Available online at www.sciencedirect.com
ScienceDirect
Cognition 106 (2008) 940–951
COGNITION
www.elsevier.com/locate/COGNIT

Discussion

Do children learn the integers by induction?

Lance J. Rips *, Jennifer Asmuth, Amber Bloomfield

Psychology Department, Northwestern University, 2029 Sheridan Road, Evanston, IL 60208, USA

Received 11 April 2007; revised 11 July 2007; accepted 21 July 2007

Abstract

According to one theory about how children learn the meaning of the words for the positive integers, they first learn that "one," "two," and "three" stand for appropriately sized sets. They then conclude by inductive inference that the next numeral in the count sequence denotes the size of sets containing one more object than the size denoted by the preceding numeral. We have previously argued, however, that the conclusion of this Induction does not distinguish the standard meaning of the integers from nonstandard meanings in which, for example, "ten" could mean set sizes of 10, 20, 30,… elements. Margolis and Laurence [Margolis, E., & Laurence, S. (2008). How to learn the natural numbers: Inductive inference and the acquisition of number concepts. *Cognition, 106*, 924–939] believe that our argument depends on attributing to children "radically indeterminate" concepts. We show, first, that our conclusion is compatible with perfectly determinate meanings for "one" through "three." Second, although the inductive inference is indeed indeterminate – which is why it is consistent with nonstandard meanings – making it determinate presupposes the constraints that the inference is supposed to produce.

Keywords: Number learning; Numerical cognition; Number concepts; Natural number; Induction; Integers

DOI of original article: 10.1016/j.cognition.2007.03.003
* Corresponding author. Tel.: +1 847 491 5947; fax: +1 847 491 7859.
E-mail address: rips@northwestern.edu (L.J. Rips).

0010-0277/$ - see front matter
doi:10.1016/j.cognition.2007.07.011

> D.S. Age: 6 years, 2 months:
> D.S.: The numbers only go to million and ninety-nine.
> Experimenter: What happens after million and ninety-nine?
> D.S.: You go back to zero.
> E: I start all over again? So, the numbers do have an end? Or do the numbers go on and on?
> D.S.: Well, everybody says numbers go on and on because you start over again with million and ninety-nine.
> E: ...you start all over again.
> D.S.: Yeah, you go zero, one, two, three, four – all the way up to million and ninety-nine, and then you start all over again.
> E: How about if I tell you that there is a number after that? A million one hundred.
> D.S.: Well, I wish there was a million and one hundred, but there isn't.
>
> Hartnett (1991, pp. 115–116).

Many investigators believe that children learn the meaning of the positive integers ("1," "2," "3," ...) by gradually connecting the first three or four number terms with sizes of sets. "One" comes to denote sets containing exactly one object. Awhile later, they learn that "two" denotes sets containing exactly two objects, and so on, for "three" and, perhaps, "four." At this point, however, children arrive at a key insight: the next term in the count sequence refers to the size of sets containing one more object than the size denoted by the preceding term. Carey (2004) and Hurford (1987) have specific proposals along these lines, but others have made similar suggestions (see Rips, Asmuth, & Bloomfield, 2006, Footnote 4, for examples). It seems likely that some such inference takes place, and this *Induction* is a reasonable candidate. It's a correct generalization about the positive integers.[1]

But in an earlier paper (Rips et al., 2006), we argued that the Induction cannot help the child distinguish between the usual numeral-to-cardinality mapping and vastly different systems. In our illustration, two twins, Jan and Fran, both follow the ordinary learning pattern: they learn to recite the integer sequence to some small numeral ("one" through "nine" in our example), learn the set sizes associated with "one" through "three," and perform the Induction. Yet we showed that the Induc-

[1] In our earlier paper (Rips et al., 2006), we called this inference "the Bootstrap," following Carey (2004). Here, we will use Margolis and Laurence's (2008) term, "the Induction," to avoid singling out a particular version of this inference. Readers should bear in mind, however, that the induction at stake is ordinary empirical induction, not mathematical induction. See Rips and Asmuth (in press) for a discussion of this difference. We also use the terms "numerosity," "cardinality," and "set size" interchangeably. As a final point of terminology, we discuss "positive integers" (1, 2, 3, ...) here rather than "natural numbers" to avoid confusion about the role of zero. (In some systems, the natural numbers begin with 0; in others, they begin with 1.)

tion is equally consistent with the possibility that Jan goes on to learn the mod_{10} system, while Fran learns the standard system.

In their comments on our paper, Margolis and Laurence (M&L, 2008) argue that our description of Jan and Fran's plight depends on tacitly describing them as having indeterminate concepts for "one," "two," and "three" prior to performing the Induction. Similarly, they believe that Jan's and Fran's diverging number systems are only possible because we have incorrectly portrayed the Induction in an indeterminate way. Given both (a) determinate meanings for the small integers, and (b) a revised formulation of the Induction's conclusion ("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 will converge on the correct understanding of the positive integers. We argue here, however, that these points are mistaken. The puzzle does not depend on the determinateness of the concepts for "one" through "three." And although the Induction is indeed indeterminate – that's precisely what's wrong with it, since it can't distinguish the standard system of positive integers from alternatives – any way of making it determinate begs the question that the Induction is supposed to solve. We won't, however, be dealing with those portions of M&L's article that don't concern our anti-Induction argument. They believe, for example, that although the Induction is sufficient to fix the meaning of the integers, it isn't sufficient to transform a representation based on a continuous measure into a countable one. We have expressed our own doubts about this possibility elsewhere (Rips, Bloomfield, & Asmuth, submitted for publication), and our brief remark about it in Rips et al. (2006) was simply meant to convey that our anti-Induction argument doesn't bear on it.[2]

Here's the argument in outline: first, for empirical reasons, children are unlikely to attach meanings to the words for small integers that could withstand further learning of the sort that Jan receives. But even if we grant (a) – assuming that the meanings of "one," "two," and "three" are fixed – the Induction is still not sufficient to produce the correct meanings for the rest of the positive integers. Formulating the Induction in the way M&L favor does nothing to keep children from learning nonstandard meanings. Hence, (a) and (b), taken together, are not sufficient to yield the correct concept of the positive integers. Of course, there are ways of constraining the Induction so that it does produce the standard meanings, for example, by adding the Dedekind-Peano axioms for natural numbers. But then it is obvious that it's the axioms, not the Induction, that are doing the heavy lifting.

[2] Similarly, we take no stand on the issue of whether a qualitative shift occurs in children's representation of cardinalities. We should also add that we are doubtful both about the idea that the meaning of a number term is a cardinality and about the idea that children can understand the meanings of the terms for small integers in isolation from the rest of the system (see Rips et al., submitted for publication, for the reasons behind these doubts). However, we grant these possibilities here for the sake of the argument.

1. How much can the Induction accomplish?

One way to see the problem with the Induction is to imagine a case in which children make a similar inference about a dissimilar structure. Suppose that a child – we'll call him Wolfgang – is learning the tones in the C major scale (C, D, E,...). So far, Wolfgang has memorized the first four letter names in the sequence and knows their order – "C" then "D" then "E" then "F." He has also noticed that when his teacher says "C," she presses a particular key on the piano (middle C), which causes a specific tone to occur. After a number of trials, Wolfgang has learned that "C" denotes the tone sounded by this key. Wolfgang then learns that "D" denotes the tone sounded by the white key immediately to the right of the first one and, still later, that "E" denotes the tone sounded by the white key immediately to the right of the second. At this point, then, Wolfgang can make an inductive inference: the next term in the sequence "C," "D," "E,"... denotes the tone that is sounded by the white key immediately to the right of the key that produces the tone for the previous term.

Wolfgang's generalization is correct, and his teacher can be proud of him. But has Wolfgang learned the meaning of the terms in the C major scale? This simply isn't credible. For starters, Wolfgang doesn't know the entire series. "G," "A," and "B" aren't part of his vocabulary at this point, so he attaches no meaning whatever to these terms. More important, Wolfgang doesn't know the structure of the sequence of terms. He doesn't know that the structure is circular (..., "C," "D," "E," "F," "G," "A," "B," "C," "D,"...) and, therefore, doesn't know how to use the sequence to name lower or higher tones. So as far as he knows, the sequence could continue with many distinct further terms (e.g., "H" for A, "I" for B, "J" for C above middle C, etc.) with no doubling back to "C." He has no idea that there are just seven tones in the C major scale rather than infinitely many.[3]

Much the same state of affairs holds for the twins in our number-learning story. Because the twins don't know how to extend the sequence of count terms (i.e., don't know the structure of these terms) beyond the ones they have memorized, they have no idea whether later terms double back (as in the case of the musical scale) or keep on going. The twins don't know that there are an infinite number of positive integers, and, according to current evidence, won't find out about this for another two or three years (Hartnett, 1991). In Wolfgang's case, the cyclical meaning is the correct one for adults. In the twins' case, the cyclical meaning is incorrect for adults (outside the modular systems), and the linear meaning is correct. But the Induction cannot determine who is right and who is wrong about the meanings of the terms, since it is consistent with both the cyclical and the linear patterns. That was the point we intended our original fable to convey. As we emphasized in our earlier paper (Rips

[3] We are not assuming that Wolfgang is confused between terms for the alphabet and terms for the C major scale; only that he does not know how the sequence of terms for the scale should continue. The same problem would occur if Wolfgang were learning the terms in the sol-fa sequence (doh, ray, me,...) instead of the letter terms (C, D, E,...). (We would have used the sol-fa terms in this illustration but for the complication that they shift with the governing key in some systems.)

et al., 2006), the problem with the Induction differs from the traditional problems with justifying inductive inference (e.g., Goodman's, 1955, new riddle). The usual problem is justifying one possible inductive conclusion over another (e.g., linear extrapolation of data vs. polynomial extrapolation of the same data). By contrast, the problem with the Induction is that its conclusion is entirely vague about the continuation and hence consistent with many different correlations between numerals and set sizes. It would take a separate inference of a different kind to decide among these possibilities.[4]

2. The problem with the Induction does not depend on ambiguity about the meanings of "one," "two," and "three"

In our story about Jan and Fran, we supposed that, before performing the Induction, both girls understand the meanings of "one," "two," and "three" in terms of set sizes. Because there are several competing proposals about the mental representations that mediate number terms and their meanings, we deliberately took no stand on which representation Jan and Fran employed. For example, the representations could be mental magnitudes, as Gallistel and Gelman (1992) and Wynn (1992) propose, or they could be internal sets, as Carey (2004) proposes. We didn't, however, intend to attribute to children meanings for the small integers that are "radically indeterminate" (M&L, p. 7). In our formulation, "one" was said to refer to a property of collections that contain one object, and similarly, for "two" and "three." According to M&L, this formulation is compatible with a range of denotations, since there are infinitely many properties that apply to collections containing one object. M&L believe this ambiguity allowed us a sleight-of-hand. But that was certainly not our aim, and nothing about the puzzle trades on the ambiguity.

Let's suppose, along with M&L, that children's initial meanings for "one," "two," and "three" are perfectly definite: "One" denotes the size of sets containing exactly one object, and similarly for "two" and "three." We argue, first, that realistic assumptions about language learning must allow these meanings to be extendible in a way that is compatible with our original scenario about Jan and Fran. But second, even if these meanings were fixed once and for all, they still wouldn't be able to prevent children from learning nonstandard meanings for the integers. In other words, even if we assume (a) the meanings of "one" through "three" are perfectly definite, (b) the meanings of "one" through "three" never change after initial learning, and (c) the Induction occurs in the way M&L formulate, the Induction won't prevent children from learning the wrong meanings for the integers.

[4] M&L correctly note that we try to distance ourselves from earlier problems about induction, but they believe that we do so by stressing that Jan's cyclical system is no more unnatural than Fran's. We agree about the relative naturalness. As the epigraph to this article shows, some children discover a cyclical number scheme by themselves; in this case, the scheme is $\text{mod}_{1,000,099}$. We've also heard informally about similar cases from others. But this is not the source of the difference between traditional problems of induction and the problem with the Induction.

2.1. Troubles with the Induction's premises

Even if children initially have perfectly definite denotations for the small positive integers, this does not make these meanings unrevisable. One- and 2-year-olds quite commonly assign denotations that are either too narrow or too broad. For example, a child might initially understand "shoe" to refer to shoes that aren't currently being worn and then later extend the term to all shoes (Reich, 1976). By one estimate (Clark, 1997), these under- or overextended terms account for up to a third of the child's vocabulary at this age. We see no reason to think that number terms don't operate like other terms in this respect. Even if Jan at first denies that "one" could apply to a collection of eleven objects, she could later include it in the denotation when she learns the mod_{10} system (just as a child who denies that "shoe" applies to a foot-filled shoe could later extend the meaning to include it). In other words, what M&L have to suppose is not just that the terms for the small integers are initially unambiguous but that their meaning never changes during further learning. But this supposition does not characterize real language learners. What's more, this assumption would make it impossible for children to learn truly cyclical systems (e.g., the C major scale, the months of the year, and the days of the week) by expanding initially narrow meanings. For example, it would prohibit Wolfgang from revising his concept of C when he learns that C can stand for distinct new frequencies above and below middle C.

One objection to this point of view is that allowing children to revise the meanings of "one," "two," and "three," in the way Jan does, makes learning of number terms an arbitrary matter. Theories of how children acquire word meanings have to assume that learning conditions don't change radically over time, since otherwise no hypothesis that children entertain could possibly be reliable. For example, it would be unfair to a theory about how children learn names for animals to insist that the theory be able to survive the appearance of a deceptive teacher who indoctrinates children with the idea that animal names refer to car parts. Isn't it similarly unfair to criticize the Induction for not being able to survive Jan's parent's insistence that "one" refers to one, eleven, twenty-one, and so on? Another way to state this objection is to say that the Induction takes it *as given* that "one" means one; hence, to suppose that Jan can later adopt a broader meaning is simply to reject the terms of the Induction theory.

But although theories of word learning can't contend with an arbitrary landscape, they do have to cope with the environment in which children actually find themselves. The Induction is an empirical theory of the acquisition of number concepts, and so it is not in a position to stipulate assumptions that are too strong to apply to normal learners. In particular, it can't simply assume that children never alter the initial meanings they assign to words, since this is wildly contrary to fact. Of course, there are limits to how much change word meanings can plausibly undergo in normal learning. However, acquisition theories have to allow for shifts that are natural extensions of initial hypotheses, for otherwise a child would never be able to correct his or her overly narrow or overly broad concepts. Although we framed the story about Jan in terms of a parent who teaches her the mod_{10} system, there is no evidence that this system is a drastic departure from her initial ideas about number words. In fact, there is evidence that children sometimes adopt a modular hypothesis about the natural numbers without external prompting, as the transcript at the beginning of this article makes clear.

2.2. Troubles with the Induction's conclusion

Let's grant, though, for purposes of the argument that children's concepts of "one" through "three" are fixed once and for all. The Induction would still not produce the right meanings for the remaining positive integers. Suppose, once again, that Jan memorizes the sequence "one" through "nine," acquires the usual meanings for "one" through "three," and then performs the Induction. But this time, instead of teaching her the mod_{10} system, her dodgy parent teaches her a numeral sequence that is the usual one for the integers less than "ten" but diverges thereafter. There are an infinite number of these, but Fig. 1 shows one example. The meanings of Jan's "one" through "three" remain intact, but the number terms loop beginning with "ten." "Ten" refers to ten, twenty, thirty,...; "eleven" refers to eleven, twenty-one, thirty-one,...; and so on. As in our original story, this result is consistent with the conclusion of the Induction. Moreover, each number term has an unambiguous meaning. But no one would confuse the structure in Fig. 1 with that of the positive integers.

It may be helpful to describe the steps leading to Fig. 1 in slow motion to show that they do not violate M&L's formulation of either the premises of the Induction ("one" refers to one; "two" refers to two; "three" refers to three) or the conclusion

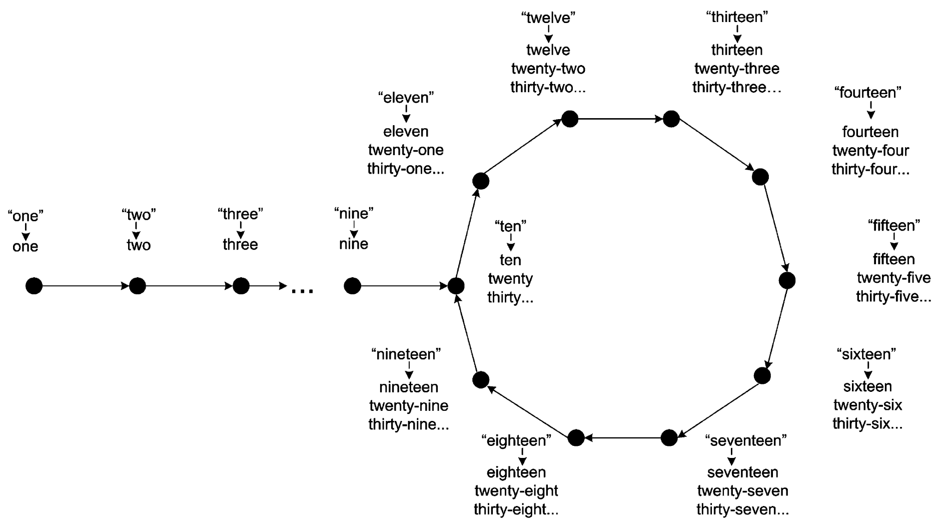

Fig. 1. The structure of an alternative number system in which numerals "one" through "nine" refer to their normal cardinalities but "ten" through "nineteen" have broader meanings.

of the Induction ("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*"):

(a) Jan starts by learning the usual denotations for "one" through "three." We assume these meanings stay fixed for the rest of her number learning.
(b) Jan performs the Induction and concludes that the next word in her count sequence, "four," refers to one more than three. She continues in this way to fill in the meanings for the rest of the terms in her count list, "five" through "nine."
(c) At this point, Jan's parent informs her that the next word in the count sequence is "ten" and that "ten" follows two different count terms, both "nine" and another called "nineteen." Jan infers that "ten" denotes one more than both "nine" and "nineteen."
(d) Jan's parent then teaches her that the count word following "ten" is "eleven," and Jan infers that "eleven" refers to one more than "ten." Exactly the same process is repeated for "twelve" through "nineteen," at which point Jan has learned the meaning of all the number words in her system, the system of Fig. 1.

The only step in this sequence that is out of the ordinary is (c), since it is the only place at which Jan's lessons would have to depart from Fran's. But there is nothing about this step that is contrary to the Induction. Jan learns at this stage that "ten" follows two terms in the count sequence; nevertheless, "ten" refers to one more than these predecessors. Thus, (a)–(d) conform exactly to the Induction as M&L have restated it, both the premises [M&L's (1*)] and the conclusion [M&L's (2*)], but they don't generate the natural numbers.

3. The problem with the Induction is that it presupposes the successor function

Imagine that children at three or four already knew how to get from any number term in the series of positive integers to the next term in that sequence. That is, suppose they knew the successor function for the numeral sequence, which we call *advanced counting* (Rips et al., submitted for publication). Let's assume, also, that they know that one can always add an element to a set to get a new set with a size one larger than that with which they started. The conclusion of the Induction would then latch these two series and would successfully provide a denotation for all the terms for positive integers. But the problem with this happy ending is that neither assumption is true. Without both pieces of information, however, correlating them through the Induction does not produce the right meanings for the integers. To see this, let's look at what children of this age know about the numerals and about increasing set size.

3.1. The successor function for numerals

One trouble with the Induction when applied to the knowledge of real 3- or 4-year-olds is that they *don't* know the successor function for numerals, and, therefore, the conclusion of the Induction doesn't settle the question of their meaning. To our knowledge, no developmental psychologist credits children of this age with knowledge of advanced counting. Correlating their knowledge of the sequence to the operation of adding one object is, therefore, of no help to them, since it can, at best, get them to understand the meanings of the short list of terms they have memorized (e.g., "one" through "nine" in Jan's case). Hence, the Induction can convey no knowledge to them of the meaning of most of the numerals – for example, no knowledge of how many positive integers there are, no knowledge of whether "fourteen" refers to fourteen or, instead, to fourteen, twenty-four, and so on.

Children at the critical age obviously don't know how to get from the numeral "*n*" to the numeral "successor(*n*)" for any *n*. But do they know enough about the structure of this sequence to constrain it in ways that would eliminate looping like that in Fig. 1? How much children know about this structure is the topic of an important line of current research, but a survey of the results to date suggests that children's knowledge of this structure is impoverished at (and just after) the point at which they would be making the Induction (Condry & Spelke, submitted for publication; Le Corre & Carey, 2007; Sarnecka & Gelman, 2004). For example, for a time after they have learned to count in order to determine the set size of collections – and so presumably after they have made the Induction – children will use the same number term (e.g., "five") to label sets of six, eight, or ten objects in situations in which they must estimate (rather than count) the objects (Le Corre & Carey, 2007). In short, children know that different number terms have different meanings, but have little notion of what constitutes this meaning and, consequently, too little information to rule out alternative number systems like that of Fig. 1.

Our understanding of the Induction proposal is that once children make the Induction they thereby attain the meaning of the full set of positive integers. M&L, however, may have a different conception in which children retain the conclusion of the Induction, but they don't fix the meaning of the larger positive integers until they encounter them sometime later. Is it possible that children could acquire the positive integers by extending the conclusion to a new set of numerals (e.g., "four" to "nine"), freezing these new meanings once and for all, repeating the same process when they learn the next batch of numerals (e.g., "ten" through "one hundred"), and so on? Obviously not, if the individual batches of numerals are finite, since the learner would never reach the end. Would such an iterative process work if the final batch of numerals is infinite? Perhaps the iterative process could continue until the children finally learn advanced counting, at which point they would acquire the full set of positive integers. Although the assumption about freezing remains psychologically unrealistic and prohibits learning the C major scale and other circular lists, this variation does avoid the difficulty with looping, such as that in Fig. 1. It does so, however, simply by taking advantage of the fact that the structure of the numerals in advanced counting *is* the structure of the positive integers. This maneuver just shifts the question of how children learn the positive integers to the equivalently difficult one of how children learn the numerals for the positive integers.

3.2. The successor function for numbers

M&L believe that the fact that Jan ends up with the wrong meanings for the numerals is due to a failure on our part to formulate unambiguously the conclusion of the Induction. Instead of supposing that what the twins learn is that the numeral "successor(n)" refers to the property of collections containing one more than those "n" refers to, we should have specified that what the twins learn is that the numeral "successor(n)" refers to "one more than n" (M&L, p. 8). The idea is to rule out the possibility that, for example, the successor to "three" could refer to four, fourteen, . . ., rather than to plain four. But this reformulation fails to do the work M&L want it to. For even in the system of Fig. 1, "successor(n)" *does* refer to one more than "n" refers to: "two" refers to precisely one more than "one," "three" refers to precisely one more than "two,". . ., and "ten" refers to precisely one more than "nine" and one more than "nineteen." Addition by 1 is a well-defined function in this system.

What M&L must have in mind is that "one more than n" should be specified independently of addition. But M&L's formulation is unclear as to how children could identify and represent the right referent. The child must see, not just that you can get a set of two by adding one object to a set of one, and a set of three by adding one to a set of two, but that one object can *always* be added to a set to produce a new set with a novel cardinality. But why would the child think this is so rather than – what may well seem a more plausible alternative – that there's a point at which we run out of things to add?

If the sets that provide the meanings for the integers are sets of ordinary physical entities (pajamas, trees, and toasters), then there is no guarantee that these sets will support the standard system. Although there are an infinite number of integers, it is uncertain whether there are an infinite number of physical objects. If there are only k objects, then the meanings of the numerals for "$k + 1$," "$k + 2$,". . ., will be undefined. Even if there happen to be an infinite number of physical objects, it seems odd to make the positive integers hostage to an empirical fact like this one. This is why formal attempts to define the integers in terms of sets have to posit axioms (e.g., the axiom of infinity) to ensure the existence of large enough sets (see Shapiro, 2000, chapter 5, for a discussion). Moreover, if producing "one more than n" means physically grouping objects, there is no guarantee that the operation will result in the right cardinalities, since physical grouping sometimes destroys objects in the process. (For related remarks about the difference between addition and object grouping, see Leslie, Gallistel, & Gelman, in press.)

This is particularly a problem for theories like M&L's in which the Induction requires a switch from a weak to a much stronger representational system. According to M&L, children start out with single nodes in a connectionist network representing each number: one node each for one through four (see M&L's Fig. 2). But no such representation could possibly do for the full set of positive integers, since humans can't represent the integers with an infinite number of nodes. As a consequence of performing the Induction, children must therefore graduate to an entirely different representational system, a system that represents the integers in a potential or generative way. To attain the right structure, the representation has to depict the integers as having a unique starting point, having no loops, and having no extraneous elements. Where do these constraints come from? Not from the preliminary representation or from external objects (as we've just noted), and not from the sequence of numerals (for reasons discussed in Section 3.1). The Induction's deficiency is that it provides no source for these restrictions.

4. Concluding comments

Children must learn that number terms in phrases like *three bears* denote the numerosity of a collection. Since they can't learn this connection one-by-one for all the positive integers, they must at some point come to recognize a general connection between the sequences of numerals and numerosities. In this respect, children's learning of numerals seems to differ from that of chimps, who never manage the generalization (Matsuzawa, 1985). The Induction is one way to implement this recognition. However, we question how much the Induction can do in providing children with the detailed meaning of the terms for the positive integers. Recognizing that there is some general connection or other is not enough to dictate a specific referent for the infinity of numbers.[5]

Number terms have uses other than in phrases denoting numerosities; so it seems likely that their meanings are much more general (Rips et al., submitted for publication). Even within the realm of numerosities, however, the Induction cannot settle the question of which numerosities go with which integers. With no understanding of the function that yields ''successor(n)'' and with no understanding of the structure of the numerosities, children can't use the Induction to secure the correct definition of the integers. M&L believe that to attain these meanings ''there is no need to represent the axioms of arithmetic, since children's representations will conform to them automatically,'' given the Induction as M&L have formulated it (p. 10). If this were correct, M&L should be able to produce a *proof* that the Induction yields the Dedekind-Peano axioms. But no such proof will be forthcoming, since structures like that of Fig. 1 provide counterexamples. Although children need not represent the axioms as such (and certainly need not represent them consciously), they do need to represent equivalent constraints in order to make the Induction work. Once these constraints are in place, however, the Induction is redundant, since the constraints yield the structure of the integers.

Acknowledgments

We thank Susan Carey, Rob Chametzky, Winston Chang, Stella Christie, Dedre Gentner, Susan Hespos, Eyal Sagi, and an anonymous reviewer for comments about

[5] An indication of just how difficult it is to pin down all and only the natural numbers, even with all the resources of first-order logic, comes from work on nonstandard models of arithmetic (see e.g., Boolos & Jeffrey, 1974, chapter 17).

the issues discussed in this paper. We also thank Eric Margolis and Stephen Laurence for comments that caused us to clarify our position.

References

Boolos, G. S., & Jeffrey, R. C. (1974). *Computability and logic* (1st ed.). Cambridge, UK: Cambridge University Press.

Carey, S. (2004). Bootstrapping and the origins of concepts. *Daedalus, 133*, 59–68.

Clark, E. V. (1997). *The lexicon in acquisition*. Cambridge, UK: Cambridge University Press.

Condry, K. F., & Spelke, E. S. (submitted for publication). Young children's understanding of number words and verbal counting.

Gallistel, C. R., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition, 44*, 43–74.

Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.

Hartnett, P. M. (1991). *The development of mathematical insight: From one, two, three to infinity*. Unpublished Ph.D. Thesis. Philadelphia: University of Pennsylvania.

Hurford, J. R. (1987). *Language and number: The emergence of a cognitive system*. Oxford, UK: Blackwell.

Le Corre, M., & Carey, S. (2007). One, two, three, four, nothing more: An investigation of the conceptual sources of the verbal counting principles. *Cognition, 105*, 395–438.

Leslie, A. M., Gallistel, C. R., & Gelman, R. (in press). Where integers come from. In: P. Carruthers (Ed.), *The innate mind: foundations and future*. Oxford, UK: Oxford University Press.

Margolis, E., & Laurence, S. (2008). How to learn the natural numbers: Inductive inference and the acquisition of number concepts. *Cognition, 106*, 924–939.

Matsuzawa, T. (1985). Use of numbers by a chimpanzee. *Nature, 315*, 57–59.

Reich, P. A. (1976). The early acquisition of word meaning. *Journal of Child Language, 3*, 117–123.

Rips, L. J., Asmuth, J., & Bloomfield, A. (2006). Giving the boot to the bootstrap: How not to learn the natural numbers. *Cognition, 101*, B51–B60.

Rips, L. J., & Asmuth, J. (in press). Mathematical induction and induction in mathematics. In: A. Feeney, & E. Heit (Eds.), *Induction*. Cambridge, UK: Cambridge University Press.

Rips, L. J., Bloomfield, A., & Asmuth, J. (submitted for publication). From numerical concepts to concepts of number.

Sarnecka, B. W., & Gelman, S. A. (2004). Six does not just mean a lot: Preschoolers see number words as specific. *Cognition, 92*, 329–352.

Shapiro, S. (2000). *Thinking about mathematics: The philosophy of mathematics*. Oxford, UK: Oxford University Press.

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