

Core Knowledge and Conceptual Change

David Barner and Andrew Scott Baron

Print publication date: 2016

Print ISBN-13: 9780190467630

Published to Oxford Scholarship Online: August 2016

DOI: 10.1093/acprof:oso/9780190467630.001.0001

Epistemic Limitations and Precise Estimates in Analog Magnitude Representation

Justin Halberda

DOI:10.1093/acprof:oso/9780190467630.003.0010

Abstract and Keywords

This chapter presents a re-understanding of the contents of our analog magnitude representations (e.g., approximate duration, distance, number). The approximate number system (ANS) is considered, which supports numerical representations that are widely described as fuzzy, noisy, and limited in their representational power. The contention is made that these characterizations are largely based on misunderstandings—that what has been called “noise” and “fuzziness” is actually an important epistemic signal of confidence in one’s estimate of the value. Rather than the ANS having noisy or fuzzy numerical content, it is suggested that the ANS has exquisitely precise numerical content that is subject to epistemic limitations. Similar considerations will arise for other analog representations. The chapter discusses how this new understanding of ANS representations recasts the learnability problem for number and the conceptual changes that children must accomplish in the number domain.

Keywords: approximate number system, conceptual change, analog magnitude representations, confidence, representational power

Introduction

A focus on conceptual change has the benefit that it motivates us to become clearer about our models of the starting states for learning, the later states for cognition, and the learning mechanisms (and input) that may be required for conceptual change to occur. In this chapter I will discuss a family of starting states—representations of analog magnitudes such as the approximate number system. I will motivate an adjustment to our understanding of these

representations, and I will briefly explore how this adjustment may change the theories we build for conceptual changes involving analog magnitudes.

Throughout, I will be concerned with magnitude representations, which include representations of approximate duration (e.g., what does 1.5 seconds feel like?), approximate distance (e.g., how far is it between me and the wall?), and approximate number (e.g., around how many people were at the party?), among many others (e.g., volume, area, brightness, loudness—e.g., Dehaene ... Brannon, 2011; Feigenson, 2007; Odic et al., 2015). The mental representations of these dimensions share a similar format and they give rise to similar behavioral signatures (e.g., Weber's law) as indexed by psychological experiments. I will describe a model for these representations, discuss a confusion we may fall prey to when theorizing about them, and present a positive proposal for replacing these notions with the notion of epistemic limitations.

(p.172) There is a long history of studying magnitude representations, and recent decades have seen an explosion of interest across multiple approaches, including animal behavior, studies of infant cognition, brain imaging, neuropsychology, and single-unit recording (Dehaene ... Brannon, 2011). Some shared beliefs about magnitudes seem to inform the work in these various disciplines. While variation exists, a common notion is that magnitude representations (e.g., representations of nonsymbolic number) are approximate or “noisy,” and this noise is a limiting factor for their expressive power. For example, numerical representations are described as subserved by, “underlying representations (that) are inherently continuous and therefore noisy and variable” (Leslie, Gelman, ... Gallistel, 2008a), and that “each numerosity is represented internally by a noisy distribution of activation on an internal continuum or mental number line” (Dehaene, 2007). Indeed, the “noisiness” of representations of number, space, and time is often taken to be their hallmark signature. In what follows I suggest a critical adjustment to this understanding. I believe that what I will argue applies to all analog magnitude representations. Throughout, I will focus my examples primarily on representations of numerical magnitude stemming from the approximate number system (ANS), but I mean for my points to apply broadly.

A Standard Model of Analog Magnitude Representations

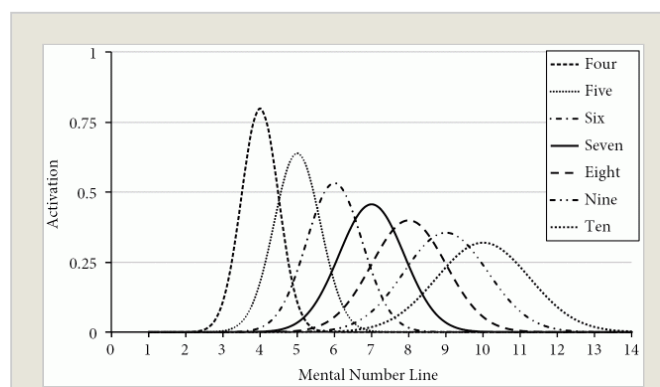
Claims about noisy, approximate, or fuzzy representations can be easier to interpret with the aid of a graphical depiction of the analog representations. Figure 10.1 offers a fairly typical depiction of the representations of the ANS. The mental number line along the bottom is often described as representing an array of number sensitive neurons, each with a preferential tuning function (e.g., Nieder ... Miller, 2003). The width of these tuning functions increases with the numerical target, with the result being a series of bell-shaped curves that linearly increase in standard deviation with increasing signal (e.g., reading from

left to right in Figure 10.1). Often each bell-shaped curve is understood to indicate a representation of a particular numerosity within the ANS. It would not be uncommon for authors to make statements such as “the curve representing 4 is sharper than the curve representing 8.” Here, a “sharper” curve is typically understood to be a more accurate representation; similarly, the representation for, say, 8 is typically taken to contain more error than the representation for, say, 4. I will argue that these notions do capture something important about analog magnitude representations, but that they also embrace a crucial misunderstanding of these representations.

The constant increase in the width of the bell-shaped activations as the to-be-represented numerosity increases is an important, law-like aspect of these representations. As one moves from the left to the right along the x-axis in Figure 10.1, the numerical value of the signal is increasing. As the bell-shaped representations take on larger and larger numerical values (e.g., as determined by the mean of their activations), the standard deviation (SD) of these curves increases (e.g., with wider spread for larger numbers). **(p.173)** This constant increase in SD provides an account for the major performance signature in behaviors that rely on the ANS—for example, the discriminability of any two numbers in the ANS is a function of their ratio. Graphically, this can be seen as the amount of overlap between two curves. Intuitively, where two curves overlap on the mental number line, their representations will tend to be confused by the observer trying to distinguish them. The more overlap there is among the curves, the more error there will be in judgments—for instance, when attempting to decide which of two collections is greater in number, the observer will increasingly struggle as the numbers become closer. For example, because there is less overlap between the ANS representations activated for five vs. six dots than eight vs. nine dots, performance will be more reliable with the former comparison. Analogous notions of bell-shaped representations underlie theorizing about all other analog magnitudes as well; for example, the spread of the curves representing durations increase as the times become longer (Odic, Im, et al., 2015).

One feature of the psychophysical model sketched graphically here is that the rate of increase in SD as number increases can be described by a single parameter (the Weber fraction) that determines the SD for each internal number representation (Halberda ...

Odic, 2014). This model also allows us to describe individual differences in performance: Individuals differ in their accuracy at a numerical discrimination task to the extent that they have



different Weber fractions and thereby different rates of increasing spread in their ANS representations, with more spread leading to more errors.

Figure 10.1 Depiction of the representations of the approximate number system (ANS).

Thus, we can understand individual differences in performance in terms of the Weber fraction. If the Weber fraction is higher, it means that the confusability of the number representations increases more rapidly as number increases (e.g., larger SDs and wider curves). If the Weber fraction is lower, it means that the confusability of the number representations increases more slowly as number increases (e.g., smaller SDs and narrower curves). Different animals (Brannon ... Roitman, 2003), and humans of different ages or abilities (**p.174**) (Halberda ... Feigenson, 2008; Piazza et al., 2010) exhibit different Weber fractions, but all participants show performance consistent with the lawful increase in SD—for example, all observers have representations similar to those depicted in Figure 10.1, just with wider or narrower bell-shaped curves.

Spread in Analog Magnitude Representations Does Not Indicate Error

How are we to understand the curved representations in Figure 10.1? What does the spread of the curves actually indicate? One understanding is that the curves depict the amount of error in each representation. For example, one might suggest that sometimes the observer's representation of ANS_8 (e.g., a representation generated by the ANS in response to a stimulus presentation of eight items) gets it correct and responds "8," but a significant amount of the time gets it a little wrong and says "7," and more rarely it gets it really wrong and says "5."

This understanding—curves as errors in the representation—can also be applied to performance in a numerical discrimination task. For example, one might expect that the observer in a numerical discrimination task will at some point be incapable of discriminating between two close numerosities. Unreliable discrimination performance is often understood as arising because the error in the two analog magnitude representations being compared makes them indistinguishable. For instance, it is not uncommon to read statements such as "because analog magnitude representations are inexact and subject to Weber fraction considerations, they fail to capture small numerical differences between large numbers" (Carey, 2009, p. 294), and "the difference between eight and nine is not experienced at all, since eight and nine, like any higher successive numerical values, cannot be discriminated" (Carey, 2009, p. 295). These appear to be claims about the representational capacities of the ANS and the thoughts and experiences it supports, not simply statements about an animal's performance in some task.

This view can also be seen in the assumption that performance in a numerical discrimination task should drop to chance (i.e., should result in an inability to discriminate) when the ratio between numbers makes them indistinguishable to the observer. For instance, it is not uncommon to read statements such as “when the ratio of two magnitudes exceeds that allowed by the subject’s Weber constant, the quantities become indiscriminable for the subject” (Beck, 2012), or “two sets can be discriminated only if they differ by a given numerical ratio, according to Weber’s law” (Piazza, 2010). Here, too, authors appear to suggest that this limitation is fundamental to the representations themselves.

The suggestion I make here, and in other work, is that such beliefs likely reflect a common misunderstanding of the psychophysical law (Halberda ... Odic, 2014). In fact, the psychophysical law predicts that the observer *will be above chance for all comparisons*, even the most difficult, barring performance or attention limitations. Indeed, even the unfortunate term “just noticeable difference,” often used to refer to the smallest difference (**p.175**) between two stimuli that an observer can reliably detect, does not actually denote a just noticeable difference—because observers’ actual performance will be above chance with discriminations that are both above and below their JND (Halberda ... Odic, 2014).

Instead, I suggest that the Weber fraction is better understood to be a kind of internal scaling factor that reflects the rate of increase in spread of the representations along a particular psychological dimension (e.g., the mental number line, area line, or duration line). The Weber fraction determines the SD for every possible representation of that dimension, not simply performance at some particular ratio (Halberda ... Odic, 2014). Because the representations along the mental dimensions of analog magnitudes are well-ordered and any two representations have some degree of non-overlap, the psychophysical model predicts that all numbers can, in principle, be distinguished (even through, in practice, with very large numbers it might take very many trials for an observer’s performance to yield enough data to show this).

To connect the current discussion back to the broader goals of understanding starting states and conceptual change, I note that these beliefs about error in ANS representations have had some impact on the theories we have built. For example, Jake Beck has reasoned from a supposed inability to discriminate close numbers in an ANS task to a hypothesis that analog magnitude representations have nonconceptual content (as opposed to conceptual content) and that they violate systematicity (Beck, 2012).¹ Susan Carey has constructed several fully worked-out proposals for how children might move from ANS representations, in combination with other primitive non-integer number concepts, to build a new conceptual understanding of the integers (Carey, 2009). These proposals, and others in the literature, assume that ANS representations do not distinguish between close numbers. This is one corollary of the broader view that the

“noisiness” of analog magnitude representations amounts to “fuzzy” or “error-prone” representations. Here I suggest that this understanding is mistaken, and offer a positive proposal for how to think about these representations.

Evidence from Human and Animal Success with Close Comparisons

In considering the suggestion that ANS representations fail to distinguish between close numbers, we should first examine the empirical evidence. This examination reveals that, typically, no such failure occurs. Figure 10.2 shows two performance patterns and two theoretical curves. The first dataset (Figure 10.2a) shows the average performance from over 10,000 people between the ages of 10 and 85 who participated in a numerical discrimination task (i.e., deciding which of two briefly flashed collections of dots was greater in number) (Halberda, Ly, Wilmer, Naiman, ... Germine, 2012). Throughout the task, the numbers of items in the two collections, and their ratio, varied so that some trials were easier and some were harder. This allows us to generate estimates of how well the observers performed across four different numerical ratios. In Figure 10.2a we see that **(p.176)** human performance (circles) does not drop to chance (50%) as the ratio becomes harder. Rather, performance shows a smooth change from better on easier trials (e.g., a ratio of 2, where there might be 20 yellow and 10 blue dots) to less good on harder trials (e.g., a ratio of 1.143, where there might be 7 yellow and 8 blue dots). Every one of the more than 10,000 participants showed performance consistent with this kind of smooth gradual change. Critically, this smooth gradual change is exactly what is predicted by Weber’s law—not an abrupt change from “above chance” to “at chance” performance (Halberda ... Odic, 2014). In many other studies across multiple labs, human participants have been tested on a wide range of numerical ratios (including ratios more difficult than those displayed in Figure 10.2a). Across these studies, it is the exception to see performance dropping down to chance at some point before a ratio of 1. Instead, we almost always see the smooth and gradual change in performance seen in Figure 10.2.

The curve in Figure 10.2a is the predicted performance based on the model of the ANS shown in Figure 10.1. It shows smooth and gradual change in performance from a ratio of 1 (where the two collections have the same number of items, the ratio is 1, and there is no correct answer) toward better and better performance as ratio becomes easier. To connect this curve of percent correct to the bell-shaped curved representations of Figure 10.1, notice that the curves in




Figure 10.2 Two performance patterns and two theoretical curves. **a:** Plotting of average performance from over 10,000 people between the ages of 10 and 85 who participated in a numerical discrimination task. **b:** Replotting of a dataset from a single pigeon who was a

Figure 10.1 have no abrupt change in their widths. There is a smooth and gradual change in SD with the increasing number signal. The mean of each curve shifts to the right, and the movement right is proportional to change in SD. This is Weber's law, and this smooth change in SD and mean is why the predicted performance in Figure 10.2a is smooth with no abrupt change in percent correct. Weber's law is a law about the smooth change in SD with increasing signal, and not a law about a "just noticeable difference" or at-chance performance at some ratio. Indeed, if participants dropped to chance before reaching a ratio of 1 (e.g., as might happen at more difficult ratios than those shown in Figure 10.2), or if participants showed some abrupt change in percent correct (e.g., from good to bad performance), than these would be *violations* of Weber's law and we would require some other mathematical model to describe such performance. **(p.177)** Note that such violations might indeed occur, as might happen with performance limitations (e.g., giving up on the task when it becomes too tiresome). If such violations occurred due to performance factors, then we may still say that Weber's law holds. Still, the data from the more than 10,000 participants in Figure 10.2a demonstrate that such violations are not the typical case.²

participant in one of the earliest papers on the ANS and numerical discrimination.

The second dataset, in Figure 10.2b, is a replotting of a dataset from a single pigeon who was a participant in one of the earliest papers on the ANS and numerical discrimination (Rilling ... McDiarmid, 1965). I include this plot to show the performance of a nonhuman animal, and to demonstrate that a belief in error-prone representations, and "at chance" performance, is also misplaced for animal data. In Figure 10.2b we see that the pigeon's choices show a smooth change with decreasing numerical ratio (never showing an abrupt change and never dropping to chance performance), and we can also see the smoothness of the predicted performance as fit by the model of Weber's law. The smoothness of this function, as it approaches the intersection with the axes at ratio = 1, reveals that the pigeon's numerical discrimination behavior is consistent with the psychophysical model that maintains that consecutive higher numbers are well-ordered all the way up and down (e.g., the pigeon's behavior is consistent with the assertion that it even knows that $47 < 50$; Rilling ... McDiarmid, 1965). The pigeon's performance suggests that it is attempting to make the correct numerical discrimination even on the most difficult trials. Note that even a numerical difference that the pigeon struggles to distinguish ($\approx 70\%$ correct at 45 vs. 50, ratio 1.11, Figure 10.2b) is nonetheless well-ordered in its competence. The evidence for this is that its smoothly decreasing performance continues to fall along the psychophysical curve; if the pigeon's ANS representations were not well-ordered, we would see a deviation from Weber's law at these harder ratios Figure 10.2b.

I believe that even human infants' performance is best understood to be consistent with Weber's law—that infants' discrimination abilities are a function of ratio, and that they have the competence to succeed with all comparisons, even the most difficult ratios, barring attention or performance limitations. Melissa Libertus and colleagues (e.g., Libertus ... Brannon, 2009) have begun to identify smooth, ratio-dependent looking-time performance in infants that is consistent with the smooth curves depicted in Figure 10.2. Showing that infants succeed at discriminating between numerical quantities that have previously been described as being "indiscriminable" at certain ages may require changes in numbers of trials and/or experimental design, but it is, I predict, quite possible.³

A Performance/Competence Distinction for Analog Magnitude Discrimination Behavior

The distinction between understanding participant behavior as involving performance errors rather than competence failures may be crucial for our theorizing about conceptual development involving analog magnitude representations. As learners, the hypotheses we **(p.178)** can form depend primarily on our competence; testing hypotheses through observation and gathering evidence depends on both our competence and our performance. For example, if human and animal behavior in numerical discrimination tasks is best understood to engage a system of well-ordered representations (the use of which is sometimes affected by performance failures), then the thought "I bet that 9 is bigger than 8" is thinkable by a creature which, upon seeing two collections of items, is only slightly better than chance at deciding which collection has more. That is, although performance errors involving the ANS may spell trouble for one's behavior when discriminating close numbers, these performance errors do not block precise thoughts about such numbers.⁴

As an analogy, consider our intuitive understanding of biology that supports thoughts such as "that cow is alive" and "that boulder is not alive." It can happen that our ability to distinguish a cow from a boulder can be disrupted by external noise. Such events happen on dark and foggy nights—one can imagine walking through a foggy field illuminated by just a sliver of moon and seeing a large, somewhat amorphous dark shape about 15 paces ahead in the field. We might not be able to discriminate much better than chance whether this object is a boulder or a sleepy cow. But this performance problem gives rise to no problems for our intuitive theories. This event doesn't lead us to a positive belief that *cows and boulders are the same*. One reason that it does not cause a problem for our theories about cows and boulders is that, standing in that field, we very much know that we can't tell whether what stands before us is a cow or a boulder. We have a clear sense of our own epistemic limitations. I believe that there is also a clear signal of epistemic limitation that plays a crucial role in how we use and understand our own analog magnitude thoughts, and I will try to

make the case here that epistemic limitation should replace our notions of error in these representations.

Analog Magnitude Representations Are Epistemically Limited Estimates with Precise Content

As a first step toward rethinking our views of magnitude representations, we might ask: Considering the curves in Figure 10.1, depicting the representations of the ANS—are these “noisy” curves spread out within the number space? Is the representation of ANS₉ more spread out, more diffuse, in the number space than the representation of ANS₅? Or, more generally, should we infer from the graded nature of the ANS activation curves in Figure 10.1 that our representation of number itself is graded? Is our approximate number content “fuzzy,” “noisy,” or somehow imprecise? I will answer “no” to all of these questions and instead suggest that our analog magnitude representations are exquisitely precise, but subject to epistemic limitations.⁵

To make my case, I will present a thought experiment from the spatial domain. Data demonstrating “fuzzy” representations of magnitude and position are well developed in the space literature (compared to more recent and smaller literatures on approximate **(p.179)** number, area, duration, etc.). And, space is a helpful choice because humans can easily manipulate our own position in space (e.g., by walking) and can manipulate our own sensory input (e.g., by closing our eyes). Our sense of distance and movement in space—being one of our most visceral daily experiences—gives rise to some of our clearest intuitions. I believe that the points I will make here should apply to all analog magnitude representations (e.g., space, time, number, brightness, loudness, felt weight).

Consider you are standing in a room, in a specific position with an unobstructed path to one of the walls. (You can imagine this situation, but I think this example works much better if you actually physically do this at some point; I urge the reader to try this activity to feel the experiences I describe here). Position yourself facing the somewhat distant wall (around 9–15 feet away is a good distance for this example) with your eyes open and looking at the unobstructed path you have to the wall. Many experiments, such as recording from neurons in humans who are preparing to undergo brain surgery (e.g., Ekstrom et al., 2003) and in rats who are freely exploring a circular arena (e.g., O’Keefe ... Nadel, 1978), show that you will have an active representation of your position in the room and the distance between you and the wall. Your representation of your position in space will be very similar to the curve-like representations for number in Figure 10.1, with a noisy curve of activation over an approximate spatial location and distance, but with the activation for position occurring within a virtual two-dimensional floor plane representation rather than the one-dimensional mental number line. This representation of approximate location forms a rough circle or oval-shaped bump of activation on a sort of two-dimensional map (Jeffery, 2007; O’Keefe ... Nadel, 1978). Such a representation

of distance and position in space can be understood as quite similar to an ANS representation of number—for example, an ANS representation is a rough number estimate, a spatial representation is a rough position estimate. I will take these spatial representations to be good examples of magnitude representations.⁶

If you continue standing in this position in the room, facing the distant wall, and close your eyes, you will still have a representation of how far the wall is away from you, and you can update your position representation by sensing your own body movement (Quirk, Muller, ... Kubie, 1990). You can also use this representation of path and distance to guide your actions. Now, with your eyes closed, you can start walking at a slow and comfortable pace toward the distant wall, and you can keep going with your eyes closed until you gently bump into the wall (it is fine to walk with your hands out in front of you). What you will experience along the way is that you gradually have a worse and worse sense of how far you've walked. You may start to feel a little tingle of nervousness as you walk. What is happening (and this is, admittedly, just my prediction based on my own sensations and my reading of the literature—because this exact experiment has not been done systematically—but see, e.g., Schwartz, 1999, Figure 5 and Durgin, Akagi, Gallistel, ... Haiken, 2009 for something close) is that the hump of activation that represents your position in space, and the hump that represents distance, are becoming more diffuse—wider humps. This type of increase in spread of activation in the position coding neurons (**p.180**) is exactly what happens in rats when running in a familiar arena with the lights turned off compared to turned on—their representational humps get wider and more diffuse (Zhang, Schönfeld, Wiskott, ... Manahan-Vaughan, 2014). Because you do not have visual sensory input as you move through the room, your representation of where you are takes on a larger and larger SD. So, as you walk toward the wall, the SD of the hump that is coding your own approximate position in space becomes wider and wider.

And here is the crucial observation: During the moments when your representations of position and distance are becoming wider and wider, when you are walking toward the wall, you do not also come to expect that your body will eventually gradually sink into the wall. You do not believe that your actual position in space has become more diffuse. You do not experience this event as if your body position has become diffuse and fuzzy. You fully expect that you will hit the wall at some point, near instantaneously, but you just don't know when you will hit it. You experience that you have one and only one position in space, but you just aren't sure which position that is.

What might this thought experiment suggest to us about how best to understand the hump-like representations of the analog magnitude systems of, say, Figure 10.1? Consistent with this thought experiment, I believe that your “noisy, hump-like” spatial representations (in this exercise, and at all times) generate two

independent signals for your mind (and that two similar signals are generated by each magnitude representation, e.g., of number, area, duration). The first signal is an estimate of the value for your position variable. This is an estimate within (what may be) a completely precise representational content. It is not a fuzzy set of values or a range of values. Your mind represents that you have one and only one position. Any “fuzziness” in this representation does not amount to a fuzzy position. Rather, the hump indicates your possible/likely positions, with only one position being the true state of affairs.^{7,8}

The second signal generated from this noisy hump-like representation is an estimate of your epistemic limitations for knowing this single precise position. This signal is generated from the width of the hump (roughly speaking), and it is independent of the position signal.⁹ For example, the same value for a distance estimate between you and the wall (e.g., around 10 feet away) can be made more or less certain by, for instance, changing the external noise in the signal (e.g., by closing your eyes, or making the room foggy, one can make the hump wider or narrower).

The width of the hump is a signal of current epistemic limitations (e.g., more or less certain) and not a signal of fuzzy content or fuzzy beliefs. Similarly, the overlap between two hump-like representations is a signal of epistemic limitation for knowing the order between the humps and not a signal of a fuzzy ordering, or of no ordering at all.

Thus, when we inspect the bell-shaped representations of number in Figure 10.1, or in the heat-map positional representations of the hippocampus, or in analog magnitude representations of area, length, and duration, we should be thinking of epistemic limitations and not errors. The “noise” in all of these representations is not a shortcoming of analog magnitudes. Rather, it is an effective solution to the challenge of tuning a sensor (**p.181**) to indicate a precise value along a dimension while simultaneously indexing the epistemic limitations for knowing this value.

Some Impacts of Epistemic Limitations and One True Value

This change—from understanding the nature of our analog magnitude representations as infused with error to understanding them as indexing the epistemic limitations for a precise estimate—can have profound implications for how such representations factor into conceptual change stories. These implications derive from the two central components of this re-understanding: 1) a commitment to one true value, and 2) a sense of epistemic limitations. Here, I gesture at some implications for a few relevant topics. For some of these topics, I think the first steps may be already coming into focus (e.g., how epistemic signals can guide learning). But for others, I think there is still much work to be done to elaborate how the re-understood analog magnitude systems may contribute to our theories (e.g., of learning the integers).

On the Increased Relevance of the ANS

A central claim from this analysis is the suggestion that analog magnitude systems engender a commitment to one true value about the outside world. That is, even when we are uncertain about our estimates (e.g., uncertain of exactly how far we are from the wall) we remain resolute that there is one true value for this distance, and we experience the internal spread in our representations as indicating our own level of uncertainty about what that one value is. A commitment to one true value may impact a host of learning stories, increasing the potential relevance of the analog magnitudes for other aspects of higher cognition, in that it bars “noise” from affecting content (e.g., it blocks vagueness from being introduced into the content of analog magnitude representations).

In the domain of number, theories involving error-prone representations, or failures to represent the numerical distinction between close numbers (e.g., see earlier discussion), may lead to the suggestion that there exists only a partial ordering among analog magnitude representations. What this means is that the sense of numerical order generated by such representations will require a notion of *no order* in addition to *greater than* and *less than*. For example, an animal whose ANS content involved error may have an ANS that cannot distinguish whether 8 or 9 is greater (e.g., see suggestions presented earlier). This animal may have a sense of numerical order that supports thoughts such as “ $9 > 5$ ” and “ $8 < 12$,” but not “ $9 > 8$.” As several authors have suggested, a partial ordering such as this causes difficulties for using such representations to support hypotheses about large exact numbers (e.g., see arguments in Carey, 2009 and Beck, 2012).

Understanding the bell-shaped representations of Figure 10.1 to depict epistemic limitations with a commitment to one true value makes a significant change to such theories. **(p.182)** This is because a commitment to one true value allows for well-ordered representations, and well-ordered representations can support more robust inferences across numbers. For example, a learner with a well-ordered ANS could form the commitment, “Every ANS estimate divides the number space into those values that are less than the estimate, those values that are greater than the estimate, and no others.” This is analogous to a commitment that “every position I can take along my path between the starting point and the wall will divide that path into the space ahead of me, the space behind me, and no others.” Such postulates can form a basis for learning and conceptual change across numbers. For theorizing, work remains to be done to provide detailed proposals of what inductive generalizations may be formed and extended to all numbers, and how these may factor into conceptual change stories (e.g., the integers, the successor function).

For empirical work, determining which numerical commitments our epistemically limited ANS supports will be an important goal towards understanding its relevance to various learning stories. For example, does the

ANS support our earliest intuitions about ordinal comparisons, and is this true for both “small” and “large” numbers? Does the ANS fail to provide an intuitive initial understanding of the integer concepts while still promoting a commitment to one true value? Do observers intuitively seek to gather more samples of evidence when attempting to discriminate between two quantities very close in number? This is merely a listing of some of the areas that may be amenable to research.

Also, an important assay of analog representation will continue to be behavioral performance that accords with the psychophysical law. Studies of early number discrimination, and ordinal judgments, may reveal behaviors consistent with the psychophysical law (e.g., curves like those in Figure 10.2 that run smoothly through the origin) or violations of the law (e.g., dropping to chance performance before reaching the origin). Where violations occur, we will need to explore the reasons for these violations (e.g., a failure of motivation on difficult trials) or consider a new model—one distinct from Weber’s law. Notice also that some results, like a drop in performance on difficult trials, may be evidence for the proposal that observers’ decisions are impacted by epistemic limitations—that is, while an observer’s ANS representations might noisily distinguish close numbers, the observer might be severely in doubt due to epistemic limitations and may thus give up on the trial. “Giving up” is related to epistemic notions while “cannot (in principle) distinguish” is a claim about content.

On Constructing the Integers

The challenge of accessing the integer concepts during development has held a particular interest in recent years. It is a difficult puzzle, and I will not attempt a positive proposal for how children solve it, but I will gesture at a few relevant points.

The first point is that several authors have drawn upon notions of error and noise when rejecting the relevance of the ANS to learning the integers (e.g., consider discussion earlier about the failure of the ANS to distinguish higher consecutive integers). As we have **(p.183)** discussed, this is an inaccurate understanding of the psychophysical law—in theory, and in practice, humans and nonhumans alike have the competence to represent the differences between consecutive integers (e.g., Figure 10.2). This competence (and the related understanding that numbers are well-ordered) has the potential to support relevant inductive inferences across all numbers, but, admittedly, the details of such proposals remain to be revealed.

For theorizing, it will be valuable to explore the generalizations that may be made possible for the learner by including a commitment to one true value within our analog systems. There is much thinking still to be done here, and

additional representations (such as object files, sets of objects, one, successor, etc.) may be necessary ingredients for building integer representations.

For empirical work, continuing to investigate children's early commitments about number before they have become sophisticated counters may be particularly revealing. For example, do children maintain a commitment that "every number divides the number space into those values that are less than the value, those values that are greater than the value, and no others?" And can it be shown that such a commitment is supported by ANS representations?

On Interfacing the Integers with the Magnitudes

One challenging problem is determining how the eventual correspondence between the integer concepts and the representations of the ANS is negotiated. It has been suggested by some that the ANS (and one might expect other analog magnitudes) might best be described using the real numbers. This may be related to the notion that analog magnitudes are continuous (Gallistel ... Gelman, 2000). In contrast, some have suggested that neither number nor time can be represented in a truly continuous, dense format (Laurence ... Margolis, 2005)—they perhaps are represented at some grain size finer than the unitizer of the integers (e.g., finer discrete bins than "1, 2, 3. . ."). If the grain size of analog magnitudes is smaller than the unit size (e.g., "1"), a puzzle remains for how our unitized notions are brought into concert with these representations. Approaches to this problem have included suggesting mechanisms that might bin ANS representations into discrete response regions (Gallistel ... Gelman, 1992; Izard ... Dehaene, 2008) or invoking notions like *one*, *successor*, or *multiplicative identity* to force the bell-shaped noise of ANS representations into discrete ordered units of numerosity (Leslie, Gallistel ... Gelman, 2008a, b). These are interesting approaches. Given that this correspondence problem requires navigating two vocabularies (e.g., discrete integers and fine-grained numerosities), even if we remove the notion of noisy content from our understanding of the ANS, I doubt that this alone will solve the challenge of how children access the integers, the successor function, and other important generalizations across numbers such as the mapping of integers to magnitudes. That is, other representations may be required, as well as a more nuanced learning story.

(p.184) On the Invitation to Learn

On a more practical side, a commitment to one true value and an experience of epistemic limitations may dramatically recast learning problems that involve the analog magnitude systems. For example, imagine that our ANS presented us with fuzzy content. Then, by this system's lights, there would be *no right answer* for the question of how many dots are before my eyes. There is just a probability density function of truthiness—a "blur on the number line", for example, *sort-of-nineish-and-also-sixish-and. . .* In contrast, for the epistemically limited learner with a commitment to one true value, there is one right answer for each

experienced numerosity, and the fuzz usefully tracks their confidence in their estimate of that answer. The epistemic limits indexed by analog representations may inspire learners rather than dissuade them. For the case of number, an experience of the epistemic limitations in the ANS may inspire a learner to ask, “What is the right answer for this number of dots that I’m having so much trouble estimating? Is there another way to find the right answer?”

I suspect that a commitment to one true value and epistemic limitations play a large role in motivating learners to approach learning problems in the first place. Indeed, a combination of epistemic limitations and a commitment to one true value may present us with an invitation to learn. For number, this drive to search for another way to understand the right answer may culminate in the creation of the integer concepts and their interface with the magnitudes.

On Building a Commitment to One True Value

Given that the dominant understanding of analog magnitude systems in recent history has promoted representations that are inherently noisy and approximate, perhaps the most unfamiliar idea in my proposal is the notion of a commitment to one true value. One might ask where such a commitment comes from. The classic answers are that it is either innate—that is, central to how the system operates and always present—or it is learned through experience via some mechanism. What might a possible learning mechanism be? I will gesture at two (though I suspect that a commitment to one true value may be inherent to how these systems operate).

When we consider how one might acquire a commitment that, say, objects have exact (i.e., non-fuzzy) positions in space, we are reminded that our distance estimates (*qua* analog magnitudes) become fuzzier as we move further away from objects (Allison, Gillam, ... Vecellio, 2009). While such a dependency between our movement and the fuzziness of our estimates will present a challenge for a learning story that begins by positing fuzzy-positional-content for things in the world,¹⁰ it might be used as a source of evidence for an epistemically limited learner along the path to constructing a commitment to one true value.

Because of such dependencies, through our own actions (e.g., walking closer, squinting our eyes or moving them, internally changing a scaling factor) we may experience that **(p.185)** our epistemic limitations are not fixed—the spread in the humps of our magnitude representations are not static. Notice that this requires us to note the sameness (e.g., the mean of our distance estimate) in the face of our changing epistemic state (e.g., the SD of our distance estimate). If this were possible, it may allow our experience, over time, to give rise to a commitment that these are merely noisy estimates of some true external value.¹¹

Another source of evidence for one true value from epistemic limitations might be multimodal representations. Having two magnitude representations each grounding out in the same external stimulus (e.g., vision and touch attempting to determine how long an object is) may provide us with some agreement in their estimate, and also some disagreement in the epistemic certainty they yield (e.g., a sharper, more certain representation of the length of the object from vision than from touch [Ernst ... Banks, 2002]). Such experiences of differing levels of epistemic limitation could give rise to the understanding that the length itself is something stable and external to the mind that these systems are attempting to estimate.¹²

Alternatively, it may be that the commitment to the existence of one true value is an immediate and permanent aspect of our magnitude representations from the get-go.¹³

On Learning across Multiple Timescales

Epistemic limitations may also be useful to observers in other ways and across multiple timescales. For example, within a single display time (e.g., 1000 msec), it may be that epistemic signals guide the learner to adjust behavior over hundreds of milliseconds in order to improve the resulting representation (e.g., moving closer to the object to estimate distance; shifting fixation to a crowded area of dots to better estimate their numerosity).

Over the course of tens of trials (or learning problems within a classroom), epistemic signals may help scaffold the learner toward better decisions and more accurate behaviors. In recent work, we have found that beginning with numerical discriminations that are easier (i.e., trials for which the observer has higher confidence about the correct answer) and building toward more difficult trials leads children to have a better sense of the distinctions between stimuli, resulting in improved performance throughout the task (Odic, Hock, ... Halberda, 2014).

At still longer timescales, we are finding that manipulating internal confidence in ANS representations can transfer to school math performance, even on problems using discrete number symbols rather than nonsymbolic arrays (Wang, Odic, Halberda, ... Feigenson, 2015). We also believe that longer-term changes in how children respond to their own internal confidence may contribute to developmental changes in ANS precision.

At the timescale of a whole life, it may be the sum total of such experiences with epistemic signals (and the successful or failed attempts to adjust behavior based on these signals) that forms the basis for our trait-like attributions (e.g., "I'm not a math person").

There remains much work to do in integrating the insights we are gaining about epistemic signals into how we understand learning problems and their solutions. Because **(p.186)** epistemic signals may be used (and have impact) across many timescales, there is also the possibility of using epistemic signals to serve as a unifying concept for both theorizing and empirical work.

Conclusions

I began this chapter by motivating a consideration of how we understand the bell-shaped representations of Figure 10.1. Do we think that the representation of ANS₉ is less precise than the representation of ANS₅? I suggested that we are at risk for misunderstanding the “noise” in analog magnitude representations as a limitation in the content of those representations. For example, some authors have suggested that the distinction between ANS₈ and ANS₉ is not represented at all. I showed that such a view is at odds with the human and animal data in discrimination tasks. Instead, I argued that the variability in our analog magnitude representations is better understood to be a signal of our epistemic limitation for our estimates and discriminations. I believe this proposal provides a satisfactory account of performance failures, and also of our deep commitment to there being one true value for any given magnitude out in the world. Finally, I closed by starting to consider how such an adjustment might affect some theoretical and empirical work involving analog magnitude representations.

In sum, my suggestion is that number, distance, time, area, volume, brightness, and many other psychological dimensions are represented by well-ordered analog magnitude representations, which provide an estimate of one true value along with a signal of the epistemic limitations for knowing this value. Progress remains to be made in understanding how these representations may play a role in some aspects of higher-level cognition (like constructing integer representations). There is also progress to be made in characterizing the foundational nature of the representations themselves. I hope this has been a useful step.

References

Bibliography references:

Allison, R. S., Gillam, B. J., ... Vecellio, E. (2009). Binocular depth discrimination and estimation beyond interaction space. *Journal of Vision*, 9(1), 10, 1-14.

Beck, J. (2012). The generality constraint and the structure of thought. *Mind*, 121, 563-600.

Berkeley, G. (1910). An essay towards a new theory of vision. In E. Rhys (Ed.), *A new theory of vision and other select philosophical writings by George Berkeley*

Bishop of Cloyne (pp. 13–86). New York: E.P. Dutton ... Co. (Original work published 1709.) **(p.189)**

Brannon, E. M., ... Roitman, J. D. (2003). Nonverbal representations of time and number in animals and human infants. In W. H. Meck (Ed.), *Functional and neural mechanisms of interval timing*. Boca Raton, FL: CRC Press.

Burge, J., Girshick, A.R., ... Banks, M.S. (2010). Visual-haptic adaptation is determined by relative reliability. *Journal of Neuroscience*, 30(22), 7714–7721.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Cheng, K., Srinivasan, M. V., ... Zhang, S. W. (1999). Error is proportional to distance measured by honeybees: Weber's law in the odometer. *Animal Cognition*, 2, 11–16.

Dehaene, S. (2007). Symbols and quantities in parietal cortex: Elements of a mathematical theory of number representation and manipulation. In P. Haggard, Y. Rossetti, ... M. Kawato (Eds.), *Sensorimotor foundations of higher cognition*, Vol. XXII, *Attention and Performance* (pp. 527–574). Cambridge, MA: Harvard University Press.

Dehaene, S., ... Brannon, E. (Eds.). (2011). *Space, time and number in the brain: Searching for the foundations of mathematical thought*. San Diego: Academic Press.

Dehaene, S., ... Changeux, J.P. (1993). Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience*, 5(4), 390–407.

Durgin, F. H., Akagi, M., Gallistel, C. R., ... Haiken, W. (2009). The precision of locomotor odometry in humans. *Experimental Brain Research*, 193(3), 429–436.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., ... Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, 425, 184–187.

Ernst, M. O., ... Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415, 429–433.

Feigenson, L. (2007). The equality of quantity. *Trends in Cognitive Sciences*, 11(5), 185–187.

Fodor, J., ... Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.

Gallistel, C. R., ... Gelman, R. (2000). Non-verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences*, 4(2), 59–65.

Gallistel, C. R., Gelman, R., ... Cordes, S. (2006). The cultural and evolutionary history of the real numbers. In S. Levinson ... P. Jaisson (Eds.), *Evolution and culture: A Fyssen Foundation Symposium* (pp. 247–274). Cambridge, MA: MIT Press.

Halberda, J., ... Feigenson, L. (2008) Developmental change in the acuity of the 'number sense': The approximate number system in 3-, 4-, 5-, and 6-year-olds and adults. *Developmental Psychology*, *44*, 1457–1465

Halberda, J., Ly, R., Wilmer, J., Naiman, D., ... Germine, L. (2012). Number sense across the lifespan as revealed by a massive Internet-based sample. *Proceedings of the National Academy of Sciences U S A*, *109*(28), 11116–11120.

Halberda, J., ... Odic, D. (2014). The precision and internal confidence of our approximate number thoughts. In D. C. Geary, D. B. Berch, ... K Mann Koepke (Eds.), *Evolutionary origins and early development of number processing: Mathematical cognition and learning*. San Diego: Academic Press.

Izard, V., ... Dehaene, S. (2008). Calibrating the mental number line. *Cognition*, *106*, 1221–1247.

James, W. (1983). *The principles of psychology*. Cambridge, MA: Harvard University Press. (Original work published 1890.)

Jeffery, K. J. (2007). Integration of the sensory inputs to place cells: What, where, why and how. *Hippocampus*, *17*(9), 775–85. **(p.190)**

Laurence, S., ... Margolis, E. (2005). Number and natural language. In P. Carruthers, S. Laurence, ... S. Stich (Eds.), *The innate mind: Structure and contents* (pp. 216–235). New York: Oxford University Press.

Leslie, A. M., Gelman, R., ... Gallistel, C. R. (2008a). Where integers come from. In P. Carruthers, S. Laurence, ... S. Stich (Eds.), *The innate mind, Vol. 3: Foundations and the future* (pp. 109–138). New York: Oxford University Press.

Leslie, A.M., Gelman, R., ... Gallistel, C.R. (2008b). The generative basis of natural number concepts. *Trends in Cognitive Sciences*, *12*(6), 213–218.

Libertus, M. E., ... Brannon, E. M. (2009). Behavioral and neural basis of number sense in infancy. *Current Directions in Psychological Science*, *18*(6), 346–351.

Meck, W. H., ... Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*(3), 320–334.

Mill, J. S. (1865). *An examination of Sir William Hamilton's philosophy*. London: Longmans, Green, and Co.

Nieder, A., ... Miller, E. K. (2003). Coding of cognitive magnitude: Compressed scaling of numerical information in the primate prefrontal cortex. *Neuron*, 37(1), 149-57.

Odic, D., Im, H. Y., Eisinger, R., Ly, R., ... Halberda, J. (2015). PsiMLE: A maximum-likelihood approach to estimating psychophysical scaling and variability more reliably, efficiently, and flexibly. *Behavior Research Review* [Epub ahead of print].

Odic, D., Hock, H., ... Halberda, J. (2014). Hysteresis affects approximate number discrimination in young children. *Journal of Experimental Psychology: General*, 143(1), 255-265.

O'Keefe, J., ... Nadel, L. (1978). *The hippocampus as a cognitive map*. New York: Oxford University Press.

Piaget, J. (1955). *The construction of reality in the child*. Translated by M. Cook. London: Routledge and Kegan Paul.

Piazza, M., Facoetti, A., Trussardi, A. N., Berteletti, I., Conte, S., Lucangeli, D., Dehaene, S., ... Zorzi M. (2010). Developmental trajectory of number acuity reveals a severe impairment in developmental dyscalculia. *Cognition*, 116, 33-41.

Piazza, M. (2010). Neurocognitive start-up tools for symbolic number representations. *Trends in Cognitive Sciences*, 14(12), 542-551.

Quirk, G. J., Muller, R. U., ... Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10(6), 2008-2017.

Rilling, M. E., ... McDiarmid, C. G. (1965). Signal detection in fixed-ratio schedules. *Science*, 148, 526-527.

Schwartz, M. (1999). Haptic perception of the distance walked when blindfolded. *Journal of Experimental Psychology: Human Perceptions and Performance*, 25(3), 852-865.

Wang, J., Odic, D., Halberda, J., ... Feigenson, L. (2015). Temporary changes to children's ANS precision affect their symbolic math performance. Poster presented at SRCD, the Society for Research on Child Development, Philadelphia, PA.

Zhang, S., Schönfeld, F., Wiskott, L., ... Manahan-Vaughan, D. (2014). Spatial representations of place cells in darkness are supported by path integration and border information. *Frontiers in Behavioral Neuroscience*, 8, 222.

Notes:

(1.) Systematicity of thought is the proposal that our ability to entertain certain thoughts (e.g., “8 is greater than 6”) is related, straightforwardly, to our ability to entertain some other thoughts (e.g., “9 is greater than 6,” and “8 is greater than 5”) and not others (e.g., “I like sailing”). Systematicity, it is argued, is related to the productivity of thought (Fodor ... Pylyshyn, 1988). The suggestion from Beck is that analog magnitudes may not be systematic in this way (Beck, 2012).

(2.) One case where this type of “drop to chance before ratio 1” violation has been reported is in a numerical discrimination task from Halberda and Feigenson (2008), in which children were tested with a wide variety of ratios. In this case, we observed that the large number of highly difficult trials led many children to become discouraged and give up on the task when a trial seemed too hard. Notice that “giving up” is a performance issue and does not reflect competence. This type of giving up can be an interesting matter to investigate in its own right. For instance, we have since found that an observer’s internal confidence can be affected by trial order and trial difficulty (Odic, Hock, ... Halberda, 2014). One implication of this is that confidence, as experienced via ANS representations, may play an important role in mediating school math performance. That is, a change in a child’s internal confidence in their ANS representations can help or hinder their later school math performance (Wang, Odic, Halberda, ... Feigenson, 2015). There is much work still to be done in this area.

(3.) This is not to say that infants will perform as well as adults do; there are genuine developmental improvements in ANS performance. What changes over development is the steepness of the curves (e.g., in Figure 10.2). I believe it is a promising area for future work to determine what underlies these age-related improvements and how we should best interpret them. The improved behavioral performance (e.g., in judging that 12 dots is more than 9) suggests that infants and children are getting better, sharper estimates with age and/or practice. Our recent work also suggests that such children experience higher internal confidence as this sharpness improves (Odic et al., 2014). Future neuroscientific evidence could be relevant as well—for example, one might find increased numbers of neurons in the intraparietal sulcus (IPS) or increased frontoparietal connectivity. All of these are consistent with the understanding that variability in ANS representations is indexing epistemic limitations rather than fuzzy content (see argument in main text). The details matter, and these questions present exciting possibilities for exploration.

(4.) Here I leave open whether magnitude representations alone can support integer thoughts like “8.”

(5.) Some authors have suggested that analog magnitude representations are the complete hump-like curves depicted in Figure 10.1, such that an instance of viewing, say, nine objects in a collection gives rise to a hump of activation centered on ANS_9 . This can occur either because the ANS representations are themselves analog transforms of earlier analog signals (Dehaene ... Changeux, 1993) that may even vary dynamically in time (e.g., imagine a vibrating wave-hump of activation that is constantly shaking a little bit, moving up and down the line, and getting thinner and fatter due to changes in incoming stimulation). Alternatively, the representations may be full humps of activation because the read-out of a discrete representation is subject to some noise (Gallistel, Gelman, ... Cordes, 2006). In contrast, other authors have suggested that activations within analog magnitudes are discrete individual samples (e.g., a bar of precise activation on the mental number line) and that the humps only become observable when the discrete samples are averaged across many instances (e.g., Meck ... Church, 1983). I believe that what I argue here applies equally well to any of these hypothesized representations, since all will need to address how noise should be understood in the system and by the animal. Personally, I like imagining fully moveable vibrating humps (all of them subject to a scaling factor) and so I will continue discussing representations that look like the humps in Figure 10.1. I will argue that even these complete humps are best understood to implement a perfectly precise number estimate with a signal of epistemic limitation, and not a fuzzy spread-out number thought.

(6.) I'll note that Weber's law may hold for our distance estimates (e.g., how far is the wall from me) but not necessarily for our spatial position estimates per se (e.g., my position is x,y). For my argument, I believe the phenomenon of interest (e.g., uncertainty of distance) emerges from an analog magnitude representation of the distance from me to the wall (cf Cheng, Srinivasan, ... Zhang, 1999). I will include references to neuronal data for position as well, partly because this is the most developed literature in the vicinity of my argument, and because the images and notions this literature has developed (e.g., place fields) are, I think, valuable for fueling our intuitions about what might be happening in the head as we move around. In fact, distance estimates are an important part of the input to the place cells (e.g., Jeffery, 2007), and so place cells may help us understand how our analog magnitude distance representations fuel later representations of position, even if position representations themselves are not bona fide analog magnitudes.

(7.) The level of granularity for each analog magnitude dimension is an open question of interest, but that granularity is not affected by a JND or the SD of the humps—by the argument I am suggesting here, these factors are about epistemic limitations, not content granularity.

(8.) I believe I can run a very similar argument if the magnitude representation in question is dense and continuous rather than organized in discrete bins. I do not wish to choose between these possibilities here. I believe an understanding of epistemic limitations will be important for whichever way we understand the underlying content.

(9.) Modulo dependencies like Weber's law. But note, such dependencies do not violate the well-orderedness of the values on the magnitude scales and so do not conflict with the point I am making here.

(10.) Indeed, a theory about analog magnitude representations that takes the humps to be humps in content (e.g., fuzzy content) may be positioning the animal to face a learning problem akin to those conceived by Piaget (1955), Mill (1865), and James' blooming buzzing confusion (1890/1983). That is, if the contents of these representations are fuzzy, diffuse, and inexact, then why don't we believe the world to be fuzzy, diffuse, and inexact? This would seem to be a straightforward extrapolation from the experience of an animal with fuzzy content. Under the model we are discussing here, the analog magnitude machines always and only generate "noisy humps." They never provide exact discrete "hash marks." So taking those fuzzy humps to indicate fuzzy content will position the animal to face the learning problem of going from *fuzzy-hump-thoughts* to *precise-hashmark-ideas*. As Piaget, Mill, James, and subsequent attempts reveal, this is a hard learning problem to say the least.

(11.) Note, this is similar to John Stuart Mill's (1865) argument for the role of our own actions in birthing a commitment to the permanent possibilities of perception, thereby giving rise to our belief in a mind-independent enduring external world.

(12.) This argument could include a learning piece that is similar to Lord Berkeley's (1709/1910) proposal for how touch could teach vision to represent distance from eye-muscle strain. However, in the case of one true value, it may be more like two limited systems (e.g., touch and vision) giving rise to a new idea about the external world (e.g., that there is one true distance for the object) rather than one "perfect" system guiding the other. If one were interested in this type of learning story, it would also be good to become acquainted with the results from studies of multimodal cue combination, which suggest that we combine evidence from both modalities (e.g., touch and vision), subject to the internal confidence in each representation (e.g., Burge, Girshick, ... Banks, 2010; Ernst ... Banks, 2002).

(13.) I must admit that I prefer to think that a commitment to one true value may be a principle that is required by the combinatorics of the language of thought; and, perhaps even earlier, may have its ancestor in the first sensory motor

transforms of living organisms. Of course, such speculations leave us with much work to do to determine how this might be so.