The Innate Mind

Volume 3: Foundations and the Future

Edited by
Peter Carruthers
Stephen Laurence
Stephen Stich

OXFORD
UNIVERSITY PRESS
2007

ALAN M. LESLIE, C. R. GALLISTEL, AND ROCHEL GELMAN

Where Integers Come From

If the primitive preverbal symbols for numbers are noisy mental magnitudes, what is the preverbal foundation for our concept of an integer? We argue that the essential problem is to answer the question where our notion of exact equality or perfect substitutability comes from. In practice, real valued variables are never exactly equal; nor is it easy to specify an algorithm for establishing exact equality (in the limit) between two random Gaussian variables. Furthermore, because number concepts must support arithmetic inference, a necessary part of the psychological foundations is the integer concept ONE. ONE is required because it is the multiplicative identity element for which no other value, approximate or exact, can be substituted. Moreover, ONE is required by the successor function, which generates all the other positive integers. We further argue that an essential constraint on any proposal for discrete (integer-valued rather than real-valued) mental symbols is computational compatibility with the real- (or rational-) valued mental magnitudes that represent continuous quantity. These constraints rule out most current proposals that postulate systems of discrete numerons or other symbols representing only very small numbers. We consider alternative proposals.

Die ganze Zahl schuf der liebe Gott, alles Übrige ist Menschenwerk.

—Leopold Kronecker

1 Introduction

Among the earliest quantitative concepts that we have language for are the first few counting numbers, {one, two, three}. They appear in development long before other types of number concepts, such as fractions, decimals, and complex numbers. When preschool children begin to count, these are the numbers they use (R. Gelman and Gallistel, 1978). Only much later, under formal instruction, and with considerable difficulty, do children learn about the mathematical concept and notation for fractions

(R. Gelman, 1991; Hartnett and Gelman, 1998), and, following that, about real numbers. People in most cultures use counting numbers, whereas the discovery of the reals appears to have required a series of historical singularities. If a language provides any explicit number words, these almost always[1] will denote at least the first few count numbers. Why is this? Where do integer concepts come from?

For us, it is critical that number concepts support arithmetic inference. This means that the concept must denote an entity over which arithmetic operations can operate. We do not require that an individual who possesses the concept must actually be able to perform a given arithmetic operation over that concept—performance and developmental constraints may prevent that. However, among the individuation conditions on numeric concepts is that they support arithmetic inference (R. Gelman, 2006). This rules out considering notions such as "a" or "few" as number concepts because they can never support arithmetic inference. Thus the sentence (a + a few) is not interpretable and it is an unnatural utterance in everyday speech.

2 Number Words

In a language like English, the words encoding natural number concepts are everyday words with mundane uses, such as counting, and languages that have such words are quite common. The natural number concepts support arithmetic operations and intuitions, and they denote *exact* integer values. The word "three" in English, for example, does not ordinarily refer to a range of real values or to a stochastic function over real values; in general, natural number words do not refer to entities such as "value(s) in the range 2.5 to 3.49" or "a Gaussian with mean 2.0". Instead, words such as "one, two, three" refer to exact values, such as 1, 2, 3, and do not mean 1-ish, 2-ish, 3-ish.[2]

This fact about adult usage presumably reflects the fact that children are disposed to learn that "one" means 1 (exactly), "two" means 2 (exactly), and so on. If children were not so disposed, but were disposed instead to learn that "one" means 1-ish, "two" means 2-ish, and so on, then presumably number words in natural languages would commonly denote ranges of real numbers. The predominant disposition, in turn, reflects the fact that children tend to entertain and settle on integer-valued hypotheses in preference to other possibilities such as "vaguely 2-ish," "Gaussian with mean 2.1," "values in the range 1.5 to 2.49," and so on. In this regard, the natural number words are strikingly unlike color terms, which do refer to ranges of values in color space or to stochastic functions over such ranges. The word "red" does mean red-ish; "green," green-ish; and so on. Adult color terms presumably have such inexact meanings just because, as children learning the meanings of those sounds, they were disposed to entertain, and settled on, hypotheses that referred to inexact regions of color space, perhaps because the brain may be incapable of remembering exact locations.

1. But see Flegg (1989) on the early widespread use of Pure-2 Counting in parts of Africa, South America, Australia, and New Guinea. In Pure-2 Counting, there are distinct words for one and two; the rest of the count words are derived by combining these words. The count list *urapon, ukasar, ukasar-urapon, ukasar-ukasar, ukasar-ukasar-urapon,* used by the Gumulgal of Australia, is but one example of such counting lists. It is not clear how many groups still use this or the more complicated version of Neo-2 systems. It is noteworthy that the system is generative. This is not so for the initial count words in English and many other languages. Nothing about the sound "one" predicts that "two" will be next.

2. However, Fox and Hackl (2006) argue that many facts about implicature imply that the mental scale to which even the counting numbers map is dense, that is, continuous, like the mental magnitude system. They argue on purely linguistic evidence that the scales underlying all mental quantification, whether of discrete or continuous quantity, are dense.

Furthermore, the count number words are not "vague" in the technical sense studied by logicians. Words whose meaning is vague in this sense are quite common. For example, *bald* applies to someone who has no hair on his head, and also to someone who has one hair on his head, and to someone with two hairs... and so on. But the number of hairs a head must have in order to stop being bald and start being hirsute cannot be specified, thus qualifying the meaning of the word as vague. Interestingly, many words for quantities are also vague in this sense. What does a rock have to weigh to be considered "heavy"? What's the cutoff point for being "tall"? How many people are required for there to be "many" people? And so on. By contrast, the meanings of the words *one, two, three,...* are not vague; indeed, they are the very paradigms of precision and exactitude when applied to discrete entities: They mean 1, 2, 3. Hurewitz, Papafragou, Gleitman, and Gelman (2006) provide evidence that the distinction between linguistic quantifiers ("some," "all," for example), which are vague, and cardinal values, which are not, is available to young language learners.

We are not claiming that words whose meanings are real values are impossible to learn, nor that integer words are mandatory in all human languages. The first claim is obviously false (for example, pi), and the second is an open question with some evidence to suggest it, too, may be false (Gordon, 2004a, 2004b; but see Gelman and Butterworth, 2005).[3] What we do wish to claim is the following: When preschool children identify that the meaning field for a given lexical item may be a *numerical* value—as they might, for example, in an activity such as *counting*—they expect that word to denote some positive integer value. To say that they expect such words to denote the natural numbers means that they draw their hypotheses regarding possible numerical values from a restricted hypothesis space, namely, the space of (the first few) positive integers. Indeed, the adult number words "one, two, three" come to refer to exactly 1, 2, 3, respectively, just because children entertain hypotheses restricted to integer-valued referents. If children supposed, for example, that real-valued referents or vague numerical referents approximately centered around 1, 2, 3, respectively, then that's what these words would (come to) mean. There is no immediately obvious reason why these are not the "correct" meanings if the hypothesis space consists of noisy reals. In this case, integer-valued hypotheses would have only an infinitesimal probability of being entertained, that is, would never be entertained. This means that no language containing words for the natural numbers would ever be learned. In fact, these are the first number words to be learned.

3. Gordon (2004a, 2004b) reports that a small, isolated group of around 300 Amazonian villagers speak a language that may lack any words for the natural numbers. There is considerable debate regarding the reliability of the Gordon data. See R. Gelman and Butterworth (2005) for discussion of comparable reports about other isolated groups.

3 Exact Equality

One use we make of integers is counting things. A fundamental intuition here is that if three things are counted, then the resulting cardinal value will be exactly equal to the cardinal value that will result from counting them again. By contrast, two measurements of the same continuous physical quantity will yield the same answer twice *only* by error (for example, rounding error), because it is impossible in principle to determine the value of a continuous (that is, real-valued) quantity such as length or duration with perfect precision (zero residual uncertainty). By contrast, counting the members of a set requires the use of integers; and this means that repeated counts should yield exact equality *unless* there is counting error, as when an item is skipped or double-counted.

Exact equality has been taken unself-consciously for granted by cognitive theorists. However, exact equality challenges most current models because the latter relate basic human numerical concepts to an underlying analog representation (Dehaene, 1997; Gallistel, 1990; Gallistel and Gelman, 1992; Wynn, 1992b, 1992c). The essential idea of the magnitude representation is that the brain represents numbers not as a series of discrete symbols such as the Arabic numeral system or the binary digits in a computer, but as a continuous quantity, such as charge in a capacitor, or water filling a test tube, or a needle moving along a linear scale (as in a speedometer). In these models, counting a set is pictured as adding successive drops of "water," "charge," or some other analog quantity, such that each drop corresponds to a distinct member of the set to be counted. The quantities accumulate in a "container," raising its "water level," or a needle is moved a regular distance along a scale, so that each rise or movement corresponds one to one with the members of the counted set. The final level or point on the scale reached thus represents the cardinal value of the set counted. It is assumed that some analogous process of accumulating physical quantities takes place in the brain as a person counts.

Dehaene (1997) discusses evidence from cognitive and neuroimaging studies of human calculation that supports the existence of an internal number continuum (mental number line). For example, the time taken by adults to compare the magnitudes of two numbers increases as the differences between the two numbers decreases. There is also impressive evidence for the existence of an analog magnitude representation in animals, suggesting a long evolutionary history of this basic numerical capacity (see Gallistel, 1990; Gallistel and Gelman, 2005, for reviews). For example, Platt and Johnson (1971) trained rats to press a lever n times before pressing a second lever to obtain a reward. Rats learned to press the first lever a mean number of times equal to n with variability (standard deviation) proportional to the mean. In other words, as the target n increased, so did the rats' bar pressing, with an error rate that was a constant proportion of the size of the target. Numerical estimation in human adults (Cordes et al. 2001), infants (Brannon, Abbott, and Lutz, 2004; Xu, 2003; Xu and Spelke, 2000), and children (Cordes and Gelman, 2005), as well as in animals, appears to respect Weber's law. This supports the existence of an underlying representation in the form of a noisy real-valued magnitude.

Dehaene (1997), after reviewing evidence on both infant and adult human numerical abilities, suggests that we are endowed with a "continuous and approximate

representation of quantities" (p. 86) and that, despite being able to "convey numbers using... digits," the brain always automatically converts into an "internal analogical magnitude" representation that he dubs the "number line" (p. 87). However, at the same time that he adduces evidence that the brain uses an analog magnitude representation, Dehaene also argues that our basic number concepts refer to integers, stating that the "number line... clearly supports a limited form of intuition about numbers [in that] it encodes only positive integers" (p. 87). Either of these claims can be argued for individually, but to say that the number line is *both* continuous (real-valued, hence dense) *and* encodes only positive integers is a mathematical contradiction. The real-number line has no special pit stops for integer values; they are just ordinary values among transfinitely many other reals. It is extremely puzzling how an analog magnitude representation could support *only* integer concepts. Indeed, given noise and considerations about the impossibility of determining the exact value of a real-valued empirical quantity, it is puzzling how analog mental magnitudes could directly represent integer concepts *at all*. Thus, it is unclear how, exactly, an analog magnitude representation could be identical with our basic integer representations. Let us look at this more closely.

3.1 *Analog Accumulations*

The best-developed model of the magnitude representation is the "accumulator" model of Meck and Church (1983), depicted in figure 7.1. The accumulator can

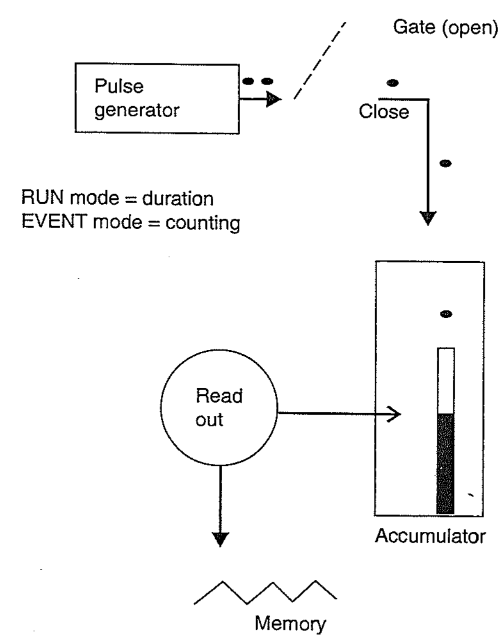

FIGURE 7.1 The accumulator model.

measure continuous time intervals (in "run" mode) or count discrete entities (in "event" mode). In run mode, pulses are gated into an accumulator at the beginning of the interval, and continue to accumulate until the gate closes at the end of the interval, preventing further accumulation. The resulting magnitude is then read out into memory. This real number represents the length of the time interval. In event mode, pulses are gated into the accumulator for a fixed amount of time for each item to be counted. In this case, the accumulator is incremented by a discrete amount for each item counted. At the end of the count, the accumulated level is again read out into memory. This number is proportionate to the number of items in the set counted, and thus represents the cardinal value of the set. Although the way gating works in event mode means that the process picks out a sequence of next magnitudes (a hallmark of the discrete), each increment is assumed to be a noisy real value. The gate is open for a duration that has continuous value, and each of the gated pulses is a continuous magnitude, so the sum of these must itself be a noisy real-valued magnitude.

The assumption that the noise (uncertainty) in the resulting magnitudes is proportional to the sum is central to this model. It is this assumption that explains why Weber's law is observed to hold in a multitude of number and magnitude estimation tasks in animals and humans. For example, in a classic study of counting in rats, Platt and Johnson (1971) showed that the probability of breaking off a sequence of lever presses when N was the number of presses reinforced was a distribution with mean N and a variance proportional to N. The coefficient of variation was therefore constant across the range of values for N. Human adults and children show the same scalar variability in rapid counting and magnitude estimation tasks (Cordes et al., 2001; Cordes and Gelman, 2004; Whalen, Gallistel, and Gelman, 1999), as well as the size and distance effects in number order judgment tasks (Buckley and Gilman, 1974; Dehaene, Dupoux, and Mehler, 1990; Holyoak, 1978; Moyer and Landauer, 1973). Scalar variability is one explanation for these effects; logarithmic compression with mean mental magnitude proportional to the logarithm of the number represented and variability-independent of it is another. In either case, the key explanans is that the underlying representations are inherently continuous, and therefore noisy and variable (for review, see Gallistel, Gelman, and Cordes, 2005). In short, there is compelling evidence for a noisy analog continuous magnitude representation underlying counting and other number tasks in animals and humans.

The strong point of the accumulator model, however, makes it hard to see why our basic number concepts—the ones picked out by language—should be integers rather than reals. One problem is that there is nothing in the account that explains why each discrete value added to the accumulator should equal exactly 1 rather than some real number (perhaps varying around 1). Similarly, accumulated values will be noisy and will never be exactly equal to integer values.[4] Furthermore, the values stored in memory will be noisy and continuously variable. Any numerical observations that a learner makes in the course of quantifying will therefore take place in a vocabulary of the reals. One plausible explanation of why the colors designated by color words are understood as imprecisely specified is the noisy (uncertain) values of remembered coordinates in color space. The accumulator account predicts a similar pattern for our basic number concepts. Therefore, nothing in the accumulator account so far explains why our basic number notions are integers. Of course, the Meck and Church model was developed to account for numerical capacity in animals, and we don't know whether the basic number concepts that *animals* may have actually include the natural numbers. One powerful reason we have for believing that humans have integer concepts is our intuitions regarding what our natural number words mean; moreover, the extensively documented history of mathematics is clearly based on those intuitions (all of number theory, for one example). So, can language learning come to the rescue of the accumulator model and show us how we get integers? Perhaps accumulator magnitudes can be translated into or constrained to integer values. Could language learning perform this trick?

4. "Never" in this context and throughout this chapter means *with an infinitesimal probability*.

4 Learning Number Words: Is This the Source of the Integer?

Gallistel and Gelman (1992) argued that children learn a bidirectional mapping between the preverbal magnitudes produced by the accumulator and the number words. They suggest that children recognize the formal similarity between the nonverbal counting process proposed by Meck and Church (1983) and the verbal counting process. Both produce a one-one correspondence between a stably ordered set of symbols (successive magnitudes in the one case, successive count words in the other) and the items in the to-be-counted set, and they both use the final symbol to represent the cardinality of the set. Gelman has repeatedly emphasized that the achievement of cardinal values is subject to arithmetic principles. In both the animal and the human cases, the meaning and use of the symbols are subject to arithmetic processing. Interestingly, young children who are still learning to count do better at counting and show clearer evidence of understanding the numerical referents of the count words when their counting is embedded in tasks that involve arithmetic processing (R. Gelman, 2006).

Learning to map number words onto magnitudes accomplishes essentially what our "speedometer" mechanism in figure 7.1a does: A labeled grid is laid alongside the magnitude representation, calibrating it in terms of integer values. The particular integer values on the "dial" gain their sound labels from the language the child is learning; if it is English, then "one," "two," "three," and so on, in that order. The underlying magnitudes, of course, remain noisy and real-valued quantities:

> A pivotal assumption about mapping from digits to preverbal magnitudes is that there is variability in the magnitudes to which a digit maps and this variability obeys Weber's law: the standard deviation of the distribution of magnitudes to which a digit maps increases in proportion to the mean magnitude. (Gallistel and Gelman, 1992, p. 55)

Attractive as this hypothesis is, it does not give us everything we need. It does not explain where the concept of the exact equality of two different instances of the same integer comes from. Learning the meaning of a word is learning to associate a sound with a meaning (concept). To do this, the learner must test hypotheses about what the concept might be that is encoded by the new sound. When a given hypothesis has been sufficiently confirmed, it is recorded in memory as the meaning of that sound. As we saw earlier, the problem with count words is how to explain why an integer hypothesis should be regularly entertained and confirmed, if the underlying space of available hypotheses is the real numbers. If the *only* underlying numerical representation the *pre* verbal child has access to is a continuous magnitude, and moreover a noisy continuous magnitude, the likelihood of entertaining a particular exact value as a hypothesis will be infinitesimal. No child would ever entertain exact integer values, and for just that reason no language would contain common words for 1, 2, 3; instead, *one*, *two*, *three* would be specialists' words, like *pi* or *e*. Yet children regularly do entertain integer hypotheses, and for just that reason, the count words (in languages that have them) have integer values as their meanings. The underlying hypothesis space for children learning count words is not a space of continuous magnitudes, but is biased toward a space of integer values.[5]

Another possibility arises from the fact that the values obtained by repeated counting of the same small numerosity would be strongly clustered. The empirically determined coefficient of variation (Weber fraction) for adult nonverbal counting is in the range .12 to .15, so the distributions of values obtained for repeated counts of a set of two objects would overlap with the distribution of values obtained for repeated counts of a set of only one object between the second and third standard deviations. Language learners might parse the values obtained by counting into clusters and assume that the words referred either to these distributions implied by the clusters or to the intervals over which given distributions dominated. However, this would not license the conclusion that two instances of valid reference for the word "two" were perfectly substitutable. Values drawn from a common continuous distribution are not substitutable, because they are never exactly the same.

There is evidence that young children who are still learning to count can reason about discrete numbers in a manner that respects substitutability. They seem to recognize the existence of additive inverses. The additive inverse of a number is the number that when added to a set whose numerosity has been altered by addition restores the numerical value of the set. "Added" and "addition" are used here in the technical sense that subsumes subtraction under addition. Thus, when a set has been reduced by adding −3 to it (that is, subtracting 3), its numerical value can be restored by adding 3 to it, because +3 and −3 are additive inverses; adding one cancels the effect of adding the other.

This evidence first appeared in the behavior of children in the magic paradigm (R. Gelman, 1972), who confronted a plate from which *one* mouse had been surreptitiously added or subtracted, making it no longer a "winner" plate. Children noticed the numerical change and saw it as central to the question of whether the plate was or was not a winner (unlike, for example, changes in item identities, which were not seen

5. This last remark applies after the point the child has figured out (somehow) that the count words refer to numerical values (that there *are* number words) and is now trying to figure out *which* values particular words refer to.

as critical). For present purposes, the most important result is that even the youngest children believed that the plate could be restored to winner status by the addition of *one* item (not some items, but *one* item). This suggests that they thought that adding 1 would cancel the effect of subtracting 1, and vice versa. It is not clear why they would think this if their reasoning rested entirely on operations with noisy magnitudes, because adding a magnitude drawn from a distribution centered on 1 will generally not exactly reverse the effect of subtracting a magnitude drawn from that same distribution. Indeed, as already noted, it is not clear how in a system that processed only noisy magnitudes, one could even confirm a restoration to the status quo ante.

The evidence for the recognition of inverse problems was extended by the arithmetized counting experiments of Zur and Gelman (2004). They had children predict the result of adding or subtracting between one and three items, and then check their predictions by counting. There was a marked tendency for the predictions to be more accurate when the children had already dealt with the inverse problem than when they encountered the problem without having already encountered its inverse. Thus, children were much more likely to say that 12 − 3 = 9 when they had already dealt with 9 + 3 = 12 than when they encountered 12 − 3 before encountering its inverse. This suggests that their reasoning assumes the existence of restorative inverses, even for numbers greater than 1.

Could learning the number words be constrained to integer values because the things that get counted with the number words are typically *physical objects*? Can the discreteness of physical objects somehow be ported over into the magnitude representation to yield integer values? For example, the child sees two cups and hears Mother say, "Here are two cups." Can the child use the fact that objects come only whole, as it were, to constrain the hypotheses about the meaning of "two" to whole-number integer values? This is so tempting, especially given how English conveniently uses the word "whole" for both cases! But notice that the "wholeness" of an object reflects how we individuate objects, *not* how we individuate numbers: Chip a little bit off a cup, and it is still the same old cup; "chip a little bit" off a number, and the result is an entirely new number. Trading on the polysemies of English is not helpful. In any case, the same problem arises as before. Whatever the discreteness of a "whole" object means, it does not disclose a number: A child still has to *count* objects to know there are *two* of them. But if the only counting-mechanism-cum-number-representation the child has *pre* verbally is the accumulator, producing noisy real values, then again the child will obtain a noninteger value, such as 1.94 or 2.[053], whenever he or she counts, and never exactly the value 2. So again, even in the co-presence of objects and the count words, children would entertain noninteger-valued hypotheses for the meaning of the count words and would never consider an integer value as a candidate referent. And again, as predicted by the continuous-magnitude-only hypothesis, we should expect that adult language communities would have count words (the ones young children learn) that refer to noninteger values, such as "Gaussian around 1.98," instead of to integers. Since there appear to be few or no such communities, it argues against the idea that magnitude representations underlie—form the hypothesis space for—number word learning. This in turn undermines the suggestion that such representations are the source of the natural number concepts.

To summarize so far, the evidence for the existence of an accumulator mechanism and its associated continuous magnitude representation seems good for both animals and humans. But, because this mechanism cannot produce integer values on demand, it *alone* cannot be the source of the natural number concepts.

5 Can "Object Files" Do the Trick?

According to one current suggestion, the relation between the origin of natural numbers and verbal counting is quite different from the above. The basic idea is that, in addition to an accumulator, young children are equipped with a second system with quite different properties. As noted earlier, infants can estimate the number of objects in sets as long as the sets contain four or more objects. But infants are also able to track the numerosity of objects in sets of three or fewer. Furthermore, whereas discrimination of large sets shows a classic Weber fraction, infant performance with set sizes of three or fewer apparently does not; with small set sizes, numerosity appears to be tracked exactly (although we note that the evidence for this claim of exactness is not as strong as one could wish). This has led to the suggestion of two independent systems (Carey, 2004; Feigenson, Carey, and Hauser, 2002; Feigenson, Dehaene, and Spelke, 2004; Xu, 2003; Xu and Spelke, 2000). These authors, however, do not propose that infants deploy natural number representations and count the number of objects in small sets. Instead, following suggestions by T. Simon (1997), Leslie, Xu, Tremoulet, and Scholl (1998), and Scholl and Leslie (1999a), they argue that infants represent object numerosity via an attentional mechanism that concurrently individuates multiple objects. This mechanism allows even young infants to track the exact numerosity of sets of objects.

According to Feigenson, Dehaene, and Spelke (2004), because of the accumulator, "humans are attuned to the cardinal values of arrays from the beginning of life" and, because of the second system for representing numerically distinct individuals, "concepts of 'enumerable individual' and 'adding one' are accessible throughout our lifetimes" (pp. 312–13). According to Carey (2004), the accumulator plays little role in the origins of integer concepts. Instead, she argues that the concurrent individuation of small sets of objects allows children to come to recognize sets of one, and then sets of two, and then sets of three objects. At this point, following Wynn (1992b), a crucial role is assigned to learning language—specifically, learning the meanings of what Carey calls the "count list," the words, *one*, *two*, *three*,...

Carey's proposal is that learning the count list focuses the child upon the difference between sets of one object and sets of two objects. The mechanism for concurrent individuation supplies the answer: The difference is *adding (another) one*. After this the child not only recognizes sets of one and sets of two objects, but can grasp what the relation between them is, namely, the successor relation. Subsequently, the child works on the meaning of the next word in the list, *three*, and figures out (thanks to the concurrent individuation mechanism) that the difference between sets of two objects and sets of three objects is, again, *adding another one*. Following this second discovery, children go on to generalize their induction to the entire count list: Each successive word in the list differs from the one before it by the addition of *another* one. Carey describes the way concurrent individuation and language learning engage each other as a bootstrapping process. She then goes on to claim that "coming to understand how the count list represents number... does nothing less than create a representation of the positive integers *where none was available before*" (Carey, 2004, p. 65; italics added).

Both Feigenson et al.'s and Carey's proposals afford a fundamental role to the mechanism for concurrent individuation of physical objects. We therefore need to examine this mechanism more closely. We turn to do this now before going on to evaluate whether it can bear the weight assigned to it by these accounts.

5.1 *Individuating Objects: Files, Bundles, and Indexes*

Feigenson, Carey, and their colleagues have made extensive use of the idea of concurrent individuation by way of representations called *object files*. Kahneman and Treisman (1984; Kahneman, Treisman, and Gibbs, 1992) introduced the idea of an object file because they perceived a missing link in traditional accounts of object perception. In traditional accounts, bottom-up sensory information is thought to directly activate long-term semantic memory traces; once the appropriate semantic categories have been thus activated—and only then—can the objects in the scene be identified and tracked. Traditionally, the task of keeping track of objects that change location was conceived of essentially as a search task. Initial contact with an object results in a memory description combining the sensory information and the semantic information it activates. When the object moves, the scene must be searched to discover which item in the scene matches this object description. When a matching item is found, then it must be the same object. Object representations, in the traditional view, are essentially feature bundles of one sort or another, including perhaps a semantic category label or a word, activated bottom up but imposed top down on sensory input. For Kahneman and Treisman this view missed important phenomena. For example, objects can be tracked through space without being identified (described); the same object can be tracked through changes in its identification ("It's a bird! It's a plane!"); and two "identical looking" objects can be perceived as distinct if there is a minute spatiotemporal gap between them, while two radically different-looking objects can readily be seen as a single transforming object (frog changes into a prince).

To accommodate such phenomena, Kahneman and Treisman introduced an intermediate level of object representation, which they called the *object file*. Object files are temporary object representations that interface between sensory information and long-term semantic information. There are two basic functional parts to the object file. The first, and in many ways more important, part of the object file is a continuously updated spatiotemporal code which locates the object corresponding to the object file. This is the *indexing* function of the object file; the file *points at* the object it refers to. We can think of this function as the file's "folder"—a container with only (continuously updated) spatiotemporal coordinates written on the folder's tab.

The second basic function of an object file is that the folder can have further information added, taken away, or changed. We can think of this as the sheets of paper that a folder might contain, each sheet having some property written upon it, either from sensory input or from long- term semantic storage. Together, the folder, plus any information it may contain, form an object file. In thinking about object files, we need to keep clear these two distinct functions. The folder may be empty, but it can still index and track an object *without* describing that object. In this regard, object file theory distinguishes itself radically from traditional theory. In traditional theory, an object representation just *was* a bundle of features; it consisted of nothing but a sheaf of papers, as it were. Without features, there is no feature bundle; without a bundle, there is no object representation. But an object file can represent and track an object, even if its folder is empty.

The way that Kahneman and Treisman thought of an empty object file as tracking an object was analogous to the way that a finger might track a moving object. One can pick out a particular object in a scene by touching it with one's index finger. Notice that the finger identifies the object *without describing it*. If you see only the *finger*, you have no idea whether it is touching something red or round or whatever. Instead, the touching finger helps you find the object without searching the scene because it indexes the object's location. Now imagine: When the object moves, the finger *sticks* to the object and moves with it.

The concept of the *sticky index* was highlighted and developed in Pylyshyn's FINST (Fingers of INSTantiation) theory (Pylyshyn, 1989, 2000). Pylyshyn argued that even spatiotemporal information does not have to be added to a folder; a coordinate code does not have to be written on the folder's tab. We can do without even that much descriptive information. Instead, a simple winner-takes-all network can solve the correspondence problem—matching the mental index to an item in the visual world—without explicitly representing coordinates in the object file and without requiring a top-down search (Pylyshyn, 2003).

5.2 *Indexing Objects and Number*

Howsoever it is implemented in the brain, indexing is an important and necessary function for any organism that tracks objects in real time. Leslie and colleagues applied these ideas to the long-studied problem of how infants come to individuate and track physical objects as they move and become occluded (Leslie, Xu, Tremoulet, and Scholl, 1998; Scholl and Leslie, 1999a). They chose to use the term *object indexes* in developing their approach to the infant object concept in order to emphasize this crucial and novel aspect of both object file and FINST theories. An object file may or may not contain a feature bundle, but it must minimally contain an index. So, how can object indexes represent numerosity? For each object that is attended in a set, there is a corresponding object file actively indexing its location. If all the objects in the set are thus indexed, then the numerosity of the set of objects will be mirrored mentally in the numerosity of the set of active object files. Not surprisingly, there is a limit to the size of sets that can be so represented. The evidence for a limit comes from multiple object tracking in adults where the limit is usually around four (Pylyshyn and Storm 1988; Scholl and Pylyshyn 1999; Trick and Pylyshyn 1994a, 1994b; but see Trick, Jaspers-Fayer, and Sethi, 2005) and from Feigenson's studies with infants (Feigenson and Carey, 2003; Feigenson, Carey and Hauser, 2002; Feigenson, Carey and Spelke, 2002; reviewed in Feigenson, in press), where the limit appears to be three.

One conclusion from the idea of object indexing under a limit is that infants may track small sets of physical objects and detect numerosity changes (Wynn, 1992a) without actually counting, and without having any symbol that represents or refers to the numerosity of the set of tracked objects. That is, they may detect numerosity changes simply by distinctly remembering each individual in the set, in which case all that is demonstrated is their commitment to object permanence. This raises an alternative interpretation to the explicit representation of numerical value that Wynn originally proposed. But that is all it does. Specifically, it provides no help whatsoever in understanding where integers come from. Because our claim goes against the claims of Feigenson and colleagues, and especially against the proposals of Carey, let us look further into this issue.

6 Implicit and Explicit Representation

Bootstrapping is an account of how a concept that, prior to the bootstrap, was not available to the learner can become available. Bootstrapping accounts should have two properties. First, the concept that was not formerly available should not be *expressible* by any combination of formerly available concepts. The new concept really should be *new* (not just more accessible). Second, the account of the bootstrap must specify a computational process that will take a combination of available concepts and yield a *new* concept, in this sense. Without specifying the bootstrap process itself, we simply have a claim that such an (unspecified) process is possible, but no way of evaluating that claim. For the skeptic, an unspecified bootstrap is akin to the magician's trick of turning a glass of milk and a few cards into a rabbit; no matter how much it *looks like* a rabbit has been pulled from a hat, without being able to imagine the natural process, it is not believable. As will become clear, we don't think that Carey has so far provided, even in outline, a computational account of a bootstrap for number.

It will help us to be clear if we establish some terminology. Marr (1982) introduced a useful distinction between *implicit* and *explicit* representation. This distinction is not what has become the more common usage of these terms as synonyms for unconscious and conscious, respectively. The popularity of the latter usage strikes us as unfortunate because we already have perfectly good words for *those* senses: unconscious, tacit, versus conscious, verbalized, and so on. By contrast, Marr's distinction revolved around whether or not a given representation made a certain piece of information available to other processes *directly*—without further inference being necessary—in which case that representation represented that piece of information *explicitly*. If a given piece of information could be recovered from a given representation only by processing that representation further—for example, by drawing inferences from it—then that piece of information is represented (only) *implicitly* by that representation. Marr provides us with a simple terminology for a fundamental property of computational systems.

It will also help if we are careful to mark when we refer to situations in the world (that the child may be thinking about) versus when we refer to concepts or strings of concepts (thoughts) that the child may possess and use. Let us use italics for when *we* describe situations in the world (that the child may in *some* manner be thinking about). And let us use small capitals when we refer to the concepts or thoughts that we believe the *child* uses to describe that situation. So, for example, when a child sees *a quantity of* H_2O (situation as we describe it), he or she may think, THERE IS WATER (situation as child's thought describes it).

6.1 Can Number Be Bootstrapped from Nonnumerical Concepts?

With these distinctions and practices in mind, let us return to the claims that have been made regarding object indexing and the origin of number concepts. Feigenson, Dehaene, and Spelke (2004) conclude that the object indexing system

> serves to represent numerically distinct individuals...and allows multiple computations over these representations. These computations include...representing the number of individuals in an array. Because this second system is also active in infancy, concepts of *"enumerable individual"* and *"adding one"* are accessible throughout our lifetimes. (pp. 312–13; italics added)

In a similar vein, Carey (2004) describes the object indexing system as a "system of representations with *numerical content*" (p. 61; italics added). She then develops proposals regarding how the indexing system's *numerical content* can, together with learning the list of count words in a language like English, allow the child to "bootstrap" his or her way to the integer concepts. As we saw earlier, the proposed bootstrap hinges on using collections of active indexes to represent sets of different sizes, then to have the child observe that each of the sets can be ordered under the *add another one* relation, and finally to see what Carey calls the "wild analogy" between this ordering and the ordering of the quantities designated by the words in the ordered count list (p. 67).

The first problem with this hypothesis is the assumption that the object-indexing system has numerical content. This assumption confuses properties of the symbols themselves—the oneness, twoness or threeness of a set of object files in the mind of the infant—with what those symbols refer to or represent. Having two object files pointing to two perceived objects is not the same as having a symbol (or symbol string) that refers to the numerosity of the set composed of the objects to which those two object files individually point. The *twoness* of the set of object files does not make that set a TWO symbol any more than the twoness of the symbol string "12" (the fact that the string is composed of two numerals) makes it a symbol for *two.* If the child assumed that the word "two" had the same referent as the two objects to which a particular set of two object files pointed, it would assume that the word was a name for that particular pair of objects. This would rapidly lead to massive confusion about what "two" could possibly refer to, because the child would hear it used to refer to many different pairs of objects having nothing in common. A child (or any other symbolic system) that lacked a symbol (or symbol string) that referred to *twoness* could not entertain the hypothesis that what all those sets had in common was their *twoness.* If the child has no symbol that refers to *twoness,* how can it learn that that is what "two" refers to?

The problem of inducing a word's reference is hard enough, even when we stipulate that the language learning system has symbolic resources that enable it to refer to whatever it is that the to-be-learned word refers to. On the face of it, it seems impossible when the system lacks the symbolic resources to refer to that which the word refers to. If one thinks, nevertheless, that bootstrapping can do this, then the process must be specified. And, to repeat, sets of object files do not—and cannot—refer to the numerosities that they instantiate. Conversely, if a set of two object files referred to *twoness,* then it could not refer to two particular objects. The essential feature of twoness is that it is a property of any set of two objects; in short, twoness is the cardinal value of a set, not something that refers to particular objects. Certainly, twoness is a property both of the set of objects to which the object files point and of the set of object files that point to those two objects; but the set of object files no more refers to *twoness* than does the set of the objects themselves.

The very first step in the proposed bootstrap also seems to us to be deeply flawed. In order to work, the bootstrap needs to assume what it sets out to explain, namely, how the child thinks thoughts such as ONE ADD ONE EQUALS TWO, especially when the most reasonable gloss for *"one , add one"* is *"add another one."* Notice, in regard to Feigenson and colleagues' proposal, that there's a big difference between an infant thinking about *two enumerable individuals* and an infant thinking THOSE ARE TWO ENUMERABLE INDIVIDUALS. The first might plausibly be true, for example, of an infant who has indexed, say, two apples sitting on a table. It has object files pointing to those two objects. But the second claims that the infant actually internally *describes* the apples as "enumerable individuals"—in *those* very terms. Fortunately, the ambiguity is fairly harmless in this case, because the second reading is presumably so implausible. But in the case of "add one" or "add another one" the ambiguity is quite pernicious and leads to a fatal question-begging. Again, it is one thing for an infant to be thinking about *one individual added to another* (situation as we describe it)—for example, a situation in which one apple is placed in a location nearby another. So far, nothing has been said about *how* the infant is thinking about that situation. It is an entirely different thing to say that the infant is thinking ONE ADD ONE EQUALS TWO. Of course, as soon as we do say this, we uncover another ambiguity, this time the ambiguities in the English phrase "add one." This phrase is commonly used to describe the physical event of placing an individual ("one") in some location alongside other individuals. Indeed, the verb "add" is used for all sorts of events that are not in the least arithmetic (e.g., "italics added"). But none of these other meanings of the phrase are relevant to the issue of who thinks thoughts like ONE ADD ONE, with the arithmetic reading 1 + 1. Thoughts like ONE ADD ONE (1 + 1) make numerical and arithmetic information *explicit.* By contrast, thoughts like PLACE OBJECT_i IN LOCATION X NEAR OBJECT_j make spatial information explicit, but leave numerical information at best *implicit.* Granted, a special class of inference process can operate on spatial representations like this to make the numerical information explicit. For example, applying an inference procedure such as *counting* to the objects in question will result in a representation that makes the cardinality of the set explicit. But then again, this assumes exactly what the bootstrap sets out to explain.

Even the English word "one" is multiply ambiguous. In the present context, the ambiguity of "one" meaning individual and "one" meaning 1 is particularly troublesome.

We must be careful about what we attribute to the child as explicit representations (as always, in Marr's sense). And indeed, we must extend this care to the notion of object files, too. If one had to translate the indexical function of object files into English, it would translate as a bare demonstrative, such as the word "that" when used on its own simply to point at some individual object, event, or property (Leslie and Káldy, 2001; Pylyshyn, 2001). Using such a concept will provide something that *could* be counted (by someone who can count), but "that" does not mean 1, and "that and that" or "that and another one" does not mean cardinal number 2. When Carey describes the object indexing system as a "system of representations with numerical content," this is true only in the same and unhelpful sense that the concepts APPLE and DOG have numerical content. One can use APPLE to refer to a situation with an *apple* and one could, if one is able and so inclined, count the *apple* to discover that one apple is present. Likewise, DOG can be used to refer to *dog* situations that provide opportunities for counting, again if one is able and inclined to count. But the fact that a *dog* may have *one tail*, *two ears*, *three toes*, and *four legs* does not mean that the concept DOG has numerical content. Likewise, thinking THAT$_i$ (as referring to a specific object$_i$) is not at all the same as thinking ONE (1); nor is thinking PLACE THAT$_i$ NEAR THAT$_j$ the same as thinking ONE ADD ONE EQUALS TWO. A theory of bootstrapping cannot rest peacefully on the polysemies of the English phrase "add one."

Strip away the ambiguities of the English phrase "add one," and neither Feigenson nor Carey, we believe, offers an account of a process that will move the child from noticing that PLACE THAT$_i$ NEAR THAT$_j$ CHANGES {THAT$_i$} INTO {THAT$_i$, THAT$_j$} to the conclusion that ONE ADD ONE EQUALS TWO. Of course, if the child could *already* think (grasp) ONE ADD ONE EQUALS TWO *in the arithmetic sense*, then the problem becomes highly tractable. Already grasping the arithmetic concept ADD ONE, the child could entertain and confirm the hypotheses regarding set relations and the meaning of the count words that Carey proposes. We could then imagine, in outline at least, a computational process underlying the "wild analogy" between physically placing objects together in a scene and integer addition. According to this account, at the moment the child entertains the hypothesis of an isomorphism (parallel or analogy) between placing objects together and arithmetic addition (or counting), the child already has the successor function available; indeed, this is what allows him or her to formulate hypotheses that mention ADDITION or ONE. But now there is no need for a bootstrap (cf. Rips, Asmuth, and Bloomfield, 2006, for related arguments). Absent this crucial assumption, however, accounts like Carey's and related ones appear to us to pull a conceptual rabbit out of a hat.

We conclude that there is no account on hand which shows how the young child can inductively construct integer representations *where none were available before*. We are skeptical as to whether there ever will be such an account. The reader interested in a critique of the bootstrapping hypothesis is urged to consult Rips, Asmuth, and Bloomfield (2006).

7 Computational Compatibility

It is tempting to think of sets of object files as analogous to sets of *hash marks*, that is, a counting notation in which the cardinality of a set is represented by the cardinality of

a set of lines or other marks. We suspect that this analogy in part motivated the Carey hypothesis. Sets of hash marks do refer to numerosities and, indeed, to the numerosity that is instantiated by the number of symbols (marks) in the set. Thus I refers to oneness, II to twoness, III to threeness, and so on. Moreover, sets of hash marks support some arithmetic operations—ordination, addition, subtraction, and, arguably, even multiplication—in an intuitively obvious and physically simple way. One constructs the symbol string that refers to the next larger numerosity simply by adding one more mark to the set that refers to a given numerosity. However, a set of hash marks refers to the numerosity of a set precisely because the individual marks, unlike object files, do not refer to particular objects in the set. Which mark was paired with which object in the construction of a set of hash marks is irrelevant once the set has been constructed, which is not in general the case when the set is composed of object files.

One could salvage a part of the Carey hypothesis by abandoning the assumption that it is sets of object files that have numerical content, and by simply assuming that there is a hash mark system for representing small numbers, an assumption that seems implicit in the hypothesis (cf., LeCorre and Carey, in press). This, however, separates the hypothesis from one of its principal empirical motivations. Because it is empirically well established that human adults can track only about four objects at any one time, the assumption that sets of object files could refer to the numerosity of the set of objects to which they pointed, explained the fact that the numbers between 1 and 4 appear somehow privileged in a variety of behavioral tests of human babies and monkeys. If we abandon the assumption that sets of objects files can somehow do double referential duty—both pointing at particular objects and referring to the numerosity of the set of objects pointed to—then we can no longer link the explanation of the privileged nature of numbers between 1 and 4 to the limits on the number of object files that can be open (pointing) at any one time.

There is a further problem with the hash mark hypothesis, which applies with equal force to the widely entertained hypothesis that numerosities between 1 and 4 are apprehended through perceptual *subitizing*, a process supposed to be analogous to the processes that form our percepts of things such as cows and trees (R. Gelman and Gallistel, 1978). In these models, oneness, twoness, threeness, and (perhaps) fourness generate discrete percepts (e.g., twoness looks like a line, threeness like a triangle, . . .). In such a model, the child learns that "one" is coreferential with its percept of oneness; "two," with its percept of twoness; and so on. The seemingly special status of the numbers between 1 and 4 arises because, by assumption, only these numerosities give rise to simple percepts. As with hash marks, these small number percepts are inherently discrete. Unlike hash marks, they have no inherent numerical content. That is, there is nothing in the percept of oneness that indicates it is a proper subset of the percept of twoness or that it can be added to the percept of twoness to get the percept of threeness. Just as there is nothing about "cowness" and "treeness" that renders them numerically ordered percepts, so there is nothing about any percept of "twoness" that dictates that "threeness" stands for one more than "twoness" (R. Gelman and Gallistel, 1978). Like the Arabic single-digit numerals 2 and 3, they are arbitrary discrete symbols for numerosities.

A problem with any hypothesis that posits a special discrete representation for the integers, a representation that is fundamentally different from and unrelated to the representation of continuous quantity, is the problem of *computational compatibility*, which we take to be a fundamental consideration in any model of the mind's representation of discrete numerical quantity. Whether potential symbols for number have inherent numerical properties or not, they are not in fact numerical symbols unless they enter into arithmetic processing (Gallistel, 2001; Gallistel and Gelman 2005; R. Gelman, 1990; R. Gelman and Gallistel, 1978; R. Gelman, 2006). Absent computational processes that exploit the subset structure of hash marks to draw conclusions about the numerical ordering of the sets whose numerosity is said to be represented by those sets of hash marks, the latter are not in fact numerical symbols. Similarly, absent a list that orders our putative percepts of oneness and twoness in accord with the numerical ordering of the sets that generate those percepts and support order inference about those sets, those percepts do not constitute a numerical representation; hence, they are not in fact numerical symbols. The same definitional consideration applies to putative representatives of continuous quantity: Mental magnitudes that are said to represent quantities such as length, weight, and duration, are not symbols for quantity unless they enter at least to some extent into arithmetic processing. We therefore state the following principle: The computational compatibility constraint on putative representatives (symbols) for *discrete* quantity is that they should be able to enter into the same arithmetic processes that operate on and produce symbols for *continuous* quantity.

As indicated earlier, the Gallistel and Gelman accumulator model is consistent with this requirement. Although different generative procedures serve the calculation of natural number and continuous number, both are stored as quantities. It would be problematic to have computationally incompatible symbols for discrete and continuous quantities. There are many occasions—such as the computation of rates by foraging animals and decision-making humans—where rates must be computed. The computation of a rate requires dividing a symbol that represents a discrete quantity (number) by a symbol that represents a continuous quantity (duration) to obtain a symbol that represents a different continuous quantity (rate). Moreover, the arithmetic processing of symbols for discrete quantities leads to symbols for what are (in effect) continuous quantities, namely, the proportions between numbers (1/2, 3/2, etc.).[6] Both adult humans and animals do represent rates (number per unit time; Leon and Gallistel, 1998; Gallistel, 2001), proportions between durations (Fetterman, Dreyfus, and Stubbs, 1993), and proportions between numbers (Balci and Gallistel, 2006; Meck, Church, and Gibbon, 1985). Humans and nonhuman animals recognize the equivalence between a proportion instantiated by two durations and the same proportion instantiated by two numbers (Balci and Gallistel, 2006; Meck, Church, and Gibbon, 1985). Thus, a central consideration for any proposal about how the mind represents discrete quantity (number) is that the proposed system *also* has symbols for continuous quantity (hence proportions), and that the symbols for discrete and continuous quantity are computationally compatible. It must be possible for the system to decide that the symbol for 7/3 represents a quantity (e.g., a rate) that is greater than the quantity represented by the symbol for 2. It must be possible to add the symbols for 7/3 and 2 to get the symbol for 13/3, and so on.

6. The symbols that refer to the proportions between integers are the symbols for rational numbers. From a purely formal standpoint, these do not form a system capable of fully representing a continuous variable such as length, because there are lengths (e.g., the length of the diagonal of the unit square or the length of the circumference of the unit circle) that cannot be so represented; they require symbols for irrational numbers ($\sqrt{2}$ and π, respectively). As a practical matter, however, most irrational proportions are uncomputable (they cannot be physically represented with perfect accuracy). They must be approximated by symbols for rational proportions, which can be done to whatever level of precision is required (short of perfect precision).

The problem of computational compatibility arises in a particularly acute form when it is suggested that the symbols for the small numerosities between 1 and 4 are discrete and noiseless, while the symbols for the large numerosities are continuous and noisy. This is equivalent to suggesting that there is a computer that represents the numbers 1 through 4 by bit patterns (00, 01, 10, 11) while representing larger numbers by voltage levels (an analog representation). How could such a device determine that $7 - 5 = 2$ (the difference between two noisy voltages) somehow becomes the bit pattern 10? How could it compute $5 + 2$ (the sum of a voltage and a bit pattern)? It is possible to add and subtract voltages or to add and subtract bit patterns, but it is not possible to subtract a bit pattern from a voltage. Bit patterns and voltages are computationally incompatible.

We view with skepticism any proposal that makes the preverbal representatives of discrete quantity—or of small discrete quantities—computationally incompatible with the preverbal representatives of larger discrete quantities and the representatives of continuous quantities. The proposal that noisy mental magnitudes are the symbols for continuous quantities of all kinds and for discrete quantity regardless of magnitude avoids the problem of computational incompatibility. However, as we have repeatedly noted, it does not account for human adults' unthinking and unshakable commitment to the principle that $i = i$, for all i, where i is an integer (that is, the symbol for a discrete quantity). So basic is this principle that in formal mathematics, it is taken to be true not just for all integers but also for all real numbers, for instance, $\pi = \pi$, an assumption for which it is hard to conceive of an empirical (inductive) basis.

8 Introducing Integers

We now consider a proposal that, in addition to accumulator magnitudes and object indexes, we should assume the existence of a third representational system which represents only integer values. We will propose that the main role of the accumulator in the development of human cognition is not as the ultimate source of the count number concepts, but instead as the principal mechanism for rapid tacit numerical calculation and estimation (Dehaene et al., 1999). For this purpose, the integer representation and the continuous magnitude representation have to be calibrated with one another. Both the accumulator magnitude representation and the natural number representation are innately specified.

8.1 *Next Number, Discrete Ordering, Exactness, and the Count Numbers*

Consider a common analog accumulator magnitude representation, the speedometer (figure 7.2a). The typical speedometer combines two distinct representations:

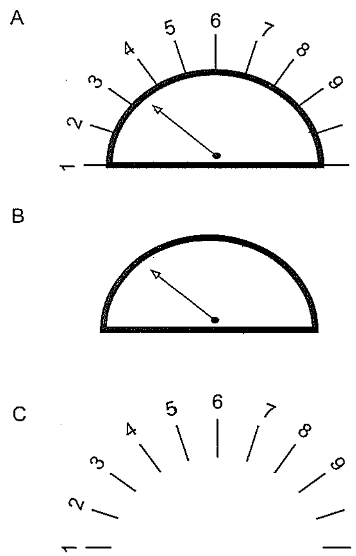

FIGURE 7.2 The typical speedometer (A) combines two distinct representations, a continuous analog representation "meter" (B) and a discrete digital representation "dial" (C). The two representations need to be carefully aligned and calibrated if the device is to be useful under normal circumstances.

one continuous and analogical (figure 7.2b), the other digital (figure 7.2c). These two representations are aligned and calibrated so that the position of the needle in the meter can be aligned with a digit on the dial and a digital reading can be taken. When these two representations are separated, as depicted in figure 7.2, it becomes clear that they have very different properties.

The "meter" returns continuous values and thus represents real numbers. The dial, by contrast, explicitly represents only discrete integer values. The dial lays out these values in a certain spatial arrangement in order that they align with and calibrate the behavior of the meter needle. But, aside from calibrating the "meter," this spatial property of the dial is inessential, as long as the ordinal structure remains. Absent the calibration, this structure could be represented quite abstractly (e.g., <1, 2, 3, . . . >), without any spatial structure.

Both the reals and the integers have ordinal structure, but only the integers support a well-defined notion of NEXT NUMBER. A well-defined NEXT NUMBER seems to be part and parcel of our basic number intuitions. When we count, we pick out a *first*, then a *second*, . . . , member of the set to be counted; intuitively it makes no sense to pick out a "second point one-th" member or that there is a place *between* second and third. Gallistel and Gelman (1992) point out that the accumulator model provides an effective procedure for picking out a next mental magnitude. This is the mental magnitude you get when you add "1" (the unit magnitude) to the mental magnitude in question. The resulting next magnitude, like all mental analog magnitudes, is,

from a psychological point of view, only stochastically differentiated from the other mental magnitudes that represent numerosity. Thus, any two samplings from the distribution of mental magnitudes for the next integer will never be exactly equal. However, it was argued that the mechanism for deciding whether one mental magnitude is greater than another should be assumed to have provision for giving "effectively equal" (~) as a result. Such a mechanism will return "~" just as often in the case of a comparison between the next integer and a separately generated mental magnitude (intended to refer to the same next numerosity) as it will in the case of two mental magnitudes generated by two counts of the same set. This brings us back to the question of exact equality and away from the issue of discrete ordering. An accumulator-continuous-magnitude counting mechanism with the assumption of an ~ operation can discretely order the magnitudes it generates.

However, there is more at stake in the notion NEXT NUMBER than simply discrete ordering and exactness. Recognizing what more is at stake, beyond order and exactness, is critical to understanding the nature and origins of our number concepts. Counting is not simply a matter of identifying *some* discrete value "minimally" greater than the current accumulator-counter magnitude or *some* discrete value "minimally" greater than a value "effectively equal" to the current accumulator value. *Some* value will just not cut it. The NEXT count value can be obtained only by adding the integer value1.

The accumulator account given by Gallistel and Gelman (1992) of course had to stipulate that the count value to be added is (effectively) equal to 1. But in a continuous magnitude representation this value is not only unobtainable with exactness, it is also ad hoc. Why should the "unit" magnitude in an accumulator count be ~ 1? Why couldn't it happen to be ~ 0.67, say, or ~ 1.134, or any other real value that would discretely order the magnitudes the accumulator generates? Such values would give you the NEXT stochastic magnitude, nicely ordered; but they wouldn't give you the *count* values, which "happen" to be integers. Moreover, no other value than exactly 1 will function as the identity element in multiplication. Support for arithmetic inference imposes heavy constraints on number representation.

This line of thought leads to the realization that the mapping from discrete quantities (numbers) to the mental magnitudes that represent them is constrained by formal considerations in a way that the mapping from continuous quantities to the mental magnitudes that represent them is not. This constraint, together with the necessity of computational compatibility, imposes a system of natural units on mental magnitudes. Suppose that we knew what the physical (neurobiological) implementation of mental magnitudes was, and could therefore measure mental magnitudes in physical units. For the sake of concreteness, suppose that mental magnitudes are realized by amounts of some intraneuronal substance, which we will call *numerin*. Thus, a particular mental magnitude would be physically realized by the synthesis of n picograms of numerin in some neuron. We could then ask what the constant k is relating the n to, for example, D, where D is duration measured in, say, seconds. There is, so far as we can see, no constraint on k other than that it be small enough so that the number of picograms of n required for even a very long duration could be comfortably contained in one neuron. We could determine k empirically only by manipulating D while measuring n.

This lack of constraint on k does not apply when we ask how much numerin represents the number (discrete quantity) 1. If quantities of numerin are really the physical realization of mental magnitudes, then they must enter into the arithmetic processes, including multiplication. Mental magnitudes are, by definition, those things in the brain that (a) are causally connected to the quantities they refer to and that (b) mediate arithmetic reasoning about those quantities. In that processing, there will be a unique quantity of numerin that corresponds to the multiplicative identity element. That is, there must be a quantity, n_1 that, when entered into one "side" of the multiplier (one functional slot in the multiplication process) together with any other quantity, n_2, entered into the other side, gives, as a result of the multiplication, the exact same quantity as was entered into the other side. That is, there must be a quantity, n_1, such that $n_1 \times n_2 = n_2$, for arbitrary n_2. Put another way, multiplication by any quantity less than n_1 will diminish the other quantity (for $n_L < n_1$, $n_L \times n_2 < n_2$), and multiplication by any quantity greater than n_1 will augment the other quantity (for $n_M > n_1$, $n_M \times n_2 > n_2$). Thus, to determine the k for the representation of discrete quantity, we do not have to manipulate or measure anything outside the brain itself. In particular, we do not have to manipulate N, the numerosity of a set represented by some amount of numerin. All we have to do is study the process that combines two amounts of numerin multiplicatively and determine the amount of numerin that functions as the identity amount in this process. That amount of numerin must be the amount that represents the numerosity of a set with only one member.

Moreover, knowing that amount would establish natural units for all of the brain's systems of mensuration—that is, the neurobiological mechanisms that causally connect objective quantities to mental magnitudes, thereby determining the amount of numerin that refers to a given amount of some objective quantity. Thus, the natural unit for the mental magnitudes (amounts of numerin) that represent, for example, duration would be the amount of numerin that functions as the multiplicative identity. However many seconds of duration that amount of numerin represented would be the mentally natural unit of physical duration, which we might, somewhat whimsically, call the *mentsec*.

In short, the intervals on the mental number line that correspond to successive increments in the counts that map discrete quantity (numerosity) to mental magnitudes are determined by the formal consideration that these intervals must be exactly equal to the interval that functions as the multiplicative identity. Otherwise, the whole system of arithmetic reasoning will not work.

8.2 *A Minimal Innate Basis for the Natural Numbers*

We have argued that basic number representation in humans is not limited to the reals but primitively includes the natural numbers. The natural numbers are exact values, the representation of which poses major difficulties for any system whose representations are inherently noisy, vague, or "fuzzy." The natural numbers are ordered values in which the notion NEXT NUMBER is well defined. Over and above all this, however, the natural numbers are not simply any sequence of well-ordered exact values, such as 0.67, 1.34, 2.01,...; they are integer values. Finally, children access integer values when they entertain hypotheses regarding the meaning of the count words.

What would the required learning mechanism look like? What properties should a mechanism have in order that it learn meanings for count words that designate integer values (as opposed to discrete reals, stochastic functions over reals, vague values, etc.), so as to order integer values according to a *next* relation, support arithmetic reasoning, and allow related magnitude estimation judgments? What properties should such a learning mechanism have in order that it can be guaranteed to complete its task in real time, that is, within the finite learning trial opportunities available to real learners? What is the *minimal* structure such a mechanism could have?

To reiterate, we take it as fundamental that numerical concepts must support a system of arithmetic reasoning. From this perspective, it might appear that the idea that the concept of any particular number may develop before the concept of other numbers of its class makes no sense. For example, one might argue that this makes no more sense than saying that the third pawn from the right is the first concept to develop in chess. However, whereas we agree that the concept ONE depends upon a system of arithmetic reasoning, and more generally that the meaning of mathematical concepts depends upon the formal system they are embedded in, it remains an open question what this means for the *psychological conditions for concept possession*. For example, how much of the system of arithmetic needs to be internalized for the system to possess the concept ONE? How high should one set the knowledge bar, and how does one motivate whatever bar setting one proposes? If mathematics is a closed deductive system, then ultimately every part is related logically to every other part. Does this mean that one must possess knowledge of *all of* mathematics to possess *any* of it, including (say) the concept ONE? If so, then no one possesses the concept ONE. We reject this precious view as prejudicial to the existence of an empirical science of concepts. However, if something less is required for concept possession, then what is that and what is the principle for determining what is required? If numbers are mind-independent properties of the world (objective properties of a formal system or of sets), and number concepts are mental symbols that refer to these properties, then what is minimally required for concept possession is a mental mechanism that can reliably lock the reference of a given number concept to its referent number. The general answer to these larger questions remains unclear. What we propose in the present case is a minimal mechanism that will generate the entire integer series and support arithmetic inference.

One part of our proposal is the distinction between a generative system and a realized system. A generative rule system specifies the derivation of an infinite set of symbols in the present context, a notation for denoting numbers. A realized system refers to those symbols that have actually been produced by running a derivation and storing the result in memory (for a longer or shorter period). To say that a system of arithmetic reasoning and ONE are mutually supportive is not to say that realized symbols for (infinitely) many numbers must exist in memory—nor, indeed, that any symbol in the series other than ONE has, as a practical matter, actually been derived and stored in memory by a given subject. It is only to say that there must be a rule of derivation and a procedure for creating those symbols (concepts) as they are needed.

Understanding cognitive development entails understanding *both* generativity and practical realization.

The accumulator is an example of such a generative procedure; it realizes mental magnitudes to represent real-valued numerosities, as needed.

We postulate, first, an integer generator that, like the accumulator, functions as a mechanism of domain-specific learning. Second, the integer generator has the property that the values it generates can be calibrated to accumulator values. Third, it allows an unbounded set of discrete values to be represented; it either provides or learns a *notational* system with an unbounded set of symbols. Fourth, it is constrained to represent only integer values. Fifth, it must guarantee an ordering of values under the NEXT relation. These basic requirements can be met by the following assumptions:

1. There is at least one innately given symbol with an integer value, namely, ONE = 1.
2. There is an innately given recursive rule $S(x) = x +$ ONE. The above two assumptions are similar to Peano's primitives, except we have 1 where Peano had 0; the rule S is also known as the *successor function* (e.g., Boolos and Jeffrey, 1989).
3. There is a regular grid that is commensurate with, and can be calibrated to, accumulator values.

The rule generates a grid alongside the accumulator magnitude representation as in figure 7.3. The grid calibrates integer symbols to noisy magnitudes, allowing the accumulator to be used in calculations and magnitude estimates whose results can be rounded to integer values.

The grid itself could conceivably provide "detachable" symbolic objects to represent integer values. If they could be detached from the accumulator and used outside of the module, general thought processes could access these symbols. For example, the symbol | would represent 1, | | would represent 2, | | | would represent 3, and so on. However, this kind of notation has a property that severely limits its usefulness. As the *n* to be represented grows in size, the "physical size" or length of the symbols themselves grows linearly with *n*. It's as if the word for elephant had to be thousands of times bigger than the word for bacteria, not a welcome property. In fact, the accumulator magnitude representation itself has this same unwelcome property, and in both cases there is a problem of how an unbounded or even a large bounded set can be represented. This suggests that in addition to the grid there should be a *compact notation*. A compact notation is one which provides symbols whose length does not grow with the size of the integer represented.

To provide a compact notation, an unbounded set of symbols must be generated, with a one-to-one correspondence between symbols and integer values, so that each symbol functions as a unique identifier for some unique integer value. Each symbol is bound to a unique rung on the grid, the ordinal position of which determines the meaning of the symbol (that is, its specific integer value). At the same time, a given symbol is also provided with an interpretation (calibrated) in terms of an (approximate) continuous accumulator magnitude. Finally, these compact symbols can be "detached" and can compose in centrally constructed and centrally processed thoughts, namely, in thoughts involving integer values.

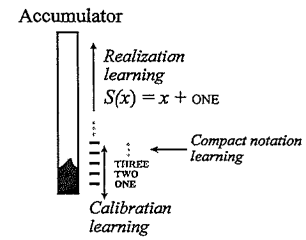

FIGURE 7.3 A model for number learning and representation, combining continuous magnitude and integer representations with three kinds of learning: that in which integer values are recursively realized by computing the function S (realization learning); that calibrated against continuous magnitudes by stretching or compressing the length of the grid relative to the accumulator magnitude (calibration learning); and that mapped to a compact notation (compact notation learning). A compact notation can be provided through learning a natural language with count words. Accumulator magnitudes are depicted as noisy. At least three variants of this model are possible: that in which only the symbol ONE is innate, or ONE and TWO, or ONE, TWO, and THREE are innate.

8.3 *Where Does the Compact Integer Notation Come From?*

Minimally, the concept of and symbol for 1 must be innately realized because the recursive rule S that generates the grid requires that concept and symbol. S also requires that the addition operation +, the identity relation =, algebraic variables, and a recursive capacity (minimal algebraic notions) also be innately realized.

Many variants on this proposal are possible; for example, variant 1: the innate integer notation also includes a realized symbol TWO (= 2); variant 2: the innate integer notation also includes a realized symbol for THREE (= 3); and so on. However, given that the set of natural numbers is unbounded, not all of them can be represented by realized symbols, and thus not all can be innately realized. Nonetheless, the entire set can be represented in the sense that it is generated recursively by S. One way to think about S is that it generates the *meanings* for the entire set of integers, using finite means. Because the means for generating the set are finite, namely S, it can be innately realized. However, for realizing an unbounded or even large bounded set of symbols, where each symbol uniquely carries an individual integer meaning, a notation is required whose symbol length does *not* grow monotonically with the magnitude of the value represented (as a grid or an accumulator representation does). Notably, the count word list in a natural language such as English has a notational system for integers with just this property: For example, English uses just two words to represent 1,000,000.

Conceivably, the brain may have an innate compact integer notation, for example, produced recursively or by a cascading notation for orders of magnitude. Alternatively, the notation for values larger than ONE may simply co-opt natural language itself and acquire *that* compact notational system. In this case, the detachable unique identifiers designating integer values larger than 1 will be drawn from a learned notation, namely, the natural language expressions of the learner's first language.

Notice that this proposal has nothing whatsoever to do with acquiring integer *concepts* by learning a count number word list. Under this proposal, natural language provides only a (compact) notation for prior existing integer concepts. Indeed, it would be impossible to learn what this lexical subsystem encodes without having the corresponding integer concepts available, since these concepts constitute the required hypothesis space for the learning process. The calibrated and ordered grid and the specific rung on the grid are realized internally by running the recursive function *S*. The ordinal position of a given rung in the grid fixes the meaning of the unique identifier bound to it. Whereas the meanings themselves do not have to be learned, the notation for the unique identifiers (except for ONE; and for TWO in variant 1; and for THREE in variant 2) is learned.

The integer grid, together with the innate unique identifier symbol, forms the language learner's hypothesis space for mapping sound forms onto meanings (i.e., integer values). An open empirical question in this account concerns the role of adult-demonstrated counting rituals. Does such a ritual itself provide the occasions on which a child will call the function *S*? If so, the adult will teach the child to count by actuating the child to realize the next step or steps in the integer grid. If a child does not belong to a language community that has count words, does the grid remain unrealized in such a child? Or will a preverbal infant routinely call function *S* in situations not involving verbal counting, for example, in spontaneously tracking the numerosity of sets of physical objects? Does verbal counting provide the only actuating circumstances or only one of a number of actuating circumstances? These questions remain entirely open.

In summary, the basis of our natural number concepts is hypothesized to be the innate representation *S* that recursively defines the positive integers and the concept *next*. The basis of these concepts cannot be a system of continuous magnitude representation, accumulator or connectionist, noisy or not, without a system that can represent exactly the value 1. Moreover, the integer representation becomes calibrated to accumulator magnitudes, allowing integer calculation and magnitude estimation. The brain may generate its own compact code for representing integer values and then learn the appropriate mapping from that internal compact code to the corresponding compact code in natural language. Alternatively, the brain may simply co-opt the compact code of a natural language. This latter account would afford an important role to language learning without embracing Whorfian claims. What natural language cannot do is determine or teach de novo the meanings of integer concepts. These meanings are known in an important sense innately: namely, as generated by *S* (and perhaps even calibrated against the magnitude representation).

8.4 *Another Proposal*

An alternative to positing a separate notation for the integers (albeit a notation calibrated to the mental number line, the system of mental magnitudes) is to assume that there is (in addition to the accumulator mechanism for generating magnitudes that refer to numerosities and enter into arithmetic operations) some innate algebraic principles that mediate or govern reasoning about discrete numerosities (cf. R. Gelman and Gallistel, 1978, pp. 227ff). This additional symbolic system has symbols that do not enter into arithmetic operations that determine numerical values. Unlike the mental magnitudes, they are not used for arithmetic computation. Rather, they are used to draw conclusions about the outcomes of computations, by licensing symbolic substitutions. In this view, the essential function of testing for equivalence is to license substituting one course of action for another, whether overt action or symbolic action.

As we have already noted, on the assumption that the arithmetic operations operate on noisy magnitudes, it is difficult to specify a mechanism that would license the conclusion that two such magnitudes are equal, and hence substitutable, one for the other. It is not hard to specify a mechanism that decides whether one magnitude is less than or greater than another. The sequential sampling diffusion model first suggested by Buckley and Gilman (1974; see also Gallistel and Gelman, 2005), which has also been proposed (and extensively tested) as a model for making nonnumerical decisions (Ratcliff and Smith, 2004), gives us a plausible model for making decisions about greater than or less than. Imagine two magnitudes that are to be compared as two speedometers, A and B, with jittering needles. The comparing mechanism takes readings, a_1 and b_1, from both speedometers; computes the difference, $a_1 - b_1$; tests whether it exceeds either a positive or a negative threshold. If it exceeds the positive threshold, the mechanism decides that $A > B$; if it exceeds the negative threshold, it decides that $A < B$; if it exceeds neither threshold, the mechanism takes another two samples, computes their difference, adds it to the previous difference, and tests whether the sum of the two differences exceeds either threshold. That is the essence of a sequential sampling decision mechanism. Proposed mechanisms of this kind include, for obvious reasons, a time limit on the sampling, at the end of which, if neither threshold has been crossed, the mechanism reports that it cannot decide. Such a report cannot, however, be taken as a decision that the two magnitudes being compared are equal, because it does not guarantee a fundamental property of the equals relation, namely, that when equals are added to equals, the results are equal. Such a mechanism is perfectly capable of reporting that it cannot decide whether $A <> B$ nor $C <> D$ and then reporting that $A + C > B + D$.

In short, it is not clear how to specify an effective procedure for determining whether two noisy magnitudes (two noisy, real-valued variables) are equal. Hints of the difficulty will be familiar to those who have run computer simulations in which integer values have been computed from floating point values and then compared. It sometimes happens that the computer decides that $1 \neq 1$, because it internally represents one of the two instances of 1 as .9999999999999999. When mental magnitudes represent estimates of continuous quantities such as duration, the inability to determine equality is arguably a feature, not a bug. The values of continuous empirical quantities cannot be known with perfect precision; therefore, the question of whether two such quantities are exactly equal is moot. But when noisy continuous quantities are used to represent numerosities, this inability is clearly a bug, at least on the realist assumption that number, as commonly conceived, is an objective property of sets, and that two sets can have exactly the same number of members; or indeed, on the assumption that $\pi = \pi$.

One solution is to posit an additional mechanism that mediates algebraic reasoning about discrete quantity. For drawing inferences about discrete quantities, there may be innate mechanisms that in effect define the relationship of exact equality and mediate reasoning in which it plays a central role. On this hypothesis, the magnitude symbols might be supplemented by an adjunctive symbol system with two symbol categories, I and n. A symbol in the category I refers to the magnitude of a cupful in the Meck and Church accumulator machine. These symbols are generated as needed and discarded as soon as they have been used. A symbol in the category n refers to any magnitude generated by using the accumulator in the count mode. Symbols in this category are also generated as needed. Neither the I symbols nor the n symbols are mental magnitudes: The relation between them and the mental magnitudes is the same as the relation between the letters in algebraic strings (x, y, k, i, etc.) and numbers; the letters refer to arbitrary instances of the numbers, but are not themselves numbers; the numerical value they refer to is left unspecified. We further postulate the existence of rewrite mechanisms or substitution licensing mechanisms operating on these two categories of symbols in accord with the principles that for arbitrary distinct instances a and b of I, $a = b$ (all instances of I refer to equal magnitudes, that is, interchangeable magnitudes), and for arbitrary instances a, b, c, and d of n, if $a = b$ and $c = d$, then $a \circ c = b \circ d$, where $\circ$; refers to any one of the arithmetic operations $+$, $-$, $\times$, and $\div$. In other words, all instances of one are interchangeable, and whenever interchangeables combine arithmetically with interchangeables, the results are interchangeable. One may also need to assume an Archimedean principle to the effect that for any instance, a, of n, $aI = a$. In words, any (natural) number may be generated by pouring cups that number of times.

In order to explain where the integers come from, this proposal blatantly posits innate mechanisms of deduction that embody defining principles of the integers. In doing so, it avoids the problem of computational incompatibility, because the symbols on which these deductive mechanisms operate are not the symbols that enter into arithmetic computations. Rather, they are symbols that enter into reasoning about the outcomes of arithmetic computations on magnitudes generated by the accumulator when operating in the counting mode. In positing these principles, we explain why the child can assume that "one," "two," and so on refer to specific mental magnitudes generated by the accumulator, and further believe that the property of a set thereby referred to may satisfy an equivalence relation.

Our solving the problem with these assumptions will remind many readers of the maxim that postulation has the advantages of theft over honest toil. In order to explain where our concept of the integers comes from, we have assumed that it is built into an innate mechanism for reasoning in the abstract about the outcomes of numerical manipulations of sets. It would be nice to be able to motivate these assumptions by considerations other than those that pose the puzzle we are trying to solve. We confess that we cannot at this time do this, which is why we put this particular proposal forward in a tentative voice. Still, it is well to keep in mind the fact that three- and four-year-old children have little difficulty switching between an approximate and an exact system, the latter being preferred when the task is an arithmetic one (Zur and Gelman, 2004).

9 Conclusions

The lengths to which we find ourselves driven serve, if for nothing else, to highlight the central features of the problem: (1) There is abundant evidence for, and considerable theoretical consensus, that discrete and continuous quantities are represented preverbally by computationally compatible symbols. (2) These symbols enter into at least some of the operations that define the system of arithmetic: namely, ordering, addition, subtraction, multiplication, and division, which is why they can be said to be numerical symbols. (3) These symbols obey Weber's law in that the confusability of two symbols or the uncertainty regarding the value to which they refer is proportional to that value. (4) This is widely assumed to imply that these symbols are analogous to noisy magnitudes. (5) The symbolic size and distance effects are generally taken to imply that judgments about the ordering of the referents of arbitrary culturally determined symbols for quantity such as "1," "2," "3," "4," and the like are mediated by order-deciding operations on the preverbal mental magnitudes that represent the quantities referred to by these symbols. (6) This implies that in learning the meaning of these symbols, verbally competent subjects take them to refer to the same properties that are referred to by the symbols in their preverbal mental magnitude system. (7) This explains why verbally competent subjects understand these arbitrary symbols to refer to properties that can be arithmetically processed. (8) However, it does not explain why subjects believe that exact equality (interchangeability) is a potentially applicable property of the quantities that these symbols refer to. In the case of continuous quantity, it is doubtful that they should believe this, and perhaps they do not. But it seems beyond argument that most adult humans believe that the positive integers (the natural numbers) represent a property of sets such that it can satisfy an equivalence relation: The numerosity of two different sets or of the same set at different times may be interchangeable; any symbol that exactly refers to the numerosity of one of the sets refers just as exactly to the numerosity of the other. (9) With discrete symbols, the determination of exact equality reduces to the determination of the identity of the symbols. (10) With noisy magnitude symbols, the determination of exact equality is much more problematic. Thus, the hypothesis that arbitrary culturally determined symbols for discrete quantity acquire their meaning from the assumption that they are coreferential with the noisy preverbal mental magnitudes that refer to those quantities fails to explain why adult humans believe in the potential exact equality of the magnitudes referred to. (11) Positing a fundamentally different discrete symbolic system that represents discrete quantity or small discrete quantities raises the problem of computational compatibility. (12) Bootstrapping models that attempt to use language to somehow create concepts that do not exist in the preverbal system for representing quantity seem always to beg the question, tacitly assuming that the concept already exists in the process of explaining how it is created (Rips, Asmuth, and Bloomfield, 2006).

In short, there does not appear to be any way to derive the integers from non-integers (reals) or from non-numerical symbols (object files, linguistic quantifiers). We have integer concepts either because there is a mental notation specific to the integers, but calibrated to the corresponding mental magnitudes, or because there is a system of algebraic reasoning about operations on discrete quantity, a system that allows the deduction of relations between computations without requiring that those computations in fact be carried out. Either way, it is hard to resist the conclusion that the generative concept of an integer is innate.

9.1 *In the End, One Proposal?*

Finally, Hartnett and Gelman (1998) report that children aged five to seven years found it surprisingly easy to articulate the idea there is a never-ending sequence of unique natural numbers, that every natural number has a successor. We say "surprisingly easy" for two quite different reasons; first, because children found this idea about infinity easier to grasp than the concept of a fraction; and second, because it is unclear what the principle of induction is that would yield this conclusion—except, of course the principle of *mathematical* induction itself, which would have to be taken as innate. Instead, the intuition of a discrete infinity surely is an intuition about the structure of the successor function itself—an intuition of its integer-closed recursion. This is evidence, then, for the psychological reality of the successor function; evidence for a little piece of intuitive algebra. When viewed this way, we can see that the two accounts we have outlined here are really one.