Published in final edited form as:
Behav Brain Sci. 2011 June ; 34(3): 152–162.

Concept Innateness, Concept Continuity, and Bootstrapping

Susan Carey
Harvard University

Abstract

The commentators raised issues relevant to all three important theses of *The Origin of Concepts* (TOOC). Some questioned the very existence of innate representational primitives, and others questioned my claims about their richness and whether they should be thought of as concepts. Some questioned the existence of conceptual discontinuity in the course of knowledge acquisition and others argued that discontinuity is much more common than portrayed in TOOC. Some raised issues with my characterization of Quinian bootstrapping, and others questioned the dual factor theory of concepts motivated by my picture of conceptual development.

I am deeply moved by the thoughtful commentaries provided by 29 colleagues, both philosophers and psychologists. Following the organization of the précis, my responses begin with preliminary issues, turn then to innateness and core cognition, and then to the issue of conceptual discontinuity and my proposal for the bootstrapping process that underlies discontinuities in conceptual development. I end with a general discussion of the nature of concepts and conceptual development.

1. Preliminary Remarks

The Origin of Concepts (TOOC) concerns mental representations—how they arise and how they come to have the meanings they do. The book is organized around case studies, such as representations of number and agency in childhood or heat and temperature in the history of science. As *Markman* points out, this approach typifies research on conceptual development at least since Piaget, and differs from much research on the cognitive psychology of adults' concepts, which tends to address conceptual representations in the abstract, independent of content. Like Markman, I believe that the psychology of concepts must follow the lead of the psychology of perception. The study of perception has long proceeded at both levels of abstraction in parallel. That is, scientists interested in perception have studied both particular cases, such as depth perception, motion perception, color perception and taste perception, as well as seeking to illuminate how sensation and perception work in general. Similarly, work on the nature of concepts in general and how human beings represent particular important concepts are mutually illuminating.

My commentator sometimes lost track of what my project in TOOC was—namely explaining how representations arise and how they come to have the meanings that they have. For example, *McClaren, Wills and Graham* discuss the pattern encoding mechanism that extracts an orientation specific prototype of stimuli such as faces that share a common configuration (e.g. Diamond & Carey, 1986; Gautier & Taar, 2002; McClaren, 1997). They ask why I didn't include this work as an excellent example of core cognition, asking how it differs from the mechanism through which indigo buntings learn to identify the north star. The reason is simple: although this mechanism plays a role in creating an iconic representation of prototypical faces (or prototypical dogs, Greebles, or meaningless checkerboard patterns), and these iconic representations in turn support better discrimination among members of those categories, the question in TOOC is how these representations come to have the meanings they do? That this mechanism works for meaningless checkerboard patterns shows just why it isn't a system of core cognition—there is no innately supported meaning assigned to the representations that are its output. What distinguishes this computational system from the mechanism through which indigo buntings identify the axis of rotation of the night sky is that the output of the latter computation has an innate conceptual role. So too do the representations of agents, objects and number that articulate systems of core cognition.

2. Innateness of representational primitives

Several commentators were not convinced that there are *any* innate representational primitives, whereas others argued for leaner ones than suggested in TOOC. Let's take these two issues in turn.

2.1 Distinguishing learning from development

The argument for innate representations I offer in my book concerns theories of learning *learning*, not theories of *development* in general. The argument is simple, and I would think, uncontroversial. Learning is a computational process that operates on representations, and so if we believe that learning plays a role in the construction of representational resources, we are committed to there being *some* innate representations: those that are the input to the initial episodes of learning. *Schlesinger* and *Amso* critique the arguments in TOOC for innate object representations on the grounds that success on some of the tasks that provide evidence for them is dependent upon experience. Remember, I am using "innate" to mean "not learned." They, instead, take "innate" to contrast with "experience dependent," and thus fail to separate learning from development more generally. We do not learn to get taller, although physical growth is certainly experience dependent (e.g., we'd better eat and move). Some changes in neural connectivity implement learning processes and some do not. Hence it is perfectly compatible with my arguments for nativism that these mechanisms are to some extent experience dependence.

Allen and Bickhard critique my argument for some innate representational primitives by pointing out that there must be processes that underlie the acquisition of some representations that do not themselves begin with representations. They ask, whatever these processes are, couldn't they underlie the developmental changes we see in infants and toddlers? It is certainly likely that some of the developmental changes we see very early in childhood are maturationally driven changes (i.e., not underlain by learning), and indeed I discuss such changes in TOOC. But noting that there are some representations whose origin is not learning is a way of acknowledging innate representations, not denying them.

2.2. Keeping one's eye on the ball: accounting for specific conceptual content

The work described by Schlesinger and Amso provides the field an important challenge, namely untangling domain general developmental changes in visual search capacities from the acquisition of specific representational schemas (e.g. of objects). Schlesinger and Amso suggest that developmental changes in visual selective attention, occulomotor skill, visual scanning, and active exploration might explain developmental changes in the capacity for object completion over the first few months of life, and I agree. In TOOC (p 58) I suggest that the development of visual selective attention might drive the process through which the input to innate computational mechanisms become becomes available to the child. Schlesinger and Amso suggest instead that these domain general changes *explain* the acquisition of object representations. But Schlesinger and Amso do not attempt to specify how representations of objects could be constructed from other primitives, such as sensori-motor ones. Absent such an account, they have not offered an alternative to the hypothesis that object representations are innate.

Both *Xu* and *Gopnik* clearly focus on learning, and suggest that the current exciting work on inductive learning in the Bayesian tradition offers a suggestion for how the structures I posit as *innate* core knowledge might instead be *learned*. However, these models have not yet been applied to explaining any domain of core cognition. Thus, this suggestion is little more than a guess that learning formulated over leaner primitives *might* yield the concepts of agent, number, and object that constitute the case studies in TOOC.

As Xu points out, we need not accept that core knowledge is innate solely on the basis of evidence for early emergence of sophisticated reasoning abilities. Of course not; there are four further arguments for the innateness of the representations in core cognition systems (see précis).

I agree with Xu that constraints on induction can be, and most typically are, learned. She might have offered her own exciting work on infants' learning of over-hypotheses as a worked example (Dewar & Xu, in press). But an over-hypothesis such as "the objects in each bag are the same color" or "the objects in each bag are the same shape" require representations of *object, bag, color and shape* to be formulated. More generally, Xu advocates abandoning the project of explaining the origin of concepts in favor of focusing on the project of understanding the processes of belief fixation. These are both important projects in cognitive science; we do not have the option of choosing between them. Belief-representations are composed of concept-representations. To decide among the hypotheses that "blicket" means *animal, dog,* or *dalmation*, (Xu & Tenenbaum, 2007) one must have these latter three concepts available. TOOC concerns how concepts arise.

Gopnik and I agree in endorsing the theory-theory of conceptual development, but we emphasize different lessons from the history of science for the understanding of human mind. Gopnik, like Xu, emphasizes the processes that underlie belief fixation (learning from evidence), whereas I emphasize the processes that result in new representational primitives that articulate abstract conceptual structure previously unrepresented. The examples in Gopnik's commentary do not address the problem I have taken on. That statistical evidence is relevant to inferring desires, or to making novel inferences in chemistry and physics, offers no explanation for the capacity to represent desires at all, or to articulate the inferences in chemistry and physics at all.

Gopnik makes an important point: Quinian bootstrapping is not the only process that results in the creation of new concepts (see below). I agree that the mechanisms in which hidden variables are posited in the course of Bayesian causal learning also may also do so. It would be very interesting indeed to show that concepts such as *goal, belief, living thing, heat, object,* or *integer* emerge through such a process. I am placing my bets on rich innate primitives and on Quinean bootstrapping, not as a hunch, but on the basis of the analyses and evidence offered in TOOC.

2.3. Leaner innate primitives

Mandler accepts the challenge of specifying what leaner primitives than those argued for in TOOC might look like, and what the learning process that creates richer representational resources from them might be. The primitives Mandler posits are spatio-temporal; the learning process she calls "perceptual meaning analysis." "A single innate analyzer (such as Perceptual Meaning Analysis) collapses attended spatio-temporal information into a small number of conceptual primitives" (Mander, 2011, p. xxx—page number from printed BBS commentary). Here's the problem. Perceptual analysis could certainly abstract the schema in common to Michottian launching, for example, or equifinal approaches by different paths to a common endpoint. And such schemata could be represented iconically. The question is what would make these schemata representations of *cause*, (or *make move*, in Mandler's words), or *goal directed action*, respectively? Where does the conceptual meaning come from?

TOOC provides an answer concerning the source of meaning: conceptual role, including innate conceptual role. It also provides an empirical argument against the suggestion that the perceptual representations through which entities in core cognition are identified are the result of perceptual meaning analysis. At least some of the representations of objects, agents, causality and number that are created in infancy are the output of *innate* input analyzers. They are present in neonates, both in humans and other animals.

TOOC's Hypothesis Space Regarding Innate Conceptual Primitives is Too Unconstrained—Mandler points out that allowing the possibility of two types of innate conceptual representations—those embedded in domain specific systems of core cognition and those embedded in central ones—greatly increases the problems of bringing data to bear on characterizing the representational repertoire of young infants. True, but unfortunately we do not get to stipulate the space of possibilities. Acknowledging both a dorsal and a ventral visual system complicates our understanding of vision, but but nobody would argue against this on the grounds of simplicity.

Poulin-Dubois raises a related worry, namely that the hypothesis that there are innate input analyzers that identify the entities in core cognition may be unfalsifiable. She points to the existence of developmental changes in the features that support core cognition. But why does this make the hypothesis of innate input analyzers unfalsifiable? An analogy with the developmental changes in depth perception might clarify what's at stake here. On some developmental accounts of depth perception, some cues to depth are innate, thus allowing infants to represent depth, which in turn allows them to then learn other cues to specifying it (Kellman & Arterrberry, 2000). The innate cues can be discovered by experiments on neonates. Similarly, that there are changes with age in the processes through which children identify agents, as well as changes in the inferences infants draw about them is to be expected on the core cognition account. If some representational capacity is innate, there must be some mechanism through which leads infants to create representations in the domain. Thus, the way to falsify this hypothesis is to counter the evidence that young infants have representations with the content in question, and also to provide a learning mechanism through which the relevant representations could be constructed from leaner primitives.

I plead not guilty to another of Poulin-Dubois' charges—namely that I see developing children as solitary agents, forming their intuitive theories in a social vacuum. It is true that scaffolding and cultural learning do not underlie the acquisition of innate representations, by definition. But my general picture of conceptual development could not be further from this view; the whole second half of TOOC concerns the cultural construction of knowledge, and takes on the challenge of showing how scaffolding and cultural learning actually works.

Finally, *Machery* accuses me of conveniently letting my intuitions decide what is the content of the representations that underlie infants' and children's performance in the tasks described in TOOC. Not guilty. During my career I have spent years diagnosing, on empirical grounds, what the Experimenters' "heat" means (Wiser and Carey, 1983), what the child's "alive" (Carey, 1985), "heavy" (TOOC) and "five" (TOOC) mean. In each of these cases I have argued that the relevant concepts are different from (and in some cases even incommensurable with) those we 21st century minimally scientifically literate adults would express with the same words. In TOOC, I brought empirical data to bear on many different interpretations of the symbolic capacities the experiments reflect (both richer and leaner than those I settle on). Machery is right that figuring out the content of any given mental representation is an extremely difficult task, and one cannot let intuitions decide the matter.

2.5.Explicit vs. Implicit Content in Core Cognition

A point raised by *Mandler* underlines a confusion I am responsible for. TOOC uses the word “explicit” in two different ways. In some places, I use it to mean “public” as when I distinguish systems of core cognition from later developing public representations—those that articulate language, mathematical notations, diagrams and the like. On this reading, no symbols in core cognition are explicit. But there is another sense of “explicit” that is important to TOOC. If we accept a representational theory of mind, we are committed to mental symbols, symbols with formats, meanings, extensions and computational roles. These are the explicit symbols. But mental representation can also be implicit. For example, within working memory models of small sets of objects (Chapter 4), the explicit symbols are the object files themselves; these are symbols that are activated in the mind, and they represent the objects that are the focus of attention. The numerical content of this whole system of representation is implicit, embodied in the processes that determine whether to open a new object file, to update working memory models when objects are added to or removed from the attended set, and to compare two models on the basis of 1-1 correspondence. There are no explicit symbols for number in this system of representation.

Mandler asks how core cognition can support the recall of event sequences that have been demonstrated in the 2nd year of life, given that core cognition contains no explicit symbols. There is no problem here; there are no public symbols in core cognition, but the explicit symbols within core cognition, in the second sense outlined above—are input to further computations, and thus underlie working memory, long term memory, action and reasoning in infancy.

3. Case studies of core cognition and other innate representational resources

3.1. The case of object representations

One issue I struggled with in writing TOOC was the senses in which core cognition representations are perceptual and the senses in which they are conceptual. Core cognition clearly goes beyond sensori-motor representations, but in many ways it resembles other clearly perceptual representations, like those of depth. I agree with several commentators who were not satisfied with the progress made on this issue; here is one place where there is much work to be done.

My strategy in TOOC is to consider many ways scholars have tried to distinguish perceptual from conceptual representations, and to argue that core cognition representations are conceptual in terms of all them, except perhaps being encoded in sentence-like propositional format. This strategy was no doubt confusing, because I do not myself accept many of those previous analyses as actually drawing the relevant distinctions between conceptual and perceptual representations. For example, Burge takes me to be endorsing Piaget’s and Quine’s characterization of how perceptual representations differ from conceptual ones, but I do not. Piaget and Quine both argued that young infants create no representations *object*, because infants are incapable of object individuation, of representing a later encountered entity as the same object as an earlier encountered one, and thus incapable of appreciating object permanence. I show that Piaget and Quine were wrong about the representational capacities of infants, and so by their analyses, young infants have a *concept* of objects. That is, infants have object representations, where *object* is glossed roughly as *bounded, coherent, separately moving, spatiotemporally continuous material entity*. Note, this gloss reflects implicit content. Infants integrate information from different modalities, and go beyond stationary snapshots of objects, in creating such representations. But I fully agree with Burge that perceptual representations go beyond sensory ones in these ways, and thus these facts do not rule out that object representations in infancy are perceptual.

Burge also seems to think that because I take object representations to be conceptual, I am denying that there are or could be perceptual representations of objects as well. Again, I do not. Adults clearly have both. The question then becomes, as he points out, what reasons there are to believe that infants have conceptual object representations. My belief that they do hangs on the central inferential role of infants' object representations; their inferential interrelations with causal representations, agent representations, and the fact that they are input into working memory models of small sets over which many different quantitative computations are defined. The quantitative capacities in questions are not the within-module computations of object individuation and numerical identity that are relevant to arguments against Piaget and Quine. Rather, TOOC reviews evidence that working memory models of small sets of objects are input into processes that compute total surface area or volume (as in deciding which bucket has more cracker stuff in it) and into computations of 1-1 correspondence (as in deciding whether all of the objects placed into a container have been removed). Chapter 4 also reviews evidence that infants can create hierarchical models of sets (e.g., a model that of a set of two sets, each containing two objects); see Feigenson and Halberda (2004), and Chapter 7 reviews evidence that prelinguistic infants have a mental symbol *plural* that applies to sets of objects. These are the quantitative computations that suggest central conceptual role to me. With respect to the integration of object representations with causal representations and agent representations, see chapters 5 and 6.

Gauker argues that by my own characterization of concepts, icons cannot be concepts, so if the representations in core cognition are iconic, they cannot be conceptual representations. He illustrates his points in relation to the concept *object*. At issue is how we draw the distinction between conceptual representations, on the one hand, and other kinds of representations, on the other. Gauker stipulates that conceptual representations must be part of a representational system with some sentence-like format; to have the concept *object* the child must be able to think the thought we can express in words "that is an object." (Thus, Gauker claims there also must be concepts of *that, is, a*, for there to be the concept *object*.) I don't agree with these claims.

By "conceptual," I do not mean "language like." I mean what Gauker says I mean—expressing content not expressable in spatio-temporal or sensori-motor vocabulary, and participating in central, conceptual, inferential roles. The question is whether a symbol with the content we express by *object* can be carried by an iconic symbol, such as "." Gauker takes it as obvious that such a symbol cannot represent what two objects have in common, or two balls, as the words "object" or "ball" do. Who says? What determines what "" represents? The position taken in TOOC is that the answer to this question, abstractly, is the same as what determines what "object" represents—on a dual factor theory, the causal connections between the symbol and the entities in its extension, and some aspects of its computational role. Chapters 2 and 3 provide an extended argument that such a symbol (an object file) for young infants does indeed have the content *object* (i.e., it represents a particular object *as such*, in a particular location relative to the baby and relative to other objects, and Chapters 4 through 7 argue for central conceptual role: objects are represented in working memory models that support action, and in particular causal relations to other objects, and in particular intentional relations to agents. Some of this content is implicit (i.e., it is not required that the child have the ability to think the thought that one object is a different object from another), but I assume that object files are explicit symbols, most probably iconic. Chapter 7 argues that such a symbol decidedly does not have the same content as does the word “ball” until late in infancy, at the developmentally earliest. In sum, the “alternative” picture Gauker offers is a good sketch of how I think about core cognition, because I do not take “conceptual” to mean “language-like.”

Hill also comments on object representations, and on my treatment of Quine’s position in particular. Hill argues that by crediting infants with representations of *object stages*, or undetached *object parts*, Quine is already committed to their representing objects, since representations of objects are presupposed by representations of stages of objects or undetached parts of objects. I understand Quine to be imagining a representational state of affairs formulated over sensory snapshots and associations among them (and I believe there is textual evidence that something like this is what Quine had in mind for his “perceptual similarity space,” see TOOC, Chapter 2, and Carey, 1994). On this exegesis of Quine’s position, infants representations of object stages does not require them to have of objects.

In TOOC, I concede one of Hill’s points; that it would be possible in principle to formulate any one of the infant’s expectancies (such as that revealed in the rotating screen experiment Hill takes as an example) in terms of statistical relations among such snapshots. I argue against this alternative on simplicity grounds.

3. 2. Representations of Causality

Butterfill raises all the right questions about infants’ causal representations. Might they be discontinuous with those of adults? Might causal representations constitute a system of core knowledge, or, alternatively, might different causal representations discussed in TOOC actually be embedded in distinct systems of core cognition?

Discontinuity exists at two levels of abstraction; representations in core cognition are discontinuous in format from later public, linguistic symbols, and representations in core cognition may express content that is qualitatively different from that expressed by later acquired concepts. In addition, the content of any system of representation (core cognition or public) is partly implicit, and conceptual development often involves creating explicit symbols that expresses content that was earlier only implicit. Butterfill speculates that infants’ causal representations, like number representations within the parallel individuation system, may be perception-like, encoded implicitly as constraints on mental models of events rather than explicitly, with a symbol that expresses the concept *cause*. I completely agree this is a possibility. However, this question is orthogonal to that of whether there is a central system of causal representations that integrates output from distinct core cognition systems, or only causal representations within the already attested core cognition systems. It is also possible that symbols in the central system, if it exists, may be language like; there may be an abstract symbol with argument slots that expresses *cause (x, y)*.

A relevant source of data that bear on deciding among these alternatives derives from the expression of causal concepts in the earliest language acquisition. There is no evidence for discontinuities in content. By contrast, the evidence from language learning strongly suggests discontinuities in number representations. Children use verbal numerals for 1 ½ years before they figure out how they represent number. “One-“knowers have “two” in their vocabulary for 6 to 9 months before they understand it to refer to the cardinality *two*. Before that, they use it as a plural marker, or referring to sets with cardinalities greater than one. Nothing like a “one”-knower stage has been discovered in regards to the expression of causal concepts in language. Children have not been observed to use a lexical item “cause” or the causative use of “make” or a lexical causative like “break” for months with non-adult meanings. Rather, children learn to form both lexical causatives (“he broke the cup”) and periphrastic causatives ("he made the cup break") early in their third year of life, and as soon as they command the periphrastic causative construction, they create novel causal alternations across wide conceptual content (e.g., they say "eat the baby," meaning "feed the baby, make the baby eat,"). Sometimes these violate adult restrictions on this construction and so could not have been in their input (Bowerman, 1974). This suggests to me an abstract representation of cause available at the outset of language learning, one that does not require conceptual change for its acquisition.

Butterfill seems to suggest that 2-year-olds' lack of explicit access to some causal mechanisms (e.g., those underlying the solidity constraint) shows discontinuities in their concept *cause*. But one should not confuse lack of knowledge of a particular causal mechanism with lack of the concept of *cause*. Uncovering the causal structure of the world is an ancient ongoing project, deeply enmeshed in culture. Kuhn (1977) argues for continuity in the concept *cause* over the history of science, in spite of profound conceptual changes within theories of causal mechanisms. Obviously children know very little about the causal mechanisms that adults around them understand, and even adults' understanding of causal mechanisms is extremely sketchy (Rozenblit & Keil, 2002). But neither fact is inconsistent with the possibility of continuity through development of the concept *cause*.

3.3 Agent Representations

Both *Kiss* and *Poulin-Dubois* suggest that the adult theory of mind might be discontinuous with infant representations of agency. As with causal representations, there are several ways this may be so: 1) the format of representation may certainly differ; infants cannot talk about their minds; 2) some of the content in external, public, representations of minds may be captured only implicitly in infants' models of the actions of agents, in the form of constraints on the computations these enter into; and 3) the content expressible, either implicitly or explicitly, in core cognition may be incommensurable with that of later developing theories of agents.

Kiss characterizes many ways infants' and young children's representations of minds may be qualitatively different from those of adults, but I know of no evidence for Kiss's specific proposals. For example, what is the evidence that the first meaning of "happy" is the behavioral manifestations of happiness, such as smiling? Contrary to Kiss' speculations, TOOC (Chapter 5) discusses the recent studies showing representation of perceptual and epistemic states by preverbal infants, suggesting continuity over development in at least *some* important conceptual content. TOOC also summarizes some of the empirical evidence for discontinuity, arguing that it is still an open possibility that developmental changes in the theory of mind involves qualitative discontinuities and requires Quinian bootstrapping.

TOOC endorses *Ristau's* observation that there is good evidence for systems of core cognition in animals (sometimes the same ones as in humans, as in analog magnitude representations of number). Non-human animals' representations of agency is a particularly important topic of active research. It is clear that non-agents satisfy the innate input analyzers that identify agents in the world (as in the experiments with computerized geometric figures attributed goals, or furry robots attributed attentional states). However, I know of no evidence for Ristau's speculation that the innate input analyzers that identify agents may lead to *wide-spread* over-attribution of agency by infants of any species, such that we should think of attributions of agency to material entities a *default* representation.

Ristau's commentary raises the question of what accounts for the profound differences between humans and other animals in their ultimate conceptual repertoire? I assume we can all agree that language is one part of the answer. Ristau qualifies this assumption with the observation that other animals also have calls with referential content, so this aspect of language cannot distinguish us from other animals. But clearly no non-human animal has a system of external symbols with the properties of human language. TOOC details the role language plays in Quinian bootstrapping, which in turn plays a role in the expansion of our representational repertoire.

3.3 Number Representations

Landy, Allen, and *Anderson* argue that number representations are neither modular nor domain-specific, because they profligately draw on representations of space. They review many fascinating phenomena that show that numerical representations share computational machinery with spatial ones. The question then arises, for each, whether it reflects mappings between the domains that occur in evolutionary or ontogenetic time.

I agree that computational machinery is recycled over evolutionary time, and that systems that evolved under selection pressures for spatial representations were later drawn upon for other purposes, including numerical representations. The worked-out example in TOOC is analog magnitude representations, which are used in representing many different dimensions of experience (brightness, loudness, duration, length, area) in addition to representing cardinal values of sets; there is even evidence of a common neural substrate for several of these representational systems (Walsh, 2003). But these facts do not undermine the claim of innate, domain-specific number representations, so long as the infant does not confuse mental symbols for number with symbols for other dimensions of experience, which infants do not, and so long as the numerical symbols have further innate numerical conceptual role, as is the case. The domain specificity in question here concerns content, not the nature of (or even the neural substrate of) the computational machinery.

Landy et al. endorse TOOC's claim that conceptual discontinuity in the course of ontogenesis also involves recycling old representations and computational capacities in new domains. Quinian bootstrapping is a specific proposal for one way this is actually accomplished. Such ontogenetic recycling also does not undermine the domain-specificity of innate systems of core number cognition.

4. On the possibility and extent of conceptual discontinuity

The commentaries express two incompatible sentiments with respect to conceptual discontinuity—some deny it (*Rips and Hespos*) and some say it is much more common and comes in many more varieties than suggested in TOOC (*Gopnik, Laurence and Margolis, Weiskopf*). I am persuaded by the commentaries that express the second sentiment.

Theories of conceptual development face two explanatory challenges; specifying how cognitive development is possible and specifying why, in some cases, it is so hard. The second half of TOOC concerns some of the hard cases, where years of input is misanalyzed, and the target conceptual system sometimes never grasped. There are many reasons conceptual development might be slow (including lack of access to relevant input), and one of them is that mastering the target system of representation requires building a representational system that is discontinuous with its input. In each case study TOOC characterizes the two successive conceptual systems, CS1 and CS2, specifying the senses in which CS2 is discontinuous with CS1—having more expressive power, being incommensurable, or both.

Clearly, if we learn or construct new representational resources, we must draw on those we already have. *Rips and Hespos* characterize discontinuity as involving the construction of "entirely new systems" of concepts. Well, yes, in some sense. Discontinuities involve creation of new representational primitives and new systems of concepts articulated in terms of those primitives. For example, the concepts, *weight, density, volume*, and *matter* are interdefined and acquired en suite. Similarly for the concepts *fraction* and *division*. But the bootstrapping processes that explain how new representational resources are constructed do draw on already existing representational resources. In that sense they are not "entirely new." Nonetheless, in cases of discontinuity, CS2 cannot be translated into the language of the CS1, and so cannot be learned by testing hypotheses stated in that language.

TOOC explicitly discusses the challenge Rips and Hespos offer to the very possibility of incommensurabilty. As Rips and Hespos put it, if bootstrapping is possible, then it must be possible to characterize a function from CS1 to CS2, and hence to translate between CS1 and CS2. The key issues here are what that function is, and whether it counts as a "translation" (Kuhn, 1982). To translate is to express a proposition stated in the language of CS2 in the language one already has (CS1), as in translating "Je suis heureux" into "I am happy." In cases of discontinuity where Quinian bootstrapping is required, this is impossible. Bootstrapping is not translation; what is involved is language construction, not translation. That is, within CS1 one constructs an incommensurable CS2 that is not translatable into CS1.

Both *Laurence and Margolis*, and *Weiskopf* point out that in addition to Quinian bootstrapping, there are also atomistic processes that create new representational primitives, Weiskopf (2008) makes a convincing case that these atomistic processes also result in increases of expressive power. I agree, and I also agree that what Weisskopf calls "coining" is always involved in the creation of new representational primitives. Discontinuities come in many different flavors, however, and I stand by my thesis that those that require bootstrapping are an important class. According to the analysis of Laurence and Margolis (2002), adding a new natural kind concept to one's repertoire depends upon the prior existence of a natural kind schema with the content "same natural kind as x" and filling in x with a kind syndrome (a prototype, a theory) for the kind x. This is clearly different process than learning that a quark is a new kind of subatomic particle, if one does not currently have the capacity to entertain thoughts about subatomic particles. It is this latter type of discontinuity, that which requires Quinian bootstrapping, that is my focus in TOOC.

5. On Bootstrapping

5.1 Bootstrapping in general

Hamin and Hernik emphasize that there is genuine conceptual development that does not require creation of new representational primitives, and that there are bootstrapping processes short of Quinian bootstrapping. I agree, and TOOC discusses syntactic and semantic bootstrapping processes in language acquisition as examples of bootstrapping processes that do not result in the creation of new representational primitives. Abstractly, all bootstrapping involves using mappings between different representations in the service of extending them in some specifiable way. Hamin and Hernik's proposal for the creation of new object kind representations, drawing on core cognition of objects, causal/functional analysis, and the capacity to create long term memory symbols would be a bootstrapping process by this characterization (and is essentially what I propose as well in Chapter 7 of TOOC). My account of the creation of long term memory models of small sets of individuals that support the meanings of "one," "two," "three" and "four" in the subset-knower stage is another bootstrapping process. Working out exactly how such processes work is an important project in the field of cognitive development. But one major thrust of TOOC is that there are episodes of conceptual development that require more.

5.2. Must the placeholder structure be articulated in public symbols?

Placeholder structures consisting of semantically impoverished symbols are the key to how Quinian bootstrapping differs from the bootstrapping processes involved in the episodes of learning described in 5.1. The placeholder symbols gain whatever initial meaning they have from their relations to each other. Several commentators (*Heinz, Weisskopf, Allen & Bickford*) wonder why it is important that the placeholder structure be encoded in the external public symbols of language, mathematical symbols, diagrams and the like.

Allen and Bickford make use of the possibility that placeholders may be mental symbols to counter my argument for innate representational primitives. They suggest that infants might create a whole suite of interrelated new mental symbols, the content of each is exhausted by its role in the mental structure. By hypothesis, none is constructed from already existing concepts. I agree this is a *logical* possibility, but what would ever lead an infant to do this?

Encountering a new external symbol from others' use of it (e.g., first hearing the word "mass" in the context of the sentence "force equals mass times acceleration," or "two" in the context of the count routine) is often the impetus for an individual's coining a new mental symbol. As Weiskopf points out, this is not the only impetus to create a new symbol in mentalese. I agree with these commentators that from the point of view of how Quinian bootstrapping generates new representational resources, it is not necessary that the symbols be public. It is certainly logically possible that a whole placeholder structure of interrelated empty symbols could be generated in the language of thought and used to model other domains of representations that already have wide content. As I said, the question is what might lead this to happen. Before I can even begin to evaluate this possibility, I'd need to see a plausible worked example.

In all of the case studies of Quinian bootstrapping in TOOC, the placeholder structures *were* external symbols, probably for several reasons. As emphasized by Heinz, conceptual change is a social process, fueled by communication. But appeals to communication, even with its assumptions of relevance, will not do all the work we need here, for what's at issue is how listeners construct the representational resources to *understand* what's being communicated. The role of placeholder structures in conceptual change may also depend upon their being public. Some public symbols may be easier to think with. Maxwell explained he used diagrammatic models that he knew captured the mathematics of Newtonian forces in a fluid medium (his placeholder structure) in his modeling of the empirical phenomena involving electricity and magnetism just this way: they are easier to think with (see Nersessian, 1992). Finally, the modeling processes involved in conceptual change unfold over years, and the placeholder structures need to be stable; external symbols may facilitate that.

Although conceptual change is a social process, its first step, contrary to Weisskopf, is not always learning *from others* a set of new symbols and their interrelations. This is usually so in development, but obviously never so in the bootstrapping episodes in cultural history (see Chapter 11 of TOOC). In the case of Maxwell, an abductive guess led him to explore the hypothesis that the forces at work in electromagnetic phenomena were similar to the Newtonian forces in a fluid medium. That is, *nobody* was using the language of electromagnetic theory; its theoretical terms and equations were the *output* of the bootstrapping episode. In the case of historical bootstrapping, the placeholders derive from conceptual structures created in other contexts and seen to be relevant to the domain at hand through an abductive leap.

5.3 The prehistory of number representations

Overmann, Wynn, and Coolidge consider the construction of explicit representations of natural number in cultural history. They assume, and I agree, that external tally systems (lines on clay tablets, notches on sticks), which are widespread both in the archeological and in the ethnographic records, were the first step in the process. Such external symbol systems constitute an example of "extended cognition," in which external objects take on symbolic functions. Note also that tally systems represent number as does parallel individuation, with the external tallies making up for the limits on working memory. That is, each tally stands in 1-1 correspondence with an individual in the set, such that the set of tallies as a whole represents the cardinality of the set. Overmann et al. speculate that strung beads, which go back in the archeological record for 100,000 years may have been an artifact co-opted for this use.

This may be, but is there any evidence that beads were ever used as external tallies? If so, understanding the invention of tally systems would involve understanding how people came to the insight that beads could serve *this* symbolic function, rather than decorative uses, or as markers of wealth, or myriad others. That is, the availability of an artifact that *could* serve as the medium of a tally system doesn't explain how it came to be one. Now that we are in the realm of speculation, I believe, contra Overmann et al., that body counting systems could well also play an extended cognition role in the cultural construction of integer representations. Bodies, before beads, may have been the basis of the first tally systems. Beads may go back 100,000 years, but fingers go back millions, and finger tally systems are also widely attested in the ethnographic literature. But even so, the question then becomes how the insight arose that fingers could be used in a tally system based on 1-1 correspondence.

5.4 Acquiring representations of natural number

Gelman, Spelke, and Gentner and Simms engaged the nitty-gritty issues concerning conceptual discontinuities in and bootstrapping of natural number representations. These debates might be hard to follow for those not steeped in the dialog, but here I try to bring out the main issues.

Gelman's and my disagreements coexist within broad agreement on the big theoretical picture, and all of my work on number starts with hers. There are two major disagreements: 1) She presupposes that there is only one innate non-verbal representational system that plays a role in the extended developmental process we both endorse ("*the* inherited…"), whereas I believe that there are other innate resources and that they are more important in the initial construction of integer representations; and 2) I argue that innate number representations are discontinuous with representations of natural number, including verbal numerals deployed in a counting algorithm, whereas she argues that children understand how counting represents number upon first learning how to count.

Gelman and I agree that analog magnitude number (AM) representations are continuous throughout development, support arithmetic computations, and compute representations of small sets as well as large ones, both in infancy and adulthood. (TOOC, Chapter 4). But, contra Gelman, I believe the parallel individuation system also has a big role to play in the construction of natural number. This system is also continuous throughout development, and also has numerical content, in the form of numerically relevant computations carried out over models of small sets. Additionally, there are other representational resources with innate support that contribute to the construction of representations of natural number. These include the logical capacities that underlie the semantics of natural language, including quantifiers. It is clear that the resources for building the culturally constructed representations of natural number include much more than those provided by the AM system. TOOC argues that none of these three systems of number representation, on their own, *contains* representations of natural number (there are no symbols for any natural numbers, for example) in any of these systems. Therein lies the discontinuity. In her recent writings Gelman acknowledges this; the AM system contains no representation of exactly one, and is not built on the successor function (Leslie, Gelman & Gallistel, 2007). Thus, AM representations are discontinuous with the numeral list representation of natural number, which has both of these properties. In creating a numeral list, or a symbolic tally system, for that matter, cultures create a representational system that can express thoughts and support arithmetic calculations that were not available to the infant or non-human animal, and when individuals master these culturally created systems, they similarly extend the expressive power of their representational repertoire.

The discontinuity at stake in Gelman's and my debate is between the numerical concepts expressed by the verbal count list and those in core cognition. What the child is creating goes beyond a way of talking about what they already represent. I stand by the evidence that learning how counting represents number is extremely difficult, and by the evidence that children who are cardinal principle-knowers are qualitatively different from those that are not. Much of the data Gelman alludes to in her commentary reflects performance by children who already are cardinal principle-knowers (i.e., many 3-year-olds are), so cannot directly bear on this debate. It would be very interesting to explore children's performance on Zur and Gelman's intuitive arithmetic tasks, for instance, as a function of knower-level rather than age. Evidence is mounting that Gelman is right that TOOC underestimates the numerical understanding of children in the subset-knower stage, For example, before they have worked out the cardinal principle, some subset-knowers do know that numerals later in the list refer to larger cardinal values than do numerals earlier in the list (e.g., Shusterman et al., 2009). How counting represents number is worked out in small steps during the subset-knower stage. Still, subsequent work has also confirmed that only CP-knowers understand, even if implicitly, how counting implements the successor function (e.g., Sarnecka and Carey, 2008; Shusterman et al., 2009). Thus, the basic bootstrapping story stands.

Spelke, like Gelman, believes I have underestimated the role of the AM system of number representation in the creation of public symbols for natural number. She offers an alternative bootstrapping proposal, in which the numerical content in the parallel individuation system is combined with that in the AM system directly, not in the service of learning verbal numerals. Notice that we still must explain how children learn the meanings of verbal numerals and how they learn how the count list represents number. However, if Spelke's bootstrapping process exists, and were to take place prior to mastering counting, it might support learning to count.

In Chapter 7 of TOOC I consider two bootstrapping proposals similar to Spelke's, although in the service of mastering counting. One involves mappings between the count list and AM representations alone, and the other involves mappings between AM representations, parallel individuation and the count list. I reject the first on the grounds that AM representations do not contain representations of "one" or the successor function, and therefore cannot support the induction that there is ample evidence children make upon becoming CP-knowers—namely that two adjacent numerals in the count list represent sets that differ in cardinal values by one. The second proposal is very similar to that outlined in Spelke's alternative. I note if children had mapped numerals to AMs and induced the "later in the list, larger AM" generalization, this could aid in working out how verbal numerals represent number. I rejected this proposal on empirical grounds—there was no evidence that subset-knowers have made this induction, and some evidence from my own lab that they had not (Le Corre and Carey, 2007). However, since then, new evidence has come to light that perhaps they have done so (e.g., Shusterman et al., 2009). Nonetheless, learning how counting works requires inducing the relations between order in the list and the successor function, not only the relation between order in the AM system and the successor function.

Thus, I stand by my bootstrapping story for how children work out how *verbal numerals* represent number.

But is it possible that children and animals might create representations of *integers* through some process that does not involve mastering a count list? I am open to that possibility. Tally systems represent exact cardinal values without counting, but of course, they do not draw on Spelke's proposed bootstrapping mechanism. Rather, they also draw on the representational resources of parallel individuation. However, I doubt that such a process occurs in the absence of the construction of some public representational system (a tally, a verbal count routine). First, one would need to specify a context in which such a mapping would be constructed. Second, one would need to understand why this construction is not achieved by humans who do not have a count list or a tally system in their public symbolic repertoire (e.g., the Piraha: Gordon, 2004; Frank et al., 2008: the Munduruku, Pica et al., 2004 and homesigners: (Spaepen, et al., under review).

Finally, although mastering counting requires understanding the relation between order in the list and numerical order, I nowhere claim that children derive representations of numerical order only from the count list. Numerical order is implicitly captured in the computations carried out in the Parallel Individuation system, and easily read off from AM representations. Once children have integrated AM representations with counting, they can and do use AM representations to support arithmetic reasoning using symbols. Thus, I do not disagree with any of Spelke's arguments that AM representations support, even are crucial to, certain adult arithmetic computations. But evidence derived from 5-year-olds and adults does not bear on the process engaged in by 2-year-olds as they create the first public representations of natural number.

Analogy underlies one of the modeling processes through which representational resources are combined during episodes of Quinian bootstrapping. I agree with *Gentner & Simm's observation* that analogy has other roles to play in conceptual development as well. But I disagree slightly with their characterization of the steps in mastering verbal numerals. It is true that the child creates a mapping between sequential order in the numeral list and ordered quantities, but it is important to notice that the quantities so modeled (captured in enriched parallel individuation) are not yet natural numbers. Gentner discusses the creation of the first natural number representations as a redescription or re-representation of the two parallel relations so that they are one and the same relations: SUCCESSOR[numeral n, numeral n + 1) <- -> SUCCESSOR (set size n, set size n +1). I doubt whether the 3-year-old child has any representation of the successor function that labels the relation, i.e., has any mental symbol *successor*. Rather, I would characterize the last step of this particular bootstrapping episode differently. The order relation in the numeral list is still merely serial order in a list, but what the child now has done is analyze "five, six, seven…" as summary symbols "n, n + 1" by analogy to "one, two, three" and "four." Following up on Gentner's and Simm's central point, many more bootstrapping episodes are yet to come before the child has a full representation of natural number and the successor function.

6. On concepts and concept acquisition, in general

6.1 On the theory-theory of narrow content

A commitment to the existence of narrow content requires distinguishing between those aspects of inferential role that determine conceptual content and those that do not. *Keil* suggests that evidence concerning what he calls the degraded nature of our knowledge militates against the possibility of succeeding at this project, as well as militating against the theory-theory of concepts. One thing Keil means by "degraded" is "sketchy." I agree our explanatory theories are sketchy and sparse in the way Keil's research has shown us, but this may be irrelevant to the existence of narrow content if, as I believe, only a very small part of conceptual role is content determining. Narrow content does not exhaust a symbol's content, which is why we have representations of entities in the world about which we have the sketchiest of knowledge. There is also wide content; the motivation of the externalist movement was to explain how we borrow reference from those from whom we hear a term. The narrow content of a natural kind symbol like *dog*, is perhaps only that it refers to a natural kind, and thus falls under the assumptions of psychological essentialism, or perhaps some placement in a framework theory (e.g., that it is a kind of animal, and thus is constrained by a skeletal vitalist biology). That we hold contradictory beliefs about dogs is part of the motivation for this guess, not a problem for the existence of narrow content. That is, the existence of contradictory beliefs requires that we must be able to distinguish concepts from conceptions, and supports the hypothesis that content determining conceptual role is a small part of everything we know about the entities that fall under our concepts. These considerations apply to metaconceptually held scientific theories as well as intuitive ones: only a few relations among concepts distinguish the source-recipient theory of thermal phenomena from the caloric theory, although those relations allow separate concepts of *heat* and *temperature* only in the latter case.

Keil raises a very interesting issue not touched upon in TOOC—how studies of the breakdown of conceptual representations with age, disease, brain damage and the like, bear on our understanding of the nature of concepts (here "degraded" means damaged). Deborah Zaitchik and colleagues (Lombrozo et al, 2007; Zaitchik & Solomon, 2008, 2009) have found that patients with Alzheimer's Disease (AD) resemble preschool children on tasks that diagnose conceptual structure—they make animistic judgments that the sun, cars, and the like are alive, and they demonstrate promiscuous teleology. These findings leave open two very different interpretations: 1) AD affects domain general processes that underlie the participant's capacity to bring relevant conceptual resources to bear on a given task and 2) AD affects the conceptual resources themselves, those that underlie the knowledge, and perhaps even the concepts, in the domain. For instance, the disease might disrupt the skeletal inferential structures that constitute a vitalist biology, such that the participants no longer have a concept of living things. Getting to the bottom of which interpretation is correct will be important to our understanding of the nature of concepts, perhaps even bearing on which aspects of conceptual role are content determining.

Contrary to *Korman's* assumption about what the theory-theory of concepts comes to, TOOC nowhere claims that conceptual changes occur only in the context of formal scientific development. Rather, the idea is that there are important commonalities between the conceptual changes in the history of science and those that occur in childhood. That is, many questions about conceptual development receive the same answers in both cases. These include: how is conceptual change distinguished from belief revision, what is an "undifferentiated concept," and what is the role of Quinian bootstrapping in episodes of conceptual change? Korman advocates the existence of smaller scale conceptual changes than those involved in the case studies in TOOC. I am open to that possibility (see above, Section 4). Before I would be able to evaluate the proposal, I would need to know in what sense a candidate developmental change is a conceptual change as opposed to a change in belief. TOOC provides such an analysis for the conceptual changes in its case studies, but Korman gives no examples of what kinds of changes she has in mind.

Although I mention the proposal that the causally deepest features of a concept are likely to be those that are content determining, I don't particularly endorse it. I agree with Korman's reading of the literature on the causal status hypothesis concerning categorization decisions, but I am skeptical about making categorization decisions the central phenomenon in our exploration of the nature of concepts. My positive proposal is quite different; namely that in at least some circumstances narrow content is determined by the conceptual role that is exhausted by the placeholder structures in which new primitives are introduced.

6.2 On the relations between the philosophers' and the psychologists' approaches to understanding concepts

Machery's commentary raises a deep issue, on which we have opposing views. He believes the traditional cognitive psychologists' project (understanding categorization, inference, conceptual combination) and the philosopher's project (understanding reference, truth, propositional attitudes) are completely distinct, whereas I do not. My appeals to a part of the philosophical literature were aimed at sketching some of the arguments for wide content. I also showed that *psychological* work supporting psychological essentialism, explanation based reasoning, and conceptual development converges with these arguments. So while he and I agree that the psychological and philosophical literatures emphasize different phenomena, many philosophers also endorse that understanding inference, categorization, conceptual combination and conceptual development (omitted from Machery's list) are part of the explananda of a theory of concepts. One goal of TOOC was to encourage psychologists to see that reference, truth, and propositional attitudes are as well. Machery suggests that I endorse information semantics, while failing to distinguish it from other approaches to wide content. Instead, TOOC endorses a dual factor theory and I draw from information semantics, the causal theory of reference, and Milliken's work as different approaches to wide content. I distinguish among them, but do not choose between them.

Machery proposes a solution to the problem of distinguishing concepts from conceptions. He suggests that the content determining aspects of the representations that underlie categorization and inference are those that are activated by default, independent of context, whenever we use a given concept. Insofar as the representational structures being considered here are prototypes, definitions, representations of exemplars, and/or intuitive theories, I am deeply skeptical of this proposal. I would bet dollars to donuts that there are no such default representations activated independent of context. Indeed, Machery's (2009) masterly review of the literature in the service of his argument that cognitive psychology's concept *concept* is not a natural kind undermines the existence of such default representations.

6.3 On the possibility of a dual factor theory of concepts

I disagree with *Rips and Hespos* that my version of dual factor theory collapses to a single factor. I offer general arguments for wide content, and I assume that the mechanisms though which it is determined will be complex and messy indeed, encompassing aspects of social/causal history, as well as the machinery envisioned by developers of information semantics. The last chapter of TOOC sees the challenge to be defending that there is any narrow content at all. Although narrow content may play some role in the processes through which wide content is determined (e.g., by constraining the nature of the causal connections between certain classes of symbols and the entities they represent), it in no way exhausts them. Nonetheless, Rips and Hespos raise an important point—any workable dual factor theory will have to specify how the two factors work together.

Mandelbaum also raises concerns about the very possibility of a dual factor theory of concepts. He points out that my proposal for specifying the narrow content of symbols in core cognition differs from my proposal for how to do so in the case of concepts that arise through episodes of Quinian bootstrapping, and he rightly points out that neither proposal applies to newly acquired atomic concepts, such as, taking his example, *fence*. I agree. I believe it is unlikely that there is one solution to the problem of specifying what aspects of conceptual role determine narrow content. Different kinds of concepts will require different solutions. The theoretical challenge is to find the principled distinctions among types of concepts, as well as to specify what aspects of conceptual role determine narrow content of each type.

I do discuss the question of how newly acquired atomic concepts might get their content, focusing on the proposal by Laurence and Margolis (2002) and of Macnamara (1986); see also Weisskopf, 2008). Their example is a natural kind term, such as *tiger*. The narrow content of a newly learned public symbol "tiger" or a newly coined mental symbol *tiger* is *same natural kind as [kind syndrome of a tiger]*, where the kind syndrome might be represented as a prototype, set of exemplars, or intuitive theory. The work done by the abstract schema for a natural kind is that the holder of this representation is open to most everything in the kind syndrome being revisable; that is, this concept falls under all of the assumptions of psychological essentialism and the division of cognitive labor of deferring to tiger experts. So the relevant aspect of conceptual role that determines content here is not the whole kind syndrome, but that which determines the nature of the function between the symbol and the world, including that which determines that the symbol is a symbol for a natural kind. An analogous solution could apply for *fence*. My guess is that there is an artifact kind schema that supports meanings like *same artifact kind as [kind syndrome for a fence]*, where the information in the kind syndrome is taken to be tentative and revisable (i.e., not itself content determining). The schema specifies the importance of function and original intent of the designers, leaving open that current representations of these may be wrong. On this view, it would be possible to have the lexical item "fence" in one's vocabulary without yet having assigned it to that minimal schema (e.g., if all one knew about fences are that they are something that make good neighbors—after all, good intentions and warm feelings make good neighbors too, but these are neither artifact kinds nor fences).

6.4 .Is concept learning "explicable by content?"

Shea raises the possibility that TOOC provides a recognizably psychological explanation of the acquisition of concepts that does not fit the dominant explanation schema in psychology. Specifically, he suggests that my accounts of concept acquisition do not rely (solely) on rational inference nor explanation-by-content. I am not sure what philosophers mean by "explanation by content," so my response is tentative. As I understand it, explanation by content is explanation in terms of non-metalinguistic inferences over mental representations alone, as when we explain an agent's opening the refrigerator by appeal to her desire for an apple and her belief that there is one in the fridge. This is explanation by content because it makes no appeal to the truth of the matter of the location of an apple.

I agree with Shea's gloss of the bootstrapping process that underlies the transition between the subset-knower stage and cardinal principle-knowers: "Carey rightly rejects the idea that it is properly described as hypothesis testing. This is doing something different: building an uninterpreted model with the counting words, and then giving that model an interpretation." My "uninterpreted" here refers to wide content; the model is not yet mapped to numbers. However, I believe that cognitive modeling is a rational process, and the cognitive model has narrow content given by inferential role, and so in one sense of "explanation by content" this is an explanation by content. Shea criticizes this proposal because it would lead to pernicious holism. I do not see why. This proposal is important to my characterization as to which aspects of inferential role determine narrow content in the case of concepts whose origin lies in episodes of Quinian bootstrapping. And the holism is not pernicious, because not every aspect of the conceptual system is potentially content-determining.

When the child induces that the each numeral represents a cardinal value one more than the preceding one, she confirms that guess against further input. That is, cognitive modeling is a rational process. Importantly, the numerical content of the count routine is partly implicit, just as the numerical content of parallel individuation is implicit. Representations of the successor function, insofar as they are attributable to the newly minted cardinal-principle knower, are carried in the counting principles the child masters. That is, learning to count is learning a skill. So the question I would ask Shea is whether skill acquisition is explainable by content? If not, it is easy to find violations of the principle that all explanation in psychology is explanation by content. If so, then learning the count-list is not a counter-example to the principle. My instinct is to favor the latter possibility.

The same questions arise with respect to Shea's gloss of how the child becomes a "one"-knower, and then a "two"-knower. What the child is learning is a procedure, and the procedure she builds draws on representational and computational resources that are part of core cognition. That is, procedurally, the child does have the resources out of which a hypothesis about one-ness could be constructed (even though she has no symbol for *one* that could articulate a hypothesis about what the word "one" means.) The enriched parallel individuation account of the meaning of "one" in the subset-knower stage is that the child makes a long term memory model of a single individual, a model that is identical in format to core cognition's working memory models formulated in terms of individual files, and acquires the following procedure for applying the word "one:" if a working memory model of a given set can be put in 1-1 correspondence with that long term-memory model, then "one" applies to that set. Similarly for "two." The process by which the child builds these procedures is rational; it does not depend only on correlations mediated by causal connections to the same external-world property.

7. Concluding Remark

I thank my commentators again for their clarifying remarks. The controversies aired in our debates confirm the interdependence of the projects of explaining the origin and nature of concepts.

References

Bowerman, M. Papers and Reports of Child Language Development. Vol. 8. Palo Alto, CA: Department of Linguistics, Stanford University; Learning the structure of causative verbs: a study in the elationship of cognitive, semantic and syntactic development; p. 142-178.

Carey, S. Conceptual Change in Childhood. Cambridge, MA: Bradford Books, MIT Press; 1985.

Carey S. Does learning a language require the child to reconceptualize the world? Lingua. 1994; 92:143–167.

Dewar K, Xu F. Induction, overhypothesis, and the origin of abstract knowledge: evidence from 9-month-old infants. Psychological Science. in press.

Diamond R, Carey S. Why faces are and are not special: An effect of expertise. Journal of Experimental Psychology: General. 1986; 115:107–117. [PubMed: 2940312]

Feigenson L, Halberda J. Infants chunk object arrays into sets of individuals. Cognition. 2004; 91:173–190. [PubMed: 14738772]

Frank MC, Everett DL, Fedorenko E, Gibson E. Number as a cognitive technology: Evidence from Piraha language and cognition. Cognition. 2008; 108:819–824. [PubMed: 18547557]

Gauthier I, Tarr MJ. Unraveling mechanisms for expert object recognition: Bridging brain activity and behavior. Journal of Experimental Psychology: Human Perception and Performance. 2002; 28(2): 431–446. [PubMed: 11999864]

Gordon P. Numerical cognition without words: Evidence from Amazonia. Science. 2004; 306(5695): 496–499. [PubMed: 15319490]

Kellman, PJ.; Arterberry, ME. The Cradle of Knowledge: Development of Perception in Infancy. Cambridge, MA: MIT Press; 2000.

Kuhn, T. The Essential Tension: Selected Studies in Scientific Tradition and Change. Chicago: University of Chicago Press; 1977.

Kuhn, T. Proceedings of the biennial meeting of the Philosophy of Science Association, 2: Symposia and invited papers. East Lansing, MI: University of Chicago Press; 1982. Commensurability, comparability, communicatability; p. 669-688.

Laurence S, Margolis E. Radical concept nativism. Cognition. 2002; 86:25–55. [PubMed: 12208650]

Le Corre M, Carey S. One, two, three, four, nothing more: An investigation of the conceptual sources of the counting principles. Cognition. 2007; 105:395–438. [PubMed: 17208214]

Leslie, AM.; Gelman, R.; Gallistel, CR. Where integers come from. In: Carruthers, P.; Laurence, S.; Stich, S., editors. The Innate Mind, Vol. 3: Foundations and the Future. Vol. 7. New York: Oxford University Press; 2007. p. 109-138.

Lombrozo T, Kelemen D, Zaitchik D. Inferring Design: Evidence of a Preference for Teleological Explanations in Patients With Alzheimer's Disease. Psychological Science. 2007; 18(11):999–1006. [PubMed: 17958715]

Machery, E. Doing Without Concepts. Oxford: Oxford University Press; 2009.

Macnamara, J. Border dispute: The place of logic in psychology. Cambridge, MA: MIT Press; 1986.

Mandler J. A leaner nativist colution to the origin of concepts. REF TO Commentary in the present volume (don't know how to do that.). 2011

McLaren IPL. Categorisation and perceptual learning: an analogue of the face inversion effect. Quarterly Journal of Experimental Psychology. 1997; 50A:257–273. [PubMed: 9225623]

Nersessian, N. How do scientists think? Capturing he dynamics of conceptual change in science. In: Giere, R., editor. Cognitive Models of Science. Minneapolics: University of Minnesota Press; 1992. p. 3-44.

Pica P, Lemer C, Izard V, Dehaene S. Exact and approximate arithmetic in an Amazonian indigene group. Science. 2004; 306(5695):499–503. [PubMed: 15486303]

Rozenblit L, Keil F. The misunderstood limits of folk science: An illusion of explanatory depth. Cognitive Science. 2002; 26:521–562. [PubMed: 21442007]

Sarnecka B, Carey S. How counting represents number: What children must learn and when they learn it. Cognition. 2008; 108(3):662–674. [PubMed: 18572155]

Sarnecka B, Gelman S. Six does not just mean a lot: preschoolers see number words as specific. Cognition. 2004; 92:329–352. [PubMed: 15019554]

Shusterman, A.; Carey, S.; Spelke, ES. Two Paths in the Development of Counting and Cardinality. Symposium talk given at the Biennial Meeting of the Society for Research on Child Development; Denver, CO. 2009.

Tenenbaum, JB.; Griffiths, TL.; Niuyogi, S. Intuitive theories as grammars for causal inference. In: Gopnik, A.; Schulz, L., editors. Causal Learning. New York: Oxford University Press; 2007. p. 301-322.

Walsh V. A theory of magnitude: common cortical metrics of time space and quantity. Trends in Cognitive Sciences. 2003; 7(11):483–488. [PubMed: 14585444]

Weiskopf D. The origins of concepts. Philosophical Studies. 2008; 140:359–384.

Wiser, M.; Carey, S. When heat and temperature were one. In: Genter, D.; Stevens, A., editors. Mental Models. Hillsdale, NJ: Erlbaum; 1983. p. 267-297.

Xu F, Tenenbaum JB. Word learning as Bayesian inference. Psychological Review. 2007; 114:245–272. [PubMed: 17500627]

Zaitchik D, Solomon GEA. Animist thinking in the elderly and in patients with Alzheimer's disease. Cognitive Neuropsychology. 2008; 25(1):27–37. [PubMed: 18340602]

Zaitchik D, Solomon GEA. Conservation of species, volume, and belief in patients with Alzheimer's disease: The issue of domain specificity and conceptual impairment. Cognitive Neuropsychology. 2009; 26(6):511–526. [PubMed: 20043252]