King's Research Portal

DOI:

[10.1007/s10539-009-9187-5](https://doi.org/10.1007/s10539-009-9187-5)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Shea, N. (2011). New concepts can be learned. *BIOLOGY AND PHILOSOPHY*, 26(1), 129-139.
<https://doi.org/10.1007/s10539-009-9187-5>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

**Open Access document
downloaded from King's Research Portal
<https://kclpure.kcl.ac.uk/portal>**

Citation to published version:

Shea, N. (2011). New concepts can be learned. *BIOLOGY AND PHILOSOPHY*, 26(1), 129-139, doi: 10.1007/s10539-009-9187-5

The published version is available at:

DOI: 10.1007/s10539-009-9187-5

This version: Preprint

URL identifying the publication in the King's Portal:

<https://kclpure.kcl.ac.uk/portal/en/publications/new-concepts-can-be-learned%28b6022507-c28d-4bf2-b043-68727fe9a89c%29.html>

The original publication is available at www.springerlink.com

The copyright in the published version resides with the publisher.

When referring to this paper, please check the page numbers in the published version and cite these.

General rights

Copyright and moral rights for the publications made accessible in King's Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications in King's Research Portal that users recognise and abide by the legal requirements associated with these rights.'

- Users may download and print one copy of any publication from King's Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the King's Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

New Concepts Can Be Learned

Review of *The Origin of Concepts*, Susan Carey

Oxford: Oxford University Press, 2009

Biology & Philosophy, doi:10.1007/s10539-009-9187-5

Nicholas Shea

Faculty of Philosophy

University of Oxford

OX1 4JJ

Abstract

Many have doubted whether the transition to genuinely new representational resources is susceptible to psychological explanation. In *The Origin of Concepts* (O.U.P. 2009), Susan Carey makes a strong empirical case for the existence of discontinuities in conceptual development. Carey also offers a plausible psychological explanation of some of these transitions, in particular of the child's acquisition of the ability to represent natural numbers. The combination amounts to a forceful answer to puzzles about the learnability of new representations.

Keywords

Carey; *The Origin of Concepts*; conceptual development; learnability; representation; Fodor

The philosopher of mind who stands convinced of the relevance of empirical results soon hits a problem. Lured in by a few interesting studies, the door opens on a cacophony of data, like a frenetic party in full swing. There are just so many studies out there. How do they fit together? And what do they all add up to? The young Darwin faced the same problem.¹ In *The Origin of Concepts*, Susan Carey comes to our rescue. In the tradition of the other *Origin*, Carey's book offers a theory which turns a mass of facts into a coherent story.

Much of this substantial book is devoted an array of fascinating studies, from Carey's own lab and from others in the discipline. Each is explained vividly and deployed to good dialectical effect. The local arguments are judicious, measured and unpolemical, and they add up to a bold and interesting narrative. For Carey has an answer to a conundrum that lies at the heart of cognitive science, one which made trouble for the discipline almost as soon as an alliance of psychologists, linguists, philosophers and computer scientists decided to work together to make good on the promise of the cognitive revolution. It is the puzzle of how humans could ever acquire new representations. The dreadful fear is that these allied disciplines lack the means to explain how people develop genuinely new representational

¹ 'I had at last got a theory by which to work', he exclaimed, after many years collecting data on the origin of species (Darwin 1887/1958, p. 120).

resources, with expressive power that outstrips the representations they start with. Is there a computational or psychological story that will explain such transitions?

Philosophers will be familiar with the worry in the guise of Fodor's argument for radical concept nativism. As a developmental psychologist, Carey is part of a tradition going back to Piaget that is invested in cataloguing the striking and sometimes strange conceptual changes through which the child progresses on its way to the familiar representational repertoire of adulthood. The project is a lot less exciting if such changes must for ever remain unexplained by the psychologist, fun as it is to do experiments with kids. If Fodor is right, developmentalists must be content with the modest project of documenting the order in which various innate representational capacities mature and come on stream.

Fodor's argument has a theoretical part and an empirical part. The theoretical part claims that all computational or psychological accounts of the learning of a new representation must start with some representations as input, and that the most such accounts can deliver as output are representations constructed out of those initial resources. The empirical part observes that most lexical concepts are not structured or constructed out of any other representations. They are atomic. So they must be innate – accounting for their acquisition is outside the remit of cognitive science, beyond the explanatory power even of the broad array of disciplines that have joined in the project.

It is a deep challenge, and Carey's riposte has not been made in haste. It too has a theoretical part and an empirical part. The empirical part consists of decades of careful studies, each of which taken alone may be debateable, but the combination of which adds up to a decisive demonstration that there really are discontinuities in the development of the human representational repertoire. The theoretical part is an ingenious argument to show that psychology can explain such transitions. Carey offers us a mechanism, Quinian bootstrapping, which gives rise to concepts with genuinely new expressive powers. By giving a psychological, quasi-computational account of the operation of this mechanism, Carey demonstrates that such explanations are firmly within the remit of cognitive science. The combination of theoretical insight and empirical support makes *The Origin of Concepts* a very satisfying whole. Taken just as a work of psychology, it is a valuable compendium. Taken just as a work of philosophy, it is an important contribution to a central debate. The book's remarkable achievement is to be both.

The Concepts of Core Cognition

Taking a tour of Carey's psychological garden there is much that is familiar. There are modules that transduce ambient energy into sensory and perceptual representations. These modules have been designed by evolution to output representations that covary with various sensory and perceptual properties in the world: objects are that are hot, red, square and over there. Carey's view is that their content derives from the causal connections between representations and entities in the world sustained by such modules. And they are innate: there is no learning story to be told about how the modules come to be configured as they are.

Also familiar are concepts, representations deployed in thought, applied by inference from sensory and perceptual evidence, and themselves embedded in rich networks of further inferences. Less familiar will be the representations of 'core cognition'. A central example is

the system for representing small arrays of objects in parallel and keeping track of their numerical identity as they move (object files). Another number-related example is the representation of the approximate number of items in a larger set of objects. Core cognition also encompasses the way we represent some entities as agents and represent their actions as goal-directed.

Carey argues that these kinds of representations deserve an entry of their own in the psychological inventory. They are like perceptual representations in being produced by modules. They take perceptual representations as input and produce conceptual representations as output – of objects, agents, etc. – which are then deployed further in our thinking. Their online operation is informationally encapsulated from the rest of cognition and they carry on working in the same way continuously throughout development. Carey also thinks that such modules are innate. Implicit in the computations performed by such modules are various assumptions, for example in the way they jump to the conclusion that an entity is an agent when it has such-and-such perceptible features. But they do not arrive at those assumptions through learning. When the module itself develops, it does not need to extract from the environment the information which will be implicit in the way it operates. That information is innate. As with perceptual modules, there is no psychological story to be told about how the modules of core cognition themselves develop.

Core cognition is also supposed to be like perceptual representation in being iconic (roughly: imagistic). That claim is open to doubt. Iconic representations are analogue: parts represent parts of the entity represented. If analogue magnitude representations are realised by some quantity in the brain, then they will indeed be iconic. Parts of that quantity will represent parts of the array represented (if _____ represents there being six objects in the array, then the part ____ represents three). Carey admits that the extension to representations of object and agent is more speculative.

Carey's idea here is that something like a perceptual image is deployed in a set of conceptual roles that gives it a content like *object* or *agent* from core cognition. The object file system may work like that. Each object file may consist of an imagistic representation of an object and its perceptible properties (size, shape, colour, etc.). It is operations on these icons – tracking, adding, subtracting and comparing sets by 1-1 correspondence – which makes them into representations of *objects*. If that is right, then the representations are indeed icons, representing properties like size, shape and colour iconically, but strictly speaking neither objecthood nor numerosity is being represented iconically, since those properties are represented only implicitly, in the inferences made using those icons.

Early representations of agency might work in the same way: agency could be represented implicitly in the kinds of inferences in which picture-like representations of objects are deployed (e.g. how such objects are likely to move, where they will look, etc.). If *agent* were instead represented explicitly, say by somehow predicating AGENT of an imagistic representation of a person, then the representation would be analogue with respect to the properties represented by the perceptual image of the person (e.g. size), but would not be analogue with respect to representing agency. It is hard to see how there are parts of the

property of *being an agent*, so that parts of a representation could represent parts of the property (in the way that ____ represents part of the analogue magnitude represented by ____).

One great puzzle in developmental research is why children fail to pass explicit theory of mind tasks until 4 years old, when in various experimental settings at much younger ages they betray implicit knowledge that in fact tracks others' mental states (e.g. in where infants look) (Apperly & Butterfill, forthcoming). Carey suggests that this puzzling disconnect could be explained by the representations of agency having a different format at the two stages. That point is well-taken, but what exactly would the difference have to be? Early representations of agency may be implicit in the inferences made with icons, as with numerosity in the object file system. As Carey helpfully makes clear, only explicit representations can act as input to further computations. That would explain why children fail many theory of mind tasks until they can represent agents and their goals explicitly. So rather than agency initially being represented explicitly and iconically, a rival hypothesis is that agency, goals, etc. are at first represented only implicitly, in the inferences made with iconic representations of objects.

Nevertheless, it is reasonably clear that at some stage children begin to have access to explicit representations of agents and their goals, i.e. representations that they can deploy for making further inferences. Why shouldn't we think of these as just being further properties that we can perceive? According to Carey, what distinguishes the representations of core cognition from perceptual representations is their richer conceptual roles. Representing an object as an agent gives rise to a richer array of inferential consequences than representing it as being red, although not as rich as the inferential roles of the concepts of explicit theories. There is a worry here about whether the inferential richness is really a matter of content, rather than mode of representation. Leaving aside that and other worries about how best to distinguish perception from core cognition, Carey makes a good case for the existence of an important class of psychological processes whose distinctive features are under-explored.

I would add other examples. Hearing the grammar of a sentence is one. The output is fast and mandatory. Although less encapsulated than other cases, it is still possible to hear a sentence as ungrammatical while simultaneously knowing that it is (think of garden path sentences). Perceptually-based classifications may be another case. I can see something as a duck, although I know it is just a line drawing. The mechanisms generating these kinds of representations are modular in some ways, but their deliverances also differ from those of perception. Phenomenologically, they seem to be in a different category, perhaps being less vivid than representations of colours, shapes, sounds, etc. Epistemically, too, they seem to have a different status. The deliverances of perception are epistemically unmediated. There is nothing for me to say at the personal level as to why I see a diagonal bar, say, although there is a rich information processing story to be told at the subpersonal level. Contrast seeing an object as a duck, where I can point to various perceptual features which seem to provide the basis of my perceptual judgement. Similarly, it seems explicable at the personal level why I see some shape as a single object (e.g. because the parts move together).

My claim that Carey is pointing out a distinctive and important psychological category raises two important open questions here. First, is consciousness required for the operation of

these mid-level representational systems? Do they take only consciously-represented perceptual properties as input? Second, if not, is it right to characterise their mode of operation (and not just their outputs) as being at the personal level, perhaps because facts about their operation can always be made conscious on reflection (unlike the inner workings of edge detection in early vision, say)?

Carey implicitly rejects what may be the best way of distinguishing the representations of core cognition from perception. Focusing on the idea that core cognition is conceptual, we can distinguish conceptual from non-conceptual representations in terms of constituent structure. Conceptual representations have a compositional semantics of something like predication. AGENT does not make claims on its own, but only when predicated of an object. If non-conceptual representations also have semantically-significant constituent structure, it is of a different sort. Icons obey the picture principle, so will have parts that represent parts of the thing represented. Each of these parts in turn has truth conditions of its own. Even when we get down to the smallest grain at which anything is represented, we still have vehicles with truth-evaluable contents. There is nothing like predication, no compositional semantics with different syntactic types playing different semantic roles.

If the deliverances of core cognition are indeed concepts – sub-propositional components of compositional structures – then we would have an explanation of why they have richer inferential roles than perceptual representations. It is because they act as a middle term in mediate inference. That way of drawing the distinction means giving up on the idea that the representations of core cognition are iconic (agency would be represented non-iconically, once it is represented explicitly at all). But it would show clearly how they differ from sensory and perceptual representations, while holding onto the idea that the mechanisms that produce them share with perception some aspects of modularity. It might also explain the phenomenological and epistemic differences suggested above.

Quinian Bootstrapping

Much of the book is occupied with spelling out the evidence for the existence of some innate representations with conceptual content: parallel individuation of small sets of objects or events; analogue magnitude representations of arrays of particulars; representations of agency, goal-directedness, etc.; and representations of causal relations. By characterising these systems so carefully, Carey is able to substantiate her claim that they differ from corresponding concepts in the mature adult repertoire. She can identify precisely where those differences lie. It adds up to a compelling case that there are indeed discontinuities in representational development. Which raises the question of whether such changes can be explained psychologically.

Carey's solution to Fodor's puzzle adopts a strategy that has been suggested in the literature (Laurence & Margolis 2002, Macnamara 1986). Assume that what makes a representation have the content it does is wholly or partly determined by its causal relations with things in the world. Such sustaining mechanisms for a representation R may depend, causally, on other representations R* without the content of R being determined by the content of R* – R's content is fixed more directly by its causal relations with things in the world. R can then be atomic, neither structured nor constructed out of R*. The anti-nativist tactic is to give

a psychological story in which existing representations R^* come to form the sustaining mechanism for a new representation R , where the process of forming the new representation type R is described merely causally, not as a content-driven process like inference.

To assess whether that tactic works, we need to see it instantiated less abstractly in a theory. Carey has one: Quinian bootstrapping as a way of acquiring the concepts of natural numbers. On one side of the conceptual discontinuity are various core cognition systems with number-related content. On the other side is a system of representations of natural number which are characterised partly by wide content and partly by conceptual role. The wide content of the second-stage representations is given by the fact that they correlate with the numerosity of sets of objects or other particulars. So there is a representation whose tokening correlates reliably with there being five objects present to perception (not tracking total surface area or density, say). Similarly, there is a representation for singleton sets (ONE). The narrow content is given by a set of conceptual roles linking these representations: adding two sets (PLUS), finding the difference between two sets, and judging equal numerosity by means of 1-1 correspondence. A crucial conceptual role is the successor relation, which orders the representations so that performing the operation PLUS ONE on any representation outputs the next representation in the sequence.

The set of conceptual roles that, according to Carey, defines the narrow content of the concept of a natural number outstrips anything that is available in the number-related systems of core cognition. For example, neither analogue magnitudes nor parallel individuation has the successor relation. The starting point for Carey's developmental transition is the parallel individuation system ('object files' for short, although the system can also individuate events). That system allows infants to judge whether two arrays of up to four items have the same numerosity by representing them in parallel and checking for 1-1 correspondence. For example, the infant can tell that an array of four objects differs in numerosity from the series of sets of threes they have just been shown. The system also supports addition and subtraction up to the same limit: if one object is added to a set of two that the infant has previously seen hidden behind the screen, they will behave differently if the screen is removed to reveal four or two objects than if the expected threesome is revealed.

To this Carey adds grasp of the grammatical distinction between singular and plural, mapped onto a semantic distinction between arrays with one vs. many items. Finally, Carey adds the crucial ingredient, which is a cultural creation: the sequence of counting words ("one", "two", ...). Children learn this list initially as an uninterpreted series of sounds, much like a nonsense rhyme or the alphabet. (Learning the alphabet, many children treat 'elemeno' as a single item.) The counting sequence provides a route into the successor relation. Carey's insight is that it can be acquired as a series of non-representational symbols. That is Quinian bootstrapping. A series of symbols are put into a set of causal relations by rote, habit or some other causal device. Initially the causal transitions are not inferences and the symbols are not representations. Only when the symbols get wired up to the world do they acquire a content and become representations. Pre-existing representations play a causal role in setting up the sustaining mechanisms. But pre-existing symbols cannot exhaust the content of the new

representations, since their content is in part determined by what have now become inferential relations between them.

An analogy may help. As a child I read the *Swallows and Amazon* books before I had even been near a sailing boat. I picked up a lot of sailing terminology: tiller, mainsheet, halyard, and so on. And I learnt some of the inferential relations: when you're luffing, pull on the mainsheet or pull the tiller towards you; when you're in irons hold the boom out and the tiller away; etc. But in a sense these words were all uninterpreted until I got into a boat and learnt what to apply them to. Once interpreted, the various inferential relations were already of some use: the first time I found myself in irons I already knew what to do. At the heart of Carey's theory is the idea of memorising a set of relations amongst uninterpreted symbols.

The trick is then to get the words to correlate with things in the world in the right way. The syntactic singular is a step onto the first mapping. Syntactic singular is already associated with object files with one element. The child only has to notice that "one" gets used in the same contexts as "a" or the singular in order to associate "one" with singleton object files. This gives an 'interpretation' to the first word in the counting sequence: "one" now gets applied when and only when parallel individuation contains a singleton. Then the child has to notice that "two" gets used exclusively for sets of 2 things. Then the word "two" will come to be activated whenever the child has just 2 object files open in a set in working memory ("two" gets used interchangeably with $\{i\ j\}$). Carey adduces evidence that children learn meanings for the first four words in the count sequence in this laborious way. (In languages that have grammatical dual or triple markers, those syntactic distinctions help too.) Children start as 'one-knowers' who can give you "one" object, but respond at random with several objects when asked for "two", "three", etc. Then they learn progressively what "two" and "three" should correlate with. Only once they have become 'four-knowers' do they generalise and thereby learn the successor relation.

The enriched parallel individuation system now has words that get used interchangeably with object files of numerosity 1 to 4 respectively (e.g. "one"/ $\{i\}$). Let's use the following neutral symbols for the mental symbols of this enriched system:

α : "one" / $\{i\}$
 β : "two" / $\{i\ j\}$
 γ : "three" / $\{i\ j\ k\}$
 δ : "four" / $\{i\ j\ k\ l\}$

Recall that object files support addition and 1-1 correspondence. So the following knowledge is implicit in the child's enriched parallel system:

β follows in counting from α γ follows in counting from β δ follows in counting from γ
 β corresponds 1-1 with $\alpha + \alpha$ γ corresponds 1-1 with $\alpha + \beta$ δ corresponds 1-1 with $\alpha + \gamma$

The child then has a chance to notice an analogy between counting on and adding α : if she adds α to any one of these symbols x (forming $x + \{i\}$) and compares the result with moving to the

next symbol in the count sequence after x , the two sets correspond 1-1.² If the child then comes to treat these two operations interchangeably (following in the counting sequence $\equiv$ adding $\{i\}$ / “one”), words further on in the counting sequence thereby become interpreted. The child now has a procedure which puts the symbols “five”, “six”, “seven” and “eight” into reliable causal correlation with sets of 5, 6, 7 and 8 elements in the world. It is only when moving forward in the count sequence is treated as an instance of the operation $+ \{i\}$ / “one” (adding 1) that these later symbols acquire a numerical meaning.

Striking evidence in support of this view is the fact that it is apparently only after children have passed the stage of being four-knowers that they reliably deploy the seemingly simple counting rule: to answer ‘how many’ about a set that you can’t judge perceptually (by subitising), count them one by one and give the last counting word as the answer. Carey calls children who have mastered this seemingly basic rule ‘cardinal principle-knowers’ and argues that children first attain representations of numbers at this stage, i.e. when they have a set of representations reliably tracking numerosity in this way and with narrow contents consisting of the set of conceptual roles just described.

There is good empirical support for several aspects of Carey’s Quinian bootstrapping model, especially the central role played by learning the counting sequence by rote. The available data do not yet speak to all the details of Carey’s model, so the story may change in the future. However, to doubt that Carey’s particular account is the last word on the issue is to miss the most important feature of the theoretical portion of her book – that she has a plausible, detailed story of how genuinely new atomic representations can develop out of representational resources with more limited expressive power, an account that manifestly proceeds at the psychological level. The new representations are not formed by inference from existing representations. At some stage there must be a leap, a step that is not a piece of computation-in-virtue-of-content, but one that psychology can describe causally nevertheless. Even if it transpires that Carey is wrong about the details, her achievement is to give us a worked-out answer to an instance of Fodor’s innateness puzzle.

A Fodorian Reply

A response in Fodorian spirit would doubt that Carey has described a case of conceptual change at all, by rejecting the idea that there is such a thing as narrow content. In this section I argue that the Fodorian reply misses its mark.

If the cardinal knower system is distinguished from earlier representations only by its having a proprietary set of conceptual roles, then if conceptual roles are not an aspect of content, no change in content has occurred. The parallel individuation system represents number, Fodor can say, albeit a very incomplete grasp of number, and children come to develop increasingly sophisticated knowledge about those numbers, which is encoded implicitly in the richer inferential roles in which they are then deployed. Similarly, I might have started

² Carey suggests this can be performed by comparing two sets held simultaneously in working memory, but I can’t see how the child could make the induction concerning the common role of the singleton set $\{i\}$ unless it is able to hold three sets simultaneously in working memory.

with a very impoverished conception of *elms*, which I then enriched into an expert body of knowledge about elm identification and ecology; but I had the same concept ELM throughout these changes. Why think that the process Carey describes is representational change rather than just knowledge acquisition (from a thin starting point)?

One rejoinder takes Fodor on his own terms. Even if there is no narrow content, the natural number system differs in *wide* content from the parallel individuation system, since numerosity is not explicitly represented in the parallel system at all. If inferential roles are no part of content, then the parallel individuation system just represents small sets of objects (or events) and their properties. Its computations encode various information about such objects, but only explicit representations can act as input to further computations. Those computations (e.g. addition) may make sense because of facts about numbers, but information that is implicit in the computations that are performed cannot itself act as input to further computations; so can't act as input to the process of developing new representations. Numerical representations don't enter the story themselves until after the Quinian bootstrapping process is complete. So numbers do not figure in the wide contents available at the first stage of Carey's story.

Fodor might point to the intermediate stage of enriched parallel individuation, when "one"/{i} causally covaries with singletons, "two"/{i j} with pairs, etc., and claim that these enriched symbols (the α , β , γ , δ above) represent numbers (as their wide contents) – a representation of number that is in place before the bootstrapping takes place. That rejoinder is little help to Fodor, though, since the transition from parallel individuation to enriched parallel individuation would then have introduced new representational resources (representations of number). And even if the formation of these enriched symbols could be explained in terms of hypothesis testing and construction out of pre-existing symbols, that story can't be carried over to "six". On Carey's account, the uninterpreted count word "six" only gets any wide content as a result of the analogical leap made by the child. Before that leap, if applications of "six" causally covary with anything at all, it is with the also meaningless preceding symbol "five" (and we might doubt that this causal connection is the application of a concept in any event). Only after the leap does the child have a sustaining mechanism that connects "six" reliably with sets of 6 items, giving it wide content. So here at least there must be genuine representational change, even in wide content.

A second response to Fodorian insists on the reality of narrow content. Carey has some strong empirical arguments for the existence of narrow contents. In the systems of core cognition, the inferential roles she has described are the same in everyone, and remain the same throughout development. That offers a stable target to act as the narrow content of such representations. Carey also documents considerable within-child consistency on either side of the bootstrapping step. Cardinal principle-knowers come at once to pass a whole group of tests that children before that stage fail; and Carey also claims considerable within-child consistency in performance of tasks by one-knowers, two-knowers, etc., although that is less clear. Carey also documents striking within-child consistency in other cases, for example a surprising correlation between seemingly different abilities at the stage when a child learns that 1 can be divided into fractions and that however many times you divide, you never get to 0.

There may be extraneous explanations for the apparent coherence of some of these sets of beliefs, and the data may be disputable, but in the realm of numbers, at least, Carey has prima facie evidence of a stable phenomenon at the level of the beliefs associated with a concept. And if there is a set of intra-personally stable beliefs that is widely shared inter-personally, there will be an explanatory role for such a coherent set – that is, our psychological explanations will appeal to narrow contents.³ If there are narrow contents, then a transition from one narrow content in the core cognition system to a different narrow content in cardinal principle-knowers would constitute the development of a new representational resource, even if the wide contents stay the same.

Extension to Other Cases

Once Carey has demonstrated that she can make good on Laurence and Margolis' and Macnamara's recipe for answering Fodor's puzzle, a question arises: can the tactic be deployed in other areas? The book has several other detailed cases, both in children (e.g. differentiation of density from weight) and in the history of science (e.g. differentiation of heat from temperature). I want to end by suggesting that the same tactic may work even more widely.

Carey defines learning (with respect to new representations) as any way of building representations on the basis of input that is itself representational. But notice that the words in the counting sequence start out without being representational at all. At first they are just a meaningless sequence of sounds. And although the object files are representational, their representational content does not enter into the acquisition story. They do not start out representing numbers. Furthermore, since the acquisition story proceeds causally, and not in virtue of content, the antecedent content of these symbols is entirely incidental. Non-representational resources could serve the purpose just as well, if they had the right causal profile to form an appropriate sustaining mechanism (just as with the non-representational counting words). In short, this tactic could in principle allow us to get from entirely non-representational resources to new representations, via a recognisably psychological story.

What could such non-representational resources be? Bare carriers of information, in the correlational, Shannon sense, are candidates. Theorists differ sharply on what it takes for some symbol to be a representation, but they are almost unanimous in agreeing that just causally covarying with some entity in the world is not sufficient. By contrast, causally covarying with an appropriate entity in the world looks like a very useful property if you are building a sustaining mechanism for a new representation. Correlational information is causal but not representational. A psychological account of how a new representation is built can appeal such information-carriers and facts about what it is they covary with. That gives us a hint that processes like connectionist learning which extract information from statistics of the environment might be capable of producing new representations from entirely non-representational resources.

³ On pain of regress or holism, the inferences that make up the narrow content of a concept should be individuated in terms of the wide contents of the concepts which figure in those inferences.

My suggestion is that the mechanisms of core cognition may not be innate. Carey worries a little about why they seem to come on-stream comparatively late in development. It is very plausible that interaction with the relevant domain is part of the aetiology of their development. If so, it would be odd if some of that causal intercourse did not involve extracting information from the environment. Doubtless, the infant does not come to the learning problem as a blank slate, but brings a whole suite of biases and constraints. But equally it would be odd if experience watching and manipulating objects did not help infants shape their expectations about how objects will move; that is, if information did not help to shape the computations which form sustaining mechanisms for representations of those objects. Carey has several good arguments for the innateness of the various systems of core cognition, which there is not space to canvass here, but learnability arguments on their own do not hold much weight.

Conclusion

There are many other riches in the book's nearly 600 pages, for example in a thorough examination of various cases of conceptual change in the history of science. Readers from any branch of cognitive science will find many topics to engage their interest. It is impossible even to catalogue them all here. It should be clear, however, that Carey has succeeded admirably in her main aim – to show that there are plausible psychological stories to be told about the development of novel concepts.

References

- Apperly, I., Butterfill, S. (forthcoming). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*.
- Darwin, Charles (1887/1958), Barlow, Nora, ed., *The Autobiography of Charles Darwin 1809–1882. With the original omissions restored*, N. Barlow (ed.) (London: Collins).
- Laurence, S., Margolis, E. (2002). Radical Concept Nativism. *Cognition*, 86, 25-55.
- Macnamara, J. (1986). *Border dispute: the place of logic in psychology*. Cambridge, MA: MIT Press.