Draft—Please do not quote or discuss in print. Critical comments most welcome!

THE ORIGIN OF CONCEPTS

Susan Carey

For Fodor/Peacocke Class, relevant for the discussion: Chapter 3 is the most important, Chapter 2 informs Chapter 3 and Chapter 1 simply gives an overview of the book from which these chapters come

Chapter 3: pp 1-13, up to section 3.9; 17-28, sections 3.10 to 3.16
Chapter 2: pp 1-19, up to section 2.10; 23-28, beginning with section 2.11
Chapter 1: pp 2-7

Chapter 1. Some preliminaries

Human beings, alone among animals, create rich conceptual understanding, representing concepts such as *evolution*, *electron*, *cancer*, *infinity*, *galaxy*… What accounts for the human capacity for conceptual representations? Any concept we may focus on has two ultimate sources—evolution and culture. Some concepts, such as *object* and *number,* arise in part through natural selection, in the course of hominid evolution or in some ancestor common to the apes, or to the primates, or even earlier. Other concepts, such as *kayak fraction*, and *gene*, are cultural constructions. Humans create complex artifacts, and religious, political, and scientific institutions, and in turn create new representational resources.

Although it is obvious that we must look both to evolutionary and historical processes to account for the origins of concepts, the problem for cognitive science is to provide a precise, explanatory, account of the origin of any particular concept in which we may be interested. When it comes to the concept of *integer,* or of *animal*, for examples, what are the relevant innate representational resources bequeathed to us by evolution? In creating such concepts, must we go beyond innately given representations? In what ways, and by what processes, do individuals and groups of individuals construct new representational resources? How is knowledge culturally constructed and maintained? These questions must be answered case by case--there is no reason, at the outset, to expect the answers for spatial knowledge, for example, to be the same as the answers for mathematical or biological knowledge or for knowledge of language or chess.

This book develops five case studies in detail: the concepts *object*, *intentional agent, number, animal* and *living thing,* and *matter, weight* and *density.* In the course of exploring these cases many others are touched, including *artifact, contagion*, and *causality.* But what really matters to me is not the cases (although I do admit finding each one intrinsically fascinating), but the lessons to be drawn from them. My goal in this book, as in the Nicod lectures, is to demonstrate that the disciplines of cognitive science now have the empirical and theoretical tools to turn age-old philosophical dilemmas into relatively straightforward scientific problems. I shall illustrate the progress science has made in resolving debates about the existence, nature, content and format of innate knowledge, about the thesis that conceptual resources are continuous throughout the life space, about the nature of concepts and intuitive theories, about the distinction between conceptual change and belief revision, and about controversies concerning the relations between language and thought.

1.1 Concepts and mental representations

Concepts are units of thought, the constituents of beliefs and theories, and those that interest me here are roughly the grain of single lexical items. Indeed, the meanings of words are paradigm examples of concepts. I am concerned with the mental representation of concepts; I use phrases such as “the infant’s concept *animal*” to mean the infant’s representation of animals. I assume representations are states of the nervous system that have *content,* that *refer* to concrete or abstract entities, to properties, to events. I do not attempt a philosophical analysis of mental representations; I will not try to say how it is that some states of the nervous system have symbolic content. Such a theory would explain how the extension of a given representation is determined, as well as providing a computational account of how that representation fulfills its particular inferential role, how it functions in thought.[1] Here I merely assume that such a theory will be forthcoming. In the pages to come, I work backwards from behavioral evidence for some concept's extension and inferential role to characterize that concept's content and to specify something of its nature and format of representation.

There are many different types of mental representations and one challenge to cognitive science is to find the principled distinctions among them. Different types of representations may well have theoretically important differences in origins, developmental trajectories, types of conceptual roles, and relations to their extensions. Also, some theories of conceptual development posit shifts in kinds of mental representations available to children of different ages—from a perceptual similarity space to natural kind concepts (Quine, 1977), from sensori-motor to symbolic representations (Piaget, 1954), from implicit to explicit representations (Karmiloff-Smith, 1990), for examples. Such theories depend, of course, on defensible distinctions among types of mental representations.

I will join forces with the many writers who draw a distinction between perceptual representations, on the one hand, and conceptual representations, on the other. Chapter 2 examines thesis that infants begin with perceptual representations and only construct conceptual representations later in development. Differentiating the perceptual from the conceptual is difficult. There are probably many different distinctions at work here, and most are probably ends of continua rather than categorical. An intuitive characterization of perceptual representations as what things in the world *look like, sound like, feel like, taste like*, contrasts these with conceptual representations as what things is the world *are.*. Distinctive properties of perceptual representations include, first of all, that their extensions are fixed by virtue of innate, modular, sensory input analyzers. There are innate shape analyzers, phoneme detectors, color detectors, motion detectors, and so forth. That representations of red have the content *red* is ensured by evolution, by how color vision works. Second, perceptual representations have very little in the way of inferential role. Almost nothing else follows from the fact that something is red. Third, and related to the above two points, perceptual representations are inferentially close the output of sensori-analyzers. Consider the difference between the representation of *red* or *loud*, on the one hand, and the representation of *electron* or *life*, on the other. Although we certainly can sometimes identify electrons or living things perceptual evidence, there is a long inferential chain between a path in a cloud chamber to the presence of an electron, or from what a bacteria colony on a petri dish looks like to the fact that it contains living things.

Natural kind concepts, paradigm conceptual representations, are at the other end of the continuum, contrasting with perceptual representations in all three respects. There are no innate input analyzers for tigers or electrons, natural kind concepts have rich conceptual roles, and there is a long inferential chain between the perceptible properties of natural kinds and the content of concepts of natural kinds. According to the Kripke/Putnam (Kripke, 1972; Putnam, 1975) analyses of natural kind concepts, their extensions are fixed not by the mind but by some social process of ostensive definition

---

[1] An excellent overview of competing accounts of concepts with the philosophical and psychological literatures is provided in the collection of classic papers assembled by Margolis and Laurence (1999). See also their own excellent introductory essay.

and by the essential nature (a metaphysical matter, not an epistemelogical one) of the entities so dubbed. The discovery of the extension of *gold* or of *wolf* is a matter for science, not for philosophy, linguistics, or psychology. As for the psychology of natural kind concepts, they fall under the assumption of "psychological essentialism," (Medin and Ortony, 1989). It is a fact about our mind that we assume (usually correctly, as it turns out, but it needn't be) that individuals of a given natural kind have hidden essences which both determine their kind and their surface properties. Often we have no fleshed-out guess as to a kind's essential properties.

A natural kind concept's features fall along a continuum from core to periphery, a continuum determined by explanatory depth (Ahn, Kim, Lassaline & Dennis, 2000; Keil, 1989). Its core, its essence, consists of its inferentially deepest features, and for natural kinds, these are its causally deepest features. Thus, the analysis of concepts of natural kinds is deeply intertwined with the analysis of the conceptual structures that represent causal/explanatory knowledge: intuitive theories.

Some writers deny a principled distinction between perceptual representations and conceptual representations, claiming that all mental representations at root perceptual representations (e.g., Thelan, Schoner, Scheir & Smith, 2000). Others (e.g., Quine, 1977; Piaget, 1954) grant the distinction and believe that conceptual development in the first few years of life involves a transition from perceptual representations alone to a representational repetoire that contains both types. These positions are considered in Chapter 2.

This book's first major thesis is that there is a third type of conceptual structure, called "core knowledge" by Spelke, that differs systematically from both perceptual domains of representation and from theoretical conceptual knowledge. I shall argue that core knowledge is the developmental foundation of human conceptual understanding. Like perceptual domains, the entities in core domains of knowledge are identified by modular innate perceptual input devices, but the representations are conceptual, not perceptual. Unlike perceptual representations, they have relatively rich inferential roles in thought, and there is a longer inferential chain between the output of sensori-analyzers and the content of the representations that articulate core knowledge. However the conceptual role of the concepts that articulate core knowledge is vastly less rich than that of the concepts embedded in intuitive theories, and the inferential depth between perceptual properties and the content of core knowledge is vastly less than in the case of intuitive theories. Finally, knowledge acquisition in core domains is supported by innate domain-specific learning devices, whereas that in intuitive theories is not. Chapters 3 and 4 characterize core knowledge more fully and summarize evidence for human core knowledge of objects, contact causality, intentional causality, number and emotion.

1.2. On Conceptual Development

I take it all would agree that any theory of conceptual development must specify the stock of innate representations the mechanisms (both maturational and learning mechanisms) that underlie change. Framing the problem this way leaves all the room in the world for cantankerous disagreement. Are the initial representational primitives limited to sense data, as the British empiricists would have had it (e.g., Locke, 1690/1975), to sensori-motor schemes, as Piaget (1954) would have had it, to a perceptual similarity space, as Quine (1960, 1977) would have had it, or is it theories all the way down, as Gopnik and Meltzoff (1997) would have it? Insofar as learning underlies conceptual development, are the learning mechanisms domain general or specifically tailored to particular conceptual content? If domain-specific, what are the domains, and what are the mechanisms of learning within each? If domain general, are they powerful pattern abstractors, such as those well modeled in connectionist systems, determinate computational algorithms formulated over symbolic machinery, or bootstrapping processes that draw on non-determinate processes such as analogical mapping, inference to best explanation, and inductive guesses?

To some extent, the proposals for the initial stock of representational primitives and for the types of learning mechanisms that underlie cognitive development are logically independent. Typically, though, researchers who believe that the initial state of the baby consists of perceptual or sensori-motor representations also believe that domain-general learning mechanisms of various sorts subserve conceptual development. This position (innate perceptual representations/domain general learning mechanisms) is discussed in Chapter 2. A second view (innate conceptual representations/domain general learning mechanisms) held most forcefully by Gopnik and Meltzoff (1977), is that the initial stock of representational resources includes conceptual representations embedded in intuitive theories, and that development is driven by domain general bootstrapping mechanisms of the sort that underlie theory development in science. Chapter 3 contrasts this position with a third position (innate conceptual representations/domain specific learning mechanisms; the core knowledge position), which holds that the initial stock of representational resources includes conceptual representations embedded in core domains, and that development is driven, at least in part, by domains specific learning mechanisms.

As debates about conceptual development become heated, a simple fact is often overlooked: These possibilities are not mutually exclusive, unless each is framed with an "only." It is possible that there are innately specified capacities to form perceptual representations (of red, of motion) and conceptual representations (of people, of objects, of causes) and it is possible that there are many different innately specified learning mechanisms. Building rich and accurate representations of the physical, social, and biological worlds is so important, to humans especially, but to all animals really, that many distinct representational and learning mechanisms are likely to have been selected for. As Gallistel (2000) commented, speaking of the proposal that there is a single domain general learning mechanism, "From a biological perspective, this assumption is equivalent to assuming that there is a general purpose sensory organ that solves the problem of sensing."

A variety of considerations lead individual scientists to favor one or another picture of the initial state and of the processes that yield adult conceptual capacities. Some reduce to matters of scientific taste, such as a preference for a picture that assumes the least built in, in order to explore how much of the adult state could be accounted for under any particular minimal assumptions (e.g., perceptual primitives, associationist learning mechanisms). But scientific taste must not be promoted to principles of theory choice with epistemic status, as when it is claimed that such a picture is more parsimonious than alternative pictures and thus should be considered right until proven wrong.

Other canons of scientific taste lead to different pictures. For example, scientific taste leads some scientists to bet that human learning processes will be continuous with those of other animals, an assumption that generates an expectation of highly domain specific and structured learning mechanisms. Gallistel, Brown, Carey, Gelman & Keil (1991) provided myriad examples of specialized learning mechanisms in other animals—ranging from passerine song learning, to vervet learning to identify predators, to spatial learning in rats, to migratory birds learning the azimuth of the night sky. In each of these cases, innately specified representations and domain specific learning processes are essential parts of the mechanisms that achieve the adult state.

1.3. An example of animal core knowledge.

Consider the mechanism through which indigo buntings (a species of migratory songbirds) learn to identify north from the night sky. In today's night sky, Polaris is the north pole, the star in the region of the sky that reliably indicates north. Because the stars change position through time, Polaris will not always be true north, nor will true north be predictable from the constellations of today, for the constellations themselves change over evolutionary time. Indeed, 100,000 years ago there was no big dipper, no Orion… Thus, evolution could not have built in a map of the sky on which Polaris is identified as the north pole, and yet, indigo buntings have such a map. Shown a stationary simulacrum of the current night sky in a planetarium, they take off as if to fly south as specified by the North Star, no matter how the planetarium's night sky has been oriented with respect to true north in the actual world. Indigo buntings must have learned where the north pole is; how could they have done so?

In a series of wonderful experiments, Emlen (1972) uncovered the mechanism by which indigo bunting nestlings achieve this feat. Because the earth rotates on an axis that goes through true north, the north pole is the center of rotation of the night sky. Emlen showed that nestling indigo buntings observe the rotating sky, and infer north as the center of rotation of what they see. He showed this by raising birds in a planetarium in which he could make arbitrary skies rotate around arbitrary centers, and then observed which direction they took off in the autumn, when their hormones told them it was time to fly south. There is a critical period for this learning device; if buntings do not learn the north pole while nestlings, they never do.

This is a paradigm species specific, domain specific, learning device. Clearly, not every animal will spontaneously note the center of rotation of the night sky, nor would every animal infer true north from the observations if they happened to make them. Thus, this device is species specific. With respect to domain specificity, this device is of no use learning what food is safe to eat, what indigo bunting song is like, or anything else other than where north is. This device exemplifies other properties of core domains. Innate perceptual mechanisms allow the bird to identify the sky, and to represent its rotation. And this device is a learning mechanism; innately specified computations over the representation of celestial rotation ensure that the buntings end up representing something important that they must, by logical necessity, learn—how to tell north from the night sky.

1.4. On the very notion of *innate representation*

Some authors have doubted that it would be possible for there to be innate representations, or that the notion is even coherent (e.g., Thelan et al., 2000; Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996). It is certainly possible that until the child has seen something red, she has never activated a representation of red. Thus, we might want to say that the *capacity* to represent red is innate, but that the child creates no representations of *red* until he/she has seen something red. When I ask whether some system of representations is innate I am asking whether the capacity to form the relevant representations, the capacity to activate states in the nervous system that refer to particular entities in the world and enter into relevant computations involving those entities, is innate.

Consider again the indigo buntings. Some innate perceptual analyzers must enable the nestling to identify the night sky and analyze its axis of rotation. That some state of the bunting’s nervous system has the content *night sky* is shown not only by the fact that it is activated by the night sky, but also that it plays a role in the computations that determine the azimuth and that determine the direction of flight during migration. This is an example of what I mean by taking behavioral evidence regarding the extension and inferential role of some representation as evidence regarding its content. Finally, these representations also exemplify the senses in which core knowledge is conceptual. The inferential role of the representations of a view of the night sky go beyond the specification of what it looks like (points of light distributed against a black background), nothing in the description of what the sky looks like has the content *north.*

Take another example. Ever since the pioneering work of Lorenz (1937) we have known that many birds have an innate learning mechanism that enables them to identify conspecifics, especially one important conspecific—their mother. The learning device Lorenz discovered and called “imprinting” causes chicks to attend to and seek proximity to the first large thing that moves in certain ways. Countless beginning psychology and ethology students have seen the footage of Lorenz walking through the country side with a line of goslings imprinted on him following behind. In this case, the perceptual analyzer that initially identifies *mother* or *conspecific* is movement of a certain sort; goslings and chicks will imprint on a large red shiny ball if that is the only object they see moving during a critical period after birth—or so we thought until the work of Johnson, Bolhuis & Horn (1985). These researchers showed that mother/conspecific identification is such an important problem that evolution has provide two redundant learning devices to support it. A newborn chick presented with a moving ball or person and *also* a stationary stuffed hen will huddle near the stuffed hen.

Johnson et al.’s work shows that chicks must have an innate perceptual analyzer that specifies what a conspecific *looks* like, and indeed, they specified its nature. The innate representation of conspecific appearance is very sketchy—there must be an overall bird shape, with a head with eyes and a beak on a neck on a body; if these elements are not in the right configuration, the chick does not recognize the object as a conspecific. However, a chick will huddle close to a stuffed duck or eagle or owl. Johnson and his colleagues have much to say about this fascinating story, bearing on the relations between the two different learning mechanisms, their critical periods and neural substrates. My point here is simply that there is an innate representation of *hen* that specifies roughly what hens look like, and contains the conceptual role, “stay close to that.” This representation is part of a domain specific learning device; the representation is sketchy so it can be filled in with the details of the chick’s own mother. In this way it is similar to the innate representations vervet monkeys have of three classes of predators—birds of prey, large cats, and snakes (Cheney and Seyfarth, 1987)—representations that play a role in learning the features of actual predators in their environments. It is similar, also, to the innate representations human babies have of conspecifics (two eyes above a mouth within an oval, in a certain configuration; see Morton, Johnson & Maurer, 1990), representations that play a role in their identifying their mothers.

This book argues for many controversial conclusions, but the claim that *some* representations are innate is not meant to be one of them. The controversy comes when we begin to consider the *content* and *nature* of innate representational systems. There certainly are innate perceptual representations and innate mechanisms of perceptual learning. Also, innate input analyzers, constrained to accept only some classes of stimuli (bird shaped entities, face shaped entities, the night sky), create conceptual representations that are input to further learning.

1.5 Continuity and Fodor’s Challenge

This book second major thesis is that human beings, alone among animals, have the capacity to create representational systems that transcend perceptual representations and core knowledge. That is, human beings create new representational resources, qualitatively different, in ways that may be made precise, from what they were built from. Many cognitive scientists deny the very possibility of true cognitive development in this sense (Fodor, 1975; Macnamara, 1986), endorsing instead a strong version of *continuity thesis*. The continuity thesis states that all the representational structures and inferential capacities that underlie adult belief systems either are present throughout development or arise through processes such as maturation. Fodor’s argument for the continuity thesis was a one-liner: one cannot learn what one cannot represent. All known learning mechanisms (e.g., hypothesis testing, parameter setting, correlation detection, prototype formation) involve computations over antecedently available representations (e.g., choosing between already formulated hypothesis, setting a previously specified parameter, noting a correlation between already represented states of affairs, abstracting what is common among a set of represented exemplars). These mechanisms, by their very nature, cannot yield a capacity to represent aspects of experience that were previously unrepresentable.

Fodor’s argument has a crucial weakness. All one needs to refute the claim that novel representational capacities cannot be formed through learning is a single counterexample. The best set of counterexamples were invoked by Piaget (1980) in a reply to Fodor’s argument, and come from the history of mathematics. Concepts of rational, real, and complex numbers, and the mathematical notations required to express them, cannot plausibly be attributed to infants, children, or even to adults lacking sufficient education in mathematics, and they do not plausibly develop through maturation. Rather, such concepts arise as one learns mathematics. Chapter 7 presents an example of this sort: the construction of an explicit integer list system for representing number. The integer list is a cultural construction with more representational power than any of the core representational systems on which it is built, thereby providing a genuine counterexample to Fodor’s argument.

Many cognitive scientists believe that cognitive development provides further counterexamples to the continuity thesis. In the course of acquiring intuitive theories of the world. new concepts are constructed that are incommensurable with those from which they are built. Chapters 8 and 9 develop two case studies of conceptual change in childhood, one in the domain of intuitive biology and one in the domain of intuitive physics. If such existence proofs show that Fodor’s argument is wrong, however, they do not show what is wrong with it. In particular, they do not show us how the infant, child, mathematician, or scientist can use her current representational resources to learn new ones. That is Fodor’s challenge to cognitive science, and it is still unmet.

Those who deny the continuity thesis face two distinct challenges. The first, descriptive, challenge is to find evidence for two successive points in cognitive development, the later of which contains a representational capacity that transcends what was previously available. The second, explanatory challenge is to specify a learning mechanism that yields the new representational capacity. This book takes on both challenges. Chapters 7, 8 and 9 each describe discontinuities within conceptual development. This book’s third important thesis is that the explanatory challenge is met by bootstrapping processes such as those described in the literature in history and philosophy of science (Quine, 1960; Nersessian, 1992).

1.6 On bootstrapping

The bootstrapping processes that underlie discontinuous conceptual development must be distinguished from a different kind of learning, also called "bootstrapping," that is debated in the contemporary literature on language acquisition. In the language acquisition literature, bootstrapping processes are invoked to explain how children solve a mapping problem. Suppose children know, thanks to innately supported universal grammar, that there will be nouns, noun phrases, transitive verbs, verb phrases, adjectives, prepositions, quantifiers, etc., in natural language. Suppose, moreover, that they know, thanks to innately supported conceptual capacities, that there will be words in natural language for individuals, kinds of objects, actions, stuff, etc. Children still face the formidable problem of identifying how their own particular language expresses these and other universal features of language and thought. "Semantic bootstrapping" gives the child a beginning wedge into the problem of discovering a particular language's syntactic devices. For example, the language learning mechanism might include the heuristic that representations of kinds of physical objects ought to be mapped onto count nouns. Then, if the child can figure out that a word is being used to refer to a kind of object, the child may assign it to the lexical category *count noun*, and use this assignment to further figure out what the adjectives, determiners and so forth are in the language. "Syntactic bootstrapping" gives the child a beginning wedge into the problem of discovering the particular words expressing particular concepts. For example, if the concept "give" includes a giver, a receiver, and a gift, then the child may map the concept to a verb with three arguments.

The bootstrapping metaphor is put to entirely different use in the older language acquisition literature, in which early linguistic categories were hypothesized to be only semantically interpreted, syntactic categories being created by a bootstrapping process of the sort envisioned in the history of science literature (e.g., Schlessinger, 1976). Here I appeal to a learning process like that of "bootstrapping" in its original sense: the learning process to which Quine (1960) appealed with his metaphors of Neurath's boat, climbing a ladder and then kicking it out from under, scrambling up a chimney supporting oneself against the walls as one builds it, etc. Chapters 7 through 9 seek to go beyond such metaphors and flesh out some details of how bootstrapping processes might actually work.

1.7 Intuitive theories

Bootstrapping mechanisms underlie the human capacity to create theoretical knowledge that transcends core knowledge. Although there are many knowledge structures worthy of study (scripts, schemas, prototypes, the integer list representation of number, the alphabet), many students of cognitive development assume that one kind of knowledge structure, intuitive theories, play a particularly important role in cognitive architecture (Carey, 1985b; Gopnik and Meltzoff, 1997; Keil, 1989; Wellman and Gelman, 1992). I endorse this assumption, and here I focus on a class of intuitive theories, those that Wellman and Gelman (1992) call "framework theories." These are the theories that ground the deepest ontological commitments and the most general explanatory principles in terms of which we understand our world. One task (but by no means the only task) in the study of cognitive development is accounting for the acquisition of framework theories.

It is worth stepping back and considering what is being presupposed by the choice of the term "intuitive theory" rather than the more neutral "cognitive structure." The claim of the above authors is that intuitive theories play several unique roles in mental life. These include: (1) determining a concept's core (the properties seen as essential to membership in a concept's extension) (2) representing causal and explanatory knowledge, (3) supporting explanation based inference. As Gopnik and Meltzoff (1997) emphasize, and I endorse, the mechanisms underlying theory development differ from those that underlie the acquisition of different types of conceptual structures. It is an empirical question whether children represent intuitive theories, and whether knowledge acquisition in childhood involves the process of theory change. Those who talk of "intuitive theories," and "framework theories," are explicitly committing themselves to an affirmative answer to those empirical questions. This commitment does not deny that there are important differences between children as theorizers and adult scientists (hence the qualifier, "intuitive theories.") Children are not metaconceptually aware theory builders (e.g., D. Kuhn, Amsel & O'Loughlin, 1988). In spite of these differences, the research enterprise in which this work is placed presupposes that there are a set of questions that can be asked, literally, of both scientific theories and intuitive theories, and which receive the same answer in both cases. Of course, the merit of this presupposition depends upon the fruitfulness of the research it generates. See, for example, the explicit comparison of conceptual change within thermal concepts in the history of science and conceptual change within concepts of matter in middle childhood; Carey, 1991; Smith, Carey and Wiser, 1985; Wiser and Carey, 1983. Chapters 8 and 9 present case studies framed within this research tradition.

Intuitive theories differ from core domains of knowledge in many ways. One of the goals of this book is to account for the origin and development of theory-embedded conceptual knowledge, given a beginning state of perceptual representations and core knowledge. Many researchers have blurred the distinction between core knowledge and intuitive theories. For example, Leslie is one of the most articulate advocates of the core knowledge position, and he too characterizes core knowledge as modular, encapsulated, supported by innate perceptual analyzers, and unchanging during development, accords with Confusingly, in discussing the infants' beginning state, he characterizes core knowledge, yet he dubs his modules: *Theory* of Mind Module and *Theory* of Bodies Module (Leslie, 1994). That is, he characterizes core knowledge, but he calls systems of core knowledge "theories." Gopnik and Meltzoff (1997), in contrast, deny the distinction; they explicitly characterize infants' early developing knowledge of bodies and agents as theoretical knowledge, claiming that the mechanisms that underlie acquisition of knowledge in these domains are the same as those that support theory development by adult scientists. This leads to the seemingly absurd (and false, I shall argue) conclusion that the processes by which infants achieve object permanence are the same as those through which Darwin formulated the theory of natural selection.

The confusion between core knowledge and intuitive theories is come by honestly, for two of the parade cases of each have overlapping content. Core knowledge of objects overlaps with knowledge of intuitive physics, and core knowledge of intentional agents overlaps with knowledge of intuitive theory of mind. This is because the output of the core knowledge systems is part of the input to theory building. But theory building can, and does, transcend core knowledge. Of course, it is an empirical and theoretical issue of great importance whether core knowledge and intuitive theories are different types of conceptual structures, arising in development through different mechanisms. The case studies in this book constitute an extended argument for this position.

1.8. Cultural and cognition; cultural relativity

Many writers have explored the implications of the core knowledge thesis for cross-cultural cognitive universals. Since core knowledge is innate and continues to operate throughout development, the core knowledge hypothesis entails a certain degree of universal conceptual structure. However, in the course of conceptual change and in the course of the construction of new representational formats, new concepts come into being. These are always culturally constructed; conceptual change is a social process. Understanding the human capacity to create new representational resources is continuous with understanding the human capacity for culture. New representational resources are then embodied in language or other symbolic systems, and reflected in artifacts and culturally supported institutions. Because these new representational resources transcend core knowledge, as children become acculturated, they must themselves engage in bootstrapping processes in order to construct them for themselves.

This process leaves open the possibility of a radical cultural relativity in conceptual structure. If different cultures transcend core knowledge in ways that are incommensurable not only with core knowledge but also with each other, then members of those different cultures will possess incommensurable conceptual systems. Chapter 8 explores cultural relativity in intuitive theories of biology.

1.9 Language and thought, linguistic relativity

Given the close connections between language and thought, the hypothesis of linguistic relativity of conceptual repertoires (a special case of cultural relativity) has long been debated. The logical possibilities concerning the origins of the concepts that are expressed in the open class vocabulary (e.g., nouns, verbs, adjectives), as well as those that structure sentences more intimately, the closed class vocabulary (quantifiers, bound morphemes, determiners, conjunctions) are the same as for any concepts. They may have arisen over evolution (either as part of hominid evolution or predating it), such that language evolved to express concepts that already articulated the primate view of the world. Or the conceptual repetoire that underlies human language may have evolved as part of the evolution of language. And finally, some may be cultural constructions, expressed in and maintained through natural language. I shall consider these alternatives for each of the cases developed in the following pages.

Linguistic relativity could arise in one of two ways. First, if some of the concepts expressed in natural languages are cultural constructions discontinuous with innate knowledge, and if different cultures create concepts different from each other, and if children create them in the course of learning language, then adults of different cultures will speak mutually unintelligible and untranslatable languages (in part). Chapters 5, 8 and 9 discuss putative cases of strong linguistic relativity. Because of the central role of language in Quinian bootstrapping mechanisms, and because Quine explicitly endorsed this view, I shall call it Quinian linguistic relativity. Second, if language variation is parameterized over a fixed set of universally available options, and if the parameter settings of each language affect habitual construals of the entities talked about, these habitual construals may become entrenched and in turn may affect non-linguistic conceptualizations as well. This hypothesis is weaker, for there is no claim of mutual unintelligibility or nontranslatability. Chapter 6 considers a putative case of weak linguistic relativity.

1.10 Overview of the book

The chapters that come distinguish the representations that constitute core knowledge from two other types of representations—perceptual representations and conceptual representations embedded in explicit intuitive theories. Chapters 3 and 4 characterize core knowledge and provide evidence for several domains of human core knowledge that exemplify its distinctive features. Chapters 5 and 6 consider the relations between core knowledge and language, considering the hypotheses of both weak linguistic relativity and Quinian linguistic relativity. Finally, Chapters 7, 8, and 9 take on Fodor's challenge, providing descriptions of discontinuous conceptual development in which concepts exist that were not expressible given earlier conceptual resources, and characterizing the bootstrapping mechanisms that underlie the change.