, 1991; Gelman & Markman, 1986; Keil, 1986; Massey & Gelman, 1988;). Also, living things are construed as belonging to a set of taxonomic classes (jointly exhaustive of the domain, and mutually exclusive) such that taxonomic proximity predicts non-obvious resemblance. These principles, however, constitute intuitive expectations, not an integrated and explicit "naive biology" [2].

Mentalistic psychology

From the earliest stages of cognitive development, children readily interpret the behaviour of animate beings, particularly persons, as caused by unobservable beliefs and intentions. Children from the age of three assume that thoughts and intentions are immaterial. They also assume that actual states of affairs cause perceptions that cause beliefs that cause intentions. From the age of four they can appreciate the possibility of false beliefs (see e.g. Gopnik, 1993; Perner, 1991; Wellmann, 1990).

Intuitive ontology is not a static system. It only emerges gradually in the course of conceptual development. I chose to stop at what is typical of a five- or six-year-old because the categories and principles will to a large extent be found in a similar form at the end-point of cognitive development. On the whole, intuitive ontological concepts as demonstrated in studies of adult categories, of neural correlates, and of neural pathologies show that further developments refine this ontological system and elaborate upon it, but generally do so in accordance with the fundamental principles described here.

This very short survey highlights the functional complexity of early-developed ontology. Intuitive ontological categorization requires a set of *specific perceptual cues*. These lead to identification of an object both in terms of a *kind* concept and in terms of a set of *ontological categories*. The latter activate *inference engines* that specialize in particular aspects of the object, and (1) stipulate non-obvious properties, (2) stipulate causal connections, and (3) produce specific expectations. Obviously, categories and aspect-specific, theoretical inference engines are related. Some inference engines, to do with biological properties, for instance, are preferentially activated about an object when that object is identified as a member of the ANIMAL or HUMAN categories; if something is identified as an ARTEFACT; neither psychological nor biological inferences are activated in normal circumstances. However, categorical identification and aspect-specific inferencing are functionally separate operations. Some inference engines apply in the same way to different ontological domains, e.g. intuitive physical constraints to animals and plants and humans. Also, some situations require that one switch from one inference engine to another, without changing the ontological identity of the object considered. For instance, the same animal species can be alternatively considered in terms of a species-specific essence (in situations like predation) and in terms of individual psychological features (e.g. pets, hunting dogs).

Intuitive ontology constrains the development of kind concepts. Conceptual development is more than the gradual acquisition or enrichment of foundational theories for ontological domains. During the same developmental period children acquire a vast repertoire of “basic-level” concepts that underpin the pre-schooler’s lexical explosion. It is plausible that developing minds acquire a number of particular concepts by fixing the values of variables postulated by intuitive ontology. Within ontological categories, the operation of inference engines introduces differentiation by specifying the range of features to be expected or stored or compared, thereby providing “respects for similarity”, i.e. constraints on similarity metrics. Take, for instance, concepts of animal kinds. Within the perceptually based identification of ANIMAL, inference engines specify that each exemplar belongs to one exclusive taxonomic class of ANIMAL, that this results in particular “internal” properties, that these are normally linked to a particular visual prototype, that inter-exemplar differences are irrelevant to behavioural dispositions, and so on (Atran, 1996). This is sufficient to generate all the seemingly “theoretical” inferences that count as possessing, e.g. the concept “tiger”. The point applies to artefact concepts as well. In both domains, the operation of ontological categories and inference engines provides a set of variables whose particular values identify a basic category. Obviously, variable-fixing is not all there is to acquiring new conceptual structures. A great deal of associative information is stored in an explicit format. This requires a meta-representational capacity to represent new conceptual slots for complex combinations of concepts. A vast amount of knowledge is represented in this way (Sperber, 1994). This makes it possible, not only to represent explicitly the output of intuitive ontology but also to create concepts that go beyond intuitive expectations or violate them (Boyer, 1998a). Formal and informal instruction result in the constant enrichment of the conceptual base, predominantly though not exclusively through meta-representational mechanism.

General epistemic stance or domain-specific presumptions?

Early-developed presumptions about “underlying structure” of kinds might result from a general epistemic stance or from particular presumptions about each of the ontological domains. In the general stance interpretation, for instance, "psychological essentialism" is characteristic of human conceptual development across domains.

The psychological evidence suggests multiple domain-specific "essentialist inferences", much more restricted in scope than a general epistemic stance. Essentialism in the child and the adult is not one general assumption but a variety of inferential mechanisms, not necessarily applied together or to all domains. Consider, for instance, the difference between PLANTS and ANIMALS. In both domains, we intuitively expect singular instances to be members of classes that are jointly exhaustive and mutually exclusive (Atran & Medin, 1999). However, plants and animals are not construed in the same way at all. For instance, children intuitively assume (1) that members of an animal species have the same "innards", (2) that inner structure is more stable than observable features, and (3) that inner features cause external ones. Now these aspects of essentialist assumptions are less readily applied to plants than to animals. Even in domains where children do apply straightforward, full-blown essentialist assumptions, they do it with remarkable selectivity. To take but one example, children are more willing to use "innate potential" as an explanation for dynamic behaviours (e.g. that an animal makes a particular noise) than for static features (e.g. that an animal has a straight tail). There are many other examples of this selectivity, and of the child's general tendency to assume that only *some* features count as essentially caused, regardless of their frequency in actual instances. Note, also, that children may use essentialist inferences in domains where they do not actually focus on "kinds", but rather on socially defined "pseudo-kinds" like racial and ethnic categories (Hirschfeld, 1996) as well as non-kind collections such as kinship-based social groups (Hirschfeld, 1989). All this leads Hirschfeld and Gelman to conclude that essentialism is neither an exclusively biological phenomenon nor a pervasive cognitive bias (Hirschfeld & Gelman, 1999). It is an inference mechanism that is selectively applied to particular aspects of the ANIMAL, PERSON and even SUBSTANCE domains, and in each domain only to particular features of the objects considered. Children refrain from essentializing certain aspects of natural kinds, even when there is supporting evidence.

This also seems to cast doubt on the notion of a general "conceptual tracking" capacity. Millikan argued that this general capacity comes first, and domain-specific presumptions second: The "practical ability to re-identify a substance … has to be *complemented* with another equally important ability … One … must have some grasp of what kinds of things can be learned" about the substance (Millikan, 1998, p. 58, my emphasis). There is, however, no positive evidence for this two-stage process. Inasmuch as we have evidence for "conceptual tracking", that is, evidence for *particular* kinds of tracking applied to *particular* aspects of *particular* categories of objects. Moreover, there is no evidence that this selectivity appears as a result of tabulating which features are more stable than others. Face-recognition capacities are in place from an early stage *before* the child can perceive that faces change less than clothes (Morton & Johnson, 1991). The link between the facial expressions the child detects and the ones she produces appears *before* the child can detect the effect of her own emotional displays on others (Meltzoff, 1994; Meltzoff & Moore, 1983). The infant detects some phonological features in her language *before* she can tabulate which features are actually more critical than others (Mehler *et al.*, 1988). In general, we have no evidence for general cross-domain same-substance detection capacities *before* the emergence of domain-specific similarity-with-respect-to-X capacities. That "conceptual tracking" is performed in parallel by several independent capacities is supported by the available evidence rather than by the notion that conceptual development "must" proceed in a particular way.

To sum up, then, both Millikan's "same-substance tracking" and Kornblith's "essentialist stance" are on the right track in the sense that the child's performance is *compatible* with "essentialism" and "conceptual tracking". However, that the child's cognitions are compatible with a given epistemological stance does not entail that they are caused by a cognitively implemented version of that stance. The evidence suggests a more complex system of intuitive, domain-specific expectations, that is, a natural metaphysics rather than a natural epistemology.

Does development suggest ontological plasticity?

Churchland has long maintained that conceptual structures display inherent "plasticity". This notion has become the focus of an intense debate. On one side, Jerry Fodor has maintained that activation and functioning of modular structures is impervious to higher-level influence from knowledge or expertise (Fodor, 1988). Churchland takes the opposite stance and observes that major revisions occur in knowledge but also in perception (Churchland, 1988) [3]. If Churchland is right about cognition in general, we should expect ontological concepts and inferences to be "plastic" too as a result of changing cultural circumstances. In this section I list different types of arguments or evidence concerning the potential plasticity of human cognition in general and examine to what extent they apply to intuitive ontology in particular.

Does ontology change in childhood?

Developmental psychologists have been arguing for some time between "continuity" (Eimas, 1994; Spelke, 1988) and "conceptual change" (Carey, 1985) interpretations of development. The continuity thesis states that development boils down to a gradual enrichment of fundamental intuitive principles. By contrast, Carey and other authors have demonstrated important changes in such fundamental properties as WEIGHT in intuitive physics and LIVE BEING in biology, and consider these as every bit as radical as scientific revisions of ontology. Here we must distinguish between two kinds of change. One is the emergence of new principles that cannot be accounted for in terms of previous principles. For example, children around the age of four begin to show some awareness of the behavioural consequences of false belief, which was not previously available. A second type of change is the acquisition of generic knowledge that can override some intuitive principles or refocus their application. Children gradually learn that some aspects of animal behaviour are properly understood in biological rather than psychological terms. The principles of biological understandings are not changed in this process but their domain of application certainly is. Asking whether there is "genuine conceptual change" in cognitive development *in general* blurs this important distinction. There is no disagreement in developmental psychology on the fact that categories and principles of intuitive ontology do mature. However, once these categories and principles are in place, further knowledge does not seem to revise or displace them.

Does training or instruction change intuitive ontology?

As Churchland points out, it often seems possible to retrain one's intuition or acquire what would appear to be new intuitive responses as a result of high-level knowledge. Obviously, the most salient example of *conceptual* training is the systematic acquisition of scientific knowledge. Science challenges the concepts that could be built by straightforward extensions of intuitive expectations. Science has no place for intuitive notions of FORCE in physics or ESSENCE in biology. Scientific psychology challenges intuitive understandings of mental functioning and may one day demonstrate that concepts of "belief" and "intention" are incoherent (Churchland, 1981). So this is where we should see major revisions of intuitive ontology.

As far as evidence can tell, acquiring scientific theories that conflict with intuitive understandings does not usually result in a straightforward *replacement* of those understandings. Knowledge of a theory does not create the kind of intuitive expectations that would be consistent with that theory. Indeed, this lack of congruence was the initial trigger for studies of "intuitive" or "folk-theories" of such domains as physics, statistics, logic or biology. Students trained in logic still tend to give "wrong" (but intuitively salient) answers to simple tests and show robust biases in their evaluation of the confirmation of hypotheses (Wason, 1983; Wason & Johnson-Laird, 1972). Intuitive understandings also provide us with frequency-based assumptions about the probability of events; evaluating the probability of a single event is a particularly difficult task for humans (Gigerenzer, 1991; Gigerenzer & Hoffrage, 1995; Koehler, 1996), which may explain why even "experts" sometimes fail to use Bayesian principles (Tversky & Kahnemann, 1982, see a discussion in Gigerenzer & Murray, 1987, pp. 163–167). In the domain of physics, subjects' expectations about ballistics for instance are not really congruent with Newtonian mechanics and tend to relate force to speed rather than to acceleration (Kaiser *et al.*, 1986; McCloskey, 1983) [4].

Intuitive understandings and scientific ones do not seem to conflict, simply because they do not really meet. This is partly caused by the fact that scientific concepts and intuitive ontology develop through very different routes. In particular, scientific concepts and propositions are invariably acquired in the form of *meta-representational* beliefs. Darwinians, for instance, are not people who just happen to believe that "evolution is caused by blind variation and selective retention", but people who believe that they have that belief and that it is true. Otherwise it would not be possible for them to argue for that view of evolution, evaluate it in the light of evidence, assess its explanatory power, and so on. All these operations require that the belief in variation and retention be explicitly represented: in other words, that the representations in question are meta-representational.

Are there historical–cultural differences?

That different cultures or periods have a different take on broad ontological categories and inference engines would constitute straightforward evidence for flexibility. At first sight, the anthropological record seems to provide evidence for such differences. "Other" people are described as believing that rocks fly in the air, mountains have blood circulation, trees have a soul, spirits fly right through physical obstacles, and so on. (Here "other people" in fact refers to anyone outside a tiny population of scientific rationalists.) While such representations are often taken very seriously by the people concerned, there is no evidence that intuitive ontology differs from one location to another. Developmental and adult evidence, when available, shows categories and developmental schedules similar to those of Western subjects (Walker, 1985). Moreover, it is probably the case that such supernatural representations are parasitic upon intuitive ontology, which (1) provides a background against which violations of intuitive expectations are attention-grabbing and (2) supplements those violations with inferences derived from non-violated principles. For instance, rituals explicitly directed at superhuman agents are founded on tacit intuitive understandings of agency (Lawson & McCauley, 1990). There is now experimental evidence for these processes (Barrett & Keil, 1996).

Turning to less fancy beliefs, there is no evidence for major cultural differences in the categorical distinctions or inferential engines described above. Consider intuitive biology: although we do find cultural differences in the way biological phenomena are explained, the major features of biological essentialism, as described here, are stable across cultures (Atran, 1996). In the domain of intuitive psychological inferences ("theory of mind"), we find a similar pattern. There are cultural variations in explicit constructs like a local theory of personality, of the connections between mind and body, of how minds develop (Heelas & Lock, 1981). These explicit cultural theories consist in an enrichment, never in a revision of the intuitive principles described above. Developmentally, too, one finds unexpected cultural similarity, for instance, in the emergence of "theory of mind" in Western and Pygmy children (Avis & Harris, 1991). Differences of developmental schedule remain within a narrow envelope and lead to the same end-point. Despite otherwise vastly different cultural environments, people seem to use categories and inference engines that are substantially similar.

This also applies to science as a cultural phenomenon. Churchland (1979) has provided a vivid description of a possible future culture based on scientific understandings rather than sloppy and ultimately false common-sense concepts. It is of course difficult to predict the course of cultural evolution, but the past record, inasmuch as it provides a sound basis for induction, suggests a rather different picture. The diffusion of scientific knowledge certainly has effects on generic knowledge in a population. But, as I said above, that's not the same as changing intuitive understandings. Despite more than three centuries of Newtonian physics, the theory fails to govern intuitive expectations. Even people who actually *use* the theory seem to produce routine, spontaneous physical inferences on the basis of the same intuitive principles as untutored subjects the world over.

To sum up, then: curing human history and across cultures, people have been and are faced with very diverse environments (social and natural) as well as different, culturally transmitted explicit accounts of what makes solid things solid, what makes living things live, what constitutes a mind and so on. These differences provide substantial evidence for the "plasticity" of mind, understood as the capacity to entertain a wide variety of concepts and beliefs. However, it is also striking that we find no corresponding differences in intuitive ontology. Cultural input seems to build upon intuitive ontology, to challenge it sometimes, even to build a more powerful ontology in the case of science. But it does not eliminate the intuitive expectations.

"Innate concepts": symbolic and computational descriptions

A crucial problem for developmental psychology is to account for the early emergence of non-trivial ontological distinctions. For instance, infants make distinctions between linguistic intonation and other sounds, between faces and other displays, between inanimate and animate motion, between physical and intentional causation. This is often interpreted as evidence for a rich conceptual repertoire at the initial stage of development (Spelke, 1988). A traditional argument here stems from the poverty of the stimulus. There is not enough information in the auditory environment to tell the infant what is language and what is noise. Similar arguments are invoked to explain early intuitive physics, face-recognition or expectations of social interaction. However, it is difficult to evaluate the plausibility of such arguments in the absence of any account, however speculative, of their implementation. As Churchland points out, the best naturalistic account of what it is for a cognitive system to have a theory is not a classical sentential model but the "state–space semantics" of neural networks (Churchland, 1995). Now networks are not just passive information-storage devices. They go beyond the information given in that the effects of input vectors depend on past activation now stored in the implicit form of synaptic weight patterns. So input poverty arguments may grossly underestimate the amounts of information complex networks can extract from a set of noisy stimuli. That such networks are capable of producing quasi-theoretical inferences is computationally coherent (Rumelhart & Todd, 1992), empirically plausible as a model of development (Elman *et al.*, 1996; Quinn & Johnson, 1997; Thelen & Smith, 1995), and consistent with neuro-physiological development (Quartz & Sejnowski, 1997). Network and dynamic system models would suggest that it is by no means necessary that infants are equipped with rich conceptual structures.

Resolution of this "debate" is hindered by two conceptual ambiguities.

First, discussions of whether certain ontological concepts are or not present at the initial stage are by and large informed by a "descriptionist" account of concepts.

But understanding developmental change may require that we take a leaf out of Millikan's book and construe most concepts, including ontological categories, as (specialized) sameness tracking *skills* rather than descriptive structures (Millikan, 1998).

Second, contending accounts of development differ along not just one but two dimensions. One question is whether it is plausible that infants' cognitive capacities include distinctions at the ontological level. The other is whether developmental change is better captured by high-level symbolic or low-level computational descriptions. As it happens, it seems that answers to these two questions are often correlated in the literature. The resulting situation can be summarized (caricatured perhaps) as this matrix of options:

	No prior ontological distinctions	Some prior ontological distinctions
Symbolic description	1. Classical empiricist models of cognitive development	2. Most models of domain-specificity
Computational description	3. Most connectionist models of cognitive development	4. ???

The pitfalls of classical empiricist accounts are familiar to most developmental psychologists. As a result, the only developmental accounts available either take low-level computational aspects into consideration *and* assume no initial stage ontological distinctions, or else accept such distinctions *and* ignore issues of computational implementation. But there is nothing necessary in these conjunctions. Indeed the evidence seems to require a domain-specific computational account.

To illustrate these two issues, consider for instance the evidence for animacy detection. Infants detect the difference between animate and inanimate motion, as preferential looking and dishabituation protocols have shown. To most developmental psychologists it would seem contrived to suggest that the distinction between physically and intentionally caused motion is extracted from other information. There seems to be no available lower-level information that would correlate with this difference between types of motion and orient a cognitive system toward the motion cues.

Considering first the issue of concepts-as-descriptions, is this evidence for the fact that the infant possesses *the* "animate agent' concept? The infant's distinction does amount to an ontological distinction in that it picks out—and applies different inferences to—two different kinds of stuff in the world. The kinematic cues used by infants are *among* those used by adults to distinguish between agents and non-agents in their environment. The rather crude cues used by infants really are evidence for very early biases, but they will undergo *continuous* complexification with development. If we adopt a descriptionist account of concepts, we may wonder at what point the child's ontological distinctions actually amount to possessing the "animate agent" concept. If we construe concepts as skills, this question does not make much sense. Concepts mature in the sense that a cognitive system becomes better skilled at detecting sameness and avoiding false positives. As it happens, this calibration process is precisely documented in the domain of animacy detection: (1) infants are sensitive to *actual* kinematic differences, (2) pre-schoolers are apt to make inferences about *possible* motion from static cues (Gelman *et al.*, 1983), and (3) older children and adults use all this plus information about internal structure and biological functioning to decide what can and what cannot move intentionally (Carey, 1988). Whether infants do or do not have *the* concept of "animate agent" may well be a red herring introduced by descriptionist accounts of concepts.

Turning to the second ambiguity, this example also shows that accepting initial ontological presumptions does not confine developmental accounts to symbolic models that ignore implementation issues. Specific neural capacities underpin the distinction between animate and inanimate causation (or if one prefers a less metaphysical description, the typical kinematic difference between these two kinds of events). We now have a precise description of the relevant computational aspects of perception (Ahlstrom *et al.*, 1997; Bellefeuille & Faubert, 1998; Heptulla-Chatterjee *et al.*, 1996) as well as of the neural structures involved (Bonda *et al.*, 1996). A neuro-computational account of developmental processes for this ontological distinction may be available very soon (see Johnston *et al.*, 1999, for an attempt).

One could make similar points for the domain of face-recognition and own-language distinction. Infants distinguish faces from other kinds of visual displays on the basis of specific cues the infant's visual system is biased to detect. This has prompted the question, whether these cues are evidence for a "face concept" (see Morton *et al.*, 1990; Kleiner, 1993, for different interpretations of the evidence). Again, the question makes sense only in a descriptionist account of the "face" concept. That the bias is mostly driven by particular low-level visual features that happen to be present in typical faces does not rule out that it constitutes a (primitive) face detection skill. Also, taking these low-level biases into consideration takes us closer to a computational account of initial dispositions for face-recognition. Finally, infants seem predisposed to detect their mother tongue (literally their mother's language). But they do this on the basis of very rough characteristics of linguistic information that are certainly insufficient to make accurate distinctions between languages (Jusczyk *et al.*, 1993). Again, the infant's skills in the domain are at a primitive stage. This does not entail that they are not specialized skills.

In general, most early ontological distinctions seem to have the following characteristics: (1) they are continuous with some adult distinctions (which is why they are often described as initial evidence for the adult concept); yet (2) they are much less efficient than more mature concepts in terms of sameness tracking; (3) they allow rich information pick-up given an initial bias; and (4) this bias can be easily modelled in terms of network dynamics. This last point is particularly important for biologically plausible developmental models. A strong argument against attributing a rich conceptual repertoire to infants is that such attributions are in general computationally intractable (Thelen & Smith, 1995). That is, there is no way to determine what computational processes would underpin the precise behavioural differences observed. By contrast, we see here that ontological distinctions may be founded on low-level detection that may be easily modelled in terms of network initial topology and biases. This interpretation of early-developed categories and inference engines is consistent with Churchland's notion that such structures are "theoretical" *and* reduce to dynamic prototype activation. However, it also suggests that some biases in the dynamics of prototype building are so entrenched that we should not be surprised that their long-term effects in terms of conceptual dispositions are extremely difficult to override.

Intuitive ontology as an evolved natural metaphysics

Developmental evidence suggests a mentally represented ontological repertoire with the following characteristics: (1) it consists of metaphysical presumptions about particular domains of reality rather than a generalized epistemic stance; (2) these presumptions are intuitive (best accessed via implicit measures) *and* theoretical (computing non-trivial inferences); (3) some of the presumptions are available at very early stages of cognitive development, indeed in infancy; (4) differences in cultural environment or training do not seem to effect profound changes in intuitively available categories and principles; and (5) intuitive ontology is far from optimal. Intuitive ontology just does not deliver any fast, intuitive inferential responses for some categories of objects and aspects of experience [5], and it often includes *ad hoc* "buck-stopping" assumptions or causal inferences without a description of causal mechanism [6].

These features are not too surprising if we consider intuitive ontological distinctions and inference engines as part of the cognitive equipment typical of the species, as a result of evolution by natural selection (Boyer, 1998a,b; Cosmides & Tooby, 1994). This is not the place to discuss evolutionary constraints on this or that particular capacity. However, it is relevant to "naturalized epistemology" arguments to consider some *general* properties one might expect from a biologically implemented, evolutionarily constrained cognitive system.

First, it seems plausible from an evolutionary standpoint that ontological dispositions put developing minds on the "right" track by providing them with a (relatively) sound ontology rather than a (relatively) sound epistemology. Having the right ontology, as long as your epistemology is not dramatically unsound, puts you on the right track without much cognitive expense. Having a feeble ontology and the right epistemology, on the other hand, could be disastrous. An organism equipped in this way would waste much cognitive power discovering valid generalizations about its environment that are not of any consequence for gene transmission.

Second, there are obvious computational advantages to a set of fast identification and inference processes that deliver *intuitive* expectations. Independent evolutionary considerations support the notion of cognitive evolution as an accretion and complexification of specialized problem-solving algorithms rather than as the development of an integrated general intelligence. This is supported both by inter-specific comparisons and by the computational cost and negative consequences of developing an unbiased, domain-general correlation-detector (Cosmides & Tooby, 1994) [7]. Also, the presence of such biases makes it less surprising that intuitive ontology is sub-optimal, being only as good as it takes to provide an adaptive edge over previous versions.

Third, empirically substantiated evolutionary considerations may better explain why ontological presumptions develop early and initially appear as rough biases that only partially map onto adult conceptual distinctions. The fact that a great amount of developmentally crucial information is stored in a species' typical environment implies that minimal but definite biases are sufficient to trigger adequate development. This may help us better understand the relevance of neural plasticity at the level of neural structures and neural connectivity, which is sometimes considered as evidence for the *general* plasticity of cognitive functioning. However, it is a general principle of ontogenetic development that low-level plasticity is recruited in the service of higher-level predictable design. Indeed, it makes such predictable form more likely through calibration.

Fourth, an evolutionary framework accounts for the stubborn aspects of intuitive ontology without implying that they are present in full-blown form at birth or inevitable in human development. The fact that definite biases lead to a particular ontology is a reliable feature of development under *normal* conditions. Children brought up in zero gravity or on some planet with no animals may well develop different ontological concepts and a different intuitive physics or intuitive biology. What makes some conditions "normal" in this account is not that they are inevitable or even common. It is only that such conditions prevailed in the environment where genes that control for neural structure and development were selected.

To conclude, the evidence suggests that concept acquisition most likely does not require a sound epistemic attitude, but rather some robust metaphysical prejudices, supporting quasi-theoretical inference processes that are in principle defeasible but in practice extremely stubborn. If such prejudices are a plausible outcome of evolution through natural selection, "natural metaphysics" is the outcome of a species' particular history and of the species-typical needs of its members, and is therefore natural only relative to that particular species. This certainly bodes ill for an extreme version of the "naturalizing epistemology" project, for we cannot deduce metaphysical validity from what is adaptive for a particular kind of organism with particular requirements. On the other hand, cognitive and evolutionary considerations may lead to a more "natural" version of the project, providing an empirically motivated account of the relative "fit" between our categories and some natural regularities.

Notes

[1] This can be seen as the consequence of an implicit principle that nouns probably refer to *kinds*, and "shape is often, though not always, a good source of information about object kind" (Soja *et al.*, 1992, p. 102). As a consequence, children are more likely to override shape similarity for living kinds than for other ontological categories (Becker & Ward, 1991).

[2] The property of being "alive" is poorly understood before the age of 10, and then mostly on the basis of culturally transmitted information (Carey, 1985). Also, young children explain some characteristics of animals either in terms of psychological dispositions (Inagaki & Hatano, 1987) or mechanistically (Au & Romo, 1998).

[3] A trained musician can "hear" in auditory imagery the change induced by adding a B-flat sound to

a C-major triad, for instance (Churchland, 1995); astronomical competence can make one "see" the planets orbiting on the ecliptic plane (Churchland, 1979, p. 32ff).

[4] It is perhaps the case that the so-called "experts" who participate in such experiments are not trained enough; perhaps "real" experts would perform differently. This, however, only strengthens the point that reeducating intuitions, if it happens at all, seems to require massive instructional effort.

[5] This is clear in children's expectations about plants, for instance. It is only gradually that the child assimilates PLANTS in a broad domain of LIVING THINGS. Such "holes" in the ontological landscape can be found in the adult too. Problems to do with the shape of the Earth are not the object of any intuitive expectation. They are handled by explicit, socially transmitted, meta-represented knowledge.

[6] People can routinely explain the behaviour of animals in terms of essential properties, without necessarily having a description of what these inner properties consist of or how they govern behaviour. A species has its "essence" and that's that. Scientific theories typically challenge this kind of "buck-stopping" and try to describe the processes whereby underlying entities like FORCE or ESSENCE cause observable properties. Doing this, science often shows that the postulates were vacuous or entirely *ad hoc*.

[7] This does not mean that an evolved architecture excludes general capacities, but that the latter probably ride piggyback on a host of specialized capacities and often handle meta-representations of what intuitive ontology delivers (Mithen, 1996; Rozin, 1976).

References

AHLSTROM, V., BLAKE, R. & AHLSTROM, U. (1997). Perception of biological motion. *Perception*, *26*, 1539–1548.

ANTELL, E. & KEATING, D.P. (1983). Perception of numerical invariance in neonates. *Child Development*, *54*, 695–701.

ATRAN, S. (1989). Basic conceptual domains. *Mind and Language*, *4*, 5–16.

ATRAN, S. (1996). Modes of thinking about living kinds: science, symbolism and common sense. In D. OLSON & N. TERRANCE (Eds) *Modes of thought: explorations in culture and cognition*. New York: Cambridge University Press.

ATRAN, S. & MEDIN, D. (Eds). (1999). *Folkbiology*. Cambridge, MA: MIT Press.

AU, T.K.F. & ROMO, L. (1998). Mechanical causality in children's folkbiology. In S. ATRAN & D. MEDIN (Eds) *Folkbiology*. Cambridge, MA: MIT Press.

AVIS, M. & HARRIS, P. (1991). Belief–desire reasoning among Baka children: evidence for a universal conception of mind. *Child Development*, *62*, 460–467.

BAILLARGEON, R. (1987). Young infants' reasoning about the physical and spatial characteristics of a hidden object. *Cognitive Development*, *2*, 179–200.

BAILLARGEON, R. & HANKO-SUMMERS, S. (1990). Is the top object adequately supported by the bottom object? Young infants' understanding of support relations. *Cognitive Development*, *5*, 29–53.

BARRETT, J.L. & KEIL, F.C. (1996). Conceptualizing a non-natural entity: anthropomorphism in God concepts. *Cognitive Psychology*, *31*, 219–247.

BECKER, A.H. & WARD, T.B. (1991). Children's use of shape in extending novel labels to animate objects: identity versus postural change. *Developmental Psychology*, *6*, 3–16.

BELLEFEUILLE, A. & FAUBERT, J. (1998). Independence of contour and biological-motion cues for motion-defined animal shapes. *Perception*, *27*, 225–235.

BONDA, E., PETRIDES, M., OSTRY, D. & EVANS, A. (1996). Specific involvement of human parietal systems and the amygdala in the perception of biological motion. *Journal of Neuroscience*, *16*, 3737–3744.

BOYER, P. (1998a). Cognitive tracks of cultural inheritance: how evolved intuitive ontology governs cultural transmission. *American Anthropologist*, *100*, 876–889.

Boyer, P. (1998b). If "tracking" is category-specific a "common structure" may be redundant, commentary on R. Millikan's article. *Behavioral and Brain Sciences*, *21*, 67–68.

Bullock, M. (1985). Animism in childhood thinking: a new look at an old question. *Developmental Psychology*, *21*, 217–225.

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: MIT Press.

Carey, S. (1988). Conceptual differences between children and adults. *Mind and Language*, *3*, 167–181.

Churchland, P.M. (1979). *Scientific realism and the plasticity of mind.* Cambridge: Cambridge University Press.

Churchland, P.M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, *78*, 67–90.

Churchland, P.M. (1988). Perceptual plasticity and theoretical neutrality. *Philosophy of Science*, *55*, 167–187.

Churchland, P.M. (1989). *The neurocomputational perspective.* Cambridge, MA: MIT Press.

Churchland, P.M. (1995). *The engine of reason, the seat of the soul*. Cambridge, MA: MIT Press.

Cosmides, L. & Tooby, J. (1994). Origins of domain-specificity: the evolution of functional organization. In L.A. Hirschfeld & S. Gelman (Eds) *Mapping the mind: domain-specificity in culture and cognition*. New York: Cambridge University Press.

Damasio, A.R. (1990). Category-related recognition defects as a clue to the neural substrates of knowledge. *Trends in Neuroscience*, *3*, 95–98.

Eimas, P.D. (1994). Categorization in early infancy and the continuity of development. *Cognition*, *50*, 83–93.

Elman, J.L., Bates, E.A., Johnson, M.H., Karmiloff-Smith, A., Parisi, D. & Plunkett, T. (1996). *Rethinking innateness: a connectionist perspective.*

Faulkender, P.J. (1974). Genelized habituation of concept stimuli in toddlers. *Child Development*, *45*, 1002–1010.

Fodor, J.A. (1988). A reply to Churchland's "Perceptual plasticity and theoretical neutrality'. *Philosophy of Science*, *55*, 188–198.

Gallistel, C.R. & Gelman, R. (1992). Preverbal and verbal counting and computation, *Cognition*, *44*, 79–106.

Gelman, R. (1990). First principles organize attention and learning about relevant data: number and the animate–inanimate distinction as examples. *Cognitive Science*, *14*, 79–106.

Gelman, R., Spelke, E. & Meck, E. (1983). What pre-schoolers know about animate and inanimate objects. In D.S. Rogers (Ed.) *The acquisition of symbolic skills.* London: Plenum.

Gelman, S. & Markman, E. (1986). Categories and induction in young children. *Cognition*, *23*, 183–209.

Gelman, S.A., Gottfried, G.M. & Coley, J. (1994). Essentialist beliefs in children: the acquisition of concepts and theories. In L.A. Hirschfeld & S.A. Gelman (Eds) *Mapping the mind: domain-specificity in cognition and culture*. Cambridge: Cambridge University Press.

Gigerenzer, G. (1991). How to make cognitive illusions disappear: beyond "heuristics and biases". *European Review of Social Psychology*, *2*, 83–115.

Gigerenzer, G. & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review*, *102*, 684–704.

Gigerenzer, G. & Murray, D. (1987). *Cognition as intuitive statistics.* Hillsdale, NJ: Erlbaum.

Gopnik, A. (1993). How we know our minds: the illusion of first-person knowledge of intentionality. *Brain and Behavioral Sciences*, *16*, 1–14.

Gopnik, A. & Wellmann, H. (1994). The theory theory. In L.A. Hirschfeld & S.A. Gelman (Eds) *Mapping the mind: domain-specificity in cognition and culture*. New York: Cambridge University Press.

Heelas, P. & Lock, A. (1981). *Indigenous psychologies. The anthropology of the self.* New York: Academic Press.

Heptulla-Chatterjee, S., Freyd, J.J. & Shiffrar, M. (1996). Configural processing in the perception of apparent biological motion. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 916–929.

HIRSCHFELD, L.A. (1989). Rethinking the acquisition of kinship terms. *International Journal of Behavioral Development*, *12*, 541–68.

HIRSCHFELD, L.A. (1996). *Race in the making: cognition, culture and the child's construction of human kinds*. Cambridge, MA: MIT Press.

HIRSCHFELD, L.A. & GELMAN, S.A. (Eds) (1994). *Mapping the mind: domain-specificity in culture and cognition*. New York: Cambridge University Press.

HIRSCHFELD, L.A. & GELMAN, S.A. (1999). Essentialism. In S. ATRAN & M. D. (Eds) *Folk-biology*. Cambridge, MA: MIT Press.

INAGAKI, K. & HATANO, G. (1987). Young children's spontaneous personification as analogy. *Child Development*, *58*, 1013–1020.

JOHNSTON, A., MCOWAN, P.W. & BENTON, C.P. (1999). Robust velocity computation from a biologically motivated model of motion perception. *Proceedings of the Royal Society of London B Biological Sciences*, *266*, 509–518.

JUSCZYK, P.W., FRIEDERICI, A.D., WESSELS, J., SVENKERUD, V.Y. & JUSCZYK, A.M. (1993). Infants's recognition of foreign versus native language words. *Journal of Memory and Language*, *32*, 402–420.

KAISER, M.K., JONIDES, J. & ALEXANDER, J. (1986). Intuitive reasoning about abstract and familiar physics problems. *Memory and Cognition*, *14*, 308–312.

KEIL, F.C. (1986). The acquisition of natural kind and artefact terms. In A.W.D. MARRAR (Ed.) *Conceptual change*. Norwood, NJ: Ablex.

KEIL, F.C. (1994). The birth and nurturance of concepts by domains: the origins of concepts of living things. In L.A. HIRSCHFELD & S.A. GELMAN (Eds) *Mapping the mind: domain-specificity in cognition and culture*. New York: Cambridge University Press.

KLEINER, K.A. (1993). Specific vs. non-specific face recognition device. In B. DE BOYSSON-BARDIES, S. DE SCHONEN, P. JUSCZYK, P. MACNEILAGE & J. MORTON (Eds) *Developmental neuro-cognition: speech and face processing in the first year of life*. Dordrecht: Kluwer Academic Press.

KOEHLER, J.J. (1996). The base rate fallacy reconsidered: descriptive, normative and methodological challenges. *Behavioral and Brain Sciences*, *19*, 1–53.

KORNBLITH, H. (1993). *Inductive inference and its natural ground*. Cambridge, MA: MIT Press.

LANDAU, B., SMITH, L.B. & JONES, S.S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, *3*, 299–321.

LAWSON, E.T. & MCCAULEY, R.N. (1990). *Rethinking religion: connecting cognition and culture*. Cambridge: Cambridge University Press.

LESLIE, A. (1979). *The representation of perceived causal connection*. Unpublished PhD Philosophy thesis, Oxford University.

LESLIE, A. (1988). The necessity of illusion: perception and thought in infancy. In L. WEIZKRANTZ (Ed.) *Thought without language*. Oxford: Clarendon Press.

MANDLER, J., BAUER, P. & MCDONOUGH, L. (1991). Separating the sheep from the goats: differentiating global categories. *Cognitive Psychology*, *23*, 263–298.

MARTIN, A., WIGGS, C.L., UNGERLEIDER, L.G. & HAXBY, J.V. (1996). Neural correlates of category-specific knowledge. *Nature*, *379*, 649–652.

MASSEY, C. & GELMAN, R. (1988). Pre-schoolers' ability to decide whether pictured unfamiliar objects can move themselves. *Developmental Psychology*, *24*, 307–317.

MCCLOSKEY, M. (1983). Naive theories of motion. In D. GENTNER & A.L. STEVENS, (Eds) *Mental models*. Hillsdale, NJ: Lawrence Erlbaum.

MEHLER, J., JUSCZYK, P., LAMBERTZ, G., HALSTED, N., BERTONCINI, J. & AMIEL-TISON, C. (1988). A precursor of language acquisition in young infants. *Cognition*, *29*, 143–178.

MELTZOFF, A. (1994). Imitation, memory, and the representation of persons. *Infant Behavior and Development*, *17*, 83–99.

MELTZOFF, A. & MOORE, M.K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, *54*, 702–709.

MILLIKAN, R.G. (1998). A common structure for concepts of individuals, stuffs and real kinds: more mama, more milk, more mouse. *Behavioral and Brain Sciences*, *21*, 55–100.

MITHEN, S. (1996). *The prehistory of the mind*. London: Thames & Hudson.

Morton, J. & Johnson, M. (1991). CONSPEC and CONLERN: a two-process theory of infant face-recognition. *Psychological Review, 98*, 164–181.

Morton, J., Johnson, M.H. & Maurer, D. (1990). On the reasons for newborn's responses to faces. *Infant Behavioral Development*, *13*, 99–103.

Nelson, K. (1973). Some evidence for the cognitive primacy of categorization and its functional basis. *Merrill–Palmer Quarterly*, *19*, 21–40.

Perner, J. (1991). *Understanding the representational mind.* Cambridge, MA: MIT Press.

Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition*, *36*, 1–16.

Premack, D. & James-Premack, A. (1995). Intention as psychological cause. In D. Sperber, D. Premack & A. James-Premack (Eds) *Causal cognition. A multidisciplinary debate.* Oxford: Clarendon Press.

Quartz, S.R. & Sejnowski, T.J. (1997). The neural basis of cognitive development: a constructivist manifesto. *Behavioral and Brain Sciences*, *20*, 537–596.

Quine, W.V.O. (1969). *Ontological relativity and other essays*. New York: Columbia University Press.

Quinn, P.C. & Johnson, M.H. (1997). The emergence of perceptual category representations in young infants: a connectionist analysis. *Journal of Experimental Child Psychology*, *66*, 236–263.

Richards, D.D. & Siegler, R.S. (1986). Childen's understanding of the attributes of life. *Journal of Experimental Child Psychology*, *42*, 1–22.

Ross, G. (1980). Categorization in 1- to 2-year olds. *Developmental Psychology*, *16*, 391–396.

Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. In J.M. Sprague & A.N. Epstein (Eds) *Progress in psychobiology and physiological psychology*. New York: Academic Press.

Rumelhart, D.E. & Todd, P.M. (1992). Learning and connectionist representations. In D.E. Meyer & S. Kornblum (Eds) *Attention and performance XIV: synergies in experimental psychology, artificial intelligence and cognitive neuroscience.* Cambridge, MA: MIT Press.

Sartori, G. & Job, R. (1988). The oyster with four legs: a neuropsychological study on the interaction of visual and semantic information. *Cognitive Neuropsychology*, *5*, 105–132.

Soja, N.N., Carey, S. & Spelke, E.S. (1992). Perception, ontology and word-meaning. *Cognition*, *45*, 101–107.

Spelke, E.S. (1988). The origins of physical knowledge. In L. Weizkrantz (Ed.) *Thought without language*. Oxford: Oxford University Press.

Spelke, E.S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29–56.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L.A. Hirschfeld & S.A. Gelman (Eds) *Mapping the mind: domain-specificity in cognition and culture*. New York: Cambridge University Press.

Starkey, D. (1981). The origins of concept-formation: object sorting and object preference in early infancy. *Child Development*, *52*, 489–497.

Starkey, P., Spelke, E.S. & Gelman, R. (1990). Numerical abstraction by human infants. *Cognition*, *36*, 97–127.

Thelen, E. & Smith, L.B. (1995). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: MIT Press.

Tversky, A. & Kahnemann, D. (1982). Evidential impact of base rates. In D. Kahnemann, P. Slovic & A. Tversky (Eds) *Jugdments under uncertainty: heuristics and biases.* Cambridge: Cambridge University Press.

Tversky, B. & Hemenway, K. (1984). Objects, parts and categories. *Journal of Experimental Psychology: General*, *113*, 169–193.

Walker, S.J. (1985). *Atimodemo. Semantic conceptual development among the Yoruba*. Unpublished PhD dissertation, Cornell University.

Wason, P.C. (1983). Realism and rationality in the selection task. In J.S.B.T. Evans (Ed.) *Thinking and reasoning: psychological approaches*. London: Routledge, Chapman & Hall.

Wason, P.C. & Johnson-Laird, P.N. (1972). *The psychology of reasoning: structure and content*. London: Batsford.

Wellmann, H. (1990). *The child's theory of mind.* Cambridge, MA: MIT Press.

WELLMANN, H. & GELMAN, S.A. (1992). cognitive development: foundational theories of core domains. *Annual Review of Psychology*, *43*, 337–375.

WYNN, K. (1992). Evidence against empiricist accounts of the origins of numerical knowledge. *Mind and Language*, *7*, 315–332.

YOUNGER, B. (1985). The segregation of items into categories by ten-month-old infants. *Child Development*, *56*, 1574–1583.

YOUNGER, B. (1990). Infant categorization: memory for category-level and specific item information. *Journal of Experimental Child Psychology*, *50*, 131–155.