Do numbers begin noisy?

Do children's number words begin noisy?

Katie Wagner, Junyi Chu, & David Barner

University of California, San Diego

Word Count = 7999

CONTACT:

Katie Wagner

Department of Psychology

9500 Gilman Drive, La Jolla, CA 92093-0109

kgwagner@ucsd.edu

Research Highlights

- New analyses of standard number comprehension tasks explain conflicting findings in prior number word learning literature
- Evidence of preliminary, noisy, meanings beyond children's "knower level"
- No evidence for the use of approximate number representations in children's early number word learning

Abstract

How do children acquire exact meanings for number words like *three* or *forty-seven*? In recent years, a lively debate has probed the cognitive systems that support learning, with some arguing that an evolutionarily ancient “approximate number system” drives early number word meanings, and others arguing that learning is supported chiefly by representations of small sets of discrete individuals. This debate has centered around findings generated by Wynn’s (1990, 1992) Give-a-Number task, which she used to categorize children into discrete “knower level” stages. Early reports confirmed Wynn’s analysis, and took these stages to support the “small sets” hypothesis. However, more recent studies have disputed this analysis, and have argued that Give-a-Number data reveal a strong role for approximate number representations. In the present study, we use previously collected Give-a-Number data to replicate the analyses of these past studies, and to show that differences between past studies are due to assumptions made in analyses, rather than to differences in data themselves. We also show how Give-a-Number data violate the assumptions of parametric tests used in past studies. Based on simple non-parametric tests and model simulations, we conclude that (1) before children learn exact meanings for words like *one*, *two*, *three*, and *four*, they first acquire noisy preliminary meanings for these words, (2) there is no reliable evidence of preliminary meanings for larger meanings, and (3) Give-a-Number cannot be used to readily identify signatures of the approximate number system.

Introduction

How do children acquire exact meanings for number words like *three* or *forty-seven*? In recent years, a lively debate has probed the cognitive systems that support number word learning, with some arguing that an evolutionarily ancient “approximate number system” (ANS) defines early number word meanings, and others arguing that learning is supported chiefly by representations of small sets of discrete individuals. On the first view (Gallistel & Gelman, 1992; Leslie et al., Wagner & Johnson, 2011; Wynn, 1992), number word meanings are acquired via the association of the count list (e.g., the words *one*, *two*, *three*, *four*, etc.) to states of the ANS, which spontaneously represents approximate magnitudes of sets in accordance with Weber’s law, whereby error is proportional to the size of the set being represented (Dehaene, 1997; Meck & Church, 1983; Gallistel, 1990). On the second view, only the smallest words – e.g., *one*, *two*, *three*, *four* – are learned associatively by young children, drawing on representations of individuals and sets (Carey, 2009). On this second view, meanings of larger number words are acquired via a sweeping inductive inference, which specifies a general logic for counting (Carey, 2009; Le Corre & Carey, 2006; Sarnecka & Carey, 2008).

Much of this debate has centered around findings generated by Wynn’s (1990, 1992) Give-a-Number task (Give-N), which she used to categorize children into “knower levels” by asking them to give different numbers of objects and grouping them according to their patterns of responses (Frye et al., 1989; Schaeffer, Eggleston, & Scott, 1974). Give-N has been used across many cultures (Almoammar et al., 2013; Barner, Chow, & Yang, 2009; Le Corre, Li, & Carey, 2016; Piantadosi et al., 2014; Marusic et al., 2016; Sarnecka et al., 2003), and has repeatedly found evidence for a stage-like acquisition of exact number word meanings. In the US, children generally begin by learning to recite a small subset of number words in sequence

(*one*, *two*, *three*, etc.), in tandem with learning procedures that involve pointing at objects in sequence (Gelman & Gallistel, 1986; Fuson, 1988, 1991; Wynn, 1990, 1992). However, initially they do not know the meanings of these words (Baroody & Price, 1983; Briars & Siegler, 1984; Frye et al., 1989; Fuson, 1988; Le Corre & Carey, 2007; Wynn, 1990, 1992). Based on Wynn's criteria, these children are classified as "non-knowers". Not long after, children learn an exact meaning for *one* and become "one-knowers" – they give one object when asked for *one* but not when asked for other numbers. After a long delay – around six to nine months – children learn an exact meaning for *two* and are called "two-knowers". And after another delay, they learn *three* ("three-knowers"), and sometimes *four* ("four-knowers"), in sequence. After passing through these "subset-knower" stages, children appear to discover a general principle about counting: That the count list they began with can be used to label and generate larger sets. Having noticed this, children become able to give any number within their count list, and thus can map number words to large precise cardinalities. These children are sometimes referred to as Cardinal Principle knowers or "CP-knowers", though the precise nature of their knowledge at this stage remains uncertain (Barner, 2017; Cheung, Rubenson, & Barner, 2017; Davidson, Eng, & Barner, 2012; Le Corre & Carey, 2007; Lipton & Spelke, 2005; Sarnecka & Carey, 2008; Wagner et al., 2015).

Since Wynn's initial description of these stages, several studies have debated their significance. According to Le Corre and Carey (2007), Wynn's data provide strong evidence that children construct early number word meanings from small sets, rather than from the ANS. To support this, they tested children with a speeded version of Gelman's (1993) "What's on this Card?" task, called "Fast Cards" and found that most children provided random responses when estimating the cardinality of large arrays of dots. Further, even many CP-knowers – who were

able to count and provide large sets upon request – responded randomly for large numbers.

Corroborating these findings, a later study by Sarnecka and Lee (2009) argued that subset knowers do not have knowledge of numbers beyond their knower level. Specifically, Sarnecka and Lee reasoned that if children had such meanings, then we should expect that their errors on the Give-N task would be non-random. To test this, they analyzed a large Give-N dataset and found that the numerosity of children's errors for numbers beyond their knower level did not grow larger as a function of the size of the requested number (see also Lee & Sarnecka, 2010, 2011).

Others have challenged these conclusions and have argued that the Give-N task can reveal knowledge well beyond children's Wynn knower levels. In some cases, these studies have found relatively modest knowledge, limited to the next number – i.e., “Knower Level +1” or KL+1 (Barner & Bachrach, 2010; Gunderson et al., 2015).¹ In other cases, however, more robust knowledge has been reported. For example, Wagner and Johnson (2011) tested children with an adaptation of Wynn's Give-N task and found that children's mean responses for KL+1 to 10 exhibited a significant positive slope. Also, they reported a key signature of the ANS: Scalar variability. According to their analysis, the variability in children's responses grew as a function of the mean number they provided, such that error was larger for larger numbers. Together, these findings led Wagner and Johnson to conclude “that children who have not mastered the cardinal word principle (“non-CP-knowers”) show the signatures of scalar variability directly *within* their mappings of number words beyond their proficiency level.” (p. 18). Thus, on their analysis, the ANS may indeed supply the conceptual foundation of children's early number word meanings.

¹ Gunderson et al. (2015) found that children had knowledge of KL+2 when tested with Fast Cards. However, this result was not found in the same sample using the Give-N task, nor in a second, younger, sample using Fast Cards.

Resolving the differences between these studies is challenging for a variety of reasons. First, as noted by Gunderson et al. (2015), not all studies that employ Give-N use Wynn's criteria to classify children, making comparison of datasets difficult. According to Wynn, an N-Knower must correctly give n objects $2/3$ times in response to a request for n , but also it must be the case that $2/3$ of cases in which the child gives n are in response to requests for n , thus ensuring that a bias to give, e.g., three objects, does not result in classification as a three-knower. In some studies, such as Wagner and Johnson (2011), only the first criterion is used, thus potentially overestimating children's knowledge.² Second, in some studies, like Wagner and Johnson's, no distinction is made between small numbers (1-4) and large numbers (5+) when conducting analyses, a factor which may be important to diagnosing the relative roles of the different representational systems. For example, as noted by Gunderson et al. (2015), although Wagner and Johnson found positive slopes for numbers from KL+1 to 10, these analyses included small numbers (< 5), which might be defined by small set representations. Finally, in addition to conducting different analyses of Give-N data, previous studies have also sometimes omitted certain data from consideration. For example, in their study, Sarnecka and Lee (2009) omitted all correct responses from analyses, which, as we show below, underestimates children's knowledge.

In the present study, we revisited the Give-N task and found four important results. First, we replicated previous analyses using a single dataset, and found that past studies differed because of how they analyzed their data, rather than differences in the data themselves. Second, we conducted non-parametric analyses, which are more appropriate for non-normal Give-N data,

² Similarly, Odic et al. (2015) reported better knowledge of large numbers using a tapping task, even relative to Fast Cards (though see Schaeffer et al., 1974), but did not bin children according to Wynn's levels.

and we included both correct and incorrect responses. These analyses found evidence for preliminary meanings for numbers as high as KL+2, which was present at all knower levels except three-knowers, who exhibited only knowledge of KL+1. Critically, we did not find reliable evidence for preliminary meanings in the large number range, which we tested using the logic of previous studies – i.e., by asking whether responses to large numbers like *ten* differed from responses just beyond a child’s knower level. Unlike past studies, however, we employed a variety of different analytic methods to assess number word knowledge. For example, to test whether children responded differently to *five* and *ten*, we analyzed the slope of children’s mean responses, differences in their responses from chance, and, given the non-normal nature of Give-N data, differences between median responses. Fourth, using a series of simulations, we found that typical Give-N data cannot readily assess the primary signature of the ANS (i.e., scalar variability) because of the restricted range of possible responses on the task (i.e., a lower bound of the child’s knower level, and an upper bound imposed by the total number of objects they are free to give). This, we show, is true even when modifications are made to the task to increase its ability to detect such signatures (e.g., giving the child twice as many objects to choose from).

In sum, although children acquire numbers in clear and distinct stages as described by Wynn, we find that the Give-N task is not sensitive to the full extent of children’s knowledge, and that many have preliminary knowledge of numbers beyond their Wynn knower level. At the same time, multiple measures fail to find reliable evidence that children have preliminary, approximate, meanings for larger number words.

Study 1

Study 1 combined Give-N data from two previous studies to replicate past findings, ask why they reached different conclusions, and to model data using non-parametric analyses.

Specifically, we asked whether children have knowledge beyond their Wynn knower level, and if so, whether this extends into the large number range and exhibits scalar variability, a key signature of the ANS.

Method

Participants

We combined data from a study of English monolinguals (Almoammer et al., 2013) and English data from a study of bilinguals (Wagner et al., 2015), resulting in a total sample size of 73 children. Only subset knowers were included from each study. Monolinguals were 44 children aged 2;0 to 4;3 years ($M=3;1$) while bilinguals spoke either English and French ($n=14$, $M=3;4$, range=2;4–4;0) or English and Spanish ($n=15$, $M=3;8$, range=2;10–4;11). Only English-dominant bilinguals (i.e. those with a higher knower-level in English) were included, to ensure that stronger number knowledge in a second language did not affect performance in English). Preliminary analyses revealed no differences between the bilinguals and monolinguals in our data set. Thus, these samples were combined for all reported analyses. Also, as described below, we replicate all previous analyses with these combined data, suggesting that they do not differ from other Give-N samples.

Procedure

Give-N Task. This task was adapted from Wynn (1990). In each of the studies from which our data are drawn, the experimenter began by presenting the child with a plate and ten similar objects. For each trial, the experimenter asked the child to place a quantity on the plate, “Can you put N on the plate? Put N on the plate and tell me when you’re all done.” Singular-plural marking was avoided to ensure that children’s responses were not based on this information, rather than knowledge of the number words themselves. As in Wynn’s (1990) original procedure, once the child responded, the experimenter asked, “Is that N ? Can you count

and make sure?” and encouraged the child to count. If the child recognized an error, the experimenter allowed the child to change their response.

Children completed up to twenty-one quasi-randomized trials, consisting of three trials for each of the seven number words tested (*one, two, three, four, five, eight, and ten*). Children were defined as an n -knower (e.g., three-knower) if they correctly provided n (e.g., three fish) on at least two out of the three trials that n was requested and, of those times that the child provided n , two-thirds of the times the child did so it was in response to a request for n . If n was *five* or higher, the child was classified as a CP-knower.

Highest count. After the Give-N task, the experimenter asked each child, “Can you count as high as you can?” to measure the child’s exposure to number words.

Results

Knower-Level Classification

We identified 28 non-knowers, 15 one-knowers, 16 two-knowers, 11 three-knowers, and 3 four-knowers (Table 1). There were 3 four-knowers who were excluded from analyses both because they were few in number, and more importantly because in the Give-N task they did not receive trials testing their knowledge of *six* (KL+2), which was critical to multiple analyses. Note that, in Study 1, we do present data according to knower level since such analyses require larger samples. In Study 2 we examine whether effects are found at each knower level by combining data with a newly collected dataset.

Table 1: Knower Levels of Monolingual and Bilingual Children

	Variable				
	Non-knower	1-knower	2-knower	3-knower	4-knower
Monolingual (n=44)	25 ¹	4	9	3	3
Bilingual (n=29)	3	11	7	8	0

¹Only a subset of the non-knowers were included in our main analyses. See section *Division of non-knowers* for more details.

Division of non-knowers. While many non-knowers likely know nothing about the meanings of number words, it is possible that some have begun learning meanings for small number words like *one* and *two*. However, not all non-knowers exhibit sufficient variability to allow a test of this possibility, since they often give the identical response across all trials. Therefore, to explore whether any non-knowers have knowledge beyond their knower level, we divided them into two groups. The first group ($n=16$) included non-knowers who gave the same number of objects in response to every request, and whose data could therefore not exhibit sufficient variability to allow tests of preliminary number word meanings (i.e., whether variability was random or non-random). The second group ($n=12$) included children who provided at least two unique responses – i.e., they exhibited enough variability that they might plausibly provide evidence of knowing at least one number. For the remainder of this report, we refer to the second group as “variable non-knowers”.

Table 2: Means of Incorrect Responses

Requested Number	Mean of Incorrect Responses	N children with one or more Incorrect Responses	Comparison to Target
1	6.44	14	$t(13) = 6.3, p < 0.01$
2	6.73	19	$t(18) = 6.5, p < 0.01$.
3	5.10	40	$t(39) = 4.3, p < 0.01$
4	5.20	50	$t(49) = 3.5, p = 0.01$
5	5.50	53	$t(52) = 0.49, p = 0.63$
8	4.93	54	$t(53) = 8.7, p < 0.01$
10	4.34	47	$t(46) = 17, p < 0.01$

Replication of Sarnecka and Lee (2009)

Sarnecka and Lee (2009) performed several analyses that supported the conclusion that children’s errors beyond their Wynn knower level are random. Their first analysis examined the means of children’s errors. Replicating their results, we found that the means of children’s

incorrect responses bore no reliable relationship to correct responses (Table 2 and Figure 1a). For children at all knower levels, the mean incorrect response was significantly different from the correct response for all but one target number, *five* (likely because means for all requested numbers were close to 5). Also replicating their results, we found no positive correlation between the mean of children's errors and the requested number. In fact, the relationship was negative ($r(6)=-0.84$, $p=0.018$; Figure 1a). This suggests that as requests grow larger, children give *fewer* items – a counterintuitive result, which, as we explain in more detail below, is due to Sarnecka and Lee's removal of correct responses from their dataset (since removing correct responses reduces children's mean response for large numbers and increases it for small numbers). While Sarnecka and Lee (2009) did not find a significant negative relationship between the mean incorrect response and the requested number, they did observe that the mean incorrect response for *one* was higher than for other numbers.

Figure 1: Responses for each knower level, using different summary statistics. (a) Mean incorrect responses, (b) Mean responses, and (c) Median responses. There are no data points for 3-knowers for requests of *one* or *two* because for these numbers they always provided requested amounts. All non-knowers are variable non-knowers.

Replication of Wagner & Johnson (2011)

Although our first analyses replicated Sarnecka and Lee (2009), their practice of excluding correct responses introduces systematic bias that underestimates knowledge. First, it assumes that correct responses do not provide information regarding children's knowledge. Second, it assumes that random responding should result in errors that exhibit a flat slope for all numbers beyond a child's knower level (i.e., errors should not be positively correlated with the size of the requested number). In fact, this assumption is only valid if all responses, both correct and incorrect, are included. Critically, if random responding is modelled as a flat distribution across all numbers, removing the responses that happen to be correct alters children's predicted responses, such that a *negative* slope is expected from random guesses (Figure 2). This is because the size of correct responses removed from analyses – whether due to random guessing or not – necessarily increases with the size of the requested number. For example, removing correct responses for a request for *ten* decreases the mean of the responses while removing correct responses for *one* increases the mean of the responses (see Figure 2).

Figure 2. Means predicted from random responding. This figure plots what children's responses would look like if they were independent of the requested numbers. The dashed line represents the observed mean of all responses in our data set. The solid line represents the mean when correct responses are removed for 1,2,3,4,5,8, and 10.

To further explore the effect of excluding correct responses, we next replicated the analyses of Wagner and Johnson (2011), who included both correct and incorrect responses. We first asked whether the slope of mean responses above knower level was positive, and then conducted the same analysis for large numbers, above *four* (see Table 3; see also Gunderson et al., 2015). While Wagner and Johnson ascribed knowledge of a number, n , to children who correctly gave n 2/3 times, we classified children using Wynn's more conservative knower level classification, which also requires that 2/3 of cases in which the child gives n are in response to requests for n . Importantly, because this method is more conservative, it increases the likelihood of finding knowledge beyond a child's knower level, and thus is more likely to find results compatible with those reported by Wagner and Johnson. This first analysis replicated Wagner and Johnson, while also confirming that the removal of correct responses systematically underestimates children's knowledge (Figures 1a & 1b). For subset knowers but not for variable

non-knowers, responses exhibited a positive slope from KL+1 to 10, ($M=0.13$, $p=0.0031$, $d=0.48$; see Table 3).

Thus far, we have shown that our data can be used to replicate both Sarnecka and Lee (2009) and Wagner and Johnson (2011). This suggests that these studies found different results chiefly because of how they treated correct responses. However, these results cannot yet speak to the role of the ANS in supporting number word meanings. This is because, as argued by Gunderson et al.'s (2015) discussion of Wagner and Johnson (2011), a positive slope from KL+1 to 10 can be driven entirely by knowledge of KL+1 (and an assumption that larger numbers than this must denote larger quantities – i.e., mutual exclusivity). Given this, like Gunderson et al. (2015), we asked whether children's responses exhibited positive slopes for numbers between *five* and *ten* and from KL+2 to *ten*. Replicating Gunderson et al. (2015), we found no positive slope for subset knowers from *five* to *ten* ($M=0.10$, $p=0.12$, $d=0.25$). However, unlike Gunderson et al., we did find that subset knowers (but not variable non-knowers) exhibited a positive slope from KL + 2 to *ten* ($M=0.11$, $p=0.037$, $d=0.33$).³ Thus, using parametric tests like those employed by previous studies in the literature, we replicated the patterns of results reported by past studies using their own analyses. Also, novel to this study, we found evidence that children's responses exhibit a positive slope beyond what is reported by Sarnecka and Lee or Gunderson et al., but, contrary to the conclusions of Wagner and Johnson, we found no evidence that this extends to numbers larger than *four*. Consistent with Gunderson et al., this suggests that the positive slopes reported by Wagner and Johnson were driven by knowledge of small number

³ Although Gunderson et al. (2015) did not find evidence for a positive slope from KL+2 to *ten* in Give-N, they did find such an effect in their Fast Cards task (i.e., the speeded version of What's on this Card?), but only in one of two populations tested.

words and an ability to infer that unknown, higher number words must denote larger sets than known number words.

Table 3. Study 1 mean slopes of subjects' responses for KL+1 *to ten*, KL+2 *to ten*, and *five to ten*

	N	KL+1 to Ten			KL+2 to Ten			Five to Ten		
		M slope	SD	<i>p</i>	M slope	SD	<i>p</i>	M slope	SD	<i>p</i>
Variable Non-knowers	9 ¹	-0.01	0.39	0.84	-0.07	0.35	0.58	-0.12	0.66	0.60
All Subset Knowers	42	0.13	0.26	<0.01	0.11	0.32	0.04	0.11	0.44	0.12

¹Three additional variable non-knowers were excluded for failing to provide responses to *ten*.

Non-parametric Analyses

Thus far we have found evidence that children have knowledge beyond their knower levels, but have failed to find support for the idea that this extends to larger numbers (e.g., beyond *three* or *four*). These results largely replicate the conclusions of past studies. However, there are reasons to worry that these analyses might not provide a valid estimate of children's knowledge. In almost all previous studies of the Give-N task, analyses of numbers beyond a child's knower level have comprised parametric comparisons of the means, which assume a normal distribution. This practice is potentially misleading because although means provide an excellent characterization of central tendency for normally distributed data, children's responses on the Give-N task violate normality assumptions in important ways, as shown in Figure 3, which depicts the responses of two-knowers for requests of *four* and *five*.

Figure 3: An example of non-normal data from the Give-N task (two-knowers' responses for requests for *four* and *five*).

Give-N data predictably deviate from the normal distribution for two reasons (see also Lee & Sarnecka, 2010, 2011). First, as seen in Figure 3, children frequently give all available items (e.g., ten out of ten) when they are uncertain of how to respond. Second, children's responses are both upper- and lower-bounded. They are upper-bounded by the number of objects the experimenter presents to the child: If the child is offered ten objects, then they can only give up to ten, even if they wish to give more. Responses are lower bounded by the child's knower level: If a child is a two-knower then responses to *three* necessarily have a lower bound of three items. This is because the criteria for two-knower classification requires that, 2/3 of the time, a child gives two objects when asked for *two*, and further, gives two objects only when asked for *two* but not when asked for other numbers. Thus, the lower bound is given by how we choose to classify children's knowledge.

To address these issues, we conducted two sets of non-parametric analyses. First, we analyzed the same data using children's median responses, rather than means, as a non-parametric measure of central tendency that does not assume normality. Second, we conducted analyses comparing responses to chance, in which chance was redefined in light of the truncated range of possible responses on Give-N as well as children's biases to provide certain responses more frequently than others.

Table 4. Study 1 Sign Tests of Median Responses

	<i>N</i>	KL+1 to <i>Ten</i>				KL+2 to <i>Ten</i>				<i>Five</i> to <i>Ten</i>			
		<i>S</i> ²	<i>F</i> ³	<i>T</i> ⁴	Sign test <i>p</i>	<i>S</i> ²	<i>F</i> ³	<i>T</i> ⁴	Sign test <i>p</i>	<i>S</i> ²	<i>F</i> ³	<i>T</i> ⁴	Sign test <i>p</i>
Variable Non-knowers	9 ¹	4	1	4	0.34	2	6	1	0.96	2	4	3	0.89
All Subset Knowers	42	20	5	12	<0.01	22	5	15	<0.01	16	9	17	0.11

¹Three additional variable non-knowers were excluded because they failed to provide any responses to *ten*. ²Successes: median response is less than median response for *ten*. ³Failures: median response is greater than median response for *ten*. ⁴Ties: median response is the same as median response for *ten*.

Analysis of median responses. Our first non-parametric analysis asked whether children gave larger responses for larger requests (Figure 1c). In place of an analysis comparing the slope of means to 0, we conducted sign tests comparing children's median responses for KL+1 to their responses to *ten*, KL+2 compared to *ten*, and *five* compared to *ten* (see Table 4). None of these comparisons were statistically significant for the variable non-knowers. For subset knowers, both KL+1 vs. *ten* and KL+2 vs. *ten* were significantly different (both $ps < 0.01$). However, there was no difference between responses for *five* and *ten* ($p = 0.11$). These results suggest two interim conclusions. First, they suggest that subset knowers have preliminary knowledge of number words up to at least KL+2. Second, they fail to support the hypothesis that children have noisy, preliminary meanings for large numbers. However, because these results are new to this study, we revisit them in Study 2 to ask whether they replicate in a larger sample of children.

Correct responses vs. chance. A second approach to testing whether children have knowledge beyond their Wynn knower level is to ask whether they provide correct responses more frequently than expected by chance. To test this, we modeled chance given the assumption that children give certain quantities more frequently than others. For example, two-knowers gave ten objects on 0.17 of trials for requests above *two*. Therefore, if two-knowers' responses for requests above *two* are independent of the requested number, we should expect that on 17% of requests for *ten*, two-knowers should give ten objects. We used binomial tests to compare children's actual responses (e.g., two-knowers provided ten objects on 0.31 of ten trials) to chance (e.g., 0.17) to ask whether children were more likely than chance to provide correct responses for KL+1 and higher. We did this analysis separately for small numbers (up to requests for *four*) and large numbers (requests of *five* and higher). For the analysis of small numbers, our null distribution was computed from all responses with values above the child's knower level. For our analyses of numbers *five* through *ten*, chance was computed from responses with values ranging from five to ten. Using this method, we found that among variable non-knowers, the frequency of correct responses did not differ from chance for requests of KL+1 through *four* or for requests of *five* and higher. In contrast, subset knowers were collectively more likely than chance to provide correct responses for all requests of KL+1 through *four* (e.g., two-knowers gave correct responses for *three* and *four* significantly more than expected by chance). Also, interestingly, though their performance was modest, subset knowers gave correct responses 16% of the time for *five*, *eight*, and *ten* (combined), which differed from chance levels of 11% (see Table 5). Thus, in contrast to all other analyses presented thus far, this last analysis finds evidence that subset knowers may have some knowledge of larger numbers. Again, we return to this question in Study 2, where we attempt to replicate this effect in a larger group of children to

determine if it is reliable, while also testing children with a larger set of objects (i.e., 20), to be sure that a tendency to give, e.g. 10 objects for requests of *ten* isn't purely due to the fact that 10 is the largest possible response in the task. If children have item-based associations between number words and particular ANS states, then such a change should have no effect on responses, and we should again find evidence of meanings for large numbers. However, if this effect is not reliable, or is driven by alternative strategies, then children should no longer perform better than chance for *five*, *eight*, and *ten* when the largest set available is 20.

Table 5: Study 1 Correct Responses for KL+1 to *four* vs. chance, and *five*, *eight*, and *ten* vs. chance

	KL+1 to <i>Four</i> vs. Chance				<i>Five, Eight, & Ten</i> vs. Chance			
	# Trials	Proportion Correct	Chance	<i>p</i>	# Trials	Proportion Correct	Chance	<i>p</i>
Variable Non-knowers	71	0.24	0.18	0.15	56	0.035	0.055	0.82
All Subset Knowers	257	0.31	0.24	<0.01	366	0.16	0.12	<0.01

CoV Analysis. All the evidence presented thus far suggests that children have preliminary knowledge of numbers up to KL+2. In one instance (i.e., non-parametric comparisons to chance), we also found that such knowledge may extend into the large number range. In Study 2, we conduct these same analyses for a new dataset, and apply special scrutiny to this ambiguity surrounding large numbers. For the remainder of Study 1, we focus on the possible nature of children's preliminary meanings.

In Wagner and Johnson's (2011) study, they argued for large approximate number word meanings on the basis of two findings: (1) a positive slope for KL+1 to *ten*, and (2) scalar variability in children's responses in this same range. While we have reported evidence of noisy preliminary meanings for small – and possibly larger – number words, all behaviors exhibit some degree of noise, whether driven by an inherently noisy representational system like the ANS, by weak mappings between words and precise representations, or by noisy deployment of mental

representations in behavior. In order to show that noise is due to the ANS, it is generally necessary to show that behaviors respect Weber's law, such that variability increases with the represented quantity. Scalar variability can be captured in constant Coefficient of Variation (CoV) values, such that the standard deviation of responses divided by the mean remains constant regardless of target set size. Consequently, to determine whether children respond to Give-N by recruiting mappings of number words to the ANS, researchers have asked whether children's responses on Give-N show constant CoV across their count list.

Based on this logic, Wagner and Johnson (2011) reported increasing means and constant CoVs for larger requests, which they argued were consistent with meanings rooted in the ANS (a finding which we replicate; see Figure 4). However, as we show below, because responses on the Give-N task are upper and lower bounded, a constant CoV in behaviors is *not* in fact predicted by approximate meanings. This is because responses for requests close to either the upper or lower bound should exhibit standard deviations that are more curtailed than the standard deviations for requests that are further from either bound. Further, children's bias of giving all ten items distorts the standard deviation of responses (though past reports have attempted to address this problem by eliminating give-all responses from CoV analyses – e.g., Wagner & Johnson, 2011).

Figure 4. (A) Coefficient of Variation (CoV) values in our dataset, collapsed across knower-levels. knowers. (B) CoV values for each knower-level. All non-knowers are variable non-knowers.

To address these two problems and test whether Give-N data are in fact consistent with a constant CoV, we used MATLAB to conduct three simulations – each simulating 15 children per knower level (non-knowers through three-knowers) 15,000 times. The three simulations were similar in that each assumed that children’s responses beyond their knower level center around the correct response. However, in order to examine the role of task constraints versus the variance of children’s numerical representations, each of the three simulations assumed a different relationship between the variance of children’s numerical representations and numerosity, to test whether it is indeed possible to provide definitive support for a constant CoV.

In Figure 5, the first row of figures, labeled as “Unconstrained” plot the error relative to the mean that is expected by models that assume either (1) a Flat CoV, (2) an Increasing CoV, or (3) a Decreasing CoV. None of these three consider the constraints imposed by the Give-N task. The second and third rows of figures in Figure 5 depict the output of our simulations – i.e., the pattern of error predicted by these same flat, ascending, and descending models of CoV, but given additional assumptions – namely the constraints imposed by the Give-N task.

The first set of simulations assumed that responses are bounded, as they indeed are in the Give-N task (see Figure 5, “Bounded”). The first model (Flat CoV) assumed variance that responses were drawn from a normal distribution with a constant (i.e., flat) CoV – a pattern consistent with meanings derived from the ANS. We simulated this scenario by assuming that children’s responses were drawn from a normal distribution with a constant CoV of 0.64 for each requested number using CoV data from Halberda and Feigenson (2008). The second model (Increasing CoV) assumed that responses were drawn from a normal distribution with a CoV that increased from 0.4 to 0.9 with requested numerosity from *one* to *ten* in equal steps – a model of error *not* consistent with the ANS. The third model (Decreasing CoV) assumed that responses were drawn from a normal distribution with a CoV that decreased from 0.9 to 0.4 with requests from *one* to *ten* in equal steps – also not consistent with the ANS. In this first set of simulations, all three models took into account the constraints of the upper (ten) and lower (knower-level+1) bounds on the Give-N task. Specifically, for each requested number, a response was randomly drawn from a normal distribution with a mean equal to the correct response and a standard deviation determined by the CoV. If the randomly generated response fell below the lower bound or above the upper bound, the response was replaced with knower-level + 1 or ten, respectively.

As shown in Figure 5, despite assuming very different underlying representations using Flat, Increasing, or Decreasing CoVs (see row labeled “Unconstrained”), all three models, when simulated, resulted in a similarly non-flat shaped relationship between requested numerosity and simulated CoV, with a peak simulated CoV in the middle of the count list. In other words, the upper and lower bounds on the Give-N task overwhelm any other pattern that might be observed in the data, making it impossible to evaluate whether or not children use ANS representations using a standard Give-N dataset. For this reason - i.e., because an underlying representation with

a flat CoV does not predict different behaviors relative to increasing or decreasing CoVs - it is not meaningful to compare children's behavior to these standards as a diagnostic of ANS representations.

While the above simulations accounted for the upper and lower bounds imposed by the Give-N task, they otherwise assumed that children's responses were normally distributed. However, as already noted, subset knowers are more likely to give all ten objects in a set than any other quantity, independent of the requested numerosity. In an attempt to account for this, our second set of simulations randomly set 15% of all children's responses to *ten*. The 15% level was chosen because, in our data set, subset knowers gave ten on 15% of trials beyond their knower level. The remaining 85% of responses were simulated in the same manner as above. In this case, our simulated CoVs decreased with increasing numerosity (see Figure 5, Bounded + All Bias), and we found again that the task constraints overwhelmed any underlying pattern that might indicate whether or not children's responses are supported by the ANS. This pattern is similar to the pattern of CoVs that we see in our collected data, if we separate children by

knower level (see Figure 4b).

Figure 5. Simulated CoVs of responses. Simulations were run under three conditions: (A) Constant CoV of 0.64; (B) Increasing CoV from 0.4 to 0.9; and (C) Decreasing CoV from 0.9 to 0.4. The first row depicts the CoVs assumed by each model, which would be predicted given no task constraints. The second row plots the predictions of a model which assumes that responses are normally distributed around a correct mean, subject to a lower-bound of KL and an upper-bound of 10; The third row adds an additional constraint to account for the bias of children to give all, with 15% of responses randomly set to 10.

Study 2

In Study 1 we replicated previous studies, and showed that differences in their conclusions were due to how data were analyzed, not differences in data. Also, we extended these findings by conducting non-parametric analyses, which provided preliminary evidence that

children have knowledge not only of their KL+1, but also of their KL+2. At the same time, we failed to find evidence of knowledge beyond KL+2 on a variety of measures, including slopes of mean responses and differences in median responses for large numbers. One measure – which tested whether children’s responses to *five*, *eight*, and *ten* differed from chance – found evidence of a modest effect. However, because this is the only measure that found such an effect, and because the effect was relatively small, we interpret it, like the other novel results in Study 1, with caution, pending replication in Study 2. Finally, Study 1 reported simulations which suggested that CoV analyses of Give-N data cannot readily diagnose a role for ANS representations in number word meanings.

In our second study, we sought to replicate the novel findings of Study 1 with a new, larger, dataset. Whereas Study 1 included only 42 subset knowers, in Study 2 we recruited participants until there were at 15 children at each knower level (from non-knowers to three-knowers), resulting in 96 children, including 65 who were either one-knowers, two-knowers, or three-knowers. Also, we tested children with a larger set of objects (i.e., 20), in order to decouple the largest available set (10 objects in Study 1, 20 in Study 2) from the largest request made (*ten* in both studies). This change was important for two reasons. First, it reduced upper-boundedness as a factor, thereby allowing children the chance to demonstrate normally distributed responses and allowing a stronger test of scalar variability. Second, it allowed us to ask whether any of the findings reported in Study 1 might be driven by, e.g., a bias to give “all” objects for numbers later in the count list, independent of actual knowledge of their meanings. If children have reliable associations between large number words and approximate magnitudes – and in particular between the word *ten* and ‘approximately ten objects’ – then this should be found regardless of whether the largest set available is 10 or 20.

Methods

Participants

Participants were 96 children aged 2;1 to 4;3 years ($M=3;3$). Data collection continued until we recruited at least 15 non-knowers, one-knowers, two-knowers, and three-knowers. This resulted in data for 22 non-knowers, 20 one-knowers, 30 two-knowers, 15 three-knowers, and 8 four-knowers. As in Study 1, four-knowers ($n=8$) were excluded from analyses, as were CP-knowers, and non-knowers who gave the same amount for every request ($n=6$). Six additional children were excluded because they stopped the task before it was possible to determine knower level.

Procedure

The Give-N task was conducted using the methods of Study 1 with one exception: Children were given 20 objects to choose from instead of 10.⁴

Results

⁴ Participants also received three additional tasks, conducted after Give-N, which were designed to probe children's understanding of the later-greater principle and approximate mappings to numbers smaller than 10. These tasks are not reported here.

Figure 6. Median Responses from Study 2 for Each Knower Level.

Analysis of Median Responses. The parametric tests described in Study 1 were carried out, and the results replicate the pattern of findings reported in Study 1 (Table S1 and S2 in the Supplemental Materials). Here, we focus on the novel non-parametric tests that we introduced in Study 1 since they alone have never been previously replicated. First, an analysis of median responses replicated the finding from Study 1 that subset knowers, but not variable non-knowers, give more in response to requests for *ten* than for *KL+1* or *KL+2* (Table 6). Also, as in Study 1, neither subset knowers nor variable non-knowers gave more in response to requests from *ten* than *five* (Table 6; Figure 6). Thus, we again find that although children have noisy preliminary meanings beyond their knower levels, there was no evidence of such knowledge in the large number range.

Table 6. Study 2 Sign Tests of Median Responses

	<i>N</i>	KL+1 to <i>Ten</i>				KL+2 to <i>Ten</i>				<i>Five</i> to <i>Ten</i>			
		<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign test <i>p</i>	<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign test <i>p</i>	<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign Test <i>p</i>
Variable Non-knowers	17	8	2	7	0.055	5	3	9	0.36	3	1	13	0.31
All Subset Knowers	65	41	10	14	<0.01	36	19	10	0.015	26	23	15	0.44

¹Successes: median response is less median response for *ten*. ²Failures: median response is greater than median response for *ten*. ³Ties: median response is the same as median response for *ten*.

Table 7: Study 2 Correct Responses for KL+1 *four* vs. chance, and *five*, *eight*, and *ten* vs. chance

	KL+1 to <i>Four</i> vs. Chance				<i>Five</i>, <i>Eight</i>, and <i>Ten</i> vs. Chance			
	# Trials	Proportion Correct	Chance	<i>p</i>	# Trials	Proportion Correct	Chance	<i>p</i>
Variable Non-knowers	171	0.18	0.13	0.036	122	0.0082	0.011	0.74
All Subset Knowers	390	0.19	0.16	0.04	567	0.069	0.064	0.35

Correct Responses vs. Chance. An analysis of correct responses vs. chance replicated the finding that subset knowers respond better than chance to requests for KL+1 up to *four*, and also, new to Study 2, found this pattern for variable non-knowers (Table 7). Finally, and critically, we did not replicate the finding from Study 1 that subset knowers respond better than chance to requests for *five*, *eight*, and *ten* (Table 6). This suggests that the result in Study 1 was either not reliable, or that it was perhaps driven by the fact that only 10 objects could be given (which corresponded to the largest request – i.e., *ten*).

CoV. Replicating Wagner and Johnson (2011) and our own results from Study 1, we found flat CoVs for responses to KL+1 and up (Figure 7). However, despite providing children with 20 objects to give, we again found that this pattern is compatible with not only with a constant ratio between error and mean responses, as predicted by Weber’s Law, but also with error that increases (or even decreases) relative to the mean. Thus, as in Study 1, we again conclude that patterns of variability in the Give-N task are not sufficiently robust to provide evidence for ANS

representations, and add to this that even relatively robust modifications of the task, like doubling the available objects, do not substantially change this (Figure 8).

Figure 7. (a) Coefficient of Variation (CoV) values in our dataset, collapsed across knowers. (b) CoV values for each knower-level. All Non-knowers are Variable Non-knowers.

Figure 8. Simulated CoVs of responses. Simulations were run under three conditions: (A) Constant CoV of 0.64; (B) Increasing CoV from 0.4 to 0.9; and (C) Decreasing CoV from 0.9 to 0.4. The first row depicts the CoVs assumed by each model, which would be predicted given no task constraints. The second row plots the predictions of a model which assumes that responses are normally distributed around a correct mean, subject to a lower-bound of KL and an upper-bound of 20; The third row adds an additional constraint to account for the bias of children to give all, with 15% of responses randomly set to 20.

Combined Analysis of Data from Studies 1 and 2. Across Studies 1 and 2 we have reported the first evidence that subset knowers have preliminary meanings for numbers as high as KL+2. However, neither study had sufficient power to analyze data at individual knower levels to determine whether this pattern was observable at each level. Here, we combined data from the two studies to obtain the needed power for these analyses. This allowed us to address two outstanding questions.

First, it allowed us to perform an adequately powered analysis of variable non-knowers in order to test whether they have reliable, preliminary knowledge of any number words (all previously reported analyses for subset knowers contained 3-4 times more subjects than analyses of variable non-knowers). With increased power from the combined datasets, we found that variable non-knowers gave more in response to requests for *ten* than *one* but did not give more for *ten* vs. *two* (see Table 7). This suggests that some non-knowers may have preliminary knowledge of KL+1.

Table 8. Combined Data Set Sign Tests of Median Responses by Knower Level

	<i>N</i>	KL+1 vs. <i>Ten</i>					KL+2 vs. <i>Ten</i>					<i>Five</i> vs. <i>Ten</i>				
		<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign test	<i>p</i>	<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign test	<i>p</i>	<i>S</i> ¹	<i>F</i> ²	<i>T</i> ³	Sign test	<i>p</i>
Variable Non-knowers	26	12	3	11	0.018		7	9	10	0.77		5	5	16	0.62	
All Subset Knowers	107				<0.01					<0.01					0.18	
One-knowers	35	16	8	11	0.076		22	6	7	<0.01		12	10	13	0.42	
Two-knowers	46	30	2	14	<0.01		26	9	11	<0.01		20	14	12	0.20	
Three-knowers	26	15	5	6	0.021		10	9	7	0.5		10	9	7	0.5	

¹Successes: median response is less median response for *ten*. ²Failures: median response is greater than median response for *ten*. ³Ties: median response is the same as median response for *ten*.

Second, using combined data we asked whether noisy preliminary meanings were present at each knower level, or were instead specific to, e.g., children at higher or lower levels. To ask this, we conducted sign tests comparing median responses for KL+1 to *ten* and KL+2 to *ten* for each knower level separately, as well as responses for *five* vs. *ten*. We found that all subset knowers gave more for *ten* than KL+1 but that only one- and two-knowers (but not three-knowers) gave more for *ten* than KL+2 (Table 8). Critically, children at no knower level gave more for requests of *ten* than for requests of *five*. Finally, we repeated this analysis on the subset of children who could count to at least *ten* and found a similar pattern of results: Noisy preliminary meanings for KL+1 and KL+2, but no differences between larger numbers (Table S3

in Supplemental Materials). In sum, the pattern of responses reported in Studies 1 and 2 is present across all knower levels, with the possible exception of three-knowers: Noisy preliminary meanings for KL+1 and KL+2, but no difference in responses for larger numbers.

Discussion

We investigated the nature of children's early number word meanings, and asked why previous studies reach such strikingly different conclusions. This led us to replicate and explain discrepancies from previous studies, and to discover a pattern of findings not previously reported in the literature. On one prominent view of number word learning, Wynn's Give-a-Number task (Give-N) provides strong evidence that early number words are acquired using representations of small sets, since children appear limited to learning number words up to *three* or *four* in absence of counting, and fail to exhibit signatures of the ANS for numbers beyond their Wynn "knower level" (Carey, 2009; Le Corre & Carey, 2006; Sarnecka & Lee, 2010). Against this view, others have argued that Wynn's method systematically underestimates children's knowledge, which extends beyond small numbers, and is rooted in the ANS. Here, we found a pattern of evidence that is incompatible with each of these previous claims. First, although we found that children acquire number words in stages, in support of the Knower Level approach, we also found that Give-N is quite conservative, and that children appear to have knowledge of numbers beyond what is found in previous reports. Specifically, children exhibited knowledge of words that were up to two numbers beyond their knower level (KL+2). Second, using methods like those found in past reports, as well as more appropriate non-parametric tests, we failed to find reliable evidence of preliminary meanings in the large number range. To show this, we employed a variety of different analyses. For example, to test whether children responded differently to *five* and *ten*, we analyzed the slope of children's mean responses, differences in their responses from chance, and,

given the non-normal nature of Give-N data, differences between median responses. Third, using model simulations, we found that, due to profound violations of normality, standard Give-N datasets cannot be used to identify scalar variability, a key signature of the ANS. Thus, previously reported data that appeared to provide evidence for the use of the ANS (Wagner & Johnson, 2011) were in fact also compatible with an absence of scalar variability. Overall, our data suggest that although children may acquire noisy, preliminary meanings beyond their knower level, there is no reliable evidence that such meanings exist for larger numbers.

To reach these conclusions, we consolidated studies by showing that a single data set can be used to replicate their disparate findings, and by demonstrating how these differences were due to different analytic choices and assumptions. For example, although we replicated the findings of Sarnecka and Lee (2011) using their analyses, their conclusion that children lack knowledge beyond their knower level is true only if correct responses are removed from analyses, as done in their study. Similarly, we found that previous studies overestimated children's knowledge of larger numbers (Wagner & Johnson, 2011) due to (1) analyses of slopes that included data from the small number range (see also Gunderson et al., 2015), and (2) an assumption that only a constant CoV, like that predicted by the use of the ANS, should predict a linear increase in variability in children's actual Give-N responses. Like Gunderson et al., we showed that subset knowers do not exhibit a positive slope for numbers ranging from *five* to *ten*. Further, using model simulations, we showed that when an appropriate model of chance is used, it is impossible to differentiate scalar variability that exhibits a constant ratio to mean responses from error that increases (or even decreases) relative to the mean.

Our conclusions are based on repeated replications of findings using diverse analyses. However, they are tempered by the fact that all effects reported here – and elsewhere in this

literature – are small. Also, the data presented here – as in past studies – do not provide sufficient power to test whether any knowledge that children might have in the small number range is reliably *stronger* than knowledge they have for larger numbers. Thus, any theoretical claims that hinge upon the existence of *stronger* preliminary meanings for small vs. large numbers cannot be readily tested using our dataset, nor any existing data set. However, to our knowledge, no previous theory hinges on such a difference. Instead, the logic of preceding studies was to determine whether number word meanings are bounded to small numbers, as presupposed by Wynn’s knower level classifications, or alternatively, whether Wynn’s classification schema is too conservative, such that analyses which eschew knower levels might find evidence of knowledge beyond the small number range. The critical question, given that all parties acknowledge the presence of small number knowledge, is whether more sensitive tests of Give-N data might detect knowledge of large numbers. Here, we showed that although we had sufficient power to replicate past evidence of such knowledge (Wagner & Johnson, 2011), these previous findings rested on the inclusion of small numbers. When only large numbers were included, we failed to find reliable evidence of noisy preliminary meanings in the large number range using the same types of analyses.

These caveats aside, our results are consistent with recent reports which suggest that, to the extent that children ever form fixed, item-based, associations between number words and cardinalities, these are likely limited to very small numbers. In one study, Sullivan and Barner (2014) showed that 5- to 7-year-old children only have reliable associative mappings between number words and cardinalities up to about *five* or *six*. In one task, children were told a number verbally (e.g., *ten*) and asked to pick which of two dot arrays the word labeled, where one set contained 10 items, and the other contained either 5 or 20. When the smaller of the two sets was

greater than 6, children performed randomly. In fact, even adults were hardly better than chance at this task for numbers higher than ten or twelve (Sullivan & Barner, 2013). In a second task, Sullivan and Barner (2014) asked children to estimate dot arrays one at a time, but prior to the task told children different information regarding the largest set they would see. Some were told there would be as many as 25 dots per trial, while others were told there would be as many as 75, or even 750 dots (whereas the true maximum was 350). Here, they found that children's estimates changed depending on the calibration they received, not only for the largest numbers, but for all numbers – right down to eight – but not for numbers *six* and below (for similar results in adults, see Sullivan & Barner, 2013). Based on this, they concluded that children only had strong item-based associations between number words and cardinalities up to *five* or *six*, with estimates for larger numbers based on a holistic, malleable, structure mapping to the ANS, not individual item-based associations.

Several studies provide converging evidence that, if ANS representations play a role in early number word learning, they play an ancillary role, rather than providing core meanings of integers. First, children's knower levels generally stop at 3 or 4, rather than 7 or 8, and these classifications are predictive of important developmental outcomes (Geary et al., 2017; Le Corre, 2014; Shusterman et al., 2016). Second, languages feature singular, dual, and sometimes trial number morphology, which appear to accelerate the learning of small number word meanings (Almoammar et al., 2013; Barner et al., 2009; Marusic et al., 2016; LeCorre et al., 2016; Sarnecka et al., 2007), but languages never have grammatical markers for larger cardinalities. Third, although it is possible to teach 3-knowers the meaning of *four* by associating it with sets of 4, CP-knowers appear unable to readily learn larger numbers like *ten* given the same training protocol (Carey et al., 2017).

In sum, although approximate number representations ultimately become associated to the count list in development, this link appears to be weak in early years, and probably takes the form of a holistic structure mapping, rather than of item-based associative mappings between words and cardinalities, as would be required for learning stable, exact, meanings like those present in the earliest moments of number word learning.

References

- Almoammer, A., Sullivan, J., Donlan, C., Marušič, F., Žaucer, R., O'Donnell, T., & Barner, D. (2013). Grammatical morphology as a source of early number word meanings. *Proceedings of the National Academy of Sciences*, 110, 18448-18453.
- Barner, D. & Bachrach, A. (2010). Inference and exact numerical representation in early language development. *Cognitive Psychology*, 60, 40-62.
- Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, 118, 87-96.
- Barner, D., Libenson, A., Cheung, P., & Takasaki, M. (2009). Cross-linguistic relations between quantifiers and numerals in language acquisition: Evidence from Japanese. *Journal of Experimental Child Psychology*, 103(4), 421-440.
- Barner, D. & Snedeker, J. (2005). Quantity judgments and individuation: Evidence that mass nouns count. *Cognition*, 97(1), 41--66.
- Barner, D. & Snedeker, J. (2006). Children's Early Understanding of Mass-Count Syntax: Individuation, Lexical Content, and the Number Asymmetry Hypothesis. *Language Learning and Development*, 2(3), 163--194.
- Baroody, A. J., & Price, J. (1983). The development of the number-word sequence in the counting of three-year-olds. *Journal for Research in Mathematics Education*, 14(5), 361-368.
- Bergelson, E. & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, 109(9), 3253-3258.

- Briars, D., & Siegler, R. S. (1984). A featural analysis of preschoolers' counting knowledge. *Developmental Psychology*, 20(4), 607.
- Carey, S. (2004). Bootstrapping & the origin of concepts. *Daedalus*, 133(1), 59-68.
- Carey, S. (2009). Where our number concepts come from. *Journal of Philosophy*, 106 (4), 220-254.
- Cheung, P., Rubenson, M., & Barner, D. (in press). To infinity and beyond: Children generalize the successor function to all possible numbers years after learning to count. *Cognitive Psychology*.
- Clark, E. V. & Hecht, B. F. (1983). Comprehension, production, and language acquisition. *Annual review of psychology*, 34(1), 325-349.
- Condry, K. F. & Spelke, E. S. (2008). The development of language and abstract concepts: The case of natural number. *Journal of Experimental Psychology: General*, 137, 22-38.
- Cordes, S. & Brannon, E. M. (2008). Quantitative competencies in infancy. *Developmental Science*, 11(6), 803-808.
- Dale, P. S. & Fenson, L. (1996). Lexical development norms for young children. *Behavior Research Methods, Instruments, & Computers*, 28(1), 125-127.
- Davidson, K., Eng, K., & Barner, D. (2012). Does learning to count involve a semantic induction? *Cognition*, 123, 162-173.
- Dehaene, S. (2007). *The number sense: How the mind creates mathematics*. New York: Oxford University Press.
- Feigenson, L. & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition*, 97, 295-313.

- Feigenson, L., Dehaene, S., & Spelke, E. (2004) Core systems of number. *Trends in Cognitive Sciences*, 8(7), 307-314.
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., ... & Stiles, J. (1994). Variability in early communicative development. *Monographs of the society for research in child development*, i-185.
- Frye, D., Braisby, N., Lowe, J., Maroudas, C., & Nicholls, J. (1989). Young children's understanding of counting and cardinality. *Child Development*, 60, 1158–1171.
- Fuson, K. C. (1988). *Children's counting and concepts of number*. New York: Springer-Verlag.
- Fuson, K. C. (1991). Children's early counting: Saying the number-word sequence, counting objects, and understanding cardinality. *Language in mathematical education: Research and practice*, 27-39.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gallistel, C.R. & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, 44, 43-74.
- Gathercole, V. C. (1985). More and more and more about more. *Journal of Experimental Child Psychology*, 40(1), 73-104.
- Geary, D. C., vanMarle, K., Chu, F. W., Rouder, J., Hoard, M. K., & Nugent, L. (2017). Early Conceptual Understanding of Cardinality Predicts Superior School-Entry Number-System Knowledge. *Psychological Science*, 1-15.
- Gelman, R. (1993). A rational-constructivist account of early learning about numbers and objects. *Learning and motivation*, 30, 61-96.
- Gelman, R., & Gallistel, C. R. (1986). *The child's understanding of number*. Harvard University Press.

- Gillund, G. & Shriffrin, R. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1-67.
- Gunderson, E., Spaepen, E., & Levine, S. (2015). Approximate number word knowledge before the cardinal principle. *J Exp Child Psychol*, 130, 35-55.
- 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(5), 1457-1465.
- Le Corre, M. (2014). Children acquire the later-greater principle after the cardinal principle. *British Journal of Developmental Psychology*, 32(2), 163-177.
- Le Corre, M., & Carey, S. (2007). One, two, three, four, nothing more: An investigation of the conceptual sources of the verbal counting principles. *Cognition*, 105(2), 395-438.
- Le Corre, M., Li, P., Huang, B. H., Jia, G., & Carey, S. (2016). Numerical morphology supports early number word learning: Evidence from a comparison of young Mandarin and English learners. *Cognitive Psychology*, 88, 162-186.
- Lee, M.D. & Sarnecka, B.W. (2010). A model of knower-level behavior in number concept development. *Cognitive Science*, 34, 51-67.
- Lee, M.D. & Sarnecka, B.W. (2011). Number-knower levels in young children: Insights from Bayesian modeling. *Cognition*, 120, 391-402.
- Leslie, A. M., Gelman, R., & Gallistel, C. R. (2008). The generative basis of natural number concepts. *Trends in cognitive sciences*, 12(6), 213-218.
- Lipton, J. S. & Spelke, E. (2005). Preschool children's mapping of number words to nonsymbolic numerosities. *Child Development*, 76(5), 978-988.

- Meck, W. H., & Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 320-334.
- Piantadosi, S. T., Jara-Ettinger, J., & Gibson, E. (2014). Children's learning of number words in an indigenous farming-foraging group. *Developmental Science*, 17(4), 553-563.
- Sarnecka, B. W. & Carey, S. (2008) How counting represents number: What children must learn and when they learn it. *Cognition*, 108, 662-674.
- Sarnecka, B. W., Kamenskaya, V. G., Yamana, Y., Ogura, T., & Yudovina, Y. B. (2007). From grammatical number to exact numbers: Early meanings of 'one', 'two', and 'three' in English, Russian, and Japanese. *Cognitive Psychology*, 55(2), 136-168.
- Sarnecka, B. & Lee, M. (2009). Levels of number knowledge during early childhood. *J Exp Child Psychol*, 103, 325-337.
- Schaeffer, B., Eggleston, V. H., & Scott, J. L. (1974). Number development in young children. *Cognitive Psychology*, 6, 357-379.
- Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior*, 6, 156-163.
- Shusterman, A., Slusser, E., Halberda, J., & Odic, D. (2016). Acquisition of the cardinal principle coincides with improvement in approximate number system acuity in preschoolers. *PloS One*, 11(4), 1-22.
- Carey, S., Shusterman, A., Haward, P., & Distefano, R. (2017). Do analog number representations underlie the meanings of young children's verbal numerals?. *Cognition*, 168, 243-255.
- Starkey, G. & McCandliss, B. (2014). The emergences of "groupitizing" in children's numerical cognition. *Journal of Experimental Child Psychology*, 126, 120-137.

Sullivan, J. & Barner, D. (2012). How are number words mapped to approximate magnitudes?

The Quarterly Journal of Experimental Psychology, 66(2), 389--402.

Sullivan, J. & Barner, D. (2014). Inference and association in children's early numerical

estimation. *Child Development*, 85(4), 1740-1755.

Tincoff, R. & Jusczyk, P. W. (2012). Six-month-olds comprehend words that refer to parts of the

body. *Infancy*, 17(4), 432-444.

Wagner, J. B., & Johnson, S. C. (2011). An association between understanding cardinality and

analog magnitude representations in preschoolers. *Cognition*, 119(1), 10-22.

Wagner, K., Dobkins, K., & Barner, D. (2013). Slow mapping: Color word learning as a gradual

inductive process. *Cognition*, 127, 307-317.

Wagner, K., Kimura, K., Cheung, P., & Barner, D. (2015). Why is number word learning hard?

Evidence from bilingual learners. *Cognitive Psychology*, 83, 1-21.

Wynn, K. (1990). Children's understanding of counting. *Cognition*, 36, 155-193.

Wynn, K. (1992). Children's acquisition of the number words and the counting system.

Cognitive Psychology, 24, 220-251.