

From the mouths of babes: Toddlers' early word production favors information in common ground

Lia Washington, Brian Leahy, Max Siegel, Sienna Radifera, Melissa Kline-Struhl*, Laura E. Schulz*

{bw18, maxx, radifera}@mit.edu, {leahy.b.p, melissa.e.kline, lauraelizabethschulz}@gmail.com

Department of Brain and Cognitive Sciences, MIT

Abstract

Toddlers can only say one or two words at a time. What do they choose to talk about? We report the results from a pre-registered online experiment on productive language. Toddlers (N=167; mean: 19.5 months) saw six movies. A curtain opened on an introductory scene, the parent closed their eyes, and a new event happened. The curtain closed and the child was asked what happened. On two trials the unseen event was new to the parent (Novel event); on two trials, one of two animals ate the only food in the scene (Agent ambiguous); on two trials, the only animal ate one of the two foods (Patient ambiguous). We predicted that toddlers would selectively generate informative utterances (i.e., referring to the novel event, the agent, and the patient, respectively). Toddlers' productive language was indeed sensitive to what listeners' know; however, unlike adults, they selectively referred to information in common ground.

Keywords: Productive language; Communication; Pragmatics; Rational Speech Act; Toddlers; Development

Introduction

One of the biggest discrepancies in all of development might be the gap between what young children know and what they can express in language. Decades of infancy research have revealed the surprising sophistication of infants' understanding of physical and psychological events, at ages where they can say nothing at all (Spelke & Kinzler, 2007). When children do begin to talk, they only say a single word at a time. Children only begin to produce multi-word utterances (typically only two words at a time) between 19 and 26 months (Berk and Lillo-Martin, 2012; Butcher and Goldin-Meadow, 2000; Capirci et al., 1996). Thus for a long time, children would seem to have a lot to say and very little way to say it. Here we are interested in how the very youngest speakers decide what to talk about, given that they must compress all the things they can think about into the very few words they can say.

The fact that toddlers rarely communicate linguistically does not mean that young children are uncommunicative. Pre-verbal toddlers generate, on average, one intentional communicative act (through gestures, pointing, and vocalizing) every waking minute (Wetherby et al., 1988; see also Salley et al., 2020). Gestures emerge as early as 8-12-months (Bates, 1976) and include gestures to attract adult attention, points to targets in the environment, and conventional gestures (e.g., waving bye-bye; (Goldin-Meadow & Alibali, 2013). Adults often translate these gestures into words (Goldin-Meadow et al., 2007; Golinkoff, 1986; Masur, 1982).

These early non-verbal communicative acts predict later linguistic ones, including early vocabulary (Iverson & Goldin-Meadow, 2005) and the onset of the two-word stage (Goldin-Meadow and Butcher, 2003; Iverson et al., 2008; Iverson and Goldin-Meadow, 2005; see Goldin-Meadow and Alibali, 2013 for review and discussion).

Moreover, infants and toddlers' early understanding and production of non-verbal communication is sensitive to communicative intent and attuned to what their audience does and doesn't know. Infants under a year old are surprised if someone points to one object and the listener hands over another; moreover, infants are only surprised if the person points, not if she gestures with a closed hand (Krehm et al., 2014). Toddlers spontaneously use gestures to indicate ignorance and uncertainty (Acredolo & Goodwyn, 1985) and point interrogatively to request information (Begos and Southgate, 2012; see Harris et al., 2017 for discussion and review). Toddlers also point to inform others, and selectively point when others are ignorant. Thus, for instance, two-year-olds point more often to an out-of-reach toy when the parents do not know its location than when they do (O'Neill, 1996). Moreover, older (but not younger) two-year-olds selectively verbalize when pointing is insufficiently informative. If a sticker is in one of two out-of-reach boxes in close proximity to each other, younger two-year-olds simply point, but older two-year-olds name the distinguishing picture on the outside of the box (O'Neill & Topolovec, 2001). These results are broadly consistent with other work suggesting that young language learners in communicative contexts are sensitive to referential intent and disambiguate utterances by tracking what others are attending to, what they want, and what they do and do not know (Baldwin, 1993; Brandone et al., 2014; Carpenter et al., 1998; Csibra, 2003; Phillips and Wellman, 2005; Tomasello et al., 2005).

However, it is a very big leap from being selectively informative in pointing to being selectively informative in language. Adult speakers do of course engage in cooperative communication (Grice, 1975). They take into account what their listener knows and wants to know and efficiently communicate the most informative thing to say in context (Frank and Goodman, 2012; Franke and Jäger, 2016; Gibson et al., 2019; Goodman and Frank, 2016).

This Rational Speech Act framework accounts for many aspects of adult communication (see Degen, 2023 for review) but what about children? Although as reviewed above,

even toddlers selectively communicate information to ignorant informants, the ability to identify precisely what information is helpful develops well through middle childhood. Early school-aged children, for instance, sometimes fail to understand when a speaker is under-informative (e.g., someone who says “I ate some of the cake” when they ate all of it; Noveck, 2001; Barner et al., 2009; Huang and Snedeker, 2009; Hurewitz et al., 2006; Papafragou and Tantalou, 2004). However, even much younger children (four and five-year-olds) are successful when relevant alternatives are clear in context (Barner et al., 2011; see also Foppolo et al., 2012; Katsos and Bishop, 2011; Papafragou and Musolino, 2003, Skordos and Papafragou, 2016).

In production, four and five-year-olds sometimes disambiguate referents to give useful information to their listeners (e.g., referring to ‘a blue cup’ when they know the listener can see both a blue and red cup; Nadig and Sedivy, 2002) but they do not do so reliably, often using pronouns (“They want to”) and definite noun phrases (“this one”) opaque to the listener (Brown, 1973; Sonnenschein and Whitehurst, 1984; see Kline et al., 2017). However, it is not clear whether these infelicitous utterances reflect limitations on children’s understanding of the informative thing to say or the difficulty of actually producing it. In general, more informative sentences are longer (e.g., Whitehurst et al., 1981) and thus harder to produce. A child who says ‘cup’ when there are both red and blue cups might have known ‘blue cup’ would have been more informative but still produced the easier, shorter sentence.

We are interested in the very youngest speakers, capable of producing only the very shortest sentences: toddlers at the one and two-word stage of production. At this stage, it is not clear what leads toddlers to say any given word at any given time beyond the presence of a candidate referent. Toddlers might prefer to use the words they learned most recently, the words they have known the longest, the words that are easiest for them to pronounce, the words just said to them, or they might respond idiosyncratically. Here, however, we investigate the possibility that even at the earliest word production, toddlers take into account what their listeners do and do not know. Specifically, we look at whether the earliest speakers preferentially say: a) words that refer to events unknown to their conversational partner (and thus are most informative in context) or b) words that refer only to events in common ground.

To follow we report the results from the first toddlers to complete a pre-registered experiment on productive language in the second year of life. Although there are many corpus and vocabulary checklist studies of toddlers’ word production (see e.g., Frank et al., 2017), experimental studies of early production are rare – for the good reason that it is difficult to get toddlers at the one and two-word stages to talk on demand. We addressed this by conducting the experiment as an automated study run online in parents’ homes without an experimenter present. This allowed us the possibility of



recruiting a large enough sample to generate sufficient data for an experimental design, even if production is sparse. In the experiment, we manipulate the information in common ground by having both the child and parent see a short animated scene. Then the parent is asked to “close their eyes and cover their ears”. A new event transpires. The parent is then asked to open their eyes and the child is prompted to say what happened. The question of interest is whether, if the child says anything about the scenes at all, do they refer to events the parent did not see or events in common ground?

To increase the likelihood that the child would be able to refer to any of these scenes, we designed the scenes around words produced by at least 50% of children by 16 months of age. To select these words, we looked for the earliest words that referred to a) animate agents; b) objects; c) locations, d) foods/drinks, and e) transitive verbs in Word Bank, a structured database of archived CDI data across several languages and labs (Frank et al., 2017). See Procedure below.

Children saw two kinds of trials, Novel Events trials and Agent/Patient trials. In the Novel Event trials the child and parent saw a scene (a duck on a lake for one trial and a bear in a room for the second trial). When the parent’s eyes were closed, something new happened (a ball bounced in/the bear put on shoes). All events were narrated. When the parent opened their eyes, children were asked to describe what happened.

There were two kinds of Agent/Patient trials, Agent Ambiguous and Patient Ambiguous. Agent Ambiguous trials included two agents and one patient; Patient Ambiguous trials contained one agent and two patients. The agents were a dog and/or a cat, always paired with an apple and/or banana, and a bird and/or a baby always paired with milk and/or juice. The parent and child both saw the initial scene. When the parent closed their eyes, one of the agents ate/drank the food/beverage. All events were narrated by audio recording. To avoid prosodic cueing, a single recording was made for both contrastive events in a pair. (Thus for instance, half the children who heard “Did you see that? The dog ate the apple!” heard it in the Agent Ambiguous condition and half in the Patient Ambiguous condition.) When the parent opened their eyes, the children were asked to describe what happened.

Our overarching (Hypothesis 1) was that children would be more likely to communicate informative than uninformative things across all six trials. For this hypothesis, the critical question was whether, if the child referred to anything, they referred to the novel or known information. We considered any mention of an informative entity or event (i.e., either “ball” or “bounce” would count as informative in a Novel Event trial and either “dog” or “eat” would count as informative in an Agent Ambiguous trial). Our secondary hypothesis (Hypothesis 2) looked only at the four Agent and Patient Ambiguous trials (two of each type) and predicted specifically that children should mention the Agent more than the Patient in the Agent Ambiguous trials and the Patient more than the

	Novel event 1	Agent ambiguous 1	Patient ambiguous 1
What Parent Saw			
What Child Saw			
Informative Utterance	BALL BOUNCE	BIRD DRINK	EAT APPLE
Common Ground Utterance	DUCK LAKE	JUICE BABY	DOG BANANA
Child Example of Informative Utterance	NO; OPENED; BALL BALL BALL	ITS A BIRD	APPLE
Child Example of Common Ground Utterance	[child screams] QUACK QUACK	MORE BABY YA YA	DOG
Child Example of Irrelevant Utterance	A CAR A CAR	UH OH, MUMMY; UH OH	DUCK HAPPENED


	Novel Event 2	Agent ambiguous 2	Patient ambiguous 2
What Parent Saw			
What Child Saw			
Informative Utterance	PUT-ON SHOES	DRINK MILK	CAT EAT
Common Ground Utterance	BEAR ROOM	BABY JUICE	BANANA DOG
Child Example of Informative Utterance	SHOES; THE SHOES	No Responses	BANA, unintelligible, BITE
Child Example of Common Ground Utterance	BEAR BEAR	A BABE	BYE BYE DOG
Child Example of Irrelevant Utterance	WHAT HAPPENED; SEE	WHAT HAPPENED	WHAT HAPPEN

Figure 1: Screenshots of stimuli, target informative and common ground utterances for the primary hypothesis, and examples (from different children) of informative, common ground, and irrelevant utterances. "No Responses" indicate that none of the current data contained an example of that type.

Agent in the Patient Ambiguous trials. Throughout, any synonymous descriptor counted as long as it clearly identified the referent (e.g., "cat" or "kitty" or "meow meow" counted as an Agent). See Figure for example stimuli, the target informative and utterances in common ground, and examples of children's responses in each category.

Methods

Participants

167 children (Mean: 19.5 months; 53.5% female) were recruited online via Children Helping Science, and tested asynchronously (without an experimenter present) on the platform. Participants were English-learning children between the ages of 14 to 24 months. An additional 20 participants were recruited but excluded from analysis due to technical video issues, not attending to the study at all, or failure to attend to at least one of the training trials. The study took ap-

proximately 10 minutes to complete and compensation was \$5.00.

Procedure This experiment was pre-registered on the Open Science Framework (<https://osf.io/bknwf>). At the start of the experiment, participants were shown two consecutive training trials. In these training trials, participants were presented with an image of a noun (a bunny in the first trial and a car in the other) and encouraged to say the noun. Parents were instructed to prompt the child if the child did not respond verbally, and to continue to the next trial if the child still did not respond. Participants who were not attentive for at least one of these training trials were excluded.

The training trials were followed by six test trials that fell into one of two conditions: Novel Event and Agent/Patient Ambiguous. In the Novel Event scenes, when the parent closed their eyes a new event would take place (a ball bounced by the lake or a bear put on shoes). In the Agent Ambigu-

ous scenes, two agents were featured (e.g., a cat and a dog) and one patient (e.g., an apple). When the parent closed their eyes, one of the two agents would act on the patient (e.g., the dog would eat the apple). The Patient Ambiguous scenes were similar except the scene featured two patients (e.g., juice and milk) and a single agent (e.g., baby).

Each participant was randomly assigned to one of four conditions, with each condition displaying the stimuli in a different sequence varying which agent or patient was mentioned first in the Agent/Patient trials. The first and fourth trial was always a Novel Event Trial; the remaining trials alternated Agent/Patient trials. All narrations were audio recorded. Novel Event trials were recorded in a neutral voice to minimize prosodic cues and a single recording was used for the contrastive Agent/Patient Ambiguous trials so no differentiating prosodic cues were available.

Each test trial began with the opening of red curtains, followed by a still image featuring three nouns (e.g., a cat, a dog, and an apple). Participants heard an audio recording naming the three nouns, ensuring that both the parent and child had common ground about the contents of the scene (i.e., “Look! There’s a cat, a dog, and an apple! Do you see that? A cat, a dog, and an apple!”). Afterward, parents were instructed to cover their eyes and ears to remain ignorant of the event that would transpire. A new animation began and the event was narrated to the participant (e.g., when the dog ate the apple, the child heard, “Wow! Do you see that? The dog ate the apple!”). The curtains closed and the parent was asked to uncover their eyes and ears. The child was then asked to tell the parent what happened in the scene.

Video Coding

Video data was transcribed and hand-coded. Prior to data analysis, sessions were excluded if the child was not a native speaker of English or Spanish, was not within the age range of 14 - 24 months, had technical issues that compromised the quality of the video, or failed to attend to at least one of the two training trials. Sessions were not excluded simply because the participant did not finish the study. After culling sessions based on these criteria, we were left with 177 sessions and 1,385 trials. (NB: a handful of children participated more than once, resulting in 167 children with 177 sessions; all but two of the repeated sessions were excluded on other criteria. The two children who repeated sessions never responded more than once to any given trial thus all reported results reflect unique responses to unique trials.)

We then excluded individual trials using the following criteria: the participant was fussy and/or not attending to the study; someone other than the participant said one of the nouns featured in the scene being watched; the participant responded in another language other than English or Spanish; the participant retook the study and had seen the trial before, or the response video was corrupted. Training trials were also excluded from analysis.

Results

After excluding sessions and trials using the aforementioned criteria, 832 trials remained. We then excluded trials on which the participant either failed to produce any utterance at all or only produced utterances unrelated to the stimuli. We retained only trials on which the child attempted to refer to some aspect of the events. At this point in the data analysis, we were left with 162 trials and a total of 162 utterances relevant to Hypothesis 1 (e.g., data from all six trials) and 99 utterances relevant to Hypothesis 2 (data from the four Patient/Agent trials).

Hypothesis 1

We were interested in whether children selectively referred to information the parent already knew or information the parent did not. We ran two pre-registered, confirmatory analyses. First, we ran a McNemar’s test looking at whether children were selectively informative in their responses. Of the 162 utterances relevant to the scene, 95 mentioned only referents in common ground and 44 mentioned only informative referents (McNemar’s $\chi^2 = 17.99$, $df = 1$, $p < .0001$). We also compared two general linear mixed models (GLMM’s) one predicting the probability of being selective from whether the response was informative (0,1) and the other with only an intercept. The model that included the predictor explained more variance ($\chi^2 = 46.24$, $df = 1$, $p < .0001$).

To look specifically at whether children were more likely to generate informative utterances or refer to information in common ground, we also ran an analysis where we restricted the data to trials where children only mentioned informative referents or only mentioned uninformative referents. An intercept-only GLMM predicting the probability of mentioning an informative referent found that the probability of mentioning the informative referent was less than .5 (estimated probability = .30, 95% CI [.22, .40], $p < .001$).

In sum, across all analyses, and contrary to our predictions, children selectively referred to information in common ground.

Hypothesis 1b

Because we had hypothesized that children might be selectively informative, the Novel Event stimuli were designed such that, if anything, they were biased against our initial hypothesis. Given that children might be more interested in animate agents than objects in general, the movies were designed such that the parent always knew about the agents in the scene (the bear and the duck) and never knew about the objects that appeared when the parents eyes were closed (the ball and the shoe). In fact, children selectively referred to information in common ground across all trial types.

Critically, the agents were more informative in the Ambiguous agent movies; the objects were more informative in the Patient ambiguous movies. Thus even if children preferentially talk about agents, this should not lead them to selectively refer to information in common ground on the Ambiguous trials. To see whether the preference towards talk-

ing about information in common ground could be explained only as a bias towards talking about agents, we ran an exploratory analysis identical to the one above except that we eliminated the Novel Event trials from this analysis. The results replicated; the probability of mentioning the informative referent was less than .5 (estimated probability = .27, 95% CI [.15, .44], $p = .01$); children selectively referred to information in common ground.

Hypothesis 2

Our second question was whether children selectively provided disambiguating evidence to the parent. To test this, we looked only at the ambiguous trials, and only at whether children selectively referred to the disambiguating noun or the unambiguous noun in the scene (e.g., the active agent in the agent ambiguous scenes and the acted-on-patient in the patient ambiguous scenes rather than the only patient in the agent ambiguous scenes and the only agent in the patient ambiguous scenes). We ran two pre-registered, confirmatory analyses. Of the 99 utterances relevant to Hypothesis 2, 41 mentioned only the unambiguous noun and 28 mentioned only the disambiguating noun (McNemar's $\chi^2 = 2.09$, $df = 1$, $p = 0.15$). We also compared two GLMM's one predicting the probability of being selective from whether the child included the informative noun (0,1) and the other with only an intercept. The model that included the predictor did not explain more variance ($\chi^2 = .52$, $df = 1$, $p = .47$).

We also ran an analysis where we restricted the data to trials where children only mentioned the disambiguating nouns or only mentioned the unambiguous nouns in the Ambiguous trials. An intercept-only GLMM predicting the probability of mentioning the disambiguating referent found that the probability of doing so did not differ from .5. (estimated probability .41, 95% CI [.30, .52], $p = .12$).

Thus across analyses, we found no evidence that children selectively provided informative information to disambiguate the referents.

General Discussion

The current results suggest that toddlers' earliest productive language is audience-directed and sensitive to what the listener knows but in a way that is perhaps surprising with respect to adult models of informative communication. Adults selectively communicate the most informative thing to say given what the listener already knows (Deegan, 2023). Young toddlers did not; however, they also did not respond idiosyncratically. They selectively referred to information in common ground.

Although this tendency is at odds with rational communication in adults, referring to known and established referents in the scene may be a very effective strategy for a child breaking into their first language. It is clear that children learn new words best when engaged in joint attention with a caregiver (Yum Chen & Smith, 2012). This study suggests the intriguing possibility that the relationship between word production and joint attention may be bi-directional. That is, toddlers

may not only learn words best when they are labeled in the context of joint attention but may also track caregivers' attention and avoid talking about things their caregiver does not know.

One might wonder if there were more things in the scenes overall that the child can refer to that are in common ground than that are informative however the context mitigates against this. Children can be informative in all trials with a single, high frequency word in their early vocabularies (e.g., ball, shoes) and for the ambiguous trials exactly the same words (e.g., dog, apple) were informative or not depending on context. Moreover, very few utterances across the board ever involved children saying any word other than the nouns we used to describe the scenes. Thus, it is unlikely that children's preference for saying things in common ground over informative things is driven by the relative availability of things they could talk about.

A different concern is that in two of the six trials, the information in common ground was an agent (the bear and the duck) and the information unknown to the parent was an object (the ball and the shoe). In the ambiguous trials, what was unknown to the parent varied: in agent ambiguous trials, it was the identity of the acting agent; in patient ambiguous trials, it was the identity of the acted-on object. Looking only at these trials, children continued to selectively refer the information the parent already knew (the patient in the agent ambiguous trials and the agent in the patient ambiguous trials). This suggests that the preference to talking about information in common ground cannot be explained only by a bias towards talking about agents. Nonetheless, children did refer to agents more often throughout. In real world contexts, depending on what information is known only to the child, this preference for talking about agents could either augment or trade-off with children's interest in referring to information in common ground.

In addition to offering new insights about children's early productive language, the current results speak to the exciting possibilities opened up by online developmental research (Scott & Schulz, 2017; Sheskin et al., 2020). It is only in the past few years that scientists have been able to easily and affordably recruit relatively large populations of infants and test them at home when they are awake, alert, and at ease. Indeed, the senior authors first attempted this study as an in-person laboratory experiment over a decade ago and abandoned the attempt due to the sparsity of language production in toddlers at the one and two-word stages. What was experimentally intractable then is now giving us unprecedented views of language as it first emerges.

References

- Acredolo, L. P., & Goodwyn, S. W. (1985). Symbolic gesturing in language development: A case study. *Human development*, 28(1), 40–49.

- Baldwin, D. A. (1993). Early referential understanding: Infants' ability to recognize referential acts for what they are. *Developmental psychology*, 29(5), 832.
- Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, 118(1), 84–93.
- Barner, D., Chow, K., & Yang, S.-J. (2009). Finding one's meaning: A test of the relation between quantifiers and integers in language development. *Cognitive psychology*, 58(2), 195–219.
- Bates, J. E. (1976). Effects of children's nonverbal behavior upon adults. *Child Development*, 1079–1088.
- Begus, K., & Southgate, V. (2012). Infant pointing serves an interrogative function. *Developmental science*, 15(5), 611–617.
- Berk, S., & Lillo-Martin, D. (2012). The two-word stage: Motivated by linguistic or cognitive constraints? *Cognitive psychology*, 65(1), 118–140.
- Brandone, A. C., Horwitz, S. R., Aslin, R. N., & Wellman, H. M. (2014). Infants' goal anticipation during failed and successful reaching actions. *Developmental science*, 17(1), 23–34.
- Brown, H. D. (1973). Affective variables in second language acquisition. *Language learning*, 23(2), 231–244.
- Butcher, C., & Goldin-Meadow, S. (2000). 12 gesture and the transition from one-to two-word speech: When hand and mouth come together. *Lang Gest*, 2, 235.
- Capirci, O., Iverson, J. M., Pizzuto, E., & Volterra, V. (1996). Gestures and words during the transition to two-word speech. *Journal of Child language*, 23(3), 645–673.
- Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen-through 18-month-old infants differentially imitate intentional and accidental actions. *Infant behavior and development*, 21(2), 315–330.
- Csibra, G. (2003). Teleological and referential understanding of action in infancy. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1431), 447–458.
- Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9, 519–540.
- Foppolo, F., Guasti, M. T., & Chierchia, G. (2012). Scalar implicatures in child language: Give children a chance. *Language learning and development*, 8(4), 365–394.
- 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.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M., & Jäger, G. (2016). Probabilistic pragmatics, or why bayes' rule is probably important for pragmatics. *Zeitschrift für sprachwissenschaft*, 35(1), 3–44.
- Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019). How efficiency shapes human language. *Trends in cognitive sciences*, 23(5), 389–407.
- Goldin-Meadow, S., & Alibali, M. W. (2013). Gesture's role in speaking, learning, and creating language. *Annual review of psychology*, 64, 257–283.
- Goldin-Meadow, S., & Butcher, C. (2003). Pointing toward two-word speech in young children. *Pointing: Where language, culture, and cognition meet*, 85, 107.
- Goldin-Meadow, S., Goodrich, W., Sauer, E., & Iverson, J. (2007). Young children use their hands to tell their mothers what to say. *Developmental science*, 10(6), 778–785.
- Golinkoff, R. M. (1986). 'i beg your pardon?': The preverbal negotiation of failed messages. *Journal of child language*, 13(3), 455–476.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11), 818–829.
- Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.
- Harris, P. L., Bartz, D. T., & Rowe, M. L. (2017). Young children communicate their ignorance and ask questions. *Proceedings of the National Academy of Sciences*, 114(30), 7884–7891.
- Huang, Y. T., & Snedeker, J. (2009). Semantic meaning and pragmatic interpretation in 5-year-olds: Evidence from real-time spoken language comprehension. *Developmental psychology*, 45(6), 1723.
- Hurewitz, F., Papafragou, A., Gleitman, L., & Gelman, R. (2006). Asymmetries in the acquisition of numbers and quantifiers. *Language learning and development*, 2(2), 77–96.
- Iverson, J. M., Capirci, O., Volterra, V., & Goldin-Meadow, S. (2008). Learning to talk in a gesture-rich world: Early communication in Italian vs. American children. *First language*, 28(2), 164–181.
- Iverson, J. M., & Goldin-Meadow, S. (2005). Gesture paves the way for language development. *Psychological science*, 16(5), 367–371.
- Katsos, N., & Bishop, D. V. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition*, 120(1), 67–81.
- Kline, M., Schulz, L., & Gibson, E. (2017). Partial truths: Adults choose to mention agents and patients in proportion to informativity, even if it doesn't fully disambiguate the message. *Open Mind*, 2(1), 1–13.
- Krehm, M., Onishi, K. H., & Vouloumanos, A. (2014). I see your point: Infants under 12 months understand that pointing is communicative. *Journal of Cognition and Development*, 15(4), 527–538.

- Masur, E. F. (1982). Mothers' responses to infants' object-related gestures: Influences on lexical development. *Journal of child language*, 9(1), 23–30.
- Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, 13(4), 329–336.
- Noveck, I. A. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition*, 78(2), 165–188.
- O'Neill, D. K. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child development*, 67(2), 659–677.
- O'Neill, D. K., & Topolovec, J. C. (2001). Two-year-old children's sensitivity to the referential (in) efficacy of their own pointing gestures. *Journal of Child Language*, 28(1), 1–28.
- Papafragou, A., & Musolino, J. (2003). Scalar implicatures: Experiments at the semantics–pragmatics interface. *Cognition*, 86(3), 253–282.
- Papafragou, A., & Tantalou, N. (2004). Children's computation of implicatures. *Language acquisition*, 12(1), 71–82.
- Phillips, A. T., & Wellman, H. M. (2005). Infants' understanding of object-directed action. *Cognition*, 98(2), 137–155.
- Salley, B., Brady, N. C., Hoffman, L., & Fleming, K. (2020). Preverbal communication complexity in infants. *Infancy*, 25(1), 4–21.
- Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicatures: Alternatives and relevance. *Cognition*, 153, 6–18.
- Sonnenschein, S., & Whitehurst, G. J. (1984). Developing referential communication: A hierarchy of skills. *Child Development*, 1936–1945.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental science*, 10(1), 89–96.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(5), 675–691.
- Wetherby, A. M., Cain, D. H., Yonclas, D. G., & Walker, V. G. (1988). Analysis of intentional communication of normal children from the prelinguistic to the multiword stage. *Journal of Speech, Language, and Hearing Research*, 31(2), 240–252.
- Whitehurst, G. J., Sonnenschein, S., & Ianfolla, B. (1981). Learning to communicate from models: Children confuse length with information. *Child Development*, 507–513.