

The Small Number System

Eric Margolis

Department of Philosophy, University of British Columbia

Abstract

I argue that the human mind includes an innate domain-specific system for representing precise small numerical quantities. This theory contrasts with object-tracking theories and with domain-general theories that only make use of mental models. I argue that there is a good amount of evidence for innate representations of small numerical quantities and that such a domain-specific system has explanatory advantages when infants' poor working memory is taken into account. I also show that the mental models approach requires previously unnoticed domain-specific structure and consequently that there is no domain-general alternative to an innate domain-specific small number system.

1. Introduction

Researchers who study numerical cognition commonly hold that there are two representational systems that are critical to the origins of numerical concepts—one that is approximate and capable of representing large numerical quantities and one that is precise and limited to small numbers of entities (e.g., Spelke 2003; Carey 2009; vanMarle et al. 2016). Moreover, it is often claimed that the second of these isn't fundamentally a system for representing number, despite the fact that it is often referred to as a "small number system". While it operates on small numbers of entities (up to three or four) and is incapable of dealing with larger numbers, it is generally thought to start out as a domain-general system or at least as a system that doesn't require any innate structure that is specific to the domain of number.

In this paper, I argue that the small number system should be understood, instead, to have innate domain-specific structure. I begin by showing that we need to postulate a richer system of representation than one that is confined to object tracking. I then go on to argue that the proposal of an innate system for representing a few small numerical quantities stands up well against the proposal that the small number system is fundamentally a general-purpose capacity for working

with mental models. I also argue that the mental models proposal ends up requiring an innate capacity for performing assessments of one-to-one correspondence. As a result, we are left with two broadly nativist options—one that relies on an innate system for representing small numerical quantities and one that relies on an innate system for comparing small numbers of items for numerical equivalence.

2. The Subitizing Module and Some Initial Objections to Its Innate Representations

Let's begin with an unabashedly nativist approach to the small number system. This is the proposal that there is an innate domain-specific system for representing a few precise numerical quantities—for representing precisely *one*, *two*, and *three* (cf. Hurford 1987; Margolis and Laurence 2008; Barner 2017). If you were to see a pair of shoes on the floor, it would register the twoness of the shoes. Or if you were to hear the tap-tap-tap of someone knocking on the door, it would register the threeness of the knocks. By hypothesis, this domain-specific system represents numerical quantity as such and is restricted to the small number range. I will refer to this system as the *subitizing module* (or SuM), since the term *subitizing* evokes a process for representing numerical quantity that is distinctive to the small-number range and that doesn't involve counting. In paradigmatic cases of subitizing, it's as if you can directly perceive certain numerical quantities but only in the small number range.

In what way does SuM represent its numerical quantities? My proposal, following Margolis and Laurence (2008), is that it includes a small stock of discrete representations that are causally responsive to particular numerical quantities and that have the function of responding to these quantities. Although there are a number of ways this could be implemented, the most straightforward is by means of a neural network that takes input from systems that individuate entities in different modalities and whose output nodes are selectively responsive to particular numerical quantities. To bring this about, the connections mediating the spread of activation would have to be weighted in such a way that the activation of any single input node suffices to activate the “one” output node, the activation of any two input nodes suffices to activate the “two” output node, etc., where each output node's activation also inhibits the activation of the other output nodes. Under this arrangement, n individuated entities would cause the activation of a unique symbol corresponding to that precise quantity.

More structure could be built into SuM, but this is the minimal amount that I will consider essential to its representation of *one*, *two*, and *three*. Notice that on this minimal account, the symbols for small numerical quantities needn't be inherently ordered, and there needn't be a procedure that ensures that *three* is represented as more than *two*, or *two* as more than *one* (unlike conventional counting terms). Yet this minimal structure is enough to put children in a position to learn about some of the basic relations between differing small numerical quantities. They could do this by using SuM in the context of observing addition and subtraction events, or by creating these changes themselves and attending to the numerical effects that SuM allows them to represent. For example, attending to two toys would generate the representation for *two*, and seeing one of these removed would generate the representation for *one* both for the quantity removed and for the quantity that remains. Observing changes like this might allow children to infer that the difference between *one* and *two* is itself *one* and hence that *two* is the larger numerical quantity. Another possibility, of course, is that SuM builds in some of this structure from the start. The present point is simply that the defining feature of SuM is its small stock of discrete symbols that represent numerical quantity as such.

Despite the fact that many researchers hold that there is some kind of representational system that is responsive to small numbers of items, few accept that it takes the form of SuM—that it is a domain-specific system for representing particular numerical quantities. Why is that? One of the main reasons is that a system like SuM is thought to be too speculative. According to this objection, it is one thing to suppose that there is an innate system like SuM in order to explain how children come to be in a position to acquire concepts for precise numerical concepts, but it is another matter entirely whether there is empirical evidence for SuM's innate numerical representations. If there aren't any data to suggest that infants can represent *one*, *two*, and *three*, there is no concrete reason to believe that they actually do.

Resistance to innate representations for small numerical quantities has also grown in recent years in response to what many researchers see as a problem that was endemic to earlier work in developmental psychology, namely, the failure to fully take into account the nonnumerical properties that correlate with number. Discrete number correlates with a variety of continuous properties. A group of three plums differs in number from a group of two plums, but also takes up more space, has more total surface area, and so on. While early research on infants' numerical

abilities did take steps to ensure that infants weren't merely responding to differences in these sorts of continuous properties, there have been questions about whether these measures went far enough. One landmark study that has fueled these doubts focused on the contour length of small numbers of geometrical figures (i.e., the sum of their perimeters) (Clearfield and Mix 1999). Infants were first shown different arrangements of either two or three same-size squares until they habituated to the stimuli. Notice that because the squares were all the same size, number was deliberately confounded with contour, as the three squares were guaranteed to have more contour than the two squares. The infants were then shown alternating instances of two and three squares where the stimuli with the number they had seen before had a new amount of contour, while the stimuli with the new number had the old amount of contour. The key finding was that the infants dishabituated to the change in contour but not to the change in number, suggesting that perhaps infants who had seemed to respond to numerical quantity in previous work were merely responding to continuous properties of the stimuli.

Finally, if we turn to language learning, there is a pattern of development that may appear to conflict with the proposal that there are innate representations for precise small numerical quantities. This is that learning the meanings of the first few natural language counting terms isn't easy for children and that children reliably learn them in order. One might have thought if children have innate representations for *one*, *two*, and *three*, then learning the words that pick out these quantities would be relatively easy and that there would be no intrinsic constraints on the order in which they can be learned. It would just be a matter of mapping three word forms to three independent and readily available representations. On the other hand, if representations for *one*, *two*, and *three* have to be constructed in the course of language learning, it would make sense that learning these words is challenging and that they are acquired in order, since the construction of these representations would involve increasing complexity as the numerical quantities get larger.

3. Object Tracking

In this section, I want to look at the first of two alternatives to the SuM theory's account of the small number system. According to this first theory, infants may have a system of representation that is confined to small numbers of items, but this isn't a system for representing numerical

quantity. It's simply a mechanism of attention that the visual system uses to track small numbers of objects (Leslie et al. 1998, Scholl and Leslie 1999).

Object-based approaches to visual attention differ from approaches in which attention is taken to function like a spotlight that directs limited processing resources to a focal region in the visual field. On an object-based model, what happens instead is that attention attaches to individual objects and temporarily sticks to each object regardless of whether it moves. On Leslie et al.'s model, this is implemented by the object-indexing system, which incorporates up to four symbols—i.e., four *indexes*—that act like pointers in that each picks out the object it is responsible for without necessarily relying on a representation of its features (color, texture, etc.). Rather, these indexes track their objects in the first instance on the basis of their spatial-temporal properties and can do so even when an object is briefly occluded.

Scholl and Leslie note that there is large body of evidence showing that object-based attention is an important aspect of mid-level visual processing in adults. For example, in the motion-object tracking task, subjects view a computer-generated image of a number of identical looking objects, a subset of them are briefly highlighted, and then all of the objects begin to move in quasi-random directions (Pylyshyn and Storm 1988). After a while they stop, and subjects have to say which are the ones that had been highlighted. This may sound like it is difficult to do. After all, the target objects look exactly the same as the distractors (e.g., they might all be black squares), and each moves independently of the others along its own erratic path. But people are fairly good at identifying the target items so long as they are not asked to keep track of more than about four. The object-indexing model can explain these and related results on the assumption that it has just a small number of indexes at its disposal and they operate in parallel, tracking their objects without having to identify them by their features.

For present purposes, what matters is how the object-indexing system promises to explain infants' apparent numerical abilities. The focus of Scholl and Leslie's model is Wynn's influential claim that infants can do simple arithmetic as it relates to events with small numbers of seen objects (Wynn 1992). This research employed a violation of expectation procedure in which five-month-olds were shown addition and subtraction events with correct or incorrect outcomes. Looking longer at an incorrect outcome is a sign that infants find it to be unexpected and, for Wynn, that they appreciate the numerical significance of the events they are witnessing.

In one experiment, some infants saw a single doll placed on an empty stage, which was then hidden behind a screen, followed by a hand placing a second doll behind the screen—that is, a 1+1 event. Other infants saw a similar 2-1 event. Then, in the test condition, the screen was removed to reveal either one or two dolls. The result was that infants looked longer at one doll for the 1+1 event and at two dolls for the 2-1 event, suggesting that they found these incorrect outcomes to be unexpected.

But does this mean that infants really appreciate that $1+1=2$? According to Scholl and Leslie, the infants' looking-time can be explained without postulating any numerical representation per se. They propose instead that it results from the effect that these events have on the assignment and maintenance of the indexes that track the dolls (1999, 34). Take the 1+1 event. In this case, one index is initially activated to track the doll that is visible on the stage. After the screen comes up, this first index maintains a link with the now hidden object and a second index is activated to accommodate the second doll that is placed behind the screen. As a result, by the time the screen is removed and just a single object is revealed, there is an extra index that has lost track of its object, and this causes an increase in attention to search for the missing object. In contrast, when the screen is removed and two objects are revealed, no extra attention is needed. In short, the $1+1=1$ event leads to longer looking not because infants appreciate the relevant arithmetic facts; infants look longer at the incorrect outcome because of the demands that are placed on attention given the way the object-indexing system works.

Many theorists who talk about the existence and significance of a small-number system seem to have something like this deflationary model in mind. To mention just one recent example, in a paper with the subtitle “Contributions of the Object-Tracking and Approximate Number systems”, vanMarle and colleagues point to the object-indexing system as the main alternative to the approximate number system in theories that aim to explain how children learn the meanings of natural language counting terms:

More recent work ... suggests an alternative account in which the verbal labels are mapped onto episodic object representations in another core mechanism—the object tracking system (OTS). This system consists of a set of indexes that ‘point’

to objects in the world, keeping track of them as they move through space ...
(vanMarle et al. 2016, 1-2)

As vanMarle et al. see things, the main lesson regarding the limits of the approximate number system and the object-indexing system is that neither on its own can explain how children learn the significance of counting. Instead, these systems must work together, where the object-indexing system's unique contribution is to produce "exact representations, but only for small numbers of individuals and without cardinal value" (vanMarle et al. 2016, 2). Unfortunately, it has never been clear how the object-indexing system's representations can combine with the representations of the approximate number system to produce the precise numerical representations needed to interpret the counting terms (Laurence and Margolis 2005). But even putting this puzzling theoretical matter to the side, the object-indexing system falls short in that it cannot accommodate the full range of findings associated with the representation of small numbers of items. This is because some of these include sensory conditions where visual cues are absent or impoverished, or where object tracking isn't viable or not even relevant to the situation at hand.

Consider a study that resembles Wynn's but where infants were given addition events that require the interpretation of intermodal cues (Kobayashi et al. 2004). Five-month-olds were presented with computer-generated events in which, when an object dropped from the top of the display, a tone was heard exactly at the point at which it hit the bottom. The general impression of these events is that the tone occurs as a result of the object impacting the ground. Initially the infants were familiarized with the sorts of events they would be tested on. They saw a screen conceal the bottom half of the display, and either two objects or three objects fell one at a time, so that they became hidden behind the screen, with a tone always occurring at the hidden point of impact. After the sequence was finished, the screen dropped to reveal the expected number of objects, either two or three. Next came the test trials. At the start of the test trials, rather than dropping from the top of the display, a single object moved horizontally along the bottom until it arrived at the center. Then a tall screen came up, obstructing the infants' view of the object as well as the entire center of the monitor along the vertical axis, and either one or two tones were heard. Finally, the screen came down to reveal the correct number of objects ($1+1=2$) or an

incorrect number ($1+1+1=2$). The crucial finding was that the infants looked significantly longer at the incorrect outcome, suggesting they found it to be unexpected.

Notice that this looking-time pattern cannot be explained by the different amounts of attention that are required by object-indexing. In *both* the 1+1 and the 1+1+1 events, only one item was in a position to trigger a visual index (1 seen object + n tones). Given the experimental setup, infants had to deduce that a tone indicates that an object behind the screen dropped to the ground; they couldn't actually see the object. It's also unlikely that the looking-time pattern can be explained either in terms of infants' response to a nonnumerical continuous property or their reliance on the approximate number system's representation of the events. The multimodal experimental design excluded any possibility of successfully responding solely to a continuous property like the seen amount of surface area that was placed behind the screen—again, this was identical across the two conditions. And while the approximate number system is functional in five-month-olds, it doesn't have the required discriminative capacity to succeed on this task until infants are significantly older than five months of age. It isn't until infants are nine months old that it has matured enough to distinguish between numerical quantities in a 2:3 ratio (Xu and Spelke 2000; Xu et al. 2005).

Work with animals also suggests that object-indexing offers an inadequate account of the small number system. Consider a study in which honeybees were successfully trained to select between two numbers of geometrical figures in a delayed match-to-sample task (Gross et al. 2009). Upon entering the training apparatus, a bee would see the sample stimulus with two or three elements and have to fly a meter down a tunnel before encountering two further stimuli—one with two elements, one with three—each of which marked a different exit but where only the numerical match was rewarded. Later, in the test trials, the bees were able to generalize correctly, selecting the numerical match for novel stimuli using patterns that carefully controlled for nonnumerical continuous properties. The bees were even able to select the matching number in the face of misdirecting cues, such as a sample that included a color that only appeared in the numerical mismatch. What they couldn't do was generalize correctly outside of the small-number range—they failed on 4v6, 4v5, and 5v6. To appreciate why object-indexing can't explain these results, it helps to picture things from the bees' perspective. They fly past the sample (e.g., two yellow stars) only to encounter two further stimuli in the choice chamber (e.g.,

one composed of two blue dots and the other of three). The yellow stars are a full meter in front of the blue dots—equivalent to *dozens* of bee body-lengths—and all of the stimuli are completely static. Hence there are absolutely no cues that the elements composing the stimuli in the choice chamber are one and the same as those previously encountered. All the bees have to go on are the numerical properties of the stimuli, for example, the fact that the correct match is numerically equivalent to the sample.

In sum, object-indexing can't be the whole story about the representation of small numbers of items. It may provide some of the crucial input to the small number system since it can individuate the entities that the small number system responds to, at least when these are confined to vision. But the representation of small numbers of items isn't unique to vision. And even within vision, there are responses that are limited to small numbers of items where object tracking isn't called for and that don't turn on the continuous properties of the stimuli.

4. Mental Models

We have seen that the object-indexing account of the small number system won't do. But there is another alternative to SuM that isn't restricted to visual representation or to attentional mechanisms for tracking objects. According to this alternative, the small number system is fundamentally a domain-general capacity to construct and manipulate mental models.

The core idea of this approach is that infants respond to perceived small numbers of objects with a mental model that is composed of distinct symbols—one per item—and that this can be held in working memory for a brief amount of time and can support a variety of inferences about the represented group. One particularly important use of a mental model is to interpret quantitative changes as an event unfolds, including changes in numerical quantity. This can be done by comparing the model held in working memory to what is perceived at a later time and checking for whether these correspond one-to-one. Simon (1997) has proposed an account along these lines to explain Wynn's addition and subtraction results. According to Simon, the reason infants look longer at the incorrect arithmetic outcome (e.g., $1+1=1$) isn't because they have arithmetic abilities, and it isn't because extra attention is needed to deal with an index that has lost contact with its object. Rather, they look longer because they recognize that the model in working memory contains an element for which there is no corresponding doll. There is a

mismatch between the model and the aspect of the world it is directed towards when the two are compared for one-to-one correspondence.

Le Corre and Carey offer a similar account (Le Corre and Carey 2007; Carey 2009). They refer to their proposed system as a system for *parallel individuation* to emphasize that it isn't inherently a visual system (although it often operates on visual input) and to emphasize that it represents small numbers of individuals via correspondingly distinct representations making up a mental model. For example, if an infant sees three balls, then the model contains three symbols, each one corresponding to one of the balls, with no symbol explicitly representing that there are three. Likewise, if an infant hears three honks, a similar three-symbol model may be formed, with each symbol corresponding to one of the sounds. Parallel individuation is a richer system of representation than an object-based attentional system not only because it isn't confined to any single modality, but because its symbols needn't be involved in online tracking. Carey notes, for example, that sometimes a two-symbol model will result from seeing pairs composed of different individuals and not from repeated sightings of the same individuals. In such cases, "a working-memory model of two objects must be abstracted from these ... arrays" (2009, 143). At the same time, parallel individuation has the same set-size limitation as object-indexing in that it can only create models that are formed from a small number of symbols. This limitation on model size is supposed to derive from the capacity limit on working memory and is thought to increase in the first year of life as the working memory system matures (Carey 2009, 83).

An important motivation for proponents of the mental models approach is that it promises to explain children's early numerical abilities without having to postulate innate number-specific structure, including innate numerical representations. As Simon puts it, "the earliest form of numerical behavior of which infants are capable arises from the deployment of some very general information processing characteristics of human cognitive architecture". He describes the foundational competences that underlie infants' success on numerical tasks as being "non-numerical" on the assumption that "they did not evolve specifically for the purpose of number processing" (1997, 350).

Likewise, while Carey holds that the parallel individuation system supports numerical quantitative assessments—because mental models are suited to figure in assessments of one-to-one correspondence—it is still a domain-general system:

The purpose of parallel individuation is to create working-memory models of small sets of individuals, *in order to represent spatial, causal, and intentional relations among them*. Unlike analog magnitude number representations, the parallel individuation system is not a dedicated number representation system. Far from it. The symbols in the parallel individuation system explicitly represent individuals. (Carey 2009, 151; italics added)

In this passage, Carey makes clear that working-memory models are meant to serve purposes that have nothing at all to do with numerical quantity.

Also, for Carey, the parallel individuation system supports quantitative comparisons that are *not* numerical. It does this by encoding some of an object's continuous properties (e.g., its amount of surface area) and associating this information with the symbol for the object in a working memory model. An operation can add the value for this continuous property to the value that is bound to the other model elements to represent the total amount for the represented group. Then it can compare this summed value to the summed value for another small group to determine whether they have the same amount or whether one has more than the other. This may seem somewhat convoluted. Why can't infants directly evaluate which has more surface area? Why do the continuous properties for each item have to be bound to a model element? The answer, for Carey, is that this arrangement is required by the data. It makes sense of cases where quantitative performance has the set-size signature of working memory—which is confined to a small number of items—but where performance is still driven by the continuous properties of the stimuli. For example, in one experiment, infants saw a certain number of crackers placed in one container, and another number placed in a second container, and the question was which container they would approach (Feigenson et al. 2002). Infants went for the larger quantity for 1v2 and for 2v3, but not when four or more crackers were involved (e.g., 2v4). Further, when the smaller number of crackers had the larger total amount of surface area, they chose the smaller number; and when different numbers had the same amount of total surface area, they showed no preference. So while the parallel individuation system would seem to explain why these infants' performance is capped at three crackers, the comparisons guiding these foraging choices are

nonnumerical. The crucial comparison is defined over the continuous properties that are bound to the model elements, allowing infants to compare the total amount of cracker in the two containers.

In contrast, in a related study with monkeys who saw different numbers of apple slices placed into two containers, the monkeys succeeded by choosing the larger number and not simply the larger total volume of apple (Hauser et al. 2000). Here too the parallel individuation system is supposed to be operative, since the monkeys only succeeded with smaller numbers. But because they chose the larger number even when this was arranged to produce a smaller total amount of apple, Carey and her colleagues suggest that the monkeys must have used number as a heuristic to obtain the larger amount. In other words, they constructed two mental models—one for the slices in each container—and compared them one-to-one, choosing the container that had slices with no corresponding match in the other container. Carey interprets other related work with infants to show that, under certain conditions, they use the parallel individuation system to perform assessments of numerical equivalence just like monkeys. In one study, which used small toys instead of food items, 12.5-month-olds were shown small numbers of toys placed in a box and were then given the opportunity to retrieve them (Feigenson and Carey 2003). The crucial measure was how much they would continue to search the box after a given number or amount had already been retrieved. Suppose infants saw three toys placed in the box. Would they subsequently search the box after seeing that only two were removed, or once one larger toy (equivalent in surface area to the sum of the three) had been removed? In this case, infants disregarded the size of the objects and continued to search when the number retrieved was less than the number that had been placed in the box. For Carey, this means that “the match must have been subserved by a computation of 1-1 correspondence” (Carey 2009, 142).

We have seen that Simon appeals to one-to-one correspondence in explaining how infants succeeded on Wynn’s addition and subtraction task. Presumably, proponents of the mental models approach would say much the same thing regarding the multimodal addition task mentioned in the previous section. The proposal would be that the infants succeeded on this task by constructing a mental model of the objects behind the screen, and that they introduced new model elements not only for the object that they saw but also for the objects they heard and had to infer were behind the screen. Then when the screen was removed, they compared the model in

working memory for one-to-one correspondence with the model of the visible outcome. When the two didn't correspond (in the $1+1+1=2$ condition), this mismatch was unexpected and caused them to look longer.

Finally, I will mention one last study that will be relevant when we compare the mental models approach to the SuM theory in the next section. This study was, in effect, an attempted replication of the Clearfield and Mix experiment mentioned in section 2 and hence a test for whether infants can only respond to the continuous properties of stimuli (Cordes and Brannon 2009). The basic experimental design, as before, was to habituate infants to two or three same-size squares, and then to show them both the same number with a new contour and the same contour with the old number. This time, with a larger number of subjects tested, seven-month-olds *did* respond to numerical changes as well as contour changes.¹ How would the mental models approach explain this result? The explanation would have two parts. The first is that infants constructed a mental model during the habituation phase of the experiment, and that each habituation trial reinforced the same two- or three-symbol model so that it was held in working memory. Second, in the test trials, infants constructed models for the alternating stimuli too, and compared these to the one held in memory. In the *new contour / old number* condition, they compared the sum of the continuous properties that were bound to the models' elements. This would have caused them to notice the contour change, hence the longer looking time relative to the habituation trials. In the *old contour / new number* condition, they compared the models for one-to-one correspondence. This would have caused them to notice that they don't match one-to-one, and would also have led to longer looking time relative to the habituation trials.

In this section, we have seen that the mental models approach to the small number system has a lot going for it. In contrast with the object-indexing approach, it can deal with cases where visual object tracking isn't possible or isn't relevant to the situation at hand. And from the point of view of its proponents, a further attraction of this approach is that it also doesn't postulate a domain-specific system to explain infants' numerical abilities. Later, in section 6, I will

¹ There was an important difference regarding how the two research teams analyzed their data. Clearfield and Mix limited the analyzed looking-time in each trial to ten seconds (in contrast with Cordes and Brannon's sixty seconds) and only included data from the first two test trials (in contrast with Cordes and Brannon using data from all of the test trials). As Cordes and Brannon point out, their own method of analysis is standard in the infancy literature and consequently preferable for making comparisons with other habituation studies.

challenge this assumption; I will argue that the mental models approach actually requires its own fair share of domain-specific structure. But first I want to return to the charge that the SuM theory is too speculative.

5. SuM Meets the Data

As noted earlier, many researchers hold that there is no evidence for an innate domain-specific system for representing small numerical quantities as such. In this section, I argue that, in fact, many findings are consistent with the SuM theory. As we will see, the SuM theory can explain most of the data regarding the way that infants and animals respond to small numbers of items, and it may even be more promising than the mental models approach in some cases. At the very least, it should be viewed as an open empirical question whether the small number system starts out as a system that can support assessments of one-to-one correspondence or as a system that represents particular numerical quantities.

Let's begin by looking at how the SuM theory explains the sorts of findings that were problematic for the object-indexing model. The key point, to begin with, is that the SuM theory *can* explain most of these data. Take the multimodal addition study with five-month-olds (Kobayashi et al. 2004). We saw that the mental models explanation of its key finding supposes that infants construct an abstract model that takes visual and auditory input, putting infants in a position to compare a model of the objects that are behind the screen for one-to-one correspondence with the model they form when the screen drops down. But the SuM theory also has an explanation of the disparity in looking time—in fact, a perfectly straightforward explanation. By hypothesis, SuM isn't a visual system. It can take auditory input too. When the first object is seen to be placed behind the screen, this initiates a spread of activation through SuM that triggers its “one” node (i.e., the node that functions to correspond with the presence of one item). But given the familiarization trials that help infants to recognize that these types of objects make a beep when they hit the ground, the subsequent beeps lead to further input to SuM and ultimately to the “two” or “three” node being triggered, depending on whether infants find themselves seeing/hearing 1+1 or 1+1+1. Of course, when the screen drops down, infants are in a position to see how many objects are actually there, and this would provide new input to SuM for the seen numerical quantity. If this quantity is identical to the remembered quantity, that isn't

surprising. But when the two differ—which would happen in the $1+1+1=2$ condition—it is surprising, and this would cause infants to look longer.

Next, consider Cordes and Brannon's (2009) study in which seven-month-olds discriminated between two and three squares. The main finding was that infants dishabituated to the novel number even though it had the same contour as the stimuli they had habituated to. The mental models approach explains this by claiming that infants compare a remembered model from the habituation trials to a model of the squares in each test trial and end up noticing when the two don't correspond one-to-one. But here too the SuM theory has a perfectly straightforward explanation of the looking-time pattern, or why the infants dishabituate to the novel number. Throughout the habituation trials, SuM is active and registers the presence of the numerical quantity *two* or *three*. Then when the novel number of squares is seen in the test trials, SuM registers a different numerical quantity, and it is the *numerical* difference that causes them to dishabituate.

Similar explanations can be given for the other studies, including the one where infants recognize that there are remaining toys in the box after some number of them have been removed. Recall that Carey says of this experiment that “the match *must* have been subserved by a computation of 1-1 correspondence” (italics added). She is certainly right that infants could be approaching this situation by constructing a model of the toys in the box and then comparing this model for one-to-one correspondence with a model of the toys that were removed. But again, there is a straightforward way for infants to determine that further objects remain in the box on the assumption that they can represent particular numerical quantities via SuM. As each toy is placed in the box, this provides input to SuM and ultimately leads to a representation of the numerical quantity of the concealed toys, say, the numerical quantity *three*. Later, the child can check whether the numerical value in memory is reproduced when SuM is directed to the objects as they are removed from the box, with each removed object being taken as input to SuM. If the remembered value isn't reproduced (e.g., if only two are removed, which wouldn't provide the needed input for the “three” node to be triggered), then this would indicate that the box hasn't been emptied, and infants would be motivated to continue to search the box.

The pattern we are seeing here is that the SuM theory is able to account for much the same data as the mental models theory. For most of the cases where infants or animals respond to

changes or differences in numerical quantity and where their response is confined to small numbers of items, they might be evaluating the stimuli for numerical equivalence using only assessments of one-to-one correspondence. But alternatively, they might be representing the specific numerical properties of the stimuli and noting changes or differences among these properties. There is one exception to this general rule, however. The SuM theory can't explain the cases where infants or animals respond to nonnumerical quantitative differences among the stimuli (e.g., the amount of cracker placed into a box). To the extent that nonnumerical quantitative assessments show the set-size signature of working memory, something like the parallel individuation approach will be needed to explain why performance is capped in the small-number range even when infants or animals aren't responding to numerical quantity. Still, this doesn't mean that when they do respond to numerical quantity that this can only be a matter of their representing matches and mismatches for numerical equivalence. They might represent nonnumerical quantity in certain contexts via an operation that sums the continuous properties that are bound to object representations, and numerical quantity via an operation that takes the activation of these object representations as its input and filters this information through SuM.

In short, both approaches that go beyond object-indexing can accommodate much the same data. One postulates a domain-general system that supports assessments of numerical equivalence; the other postulates a domain-specific system for representing particular small numerical quantities. Are there any further considerations that might help to tease them apart? One that may prove useful focuses on the source of infants' performance limitation—why success in these different tasks is restricted to small numbers of items. For the domain-specific proposal in which infants represent particular small numerical quantities, the performance limitation is, by hypothesis, one of SuM's design features. SuM is built to represent just a few numerical quantities. In contrast, the domain-general mental models proposal traces the performance limitation back to the capacity limit of working memory. The reason infants can only construct or work with models that have a small number of elements is that working memory imposes this constraint. Recall that Carey reports that working memory matures in the first year of life. One factor we can look at, then, is the developmental trajectory of working memory and whether this lines up with infants' performance on numerical tasks. If children are successful at discriminating between small numbers of items in a way that exceeds the immature

working memory system's capacity, this would cast doubt on the idea that their success turns on comparing working memory models for one-to-one correspondence.

This is an area where we need to proceed cautiously. Not a lot is known about the development of working memory in infancy, and there are questions about whether different researchers who study its development are studying the same thing—for example, there is a controversy about whether and how to distinguish working memory from short-term memory (Reznick 2014). Nonetheless, there is evidence of a correspondence across a variety of tasks in which infants' working memory goes through a developmental expansion in the first year of life in which, for older infants, it can accommodate three or four items, but at six months, it is limited to just one (Oakes and Luck 2014).² For example, in one study, infants were tested on whether they could remember the location of an occluded object by encoding its shape (e.g., whether they would recognize that two objects had been switched given that the shape that had been placed on the left subsequently appeared on the right, and vice versa). Nine-month-olds were able to do this for two objects, but it was found that six-month-olds were only able to do this for one object (Káldy and Leslie 2005).

Suppose we take at face value this pattern of findings and the provisional estimate that infants' working memory is severely limited half way through the first year of life. This would cast doubt on the hypothesis that such young infants are able to perform complex one-to-one comparisons between mental models with three elements. Now this isn't a problem for the mental models account of the study in which 12.5-month-olds succeeded in determining that a toy remained in the box after a certain number had been removed. These infants presumably have a working memory system that is mature enough to accommodate three or four items. But a number of the studies mentioned above had far younger subjects. The study in which infants dishabituated to changes in number as well as contour for two versus three static squares used infant subjects who were just seven months old. And the study in which infants successfully combined visual and auditory information for objects behind a screen (looking longer at the $1+1+1=2$ event than the $1+1=2$ event) were a mere *five months* old. The point is that that there

² In summarizing this research, Oakes and Luck remark that “multiple studies using very different paradigms suggest that young infants (e.g., six months) can retain only a single item in STM... Moreover, because the information is used to compare images before and after occlusion, find hidden objects, and so on, these results may reveal the nature of a WM [working memory] system” (2014, 171).

are grounds for supposing that infants at this young age should have difficulty comparing and applying models that require this many elements. On the other hand, the SuM theory has no difficulty in accounting for why these younger infants respond as they do. As each new object is registered to occur behind the screen, this provides further input to SuM, whose network adjusts which output node is active. The only thing infants have to remember when the occluder is removed is the numerical quantity that SuM has registered for the objects behind the screen. Then all they have to do is compare this one value to the current numerical quantity that SuM registers for the objects they can see. If the two values differ, this would be unexpected—hence the longer looking time.

To be clear, my claim isn't that we currently have decisive evidence against the mental models explanation. It's that this is a critical juncture where the SuM theory and the mental models theory make different predictions. The mental models theory ascribes the restriction to small numbers of items to the capacity limit on working memory. So if we can independently determine working memory's capacity limit at different ages, we might be able to show that infants have numerical abilities that exceed the cap predicted by the mental models approach. At present, there are some indications that infants *do* have numerical abilities of this kind, and this in turn is a reason to favor the SuM theory.

Of course, there were the objections to the SuM theory mentioned earlier (in section 2). One of these we have already dealt with, namely, the concern that SuM is a purely speculative theory and that there is no evidence that such a system exists. We have seen that there is actually much evidence for the theory. The problem is just that the evidence that supports the SuM theory also generally supports the mental models theory, and consequently we need to think hard about the types of findings that might favor one over the other.

The other objection was based on the observation that learning the meaning of words for small numbers takes a long time and that children invariably learn them in order. If SuM provides children with innate and perhaps unordered representations for the numerical quantities *one*, *two*, and *three*, why the difficulty and why don't children sometimes learn the meanings of these words in some other order? In contrast, the mental models approach explains the facts about language learning on the assumption that children have to construct a special stock of models that are regularly used for performing assessments of numerical equivalence and that

come to be associated with these number words. On this view, what happens when children learn the meaning of the word “two”, say, is that they figure out that it should be associated with a given model that is to be regularly used for performing one-to-one comparisons and that the word “two” applies to just those groups of items to which this special (two-membered) model stands in one-to-one correspondence. Given that these long-term memory models have to be constructed as children are confronted by the difficulty of having to interpret the meanings of number words, it isn’t surprising that this takes time. Also, since the model for “three” is more complex than the one for “two”, and the one for “two” is more complex than the one for “one”, it stands to reason that they will be learned in order, starting with the simplest model.

Do these facts about language learning discredit the SuM theory? Not at all. Even if there are innate representations for a few small numerical quantities, it doesn’t follow that learning the meanings of the number words for small numbers is a trivial matter. There is still a very challenging mapping problem in which children have to determine that number words (and certain morphological features in language) pick out numerical quantities to begin with. And once they recognize that their interpretation should focus on numerical quantity, there is a further question about which numerical quantity is the right one for a given term. This is a daunting problem even if children don’t have to construct new long-term memory models. What’s more, given that the linguistic data to which children have access vary enormously for these different terms—with “one” being far more frequent than “two”, and “two” far more frequent than “three”—any theory that recognizes the difficulty of the mapping problem that children face ought to predict that children are going to learn them in order.

To summarize, although it is widely thought that there is no evidence for innate representations of the numerical quantities *one*, *two*, and *three*, I have argued that we just have to look in the right place. In fact, there is a great deal of evidence that fits with this theory, albeit evidence that can be explained by the mental models approach as well. What’s more, the SuM theory may be in a better position to explain some of this data, since children’s successful performance on some tasks occurs at an age when there is a question about whether their working memory system would be mature enough to handle the needed complex models. But suppose that we put this last point to the side. Suppose that younger infants’ working memory is capable of dealing with these models and that the SuM theory and the mental models theory are

on a par regarding the data. Many theorists with empiricist leanings would conclude that we should reject the SuM theory in this situation on the grounds that, all things being equal, purely domain-general accounts are simpler and hence better developmental theories. I reject this principle. I don't see any reason to suppose that there are general methodological grounds for preferring domain-general theories over domain-specific theories in accounting for early developing cognitive capacities. But I won't argue for this claim here. Instead, what I propose to do in the next section is to argue that such methodological considerations are beside the point because, in the end, the mental models approach has to accept a significant amount of innate domain-specific structure too.

6. Where Does One-to-One Correspondence Come From?

According to the mental models theory, the small number system is fundamentally a domain-general system that supports assessments of one-to-one correspondence. Let's assume for the sake of argument that younger infants are able to compare the needed mental models for one-to-one correspondence and that this explains their success on the sorts of tasks reviewed above. Still, there is a question about the innate structure of the component systems that underlie this ability; it shouldn't be assumed that the mental models approach invariably vindicates a domain-general basis for the representation of small numbers of items. One way to see that a fully domain-general account is problematic is to ask *why* children so readily compare mental models for one-to-one correspondence and *how* they manage to reliably perform one-to-one comparisons.

Take the *why* question first. The point is that it is one thing for a certain form of representation to be able to support assessments of one-to-one correspondence and quite another for an agent to recognize the value of carrying out the assessment and to spontaneously perform these comparisons. If the system that underlies these operations is supposed to be a general-purpose system—one that isn't geared towards numerical representation in particular—where would children even get the idea that they can determine whether two groups of items have the same number by settling whether there is exactly one item in the first group for every item in the second? It may be obvious to numerate adults that this is a good technique for deciding whether two groups are numerically equivalent, but it isn't a self-evident procedure, one that would

necessarily occur to any agent who happens to have a general-purpose capacity for constructing mental models.

In his discussion of why his mental models approach is a “nonnumerical” theory, Simon argues that each of the fundamental competencies that it requires isn’t specific to the domain of number and that there is independent evidence that infants have the competence. For example, he mentions that infants have mechanisms for remembering what they have seen and for generalizing without particular regard to an object’s perceptual details. But throughout his discussion, he never asks about the origins of the process that checks for one-to-one correspondence, as if it should go without saying that once the other capacities are in place, his work is done. But it isn’t done because there is still the matter of what would drive infants to compare a model in working memory to what they currently see for whether they match in this way. In developmental psychology, it has often been thought that a general understanding of one-to-one correspondence must be learned and that it takes *years* to develop through observations and activities where objects are paired with one another (candies paired with containers, forks paired with napkins, etc.). Mix et al. trace a developmental trajectory in which recognition of numerical equivalence begins when children are around 2.5 years old and expands in the preschool years as children first become able to match items that have somewhat different features, followed by heterogeneous items in a single modality, followed by crossmodal matches between such disparate items as dots and sounds (2002, 39). But if five-month-olds are supposed to already have a system in place that compares abstract models for one-to-one correspondence, these types of experiences can’t be essential to the capacity to establish numerical equivalence. In fact, it’s hard to see what alternative there is to a system that incorporates innate operations for determining numerical equivalence. This would be a system that isn’t limited to creating models or to using these models to perform nonnumerical quantitative comparisons, but a system that is designed, in part, for making judgments about numerical equivalence.

The *how* question leads to the same conclusion. Assuming that younger infants do manage to use one-to-one correspondence to make judgments of numerical equivalence, how are they able to reliably map each and every item in one model to just one in another? This isn’t an easy procedure to execute. Notice that it requires a form of bookkeeping when moving back and forth between the two models, so that, among other things, no item is fed into the process multiple

times and the process stops just when it should. A similar issue comes up as children learn to count, which is a comparatively demanding process. To achieve an accurate count, children need a way of keeping track of which counting terms have been used and which items have already been tagged—something young children find to be difficult and that is learned with the aid of adult training and much cultural support. If infants as young as five months old can reliably perform one-to-one comparisons without comparable aid and support, we have reason to believe that the system that underlies this capacity includes innate operations for just this purpose. What we are looking at is no longer a general-purpose representational system, or even a system that is designed for making quantitative comparisons. It is a system that is designed for making *numerical* comparisons.³

The animal data only reinforce this way of thinking about the small number system. Recall that honeybees can be trained to discriminate between instances of two and three and that their performance is limited to the small number range.⁴ A very natural explanation of how they do this is that they detect particular small numerical quantities via SuM—for example, they learn to choose the exit with three figures if they see three other figures when they first enter the apparatus. Still, another possibility is that they succeed by determining which of the choice stimuli matches the remembered sample when they are compared for one-to-one correspondence, that is, by choosing the numerically equivalent stimulus rather than the one that also has three elements. However, if this is how they succeed on the task, it is tremendously unlikely that individual bees have learned to determine whether two stimuli are numerically equivalent by drawing on a general-purpose capacity for working with mental models. The only plausible way of developing this account is to hold that the system that implements these processes is an innate system for making judgements of numerical equivalence.⁵

³ A reviewer has questioned whether the *how* question argues for a domain-specific system, noting that a similar difficulty arises when infants compare small collections for nonnumerical quantity. For example, when comparing two groups of crackers, infants couldn't reliably choose the one with more surface area if they didn't have a way of keeping track of whether an item's surface area has already been incorporated into the overall sum. But notice that, with nonnumerical quantity comparisons, infants don't have to map items from one group to the other. They can compute each group's total surface area independently of the other and then simply compare the two values. This considerably eases the demands on bookkeeping and suggests that a different type of process underlies nonnumerical comparisons.

⁴ For related work on the representation of small numerical quantity in newborn chickens, see Rugani et al. (2008, 2010).

⁵ Of course, it is possible that bees and humans employ different types of cognitive mechanisms in their dealings with small numbers of items. As a reviewer has noted, bees might "have a domain-specific mechanism, but more

Recall that Carey characterizes the parallel individuation system as a domain-general system by emphasizing that it serves multiple functions in which mental models support inferences about spatial, causal, and intentional relations, as well as assessments about nonnumerical quantity and numerical equivalence. At the same time, the architecture she is proposing is flexible enough that the systems that drive any of these inferences can themselves be domain-general or domain-specific. What I am suggesting is that *if* one-to-one correspondence accounts for the sorts of findings that might otherwise be explained by SuM, the early developing facility with one-to-one correspondence argues for *an innate numerical comparator* that interacts with the domain-general capacity for forming mental models. If I am right, then the main choice isn't between one theory that postulates innate domain-specific structure and another that postulates only domain-general structure. It is a choice between two theories that postulate different types of innate domain-specific structure—one that is committed to an innate system for representing particular numerical quantities (SuM) and one that is committed to an innate system for performing assessments of numerical equivalence (an innate comparator). To be sure, these are different approaches to the fundamental structure of the small number system, but either way, the representation of small numerical quantity would be grounded in an innate domain-specific system. There is no domain-general alternative.

7. Conclusion

I have argued that more consideration should be given to the proposal of an innate system for representing a few precise small numerical quantities—the SuM theory. I began by showing that we need a richer form of representation than a system for visually tracking small numbers of objects. I then went on to show that the SuM theory does well when compared to the alternative proposal that the small number system starts out as a general-purpose capacity for working with mental models. Not only can the SuM theory explain the majority of the data that the mental models theory has been claimed to explain, but it also has a potential advantage over the mental

sophisticated creatures like us use a domain-general mechanism". While this is a possibility, there is a remarkable similarity across a broad range of species (including humans) regarding systems for representing core facets of space, time, and number (Dehaene and Brannon 2011). This suggests that the representation of small number is also likely to be similar in humans and other animals. In any case, my reference to the bee data and how it is best interpreted is not meant to settle the matter in favor of SuM. It's just one part of an overall inference to the best explanation regarding the nature of the human capacity for representing and responding to small numbers of items.

models approach when the capacity limit on infants' working memory is fully taken into account. Moreover, the mental models theory itself requires innate domain-specific structure that its proponents have failed to recognize or acknowledge, namely, an innate capacity for performing assessment of one-to-one correspondence. So regardless of whether the SuM theory or the mental models theory is accepted in the end, the representation of small numbers of items requires a considerable amount of innate domain-specific structure.

Acknowledgments

I would like to thank Stephen Laurence, Gerardo Viera, and two anonymous reviewers for valuable comments on earlier drafts of this article. This research was supported by the Social Sciences and Humanities Research Council of Canada.

References

- Barner, David. 2017. "Language, Procedures, and the Non-Perceptual Origin of Number Word Meanings." *Journal of Child Language* 44.3: 553-90.
- Carey, Susan. 2009. *The Origins of Concepts*. New York: Oxford University Press.
- Clearfield, Melissa W., and Kelly S. Mix. 1999. "Number Versus Contour Length in Infants' Discrimination of Small Visual Sets." *Psychological Science* 10.5: 408-11.
- Cordes, Sara, and Elizabeth M. Brannon. 2009. "The Relative Salience of Discrete and Continuous Quantity in Young Infants." *Developmental Science* 12.3: 453-63.
- Dehaene, Stanislas, and Elizabeth Brannon, eds. 2011. *Space, Time, and Number in the Brain: Searching for the Foundations of Mathematical Thought*. New York: Academic Press.
- Feigenson, Lisa, Susan Carey, and Marc Hauser. 2002. "The Representations Underlying Infants' Choice of More: Object Files Versus Analog Magnitudes." *Psychological Science* 13.2: 150-56.
- Gross, Hans J., Mario Pahl, Aung Si, Hong Zhu, Jürgen Tautz, and Shaowu Zhang. 2009. "Number-based Visual Generalisation in the Honeybee." *PLoS ONE*, 4.1: 4263.
- Hauser, Marc D., Susan Carey, and Lilan B. Hauser. 2000. "Spontaneous Number Representation in Semi-free-ranging Rhesus Monkeys." *Proceedings of the Royal Society of London B: Biological Sciences*, 267.1445: 829-33.
- Hurford, James R. 2001. "Languages Treat 1-4 Specially." *Mind and Language*, 16.1: 69-75.
- Káldy, Zsuzsa, and Alan M. Leslie. 2005. "A Memory Span of One? Object Identification in 6.5-month-old Infants." *Cognition*, 97.2: 153-77.
- Kobayashi, Tessei, Kazuo Hiraki, Ryoko Mugitani, and Toshikazu Hasegawa. (2004). "Baby Arithmetic: One Object Plus One Tone." *Cognition*, 91.2: B23-B34.

- Laurence, Stephen, and Eric Margolis. 2005. "Number and Natural Language. In *The Innate Mind, vol. 1: Structure and Contents*, ed. Peter Carruthers, Stephen Laurence, and Stephen Stich, 216-35. New York: Oxford University Press.
- Leslie, Alan M., Fei Xu, Patrice D. Tremoulet, and Brian J. Scholl. 1998. "Indexing and the Object Concept: Developing 'What' and 'Where' Systems." *Trends in Cognitive Sciences*, 2.1: 10-18.
- Margolis, Eric, and Stephen Laurence. 2008. "How to Learn the Natural Numbers: Inductive Inference and the Acquisition of Number Concepts." *Cognition*, 106.2: 924-39.
- Mix, Kelly S., Janellen Huttenlocher, and Susan C. Levine. 2002. *Quantitative Development in Infancy and Early Childhood*. New York: Oxford University Press.
- Oakes, Lisa M., and Steven J. Luck. 2014. Short-term Memory in Infancy. In *The Wiley Handbook on the Development of Children's Memory*, ed. Patricia J. Bauer and Robyn Fivush, 157-80. Oxford: John Wiley & Sons.
- Pylyshyn, Zenon W., and Ron W. Storm. 1988. "Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism." *Spatial vision*, 3.3: 179-97.
- Reznick, S. J. (2014). "Methodological Challenges in the Study of Short-Term Working Memory in Infants." In *The Wiley Handbook on the Development of Children's Memory*, ed. Patricia J. Bauer and Robyn Fivush, 181-201. Oxford: John Wiley & Sons.
- Rugani, Rosa, Lucia Regolin, and Giorgio Vallortigara. 2010. "Imprinted Numbers: Newborn Chicks' Sensitivity to Number vs. Continuous Extent of Objects They Have Been Reared With." *Developmental Science*, 13.5: 790-97.
- Rugani, Rosa, Lucia Regolin, and Giorgio Vallortigara. 2008. "Discrimination of Small Numerosities in Young Chicks." *Journal of Experimental Psychology: Animal Behavior Processes*, 34.3: 388.
- Scholl, B. J., & Leslie, A. M. 1999. "Explaining the Infant's Object Concept: Beyond the Perception/Cognition Dichotomy." In *What Is Cognitive Science?*, ed. Ernest Lepore and Zenon Pylyshyn, 26-73. Oxford: Blackwell.
- Simon, Tony J. 1997. "Reconceptualizing The Origins of Number Knowledge: A "Non-Numerical" Account." *Cognitive Development*, 12.3: 349-72.
- Spelke, E. S. (2003). "What Makes Us Smart? Core Knowledge and Natural Language." In *Language in Mind: Advances in the Study of Language and Thought*, ed. Dedre Gentner and Susan Goldin-Meadow, 277-311. Cambridge, MA: MIT Press.
- vanMarle, Kristy, Felicia W. Chu, Yi Mou, Jin H. Seok, Jeffrey Rouder, and David C. Geary. 2016. "Attaching Meaning to The Number Words: Contributions of the Object Tracking and Approximate Number Systems." *Developmental Science*, 21.1: e12495[1-17].
- Wynn, Karen. 1992. "Addition and Subtraction by Human Infants." *Nature*, 358.6389: 749.
- Xu, Fei, Elizabeth S. Spelke. 2000. "Large Number Discrimination in 6-Month-Old Infants." *Cognition*, 74.1: 1-11.
- Xu, Fei, Elizabeth S. Spelke, and Sydney Goddard. 2005. "Number Sense in Human Infants." *Developmental Science*, 8.1: 88-101.