

Margolis, E., & Laurence, S. (2008). How to Learn the Natural Numbers: Inductive Inference and the Acquisition of Number Concepts. *Cognition*, 106, pp. 924-939.
DOI: 10.1016/j.cognition.2007.03.003

This is a preprint. Please visit Cognition for the official publication.

How to learn the natural numbers: Inductive inference and the acquisition of number concepts

Eric Margolis & Stephen Laurence

Abstract. Theories of number concepts often suppose that the natural numbers are acquired as children learn to count and as they draw an induction based on their interpretation of the first few count words, viz., that each successive word expresses a quantity one greater than its predecessor. In a bold critique of this general approach, Rips et al. (2007) argue that such an inductive inference is consistent with a representational system that clearly does not express the natural numbers and that possession of the natural numbers requires further principles that make the inductive inference superfluous. We argue that their critique is unsuccessful. Provided that children have access to a suitable initial system of representation, the sort of inductive inference that Rips et al. call into question can in fact facilitate the acquisition of large integer concepts without the addition of any further principles.

1. Introduction¹

One of the distinctive features of human cognition is our facility with numerical concepts. While other animals can represent approximate quantities (Gallistel 1990, Brannon 2005), humans have a knack for developing concepts for the natural numbers, allowing us to pick out even large quantities with precision. Not surprisingly, where this ability comes from is enormously controversial, but an important core idea shared by a wide variety of theorists is that it emerges as children learn a natural language counting system (e.g., Hurford 1987, Wynn 1990, Dehaene 1997, Bloom 2000, Spelke 2003, Carey 2004). The common supposition is that children initially interpret the first few number words in a relatively direct manner and that working out the relations among these words allows children to interpret the rest of the count sequence by extracting the relevant pattern, viz., that each word picks out a quantity one greater than its predecessor. Achieving a full interpretation of the count sequence amounts to acquiring the concepts that its terms express, so the general picture is that the bulk of the number concepts are acquired on the basis of an induction. In a bold critique of this general approach, Rips, Asmuth, and Bloomfield (2007; hereafter RAB)² challenge the role of inductive inference in explaining how children

¹ This paper was fully collaborative; the order of the authors' names is arbitrary.

² All page number references are to this paper unless otherwise noted.

acquire the natural numbers. RAB argue that the standardly cited inductive inference is by itself insufficient to ground concepts for the natural numbers.³ What's needed, they claim, is an understanding of further principles that would render the induction superfluous. RAB do allow that the inductive inference may have a cognitive function in that it may cause children to see that the number concepts it ranges over are discrete, whereas children's initial representations are of continuous quantities. But if RAB are right, the inductive inference plays a marginal role in the acquisition of number concepts and the question of where we get them from must be addressed in a radically new way.

In this paper, we will defend the general approach to number concepts that RAB oppose and will argue that RAB's critique is unsuccessful. As we see it, the inductive inference *can* form an integral part of a satisfying account of large precise numbers—and this without the addition of any further principles—so long as the account postulates a sufficiently rich initial system of representation. At the same time, we argue that the inductive inference isn't able to sustain a transition from continuous quantities to discrete ones. Together these points reinforce the need to be as explicit as possible about the nature of the representations that precipitate children's learning of the natural numbers. We end our discussion by sketching our own model of what the initial system might look like. We speculate that it takes the form of a number module that incorporates a modest amount of innate number-specific representation, and we explain how this system could be implemented.

2. From the Bootstrap to the Induction

The general approach to number concepts that is at issue involves at least two stages—one that is directed to concepts of small numbers (ONE, TWO, THREE) and one that is directed to concepts of larger natural numbers.⁴ According to RAB, the first stage is satisfied when children acquire the knowledge specified in (1):

- “(1) a. ‘One’ refers to a property of collections containing one object.
 b. ‘Two’ refers to a property of collections containing two objects.
 c. ‘Three’ refers to a property of collections containing three objects.” (p. 2)

RAB want to be neutral about how the knowledge in (1) is achieved. They mention three possibilities that have been suggested in the literature. The first relies on a perceptual process (subitizing), the second on an analog system of representation (the accumulator), and the third on a mechanism of object-based attention (object-indexing). As RAB see it, however, the way that children come to interpret the first few count terms isn't so important for understanding why standard approaches to the natural numbers are flawed. What matters is simply that children are

³ RAB sometimes write as if the issue were the acquisition of the general concept NATURAL NUMBER (e.g., p. 5). However, we take it that the more fundamental issue—and certainly the one that concerns the authors RAB criticize—is how children acquire concepts for specific numbers, such as FOUR, FIVE, and SIX. We focus on these specific concepts for two reasons. First, it's doubtful that possession of concepts for specific numbers necessitates having of the general concept NATURAL NUMBER. First, it's doubtful that possession of concepts for specific numbers necessitates having of the general concept NATURAL NUMBER (any more than possession of such concepts as THREE and SEVEN, which pick out prime numbers, necessitates having the concept PRIME NUMBER). Second, since RAB's critique clearly also applies to specific natural number concepts, we see no harm in putting aside concerns about the general concept.

⁴ We use small caps for mentioned concepts, quotation marks for words, and italics for properties.

said to recruit the knowledge given in (1) in support of the second stage. This is the point at which the crucial induction occurs. RAB characterize this step as eventuating in the knowledge specified in their (2):

- “(2) If ‘ k ’ is a number word that refers to the property of collections containing n objects, then the next number word in the counting sequence ‘ $\text{next}(k)$ ’ refers to the property of collections containing one more than n objects.” (p. 3)

RAB label the inference embodied in (2) as *the Bootstrap*, alluding to Susan Carey’s (2004) metaphorical image of a person pulling himself up from his own bootstraps. For Carey, bootstrapping is a process in which a system of external symbols is used initially without understanding what its symbols mean. The symbols act as placeholders and gradually acquire meaning by bridging previously isolated representational domains through processes involving analogical reasoning and inductive inference. Carey has proposed that bootstrapping in this sense is the key to understanding how the human mind is able to acquire wholly new concepts, ones that aren’t directly constructed from the innate conceptual primitives. While Carey’s theory of conceptual development has a lot to be said for it, it isn’t fully representative of the class of theories that are subject to RABs critique. These theories make a particular claim about the acquisition of number concepts, not a claim about the nature of conceptual development in general. What these theories share is simply a commitment to the idea that children come to acquire larger number concepts by formulating an induction that takes them from the interpretation of small number words to the interpretation of larger number words. Undoubtedly, many theorists would agree to the role of this induction without committing themselves to Carey’s general theory of conceptual development. For this reason, we will refer to the disputed claim not as “the Bootstrap” but as *the Induction*.

3. RAB’s objection

RAB ask us to imagine a case of two children, Fran and Jan, acquiring the natural numbers as they learn how to count. Both children learn the count sequence from “one” to “nine” and both learn how to apply these words to collections of appropriate sizes. They then go on to tackle words for collections of arbitrarily larger size. But there’s a twist. While Fran is exposed to the ordinary counting system, Jan is taught a system in which the count sequence cycles back around from “nine” to “none” and then proceeds to “one”, “two”, etc. over and over again (see Fig. 1). So when the children are asked to report on the number of cookies on a plate (suppose there are eleven), Fran would use the word “eleven” but Jan would use the word “one”.

Figure 1. The structure of Jan's number system (from Rips et al. 2007, p. 4). Since Jan's number system is cyclical, the concept that she uses for collections with n items is the very same concept she uses for collections with n -plus-a-multiple-of-10 items (e.g., the concept she uses for 1 item is the same as the one she uses for 11, 21, 31, etc.). As a result, while Jan's system is confined to ten radically indeterminate concepts, it still conforms to the principles that Rips et al. associate with the Induction.

The problem, according to RAB, is that both children's understanding of the count sequence is entirely consistent with the Induction (as characterized by (1) and (2)), and so, by hypothesis, both children should be credited with possessing concepts for the natural numbers. While Jan's understanding is clearly nonstandard, she *has* interpreted the first few count terms in such a way that each picks out a property of appropriate collections. For example, she interprets "two" as picking out a property of collections containing two objects. It just happens to be a property of collections containing twelve objects and twenty-two objects as well. Thus her usage clearly satisfies (1). Her usage also conforms to (2). Each subsequent term in her counting sequence *does* pick out a property of collections with "one more than" those labeled by the preceding term. It just happens that the property in question is one that is shared by collections that differ by a multiple of ten. RAB observe that while Jan and Fran both succeed at making the Induction, only Fran comes to acquire the natural numbers. After all, Jan's system includes just ten radically indeterminate concepts. The conclusion RAB draw from all this is that the Induction doesn't suffice for coming to acquire the natural numbers.⁵

⁵ RAB summarize their case against the Induction in the following way: "There is no need to deny that Jan has the concepts PROPERTY OF COLLECTIONS CONTAINING ONE OBJECT, PROPERTY OF COLLECTIONS CONTAINING TWO OBJECTS, and so on. In Jan's representation, however, these properties are not related in a way that yields the concept NATURAL NUMBER any more than they are related in a way that yields the concept MERSENNE PRIME, which some of these same properties compose" (p. 5). Once again, the main issue isn't about the general concept NATURAL NUMBER so much as it is about concepts that pick out particular natural numbers (see note 3, above). But RAB's conclusion is

RAB's imaginary scenario is reminiscent of the various skeptical challenges that philosophers have identified with inductive reasoning, especially Saul Kripke's (1982) challenge to theories of meaning and understanding. Kripke has us consider an example in which we are to imagine someone learning the symbol "+" and who has provided the correct answer to a finite number of sums. Kripke notes that no matter how well this person does, there is always a question of whether he is interpreting the symbol correctly, since it's possible that his answers systematically differ from the norm for cases outside of the test range (e.g., he might interpret "+" to yield a value of five for all sums in which one of the addenda is greater than five-hundred). Of course, Kripke's puzzle generalizes to all words and concepts; the mathematical example is used for expository purposes only. To distance themselves from such general forms of skepticism, RAB note that symbol sequences sometimes pick out cyclically ordered properties (e.g., the days of the week) and so Jan's aberrant system isn't inherently any more unnatural than Fran's.

It's important to be clear about what RAB think the Induction can and cannot do. Their Jan example is supposed to show that the Induction alone is not sufficient to deliver concepts for the natural numbers. They claim that the Induction would need to be supplemented by further principles in order to generate concepts of natural numbers, and they offer the following:

- (3) a. "There is a unique first term in the numeral sequence, say, 'zero', which is never equal to 'next (k)'" (p. 9).
- b. "If 'next (k)' = 'next (j)', then 'k' = 'j'" (p. 9).
- c. "Nothing else—nothing that cannot be reached from 'zero' using 'next'—can be part of the sequence" (p. 9).

The problem, according to RAB, is that these further principles make the Induction superfluous, since the principles in (3) will generate concepts for the natural numbers all by themselves. "When Jan and Fran have learned them, they have no need for the [Induction]" (p. 9). RAB do allow, however, that the Induction "might not be irrelevant" (p. 9) in that it may help children to surmount a limitation of their initial system of representation. "Perhaps the [Induction] is successful in convincing children that the integers—at least those within their counting range—are discrete, overcoming an initial dependence on a concept of continuous magnitude that appears to be common to infants and nonhuman animals" (p. 9).

4. Where RAB's objection goes wrong

We believe that RAB's assessment of the Induction is mistaken in a number of important ways. In this section, we will focus on their critique of the Induction. Later, in section 5, we'll take up the question of whether the Induction might facilitate the transition from continuous to discrete quantity.

4.1. RAB's incorrect formulation of the Induction

easily reformulated as the claim that Jan lacks these concepts since her internal system is composed of just ten radically indeterminate concepts—concepts that are significantly broader in their application than our integer concepts.

Perhaps the best way to begin is by attending more closely to the role that inductive inference plays in standard accounts of the acquisition of the natural numbers. Though RAB's principles (1) and (2) may appear innocuous, they distort the general approach that proponents of the Induction have in mind. For this reason, we need to clarify the nature of the inductive inference that children are supposed to make and reformulate RAB's principles accordingly. We'll see that once they are properly reformulated, Jan no longer poses a problem for the Induction and its account of how the natural numbers are acquired.

RAB characterize the Induction in terms of a single inductive inference; however, proponents of the Induction typically see the acquisition of the natural numbers as depending upon at least three distinct inductive inferences.⁶ The first of these underlies the acquisition of the first few number words taken individually (i.e., not as part of the counting procedure) and results in the knowledge of the numerical quantities they pick out. In learning these words, children have to attend to how they are used and, in each case, determine what the word means via a process that involves an inductive generalization. Just as they might learn the word "cat" by noting that the word is applied to cats, they learn the words for small numbers by noting that they are applied to collections with specific quantities (e.g., "two" is applied to collections with two items). At this stage, children have a mapping between the first few number words and their respective quantities, but they needn't have any further understanding of the number system. In principle, they needn't even represent that the words "one", "two", and "three" are ordered or even that the corresponding quantities are ordered and bear distinctive quantitative relations to one another. (The fact that children learn number words in order—"one" first, "two" second, etc.—may just reflect facts about word frequency, not children's earliest understandings of their meanings. In every language that has been studied, the word for *one* is more common than the word for *two*, and the word for *two* is more common than the word for *three* (Dehaene & Mehler 1992).)

The second inductive inference involves what is known as the *cardinality principle* (Gelman & Gallistel 1978). It results in the realization that the last word in a count sequence denotes the numerical quantity (or cardinality) of the collection being counted. Initially children treat counting as a meaningless routine, a kind of game. But later they begin to notice that when they count to "two", the word "two" is surprisingly apt as it happens to be the word that picks out the quantity of the collection, and likewise for "three" and collections of three items. Children can notice the correspondence because of their prior facility with words for small numbers. At first, they may see this as just a happy coincidence. But over time and with enough exposure to the contingencies of counting small collections, children draw the inductive inference that counting is a way of enumerating a collection and that the final word in a count says how many items are in the collection.

Finally, the third inductive inference underlies children's understanding of subsequent count terms and the corresponding number concepts that they encode. Having mastered the meanings of the words for small numbers (the first induction) and having learned the significance of counting (the second induction), children are in a position to notice that each successive count term picks out a numerical quantity that is precisely one more than the term that precedes it. "Two" picks out a quantity precisely one more than "one", and "three" picks out a quantity precisely one more than "two". Children may then surmise that this pattern continues beyond the small number words and

⁶ For ease of exposition, our characterization of these inferences is centered around the natural language count system (this is the typical way of thinking about the Induction), but natural language isn't essential to the Induction. Other suitable symbol systems can serve much the same function.

that every word in the count sequence has a numerical value of precisely one more than its predecessor. It's this last inference that RAB wish to highlight.

RAB ignore many of the details that precede the final inductive inference, but we can now see that some of the details matter. Take RAB's principle (1). RAB formulate (1) as if children's interpretations of "one", "two", and "three" are radically indeterminate. For example, (1b) says that "two" refers to "*a property of collections containing two objects*" (italics added). This leaves it completely open *which property* of these collections is the proper referent. The point may be hard to see at first because one may well suppose that there really is just one property that all collections of two items share—the property *two*. But if we are generous about what counts as a property, collections with two items share infinitely many other properties as well. All collections with two items are collections with two or twelve (the property *two or twelve*), and are collections with two or one-thousand (the property *two or one-thousand*), and, to anticipate Jan, are collections with two or twelve or twenty-two... (the property *two plus a multiple of ten*). So RAB's (1) ensures that young learners are extraordinarily open-minded about the interpretation of the small number words *before they move on to the larger number words*.

We'd suggest that proponents of the Induction should insist that RAB's (1) doesn't convey their general approach to number concepts. The whole idea behind the Induction is that children start small. It's in part because they have an independent grasp of the meanings of "one", "two", and "three" that they are able to succeed in interpreting the beginning of the count sequence and then deduce the meanings of its subsequent terms. Of course, if they don't have a firm grounding for the small number terms—if they don't have the concepts ONE, TWO, and THREE—then they may very well end up with an odd interpretation of the sequence. But proponents of the Induction suppose that children aren't in that disadvantageous situation. So the proper formulation of (1) ought to be something along the lines of (1*):

- (1*) a. "One" refers to *one*.
 b. "Two" refers to *two*.
 c. "Three" refers *three*.⁷

RAB may object that we have no reason to credit children with the determinate concepts ONE, TWO, and THREE. All we really have to go on is behavioral evidence regarding how children apply their number words, and this evidence is consistent with the claim that children interpret number words differently than adults. Just because a child uses "two" for collections with two items doesn't mean he is interpreting it in a way that rules out collections with other numbers of items. But if RAB were to argue in this way, then their position would be dangerously close to the Kripkean skepticism that they are trying to distance themselves from. The concern about number words would extend to all sorts of ordinary words. For example, it's consistent with a child's—or for that matter an adult's—use of "cat" that she interprets it differently than others in her linguistic community and would apply it to non-cats in certain situations she has not yet considered. Perhaps, never having been to Australia, she would judge Australian chickens to be "cats" (along with all of the felines outside of Australia) while excluding Australian felines. Clearly the fact that behavioral evidence is unable to decide between certain alternative meanings does not show that

⁷ Some proponents of the Induction take children's initial system of numerical representation to be an analog system that picks out approximate numerical quantities, where the imprecision of the representations is governed by Weber's law. For these theorists, we would need to adapt (1*) to reflect the approximate nature of its representations, but as we argue below, this would not substantially affect our response to RAB.

a term's meaning is indeterminate between these alternative meanings. RAB might respond that number concepts are special in that children have to be able to learn cyclical sequences and that the count sequence must initially be interpreted in a way that doesn't exclude the possibility that it has a cyclical structure. But this response won't do. Given the standard view that children learn these words individually and quite apart from their role in counting, the possibility of cyclical *counting systems* is simply not relevant. Also, while it may be true that children can learn cyclical systems, this still wouldn't show that their initial understanding of the first few count terms would have to be indeterminate. Under the highly unusual circumstance of the sort that Jan faces, children might equally be said to *revise* their interpretation of these words as they discover that the community uses a count sequence that cycles—for example, they might start with the concepts specified in (1*) but switch to other concepts in light of their (highly unusual) experience. (Similarly, if it happened that children found themselves in the unusual circumstance that “cat” doesn't apply to Australian felines, then they could revise their interpretation of “cat”.)

Now consider RAB's principle (2). (2) is supposed to convey children's discovery that each successive term in the count sequence picks out the property of collections containing one more than the preceding term. But RAB's (2) incorporates the same sort of radical indeterminacy as their (1). If we instead take (1*) as our starting point, then the natural inductive principle is something like (2*):

- (2*) If a word in the counting sequence “one, two, three...” refers to n , then the next word in the counting sequence refers to one more than n .

Children who adopt RAB's (1) and (2) may very well end up with something other than the natural numbers. But children who adopt (1*) and (2*) are far better positioned to extend their counting sequence. They have the concepts ONE, TWO, and THREE, and have mapped these correctly to the terms “one”, “two”, and “three”. In addition, they have integrated their knowledge of the meaning of these words with the counting routine, having learned that counting has the significance that final word in a count applies to the collection as a whole. So what's to stop them from noticing that, for the part of the sequence that they understand, successive words exhibit a difference of one? They might do this by noting that a single object becomes a collection with two items when one item is added, and a collection with three items when another item is added. Children could then represent that “four” picks out *four* (the property of collections with one more than three), “five” picks out *five* (the property of collections with one more than four), and so on. RAB are certainly right that (1) and (2) can lead children astray, but there is no reason to accept that the Induction appeals to (1) and (2) rather than (1*) and (2*).

One further objection is worth considering. RAB might claim that proponents of the Induction can't make use of (1*) and (2*) since, by their own lights, children start out with inadequate numerical representations. While we have stated things as if children initially map the first few number words to precise numerical representations and not to RAB's radically indeterminate ones, proponents of the induction often suppose that children map them to analog representations that pick out approximate quantities (e.g., Wynn 1992, Dehaene 1997). Later, in section 5, we'll see that approximate representations do create their own difficulties, but for present purposes we need only note that even approximate representations are enough to short-circuit RAB's objection to the Induction. Suppose that we substitute approximate representations for children's interpretation of “one”, “two”, and “three” in (1*) and “approximately one more than” for “one more than” in the inductive inference (2*). Children in this situation would still settle on an interpretation of the

counting sequence that is non-cyclical, as the next number in the sequence would always be understood to be approximately one greater than the previous number. So, taking the child's initial representations to be approximate rather than precise doesn't affect our response to RAB. The purported counterexample built around Jan's situation fails on the understanding of the Induction that we suggest. The only way for it to succeed is on the assumption that children's initial representations of numerical quantities are radically indeterminate. But proponents of the Induction have no reason to grant this assumption.

4.3. RAB's three principles

Let's turn to RAB's claim that acquisition of the natural numbers requires something along the lines of the principles (3a-c) and that these render the Induction superfluous. We would argue, on the contrary, that (3a-c) are themselves inadequate but also, and more importantly, that nothing like (3a-c) is needed to supplement the Induction.

Suppose we take RAB's three principle at face value. In that case, the problem is that they only refer to symbol types and their relations and so they only place formal constraints on the symbol system. In particular, they do not place any substantive constraints on the meanings of the symbols. They generate a sequence that begins with the symbol "zero" and unfolds in such a way that each symbol in the system can't follow more than one distinct symbol (thus excluding cyclical systems). But for all they say, "zero" might refer to the Eiffel Tower, "one" to the number seventy-three, and "two" to alpha centauri; for that matter, these symbols could be treated as being entirely meaningless. To make explicit that RAB's principles have semantic content, they could be reformulated as follows:

- (3') a. There is a unique first term in the numeral sequence, say, "zero", whose value is never equal to the value of "next (k)".
- b. If the value "next (k)" = the value of "next (j)", then the value of "k" = the value of "j".
- c. Nothing else—nothing that cannot be reached from "zero" using "next"—can be part of the sequence.

Unfortunately, these revised principles won't do either. Among other things, they fall prey to a counterexample that is closely related to RAB's example involving Jan and Fran. On this variation, Jan and Fran use the same symbol sequence but they interpret it differently. Faced with a plate with eleven cookies on it, both describe the situation by using the word "eleven". But Jan, unlike Fran, would say that it also has "one" cookie on it since she interprets "eleven" as having the same value as "one", "twenty one", "thirty one", and so on. In other words, Jan has a *symbol sequence* that never cycles (consistent with 3'a-c) but she interprets it in terms of a *conceptual system* that does (one that corresponds to RAB's original Jan).

These objections are troubling but they don't get to the heart of the matter because RAB clearly have in mind something stronger—a set of principles that not only specifies the right sort of symbol sequence but also the right interpretation and that does so by embodying the axioms of arithmetic. Given that the axioms are typically formulated directly for numbers themselves, we can avoid the use-mention difficulties that RAB inadvertently stumble into. We can omit the reference to public language symbols altogether and express what RAB take children to learn in terms of a set of principles governing the numbers themselves:

- (4) “(i) 0 is a natural number.
 (ii) For any x and y , if x is a natural number and y succeeds x , then y is a natural number.
 (iii) Every natural number has a unique successor.
 (iv) 0 is not the successor of any natural number.
 (v) If x and y are natural numbers and $x \neq y$, then the successor of x is not equal to the successor of y .
 (vi) For any [property] F , if [0 has F] and F is hereditary with respect to successor⁸ and k is any natural number, then [k has F].” (George & Velleman 2002, p. 37)

RAB’s general idea seems to be that children must somehow acquire something akin to (4) to supplement the Induction and only then can children be counted upon to represent the natural numbers and not an indeterminate system like Jan’s. The crux, however, is supposed to be that once (4) is acquired, children already have the natural numbers and have no need for the Induction.

We are now in a position to see the fundamental difficulty with RAB’s proposal. It’s that it fails to sufficiently distinguish the principles that govern the natural numbers themselves from those that govern the psychological states that represent them. While it’s certainly true that the axioms of arithmetic govern the natural numbers, it doesn’t follow that the axioms have to be explicitly represented in order for us to think about the natural numbers. Water may be H₂O but it doesn’t follow that you have to be able to represent this fact to have the concept WATER. In section 4.2, we argued that (1*) and (2*) are adequate to represent the natural numbers. If that’s right, then there is no need to represent the axioms of arithmetic, since children’s representations will conform to them automatically.

5. From continuous quantity to discrete quantity

Although RAB maintain that the Induction is inadequate for acquiring the natural numbers, they allow that it might nonetheless have a role to play in explaining how children come to be able to represent discrete quantity despite an initial dependence on representations of continuous quantity. In this section, we offer reasons for thinking that this isn’t right. Though we’ve defended the Induction from RAB’s critique, we don’t believe that it can play this sort of role in cognitive development. Seeing why is important because it helps to clarify an important limitation of the Induction that often isn’t fully appreciated. The limitation is that, for the Induction to succeed, it has to have the right representations to draw upon.

Suppose for the sake of argument that children undertake the Induction using an analog system in which the numerical quantity of a collection corresponds to a mental magnitude—the bigger the magnitude, the bigger the quantity (e.g., Wynn 1992, Dehaene 1997). On such an account, the correspondence is never exact. A mental magnitude picks out an approximate quantity, where the discriminative capacity of the system becomes increasingly approximate as the quantities become larger. If a system of this sort is in place early enough in development, then children might be thought to have a huge leg up when it comes to learning words for small numbers. Presumably,

⁸ *Being hereditary with respect to successor* means that whenever a given natural number has the property F , its successor has F too. So principle (vi) says, in effect, that all natural numbers have any properties that 0 has and that are had by every successor.

the system of mental magnitudes is fairly accurate for small quantities and children can map these to the corresponding small magnitudes. The question, one might suppose, isn't how children learn "one", "two", and "three"; it's how they learn words that express large precise quantities. And the natural suggestion is that they do so in accordance with the Induction.

The problem with this general approach is that imprecise analog representations aren't up to the task of supporting the Induction. Even though the analog representations for small quantities are more precise than those for large ones, analog representations are by their nature approximate and hence incapable of expressing a difference of exactly one. Even for small collections (e.g., 2 vs. 3) an analog system will represent these quantities only variably and approximately. What's more, the difference 1 vs. 2, when it is registered, will be represented as being slightly more of a difference than the difference between 2 vs. 3, or 3 vs. 4. Susan Carey rightly notes that the latter fact makes a system of mental magnitudes particularly ill-suited as a precursor to the integers, since mental magnitudes "positively obscure the successor function" (Carey 2001, p. 43). The ability to represent a difference of exactly one is essential to the Induction. Children need to notice that the small numbers in the count sequence proceed by a difference of exactly one, and they need to be able to entertain the thought that the sequence continues this pattern. So contrary to the suggestion that RAB entertain, *the Induction* will not help children move from representing continuous quantities to representing discrete quantities. The Induction can't take place unless they *already possess* precise numerical representations.

A similar problem affects theories that credit children with an initial representation that draws primarily from a mechanism of object-based attention. On these theories, the visual system is able to track a small number of objects in parallel using object-indexes—representations that act like pointing devices. Each index sticks with its object by responding to the object's spatial-temporal properties. When there are n objects (up to four), the visual system maintains n indexes, one per object (e.g., Leslie et al. 1998). This may make it seem like object-indexes amount to an early form of numerical representation, leaving the Induction to generate concepts for the rest of the natural numbers. But the problem is that mechanisms of object-based attention don't include any specifically numerical representations or even fully abstract representations—ones that transcend the visual modality. So object indexes, like mental magnitudes, lack the needed representational content to get children to the point where they can contemplate the crucial facts that support the Induction.

These brief remarks aren't meant to show that analog representations or object-indexes have no role to play in the acquisition of the natural numbers. But they do indicate the need to give more attention to the representations that get the Induction going. Without the ability to represent small precise quantities and especially to represent a difference of one—precisely one—the Induction hasn't a chance at explaining where the rest of the natural numbers come from.⁹ In our view, this is the real problem facing many accounts that appeal to the Induction, not the problem that RAB allege.

6. Getting ahead with a number module

Children need a sufficiently rich initial system of representation if they are going to succeed with the Induction. They need the prior ability to represent small precise quantities. Where does

⁹ For related discussion, see Laurence & Margolis (2005).

this ability come from? At present, any answer to this question has to be somewhat speculative. We propose that it derives from an *innate number module*. The number module provides a way to meet the challenge that we saw in the previous section.¹⁰ Since this is not the place for a detailed elaboration and defense of our own account, we'll just present just a stripped-down version (for further details, see Laurence & Margolis (in preparation)).

The number module, as we envision it, functions to detect the numerical quantity of small collections. In its barest form, the number module needn't be taken to represent any arithmetic facts or even the basic quantitative relations among the collections to which it can respond. How does the number module represent small numbers? Our suggestion is that it takes input from the object indexing system—and also from analogous systems corresponding to other modalities—and filters this through a neural network that has three or four output nodes (see figure 2). The network's connections are weighted so that any single index taken alone is sufficient to activate a specific output node (the one-node), any two indexes are sufficient to activate another specific output node (the two-node) and to inhibit the one node, and so on.¹¹ This means that the one-node will be active just in case one and only one object index is active, and, likewise, each of the four output nodes selectively responds to collections with a given cardinal value and becomes associated with representations of the natural language count terms “one” through “four”. Numerical representations like these have their content owing to their external causal relations. What makes them represent what they do is a matter of what they are reliably connected to in the world. The core numerical system, in other words, is to be described in terms of an information-based semantics (Fodor 1998).

¹⁰ We do not claim that ours is the only account capable of meeting this challenge. See Carey (2004) for an alternative account that aims to make do without the representations of a number module. Carey's proposal is that the Induction draws upon prior quantitative representations that are associated with natural language quantifiers (e.g., singular-plural markers).

¹¹ Object indexes provide inputs to the number module via the input nodes. Each input node is connected (with a strength of 1) to the output node that represents *one*, so that activation of any of the input nodes suffices to activate this output node. Each input node is also connected (with a strength of .5) to the output node that represents *two*, so that activation of any two of the input nodes suffices to activate this output node. The two output node is also connected to the one output node (with a strength of -2), so that if the two output node is activated, the one output node is deactivated. The module works in much the same way for representing *three* and *four*, where the activation of multiple input nodes leads to the activation of the appropriate output node while deactivating all of the other output nodes.

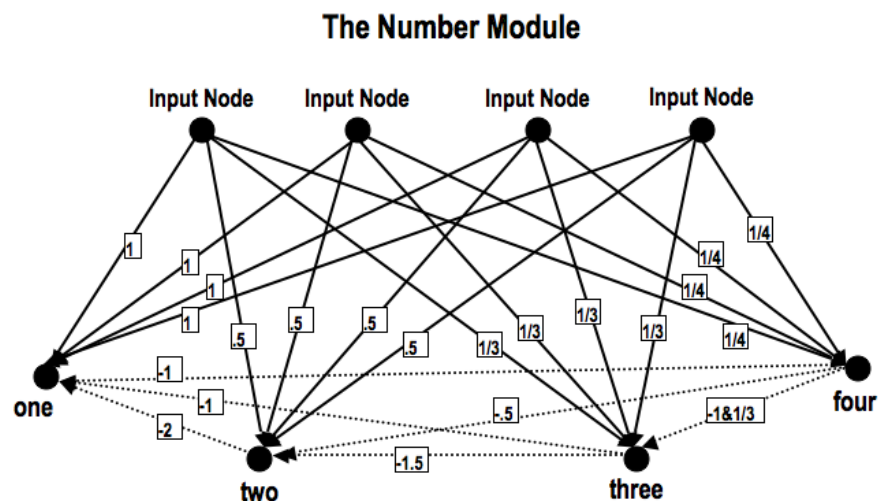


Figure 2. The Number Module. The network's input comes from the object indexing system and from comparable non-visual systems. The output nodes are selectively responsive to specific numerical quantities.

Though in many ways this proposal is rather modest, the number module offers enormous explanatory advantages. Earlier we noted that the acquisition of the natural numbers involves at least three distinct inductive inferences. It's easy to see how a number module would support these. The first inference culminates in children learning the meanings of the small number words taken individually and not as part of the count sequence. In this case a number module would provide minimal numerical representations that can be used to interpret these words. The module delivers a representation associated with "one" that selectively responds to collections that have just one item, a representation associated with "two" that selectively responds to collections with two items, and so on. In the second inference children learn the significance of the count sequence. In this case, the number module would allow children to detect the quantity of a small collection as it is being counted so that they can recognize that the final term in the count is the very term that picks out the collection's quantity. If they count to "two", the number module helps them to see that the collection has the quantity two, and having made the first inference, they know this to be the meaning of "two". Finally, in the third inference children work out the quantitative relations among the small number words in the count sequence and then infer that the pattern continues for the rest of the sequence. Here the number module allows children to represent a difference of precisely one among adjacent small terms. It does this by the activation of the representation that corresponds to a collection with just one item. (One way the process might unfold is for children to compare sets of differing sizes, e.g., 2 vs. 3, but another way is for children to note changes as they manipulate a collection by adding or subtracting items.) In short, the number module provides the precise representations that make the Induction a viable process.

7. Conclusion

Contrary to RAB, the Induction can provide a satisfying account of how to learn the natural numbers. We just have to be clear about the cognitive requirements for the Induction to take place. RAB overlook one of the most basic, that the Induction presupposes that children have the ability to correctly interpret the small number words apart from how they figure in the count sequence.

But even proponents of the Induction can sometimes overlook the equally basic requirement that children have to be able to represent small precise numerical quantities, including a difference of precisely one, for the Induction to do its job. Our suggestion is that both of these requirements can be met on the supposition that a number module gives children a small stock of numerical representations. It's on the basis of these representations that the rest of the natural numbers may be constructed.

Acknowledgements

We'd like to thank Susan Carey, Ian Dove, Richard Grandy, Bob Hale, Rosanna Keefe, and two anonymous referees for their comments.

References

- Brannon, E. (2005). What Animals Know about Numbers. In J. Campbell (ed.) *Handbook of Mathematical Cognition* (pp. 85–108). New York: Psychology Press.
- Bloom, P. (2000). *How Children Learn the Meanings of Words*. Cambridge, Mass.: MIT Press.
- Carey, S. (2001). Cognitive Foundations of Arithmetic: Evolution and Ontogenesis. *Mind & Language*, 16.1, 37-55.
- Carey, S. (2004). Bootstrapping and the Origin of Concepts. *Daedalus*, 133, 59-68.
- Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics*. Oxford: Oxford University Press.
- Dehaene, S. & Mehler, J. (1992). Cross-linguistic Regularities in the Frequency of Number Words. *Cognition*, 43, 1-29.
- Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Gallistel, C.R. (1990). *The Organization of Learning*. Cambridge, MA: MIT Press.
- George, A. & Velleman, D. (2002). *Philosophies of Mathematics*. Oxford: Blackwell.
- Hurford, J. (1987). *Language and Number: The Emergence of a Cognitive System*. Oxford: Blackwell.
- Kripke, S. (1982). *Wittgenstein: On Rules and Private Language*. Oxford: Blackwell.
- Laurence, S. & Margolis, E. (2005). Number and Natural Language. In P. Carruthers, S. Laurence, & S. Stich (eds.) *The Innate Mind: Structure and Contents*. New York: Oxford University Press, pp. 216-35.
- Laurence, S. & Margolis, E. (in preparation). *Think of a Number*.

- Leslie, A., Xu, F., Tremoulet, P., & Scholl, B. (1998). Indexing and the Object Concept: Developing 'What' and 'Where' Systems. *Trends in Cognitive Sciences*, 2.1, 10-18.
- Mix, K., Huttenlocher, J., & Levine, S. (2002). *Quantitative Development in Infancy and Early Childhood*. Oxford: Oxford University Press.
- Rips, L., Asmuth, J. & Bloomfield, A. (2007). Giving the Boot to the Bootstrap: How Not to Learn the Natural Numbers. *Cognition*, xx, xxx-xxx.
- Spelke, E. (2003). What Makes Us Smart? Core Knowledge and Natural Language. In D. Gentner & S. Goldin-Meadow (eds.) *Language in Mind: Advances in the Study of Language and Thought*. Cambridge, Mass.: MIT Press, pp. 277-311.
- Wynn, K. (1990). Children's Understanding of Counting. *Cognition*, 36, 155-93.
- Wynn, K. (1992). Evidence Against Empiricist Accounts of the Origin of Numerical Knowledge. *Mind & Language*, 7, 315-32.