1 Parents fine-tune their speech to children's vocabulary knowledge

2 Ashley Leung¹, Alexandra Tunkel¹, & Daniel Yurovsky²

3 ¹ The University of Chicago

4 ² Carnegie Mellon University

5 Author Note

6 All data and code for these analyses are available at <https://osf.io/3f8hy>. Videos of
7 experiment sessions are available on Databrary: <https://nyu.databrary.org/volume/1128>.

8 Correspondence concerning this article should be addressed to Ashley Leung, 5848 S
9 University Ave, Chicago, IL 60637. E-mail: ashleyleung@uchicago.edu

Abstract

Young children learn language at an incredible rate. While children come prepared with powerful statistical learning mechanisms, the statistics they encounter are also prepared for them: children learn from caregivers motivated to communicate with them. How precisely do parents tune their speech to their children’s individual language knowledge? To answer this question, we asked parent-child pairs ($n=41$) to play a reference game in which the parent’s goal was to guide their child to select a target animal from a set of three. Parents fine-tuned their referring expressions to their children’s knowledge at the lexical level, producing more informative references for animals they thought their children did not know. Further, parents learned about their children’s knowledge over the course of the game, and tuned their referring expressions accordingly. Child-directed speech may thus support children’s learning not because it is uniformly simplified, but because it is tuned to individual children’s language development.

Statement of Relevance

The pace at which children learn language is one of the most impressive feats of early cognitive development. One possible explanation for this rapid pace is that the language caregivers produce is tuned to children’s developing linguistic knowledge, maintaining just the right level of complexity to support rapid learning. We present the first experimental evidence of just how precise this tuning is, showing that parents tune not just to children’s holistic language development, but to their knowledge of individual words. We developed a new method in which we experimentally controlled what parents talked about, but not how they could talk or what they could say, increasing the chance that these results will generalize outside the lab. This work points to the importance of studying the parent-child dyad as a unit instead of focusing on children as isolated learners, both in the domain of language and in social learning more broadly.

³⁵ *Keywords:* parent-child interaction; language development; communication

³⁶ Word count: 1385

Parents fine-tune their speech to children's vocabulary knowledge

In just a few short years, children master their native language. Undoubtedly, a large share of the credit for this feat is due to powerful learning mechanisms that children bring to their input (Kuhl, 2004; Saffran, Aslin, & Newport, 1996; Smith, Suanda, & Yu, 2014). However, a share of the credit may also be due to the structure of linguistic input itself: individual differences in both the quantity and quality of the language children hear are associated with individual differences in those children's language learning (Hart & Risley, 1995; Huttenlocher, Waterfall, Vasilyeva, Vevea, & Hedges, 2010; Rowe, 2012). Further, associations between input and uptake are primarily driven by differences in speech directed to children. Differences in overheard speech do not predict differences in language learning, even in communities where child-directed speech is relatively rare (Romeo et al., 2018; Shneidman & Goldin-Meadow, 2012; Weisleder & Fernald, 2013). Why is child-directed speech so important to language learning?

The way we speak to children is markedly different from the way we speak to adults. For instance, child-directed speech tends to be slower, higher pitched, and exaggerated in enunciation relative to adult-directed speech (Cooper & Aslin, 1990; Grieser & Kuhl, 1988). Beyond acoustic and prosodic differences, child-directed speech is also marked by repetition, simpler syntactic structures, and higher proportions of questions (Fernald & Simon, 1984; Newport, Gleitman, & Gleitman, 1977; Snow, 1972). Children preferentially listen to child-directed speech over adult-directed speech (Cooper & Aslin, 1990; ManyBabies Consortium, 2020), and their increased attention to child-directed speech may play a part in driving language acquisition (Soderstrom, 2007).

In addition to attentional effects, structural simplifications in child-directed speech have been tied to specific benefits in children's language learning (e.g. Brent & Siskind, 2001; Ma, Golinkoff, Houston, & Hirsh-Pasek, 2011). For instance, when parents refer to a particular object, they tend to place this object in the final position of their utterances, even

in languages where this is not the canonical word order (Aslin, Woodward, LaMendola, & Bever, 1996). This structural tendency has been tied directly to ease of word segmentation and subsequent word learning (Endress, Nespor, & Mehler, 2009; Yurovsky, Yu, & Smith, 2012).

Crucially, both acoustic and structural properties of child-directed speech change over development, with sentences getting longer, more complex, and less acoustically variable (e.g. Huttenlocher et al., 2010; Liu, Tsao, & Kuhl, 2009; Phillips, 1973). The linguistic tuning hypothesis suggests that this changing nature of child-directed speech is what allows it to be such a powerful driver of language development (Snow, 1972). If parents tune their speech to children’s developmental level, increasing the complexity of input at the same rate that children are developing their linguistic knowledge, input may always be at the optimal level of complexity to support language learning (Vygotsky, 1978).

How precisely do parents tune their speech? One possibility is that tuning is *coarse*: caregivers could tune the complexity of their speech generally, using a holistic sense of their children’s developing linguistic abilities. Consistent with a coarse-tuning hypothesis, parents tune their utterance lengths, articulation of vowels, and diversity of clauses to children’s age (Bernstein Ratner, 1984; Huttenlocher et al., 2010; Moerk, 1976). Over and above this coarse-tuning, parents might *fine-tune* their speech, taking into account not only children’s global linguistic development, but their specific knowledge of smaller units of language, such as lexical items. Fine-tuning would provide a particularly powerful and efficient vehicle for scaffolding language acquisition because of its specificity. If parents could fine-tune utterances containing specific words, phrases, or constructions, they could keep each aspect of language at a desirable difficulty to support learning, retention, and generalization (Bjork & Kroll, 2015; Vlach & Sandhofer, 2014).

To date, the only evidence for fine-tuning comes from two observational studies, one showing that parents are more likely to provide their child with labels for novel as compared

to familiar toys (Masur, 1997), and the second showing that one child’s caregivers produce their shortest utterances containing a particular word just before the child first produces that word (Roy, Frank, & Roy, 2009). Here, we present the first *experimental* evidence for fine-tuning.

Children and their parents played a reference game in which the parent’s goal was to guide their child to select a target animal from a set of three. Parents tuned the amount of information in their utterances not just to the average difficulty of each animal word, but to their prior estimates of their individual child’s knowledge of that animal. Further, parents sensitively adapted over the course of the game, providing more information on subsequent trials when they discovered that their child did not know an animal. Together, these results show that parents leverage their knowledge of their children’s language development to fine-tune the linguistic information they provide.

Method

Participants

Toddlers (aged 2-2.5 years) and their parents were recruited from a database of families in the local community or approached on the floor of a local science museum in order to achieve a planned sample of 40 parent-child dyads. Because our method was novel, we chose a sample size that would give us 95% power to detect a medium-sized effect ($d=.6$) within-subjects and rounded up to the nearest multiple of 10. A total of 48 parent-child pairs were recruited, but data from 7 pairs were dropped from analysis because of failure to complete the experiment as designed. Of the 7 pairs that were dropped, 5 children fussed out, 1 had an older sibling interfering with the study, and 1 was a twin (only the twin who participated first was included). The final sample consisted of 41 children aged 24 mo.; 5 days to 29 mo.; 20 days ($M = 26$ mo.; 0 days), 21 of whom were girls.

In our recruitment, we made an effort to sample children from a variety of racial and

socio-economic groups. Our final sample was roughly representative of the racial composition of the Chicago Area and the US more broadly (56% White, 27% Black, 8% Hispanic). However, our sample was significantly more educated than the broader community (85% of mothers had a College or Graduate Degree).

Stimuli

Eighteen animal images were selected from the Rossion and Pourtois (2004) image set, a colorized version of the Snodgrass and Vanderwart (1980) object set. Animals were selected based on estimates of their age of acquisition (AoA) for American English learners. To obtain these estimates, we used two sources of information: parent-report estimates of children's age of acquisition from Wordbank (Frank, Braginsky, Yurovsky, & Marchman, 2017), and retrospective self-report estimates of AoA from adults (Kuperman, Stadthagen-Gonzalez, & Brysbaert, 2012, see Supplemental Materials for details). The AoA of the selected animals ranged from 15 to 32 months. Half of the animals were chosen to have an early AoA (15-23 months), and the other half were chosen to have a late AoA (25-32 months). Each trial featured three animals, all from either the early AoA or late AoA category. This separation was designed to lower the likelihood that children could use knowledge of early AoA animals to infer the correct target on late AoA trials.

A modified version of the MacArthur-Bates Communicative Development Inventory Short Form (CDI; Fenson et al., 2007), a parent-reported measure of children's vocabulary, was administered before the testing session via an online survey. The selected animal words were added to the standard words, producing an 85 word survey. Two of the animal words—one in the early AoA (pig) and one in the late AoA category (rooster)—were accidentally omitted, so trials for those words were not included in analyses as we could not obtain individual-level estimates of children's knowledge.

Figure 1. A parent-child dyad playing the reference game. On each trial, the parent’s goal was to use language to communicate to their child which animal to choose.

Design and Procedure

Each parent-child pair played an interactive reference game using two iPads (Figure 1). Children began with two warm-up trials in which they tapped on circles that appeared on the iPads. Following these warm-up trials, children and their parents moved on to practice and then experimental trials. On each trial, three images of animals were displayed side by side on the child’s screen, and a single word appeared on the parent’s screen. Parents were instructed to communicate as they normally would with their child, and to encourage their child to choose the object corresponding to the word on their screen. The child was instructed to listen to their parent for cues. Once the child tapped an animal, the trial ended, and a new trial began. There were a total of 36 experimental trials, such that each animal appeared as the target twice. Trials were randomized for each participant, with the constraint that the same animal could not be the target twice in a row. Practice trials followed the same format as experimental trials, with the exception that images of fruit and vegetables were shown. All sessions were videotaped for transcription and coding.

Data analysis

Our primary quantity of interest was the amount of information that parents provided in each of their utterances. To approximate this, we measured the length of parents' referring expressions—the number of words they produced on each trial before their child selected an animal. Length is an imperfect proxy for information, but it is easy to quantify and theory-agnostic. Because utterance length is highly right-skewed (i.e. most utterances are short), we log-transformed length in all analyses. However, to facilitate interpretability, we show raw utterance length in our figures. Subsequently, utterances were manually coded for the following: (1) Use of an animal's canonical label (e.g., "leopard"), (2) Use of a descriptor (e.g., "spotted"), (3) Use of a comparison (e.g., "like a cat"), (4) Use of a superordinate level category label (e.g., "bird" for peacock), and (5) Use of a subordinate level category label (e.g., "Limelight Larry," a fictional character from a children's book, for peacock). Parent utterances irrelevant to the game (e.g. asking the child to sit down) were not analyzed. Children's utterances were coded when audible, but were not analyzed. Our second source of data was the vocabulary questionnaire that parents filled out prior to participation. Parents indicated whether their child produced each of the 85 words on the survey. In addition to analyzing parents' judgments for the animals in the task, we also computed the total number of words judged to be known for each child as a proxy for total vocabulary.

All of our analyses were done using mixed-effects models. In all cases we began with maximal random effects structures and pruned random effects until the models converged. We removed interaction terms before removing main effects, and opted to keep the most theory-relevant random effects when only a subset of main effects could be kept. For clarity, we present only the key findings and statistics here, but full model details can be found in the Supplemental Materials.

Results

We begin by confirming that our a priori divisions of animals into early and late age of acquisition (AoA) in the study design were reflected in parents' survey judgments, and that children were able to follow parents' references to select the correct target animal on each trial. After this, we show that parents fine-tune their referring expressions, producing more information in their references to animals that they think their individual children do not know. Further, parents updated their tuning over the course of the experiment, producing more information on subsequent references to animals they thought their children knew after observing evidence to the contrary (i.e. children made an incorrect selection).

Target animal difficulty

We first confirm that the early AoA animals were more likely to be marked "known" by the parents of children in our studies. As predicted, parents judged that their children knew 94% of the animals in the early AoA category, and 33% of the animals in the late AoA category, which were reliably different from each other ($\beta = -6.48$, $p < .001$, $d = -3.57$ [-4.48, -2.67]). Parents' judgments for each target word are shown in the Supplemental Materials.

Selection accuracy

On the whole, parents communicated effectively with their children, such that children selected the correct target on 69.05% of trials, reliably greater than would be expected by chance (33%, $\beta = 2.07$, $p < .001$, $d = 1.14$ [0.93, 1.36]). Children were above chance both for animals that parents thought they knew ($M = 75.08\%$, $\beta = 2.61$, $p < .001$, $d = 1.44$ [1.18, 1.70]), and for animals that parents thought their children did not know ($M = 55.19\%$, $\beta = 1.23$, $p < .001$, $d = 0.68$ [0.52, 0.84]). Thus, parents successfully communicated the target referent to children, even when parents thought their children did not know the name for the animal at the start of the game.

Was this accuracy driven by children's knowledge or parents' referring expressions?

Because we did not measure children’s knowledge of each animal directly, we used parents’ estimates of children’s knowledge as a proxy to answer this question. We fit a mixed-effects logistic regression predicting children’s accuracy on each trial from children’s total estimated vocabulary, parent-reported knowledge of the target animal, and the (log) length of parents’ expressions. We found that children with bigger vocabularies were more accurate in general ($\beta = 0.40$, $p = .001$, $d = 0.22$ [0.09, 0.36]), and that children were less accurate for animals whose names parents thought they did not know ($\beta = -1.86$, $p < .001$, $d = -1.02$ [-1.46, -0.58]). Longer referring expressions were associated with greater accuracy for animals that parents thought their children did not know ($\beta = 0.46$, $p = .025$, $d = 0.25$ [0.03, 0.47]), but not for animals that parents thought their children knew ($\beta = -0.40$, $p = .007$, $d = -0.22$ [-0.38, -0.06]).

Thus, longer referring expressions were associated with more successful communication for animals that parents thought their children did not know, but not for animals that parents thought they knew. This suggests that parental tuning contributed to children’s success. We next ask whether parents tuned the lengths of their utterances appropriately, producing longer expressions for animals they believe their children do not know.

Tuning

If parents calibrate their referring expressions to their children’s linguistic knowledge, they should provide more information to children for whom a simple bare noun (e.g. “leopard”) would be insufficient to identify the target. Parents did this in a number of ways: with one or more adjectives (e.g. “the spotted, yellow leopard”), with similes (e.g. “the one that’s like a cat”), and with allusions to familiar animal exemplars of the category (e.g. “Limelight Larry”). In many of these cases, parents would be required to produce more words (see below for further qualitative analyses). Thus, we first analyzed the (log) length of parents’ referring expressions as a proxy for informativeness.

Figure 2. Parents produced longer referring expressions to communicate about animals that children were generally less likely to know. Animals in blue are those that were estimated to be earlier learned from prior norms, animals in red are those that were estimated to be later learned. Error bars show 95% confidence intervals computed by non-parametric bootstrap.

When do parents produce longer referring expressions? One possibility is that parents tune at the coarsest level, using more words when speaking to children with smaller vocabularies. This was not the case—the total number of words parents thought their children knew did not reliably affect the length of their referring expressions ($\beta = -0.02$, $p = .595$, $d = -0.17$ [-0.79, 0.45]). A second possibility is that parents have a sense for how difficult each animal is in general, and tune coarsely to this. Our analyses confirmed this coarse-tuning: parents said reliably fewer words for animals that more children were reported to know ($\beta = -0.17$, $p = .034$, $d = -1.19$ [-2.26, -0.09]); Figure 2). Finally, parents could fine-tune their referring expressions to their children’s individual knowledge, over and above the average difficulty of each animal. Our analyses supported this conclusion: parents used reliably fewer words to refer to animals that they thought their individual child knew ($\beta = 0.25$, $p = .003$, $d = 0.98$ [0.34, 1.61]); Figure 3A). Thus, parents fine-tuned the amount of information in their referring expressions, calibrating the amount of information they provided to their children’s knowledge, even after accounting for the average difficulty of the target animal.

Figure 3. (A) Parents produced longer referring expressions for animals they thought their child did not know (first appearance). (B) When children chose the correct animal, parents continued to produce longer expressions for animals they thought their children did not know (left). However, if parents thought their child knew an animal, and they chose incorrectly, parents produced longer expressions on its second appearance (right). Gray points and lines represent individual participants, black points and lines show group averaged proportions; error bars show 95% confidence intervals computed by non-parametric bootstrap.

In addition, because each animal appeared as a target twice, we asked whether parents tuned their referring expressions over successive appearances. We found that parents used fewer words on the second appearance of each animal ($\beta = -0.08$, $p = .044$, $d = -1.06$ [-2.07, -0.03]). This effect is shown in Figure 3: the mean length of referring expression is lower on the second appearance (Fig. 3B) when compared to the first (Fig. 3A). However, the difference in utterance length between animals they thought their children knew versus didn't know was smaller on their second appearance ($\beta = -0.14$, $p < .001$, $d = -0.17$ [-0.22, -0.12]; Figure 3B). Why might that be? One possibility is that parents obtain information from the first appearance of each animal: they may have thought their child knew “leopard,” but discovered from their incorrect choice that they did not. If so, they might provide more information the second time around.

To test this prediction, we fit a model predicting the (log) length of parents' referring expressions from appearance type (first, following correct, following incorrect), whether the parent thought their child knew the animal prior to the experiment, and their interaction between appearance type and prior belief. Relative to their utterances on an animal's first appearance, parents produced shorter referring expressions on an animal's second appearance following both correct responses ($\beta = -0.14$, $p = .036$, $d = -0.12$ [-0.23, -0.01]) and incorrect responses ($\beta = -0.28$, $p = < .001$, $d = -0.22$ [-0.34, -0.11]). As before, parents produced shorter utterances for animals they thought their child knew ($\beta = -0.31$, $p = < .001$, $d = -0.92$ [-1.43, -0.41]). When children were correct on an animal's first appearance, parents' referring expressions on its second appearance did not differ in length based on whether they thought their child knew the animal prior to the experiment ($\beta = -0.02$, $p = .771$, $d = -0.02$ [-0.13, 0.10]). However, when children were incorrect on an animal's first appearance, and parents thought they knew the animal prior to the experiment, they produced reliably longer referring expressions on its second appearance ($\beta = 0.43$, $p = < .001$, $d = 0.24$ [0.13, 0.35]; Figure 3B).

As we predicted, when parents thought their children knew an animal, but then observed evidence to the contrary, they provided more information in their referring expressions to help children make the correct selection the second time. However, we did not find the opposite pattern: when children were successful for animals that parents thought they did not know, parents did not update their beliefs. Why should parents update their beliefs in one direction but not the other? One likely explanation comes from parents' linguistic tuning itself. Parents' goal in this task is to produce a referring expression that allows their children to select the target animal whether or not they know its canonical label. Consequently, when children select correctly on these trials, parents cannot know how their child arrived at the correct target. For example, if a parent said "spotted leopard" and their child was correct, they could not know whether their child succeeded because they actually knew the word "leopard," or instead because the word "spotted" was sufficiently informative.

Together, these two sets of analyses suggest that parents tune their referring expressions not just coarsely to how much language their children generally know, nor their knowledge about how hard animal words are on average, but finely to their beliefs about their individual children’s knowledge of specific lexical items. Further, when parents discover that they have incorrect beliefs about their children’s knowledge of an animal, they update these beliefs in real-time and adjust subsequent references to the same animal.

Content of referring expressions

Figure 4. Proportion of trials on which parents used comparison, descriptors, superordinate category labels, and subordinate category labels in their referring expressions. Colored bars show means across parents, error bars show 95% confidence intervals computed by non-parametric bootstrap.

Parents produced reliably longer referring expressions when trying to communicate about animals that they thought their children did not know. In the analyses presented so far, we used length as a theory-agnostic, quantitative measure of information. *How* did parents successfully refer to animals that their children did not know? As a post-hoc descriptive analysis, we coded five qualitative features of referring expressions: (1) Use of the animal’s canonical label (e.g., “leopard”), (2) Use of a descriptor (e.g., “spotted”), (3) Use of a comparison (e.g., “like a cat”), (4) Use of a superordinate level category label (e.g., “bird”

for peacock), and (5) Use of a subordinate level category label (e.g., “Limelight Larry” for peacock). Because the rates of usage of each of these kinds of reference varied widely (e.g. canonical labels were used on 94.48% of trials, but subordinates were used on 1.21% of trials), we fit a logistic mixed effects model separately for each reference kind, estimating whether it would be used on each trial from whether the parent thought their child knew the animal.

Canonical labels were used on almost all trials, and did not differ in frequency between animals parents thought their children did not know ($M = 94.68\%$) and animals they thought their children knew ($M = 93.43\%$, $\beta = 0.43$, $p = .216$, $d = 0.23$ [-0.14, 0.61]). Parents thus produced canonical labels even when they thought their children did not know these labels. One plausible explanation for this is that the target animal on each trial was identified in writing for the parent, activating the canonical label and thus lowering the cost of retrieving and producing it (Wingfield, 1968). Another possibility is that this reflects parents’ general tendency to produce basic-level category labels when talking to children (e.g., Blewitt, 1983). Finally, it could have been produced for implicitly or explicitly pedagogical reasons even though it was not referentially necessary. We expand on this possibility in the Discussion below.

Comparisons were used reliably more for animals parents believed their children did not know than for animals they thought their children knew ($\beta = 2.29$, $p = .001$, $d = 1.26$ [0.49, 2.04]), as were descriptors ($\beta = 3.09$, $p < .001$, $d = 1.71$ [1.07, 2.35]) and superordinate category labels ($\beta = 3.01$, $p = .026$, $d = 1.66$ [0.20, 3.12]). Subordinates were used less for animals parents thought their children did not know than for animals they thought their children knew ($\beta = -2.19$, $p = .025$, $d = -1.21$ [-2.26, -0.15]). Thus, parents used a variety of strategies to refer to animals that they believed their children do not know, but the use of descriptors was the most prominent (Figure 4). These descriptors are particularly apt to facilitate children’s learning, connecting parents’ fine-tuning for reference with their

children’s language acquisition.

Discussion

Parents have a wealth of knowledge about their kids, including about their linguistic development (Fenson et al., 2007). In this study, we asked whether parents leverage this knowledge to communicate successfully with their children. When playing a referential communication game, parents drew on their knowledge of their children in three ways: (1) parents produced longer, more informative referring expressions for animals that children generally learn later, (2) over and above this coarse-tuning, parents fine-tuned information to their individual children’s knowledge of specific animals, and (3) when children did not know an animal that parents thought they did, parents subsequently produced longer referring expressions for that animal. Further, this tuning was associated with more successful communication: children were more likely to correctly select animals whose names parents thought they did not know if parents produced more informative referring expressions.

These data are consistent with prior evidence of coarse-tuning in child-directed speech, but importantly provide the first experimental evidence for fine-tuning at the lexical level. When communicating with their children, parents not only take into account the average difficulty of each animal word, but they also rely on (and update) their estimates of their individual child’s knowledge of those animals. Coarse-tuning and fine-tuning may be distinct adaptations that happen independently at different timescales, but our data suggest an intriguing alternative possibility: parents’ coarse-grained estimates of their children’s language development may be built hierarchically. That is, parents may use estimates of their children’s knowledge of specific lexical and syntactic items to form their general representations of their children’s overall language development. Hierarchical representations are a powerful vehicle for maximizing both speed and generalizability of learning, and they may play the same role here, allowing parents to efficiently track and use their knowledge of their children’s language development (Tenenbaum, Kemp, Griffiths, & Goodman, 2011).

While parents' speech to children is unlikely to reflect an explicit goal to teach, it is nonetheless goal-oriented: parents want to communicate successfully (Bruner, 1983). Our reference game was designed to manipulate and measure a particular communicative goal that can be instantiated in the laboratory, but similar communicative pressures structure the daily conversations between children and their parents (Tamis-LeMonda, Kuchirko, Luo, Escobar, & Bornstein, 2017). When talking about animals that they thought their children did not know, parents used referring expressions rich with descriptors and comparisons, as in previous observational studies (Blewitt, 1983; Masur, 1997; Mervis & Mervis, 1982). These strategies scaffold communication—parents use what they think their children know (e.g. color words) in order to communicate about animals they think their children do not know. Because communication and learning are intertwined, these same strategies may work in the service of language acquisition (Yurovsky, 2018). While parents produced rich descriptions to help their children select unfamiliar animals, they almost always produced the canonical label as well. These referring expressions are thus an ideal opportunity to learn the relationship between the referent, its label, and its important identifying features. We did not independently measure children's knowledge of each animal, so we cannot determine whether they learned any new animals while playing the game. The relationship between referential strategies and ultimate learning is a promising direction for future work.

Parents fine-tune language to their children's knowledge in order to communicate successfully. In the service of proximal communicative goals, they may also provide children with input that ultimately accelerates learning. Focusing on the interactive and communicative nature of language captures a more complete picture of language development: while children bring powerful learning mechanisms to language acquisition, these mechanisms are supported by an ecological niche designed for their success.

Acknowledgements

We are grateful to all of the parents and children who participated in our work, as well as to Brian MacWhinney, Kody Manke, Molly Lewis, and all of the members of the Communication and Learning lab for their feedback on our manuscript. This research was funded by a James S. McDonnell Foundation Scholar Award to DY as well as a Social Science Research Council Faculty Seed Grant from the University of Chicago.

References

- Aslin, R. N., Woodward, J. Z., LaMendola, N. P., & Bever, T. G. (1996). Models of word segmentation in fluent maternal speech to infants. *Signal to Syntax: Bootstrapping from Speech to Grammar in Early Acquisition*, 117–134.
- Bernstein Ratner, N. (1984). Patterns of vowel modification in mother–child speech. *Journal of Child Language*.
- Bjork, R. A., & Kroll, J. F. (2015). Desirable difficulties in vocabulary learning. *The American Journal of Psychology*, 128(2), 241.
- Blewitt, P. (1983). Dog versus collie: Vocabulary in speech to young children. *Developmental Psychology*, 19(4), 602.
- Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, 81(2), B33–B44.
- Bruner, J. (1983). *Child’s talk: Learning to use language* (pp. 111–114). Norton.
- Cooper, R. P., & Aslin, R. N. (1990). Preference for Infant-Directed Speech in the First Month after Birth. *Child Development*, 61(5), 1584–1595.
- Endress, A. D., Nespors, M., & Mehler, J. (2009). Perceptual and memory constraints on language acquisition. *Trends in Cognitive Sciences*, 13(8), 348–353.
- Fenson, L., Marchman, V. A., Thal, D. J., Dale, P. S., Reznick, J. S., & others. (2007). *MacArthur-bates communicative development inventories: User’s guide and technical manual*. Baltimore, MD: Brookes.
- Fernald, A., & Simon, T. (1984). Expanded intonation contours in mothers’ speech to newborns. *Developmental Psychology*, 20(1), 104.

- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3), 677–694.
- Grieser, D. A. L., & Kuhl, P. K. (1988). Maternal Speech to Infants in a Tonal Language: Support for Universal Prosodic Features in Motherese. *Developmental Psychology*. Retrieved from <http://arxiv.org/abs/0803973233>
- Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young american children*. Paul H Brookes Publishing.
- Huttenlocher, J., Waterfall, H., Vasilyeva, M., Vevea, J., & Hedges, L. V. (2010). Sources of variability in children’s language growth. *Cognitive Psychology*, 61(4), 343–365.
- Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*, 5(11), 831–843.
- Kuperman, V., Stadthagen-Gonzalez, H., & Brysbaert, M. (2012). Age-of-acquisition ratings for 30,000 english words. *Behavior Research Methods*, 44(4), 978–990.
- Liu, H.-M., Tsao, F.-M., & Kuhl, P. K. (2009). Age-related changes in acoustic modifications of mandarin maternal speech to preverbal infants and five-year-old children: A longitudinal study. *Journal of Child Language*, 36(4), 909.
- Ma, W., Golinkoff, R. M., Houston, D. M., & Hirsh-Pasek, K. (2011). Word learning in infant-and adult-directed speech. *Language Learning and Development*, 7(3), 185–201.
- ManyBabies Consortium. (2020). Quantifying sources of variability in infancy research using the infant-directed-speech preference. *Advances in Methods and Practices in Psychological Science*, 3(1), 24–52.

- Masur, E. F. (1997). Maternal labelling of novel and familiar objects: implications for children's development of lexical constraints. *Journal of Child Language*, 24, 427–439.
- Mervis, C. B., & Mervis, C. A. (1982). Leopards are kitty-cats: Object labeling by mothers for their thirteen-month-olds. *Child Development*, 267–273.
- Moerk, E. L. (1976). Processes of language teaching and training in the interactions of mother-child dyads. *Child Development*, 1064–1078.
- Newport, E. L., Gleitman, H., & Gleitman, L. R. (1977). Mother, I'd rather do it myself: Some effects and non-effects of maternal speech style. In C. A. Ferguson (Ed.), *Talking to children language input and interaction* (pp. 109–149). Cambridge University Press.
- Phillips, J. R. (1973). Syntax and vocabulary of mothers' speech to young children: Age and sex comparisons. *Child Development*, 182–185.
- Romeo, R. R., Leonard, J. A., Robinson, S. T., West, M. R., Mackey, A. P., Rowe, M. L., & Gabrieli, J. D. (2018). Beyond the 30-million-word gap: Children's conversational exposure is associated with language-related brain function. *Psychological Science*, 29(5), 700–710.
- Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception*, 33, 217–236.
- Rowe, M. L. (2012). A longitudinal investigation of the role of quantity and quality of child-directed speech in vocabulary development. *Child Development*, 83(5), 1762–1774.
- Roy, B., Frank, M., & Roy, D. (2009). Exploring word learning in a high-density longitudinal

corpus. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 31).

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928.

Shneidman, L. A., & Goldin-Meadow, S. (2012). Language input and acquisition in a Mayan village: How important is directed speech? *Developmental Science*, 15(5), 659–673.

Retrieved from <http://arxiv.org/abs/NIHMS150003>

Smith, L. B., Suanda, S. H., & Yu, C. (2014). The unrealized promise of infant statistical word–referent learning. *Trends in Cognitive Sciences*, 18(5), 251–258.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2), 174.

Snow, C. E. (1972). Mothers' Speech to Children Learning Language. *Child Development*, 43(2), 549–565.

Soderstrom, M. (2007). Beyond babytalk: Re-evaluating the nature and content of speech input to preverbal infants. *Developmental Review*, 27(4), 501–532.

Tamis-LeMonda, C. S., Kuchirko, Y., Luo, R., Escobar, K., & Bornstein, M. H. (2017). Power in methods: Language to infants in structured and naturalistic contexts. *Developmental Science*, 20(6), e12456.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.

Vlach, H. A., & Sandhofer, C. M. (2014). Retrieval dynamics and retention in cross-situational statistical word learning. *Cognitive Science*, 38(4), 757–774.

464 Vygotzky, L. S. (1978). *Mind in society: The development of higher psychological processes*.
465 Harvard university press.

466 Weisleder, A., & Fernald, A. (2013). Talking to Children Matters: Early Language
467 Experience Strengthens Processing and Builds Vocabulary. *Psychological Science*,
468 *24*(11), 2143–2152.

469 Wingfield, A. (1968). Effects of frequency on identification and naming of objects. *The*
470 *American Journal of Psychology*, *81*(2), 226–234.

471 Yurovsky, D. (2018). A communicative approach to early word learning. *New Ideas in*
472 *Psychology*, *50*, 73–79.

473 Yurovsky, D., Yu, C., & Smith, L. B. (2012). Statistical speech segmentation and word
474 learning in parallel: Scaffolding from child-directed speech. *Frontiers in Psychology*, *3*.