

Predicting Pragmatic Cue Integration in Adults' and Children's Inferences About Novel Word Meanings

Manuel Bohn¹, Michael Henry Tessler², Megan Merrick³, and Michael C. Frank³

¹ Department of Comparative Cultural Psychology, Max Planck Institute for Evolutionary Anthropology

² Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

³ Department of Psychology, Stanford University

Language is learned in complex social settings where listeners must reconstruct speakers' intended meanings from context. To navigate this challenge, children can use pragmatic reasoning to learn the meaning of unfamiliar words. A critical challenge for pragmatic reasoning is that it requires integrating multiple information sources, which have typically been studied separately. Here we study this integration process. First, we experimentally isolate two sources of pragmatic information: expectations about informative communication and common ground. Next, we use a probabilistic model of conversational reasoning to formalize how these information sources should be combined and how this process might develop. We use this model to generate quantitative predictions, which we test against new experimental data from 3- to 5-year-old children ($N = 243$) and adults ($N = 694$). Results show close alignment between model predictions and data. Furthermore, the model provided a better explanation of the data compared with simpler alternative models assuming that participants selectively ignore one information source. This work integrates distinct sets of findings regarding information sources for early language learning and suggests that pragmatic reasoning models can provide a quantitative framework for understanding developmental changes in language learning.


Keywords: language acquisition, social cognition, pragmatics, Bayesian modeling, common ground

Supplemental materials: <https://doi.org/10.1037/xge0001216.sup>

Successful communication often requires an understanding that extends beyond just the meaning of words. It takes pragmatic inference—context-sensitive reasoning about the speaker's intentions—to recover a speaker's intended meaning (Grice, 1991; Levinson, 2000; Sperber & Wilson, 2001). Contextual information

comes in many forms. On the one hand, there is information provided by the utterance itself.¹ Competent language users expect each other to communicate in a cooperative way such that speakers produce utterances that are relevant and informative. Semantic ambiguity can be resolved by reasoning about why the speaker produced these particular behaviors (H. H. Clark, 1996; Grice, 1991; Sperber & Wilson, 2001; Tomasello, 2008). On the other hand, there is information provided by common ground: Through interaction, interlocutors gradually build up a body of mutually shared knowledge and beliefs (Bohn & Köymen, 2018; E. V. Clark, 2015; H. H. Clark, 1996). Interlocutors expect each other to observe common ground and communicate in ways that are relevant to it.

Common ground and utterance-level information operate on different timelines. Utterances allow for in-the-moment inferences because they are composed of behaviors that the speaker chooses to express their intention in the here and now. On the other hand, common ground is built up over time through interaction (E. V. Clark, 2015; H. H. Clark, 1996). Nevertheless, the two

Manuel Bohn  <https://orcid.org/0000-0001-6006-1348>

We thank Jacqueline Quirke and Sabina Zacco for help with the data collection and Bria Long and Gregor Kachel for comments on an earlier version of the article. Manuel Bohn received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant 749229. Michael Henry Tessler was supported by a National Science Foundation SBE Postdoctoral Research Fellowship Grant 1911790. Michael C. Frank was supported by a Jacobs Foundation Advanced Research Fellowship and the Zhou Fund for Language and Cognition. Parts of Study 1 appeared in the proceedings of the 41st Annual Meeting of the Cognitive Science Society, Montreal, Canada, 2019. All experiments and analyses were preregistered (<https://osf.io/u7kxe/>; Bohn & Frank, 2018). Experiments, data, and analysis code are available in a public repository (<https://github.com/manuelbohn/mcc>). The authors declare no conflict of interest.

Correspondence concerning this article should be addressed to Manuel Bohn, Department of Comparative Cultural Psychology, Max Planck Institute for Evolutionary Anthropology, Deutscher Platz 6, 04103 Leipzig, Germany. Email: manuel_bohn@eva.mpg.de

¹ We use the terms utterance, utterance-level information or utterance-level cues to capture all cues that the speaker provides for their intended meaning. This includes direct referential information in the form of pointing or gazing, semantic information in the form of conventional word meanings as well as pragmatic inferences that are licensed by the particular choice of words or actions.

information sources are intimately related because utterances are embedded in common ground (Bohn & Frank, 2019; Tomasello, 2008). As a consequence, pragmatic reasoning in context always requires information integration. But how does this integration process develop? An even more important question is how this integration process develops. After all, young children have less knowledge of words and syntax than adults and cannot rely on the linguistic context to infer what a new word means. Instead, they heavily rely on pragmatic inferences that require integrating different information sources (Bohn & Frank, 2019; E. V. Clark, 2015; Tomasello, 2008).

In the current work, we try to answer these questions by formalizing information integration in a probabilistic model of pragmatic reasoning in development. In the remainder of this introduction, we describe the development of pragmatic inference and reasoning about common ground in childhood and then discuss the Rational Speech Act model, a formal framework that we use as the basis for our account of information integration.

Pragmatic Development in Childhood

Children make pragmatic inferences about intended meanings based on utterance-level information, both for language understanding and language learning (Bohn & Frank, 2019; E. V. Clark, 2009; Tomasello, 2008). Starting very early, preverbal infants expect adults to produce utterances (in the form of pointing gestures) in a cooperative way (Behne et al., 2005), and expect language to be carrying information (Vouloumanos et al., 2012). By age two, children are sensitive to the informativeness of communication (O'Neill & Topolovec, 2001). By age three, children can use this expectation to make pragmatic inferences (Stiller et al., 2015; E. J. Yoon & Frank, 2019) and to infer novel word meanings (Frank & Goodman, 2014). In this, they are not restricted to linguistic utterances: 3-year-olds also readily infer the referent of novel nonlinguistic behaviors and gestures (Bohn et al., 2019; Moore et al., 2015). And although older children continue to struggle with some complex pragmatic inferences until age five and beyond (Noveck, 2001), an emerging consensus identifies these difficulties as stemming from difficulties reasoning about the semantic scope of quantifiers rather than pragmatic deficits (Barner et al., 2011; Horowitz et al., 2018; Skordos & Papafragou, 2016). Thus, children's ability to reason about utterance-level pragmatics is present at least by ages three to five, and possibly substantially younger. In the present study, we focused on how children (and adults) make pragmatic inferences about word meanings based on the nonverbal aspects of an utterance: gaze and pointing gestures that accompany an unknown word. We adapted the procedure from Frank and Goodman (2014), in which adults and children learned a new word based on contrasting the pointing gesture a speaker produced with alternative gestures they could have produced but did not.

What is the role of common ground information in language understanding and learning? Before reviewing the developmental literature, we want to briefly clarify how we use the term *common ground* in this article. In the adult literature, common ground has traditionally been defined in recursive terms: to be part of common ground, some piece of information has to be not just known to both interlocutors but also known to both to be shared between them (H. H. Clark, 1996). Numerous studies probed the role of

sharedness of information and found that it plays a critical role in communicative interactions (Brown-Schmidt, 2009; Hanna et al., 2003; Heller et al., 2016; Mozuraitis et al., 2015). A major research focus has been to identify the processes and representations that support shared information and common ground (Galati & Brennan, 2021; Horton & Gerrig, 2016; Knutsen & Le Bigot, 2014; Wilkes-Gibbs & Clark, 1992; S. O. Yoon & Brown-Schmidt, 2014). In this article, we do *not* contribute to this discussion. Our focus is on how common ground—once established—is used and integrated with other pragmatic information sources. Thus, for the discussion that follows, we assume that the consequence of a direct interaction—with matching perspectives—is that information is shared and hence part of common ground (Bohn & Köymen, 2018). This perspective has been more prominent in the developmental literature (e.g., Matthews et al., 2006).

Construed this way, evidence for the use of common ground information by young children is strong already very early in life. For example, speaker-specific expectations guide how infants produce nonverbal gestures and interpret ambiguous utterances (Bohn et al., 2018; Saylor et al., 2011). For slightly older children, common ground also facilitates word comprehension and learning (Akhtar et al., 1996; Bohn, Le, Peloquin et al., 2021; Saylor et al., 2009; Sullivan et al., 2019).

In the present study, we will focus on two types of common ground information: discourse novelty and speaker preferences. Akhtar and colleagues (Akhtar et al., 1996; see also Diesendruck et al., 2004) showed that 2-year-olds learn a new word by reasoning about which objects are new to the speaker in the unfolding discourse—and the more likely to be referred to. Saylor and colleagues (Saylor et al., 2009) showed that 3- and 4-year-olds learn words by tracking the preference a speaker expressed during an ongoing interaction. In both studies, children's word learning was conditional on the identity of the speaker. When a different speaker produced a novel word, children did not map it onto the object that was novel to or preferred by the speaker they interacted with previously. Children formed speaker specific communicative expectations based on common ground built up during a preceding interaction (see also Graham et al., 2014; Matthews et al., 2010). In the present study, we build on this work and ask how these speaker specific expectations are integrated with additional pragmatic information provided by the utterance.

Information Integration in Pragmatic Language Learning

The work discussed so far highlights children's use of a single pragmatic information source or cue. Harnessing multiple—potentially competing—pragmatic cues poses a separate challenge. A central aspect of this integration problem is how to balance common ground information that is built up over the course of an interaction against information gleaned from the current utterance. Much less is known about whether and how children combine these types of information. Developmental studies that look at the integration of multiple information sources more generally find that children are sensitive to multiple sources from early on (Ganea & Saylor, 2007; Graham et al., 2017; Grosse et al., 2010; Khu et al., 2020; Matthews et al., 2006; Nilsen et al., 2009).

To take one example of integration processes, in a classic study, Nadig and Sedivy (2002) found that children rapidly integrate

information provided in an utterance (a particular referring expression) with the speaker's perspective (the objects the speaker can see). Integration is assumed to be occurring in that common ground constrains the later processing of language. However, how this constraining works is not specified—for example, presumably, these constraints are not absolute, implying some sort of graded combination. Furthermore, the information sources to be integrated in these studies are not all pragmatic in nature. For example, children's ability to pick out a referent following a noun reflects their linguistic knowledge and not necessarily their ability to reason about the speaker's intention in context. As a consequence, earlier work of this type—while providing important experimental evidence for information combination in childhood—still does not speak to the question of how (or even if) listeners integrate different forms of *pragmatic* information.

The Rational Speech Act Framework

Recent innovations in probabilistic models of pragmatic reasoning provide a quantitative method for addressing the problem of integrating multiple sources of contextual information. This class of computational models, which are referred to as Rational Speech Act (RSA) models (Frank & Goodman, 2012; Goodman & Frank, 2016) formalize the problem of language understanding as a special case of Bayesian social reasoning. A listener interprets an utterance by assuming it was produced by a cooperative speaker who had the goal to be informative. Being informative is defined as providing a message that would increase the probability of the listener recovering the speaker's intended meaning in context. This notion of contextual informativeness captures the Gricean idea of cooperation between speaker and listener, and provides a first approximation to what we have described above as utterance-level pragmatic information.

Within the RSA framework, one way to incorporate common ground is to treat it as a conversational prior. That is, previous social interactions result in a prior distribution over possible intended meanings for the current social interaction, in a manner specific to a particular speaker (i.e., Xs previous interactions with Y inform Xs current expectations about what Y is likely to talk about). However, RSA offers no formal description of how social interactions give rise to this speaker-specific prior distribution. These processes have been addressed computationally by Heller and colleagues (Heller et al., 2016; Mozuraitis et al., 2018; Ryskin et al., 2020), who examined how listeners identified the referent of ambiguous referring expressions. Their probabilistic model simultaneously considered the (differing) perspectives of both interlocutors and traded off between them. In principle, this model could be combined with RSA. However, as mentioned above, our focus was not on the internal structure of common ground, but on how information that is part of common ground is integrated with other pragmatic information. A natural locus for information integration within probabilistic models of pragmatic reasoning is the combination of the prior probability of a particular meaning and the likelihood of the current utterance being used to express that meaning. This feature of RSA models allows them to capture situations in which different information sources (e.g., common ground vs. utterance information) point to different meanings.

When integrated into variants of RSA that allow for uncertainty about word meaning (e.g., Frank & Goodman, 2014), this natural

weighting of prior and likelihood allows for the modeling of information integration. Despite of the broad use of probabilistic models in understanding word learning, other computational models of word learning have focused primarily on learning from cross-situational, co-occurrence statistics (Fazly et al., 2010; Frank et al., 2009) or describing generalizations about word meaning (Xu & Tenenbaum, 2007) and do not provide a clear route for pragmatic information integration.

The Current Study

We make use of this framework to study pragmatic cue integration across development. To this end, we adapt a method used in perceptual cue integration studies (Ernst & Banks, 2002): we make independent measurements of each cue's strength and then combine them using the RSA model described above to make independent predictions about conditions in which they either coincide or conflict. We preregister these quantitative predictions and test them against new data from adults and children.

We start by replicating previous findings with adults showing that listeners make pragmatic inferences based on nonlinguistic properties of utterances in isolation (Experiment 1). Then we show that adults make inferences based on common ground information (Experiments 2A and 2B). We use data from these experiments to estimate parameters and generate a priori predictions from RSA models about how utterance information and conversational priors should be integrated.

Models are most useful in comparison to one another. By examining differences in model fit as a function of different assumptions, we can make inferences about how specific choices lead to success or failure in capturing data. Here we consider three models that make different assumptions about the integration process. In the *integration model*, common ground and utterance-level information are integrated with one another as prior and likelihood (as described above). Our comparison models are lesioned models that assume that participants focus on one type of information and disregard the other whenever they are presented together. According to the *no conversational prior model*, participants focus only on the utterance information and in the *no informativeness model*, only the conversational prior is considered.² These models represent plausible alternative accounts; for example, Gagliardi et al. (2017) found that a model that selectively ignored parts of the input best captured children's use of statistical information during word learning. We compare predictions from the three models to new empirical data from experiments in which utterance and common ground information are manipulated simultaneously (Experiment 3 and 4).

After validating this approach with adults in Study 1, we apply the same model-driven experimental procedure to children (Study 2): We first show that children make pragmatic inferences based on utterance and common ground information separately (Experiments 5 and 6). Then we generate a priori model predictions and compare them to data from an experiment in which children are provided with both information sources (Experiment 7).

² We chose to refer to the alternative models by the information source they leave out (a) to highlight that they are lesioned versions of the *integration model* and (b) to avoid the impression that the *integration model* takes in qualitatively different information sources.

Taken together, this work makes three primary contributions: first, it shows that both adults and children integrate utterance-level information with common ground to make graded inferences about word meaning. Second, it provides an explicit theory of how this integration process proceeds and develops. Third, it uses Bayesian data analysis within the RSA framework to make a quantitative comparison of the evidence for competing hypotheses.

In a recent study (Bohn, Tessler, Merrick et al., 2021), we used a similar approach to study information integration in children. We used a related modeling framework, but focused on different experimental manipulations and different information sources (the previous study investigated how children's lexical knowledge integrates with discourse novelty in a mutual exclusivity context). The studies presented here conceptually extend this work in three critical ways. First, while the previous work focused exclusively on the interpretation of words, here we used multimodal utterances involving words, gestures, and gaze cues. Second, by including adults in our study, we show that the same modeling framework can be used to predict the behavior of adults and children. Finally—and perhaps most importantly—we probed the social nature of how common ground information interacts with other information sources by testing and modeling how the identity of the speaker influences the interpretation of the utterance. The prior work simply assumed that participants construe the experimental setting as a social interaction but this was neither tested nor modeled explicitly. Taken together, the current studies bring us closer to the complexity of real-world communicative interactions.

Study 1: Adults

Participants

Adult participants were recruited via Amazon Mechanical Turk (MTurk) and received payment equivalent to an hourly wage of ~\$9. Each participant contributed data to only one experiment. Experiment 1 and each manipulation of Experiment 2 had $N = 40$ participants. Sample size in Experiment 3 was $N = 121$. $N = 167$ participated in the experiments to measure the strong, medium and weak preference and novelty manipulations that went into Experiment 4. Finally, Experiment 4 had $N = 286$ participants. Sample sizes in all adult experiments were chosen to yield at least 120 data points per cell. All studies were approved by the Stanford Institutional Review Board (protocol no. 19960).

Materials

All experimental procedures were preregistered. The specific link for each experiment can be found in the online supplemental materials. Experimental stimuli are freely available in the following online repository: <https://github.com/manuelbohn/mcc>. All experiments were framed as games in which participants would learn words from animals. They were implemented in HTML/JavaScript as a website. Adults were directed to the website via MTurk and responded by clicking objects. For each animal character, we recorded a set of utterances (one native English speaker per animal) that were used to provide information and make requests. All experiments started with an introduction to the animals and two training trials in which familiar objects were requested (car and ball). Subsequent test trials in each condition were presented in a random order.

Analytic Approach

We preregistered sample sizes, inferential statistical analysis and computational models for all experiments. All deviations from the registered analysis plan are explicitly mentioned. All analyses were run in R (R Core Team, 2018). All p -values are based on two sided analysis. Cohen's d (computed via the function `cohensD`) was used as effect size for t -tests. Frequentist logistic GLMMs were fit via the function `glmer` from the package `lme4` (Bates et al., 2015) and had a maximal random effect structure conditional on model convergence. Details about GLMMs including model formulas for each experiment can be found in the online supplemental materials.

All cognitive models and model comparisons were implemented in WebPPL (Goodman & Stuhlmüller, 2014) using the R package `rwebppl` (Braginsky et al., n.d.). Probabilistic models were evaluated using Bayesian data analysis (Lee & Wagenmakers, 2014), also implemented in WebPPL. In Experiments 3, 4 and 7, we compared probabilistic models based on Bayes Factors—the ratio of the marginal likelihoods of each model given the data. Details on models, including information about priors for parameter estimation and Markov chain Monte Carlo settings can be found in the online supplemental materials available online. Code to run the models is available in the associated online repository.

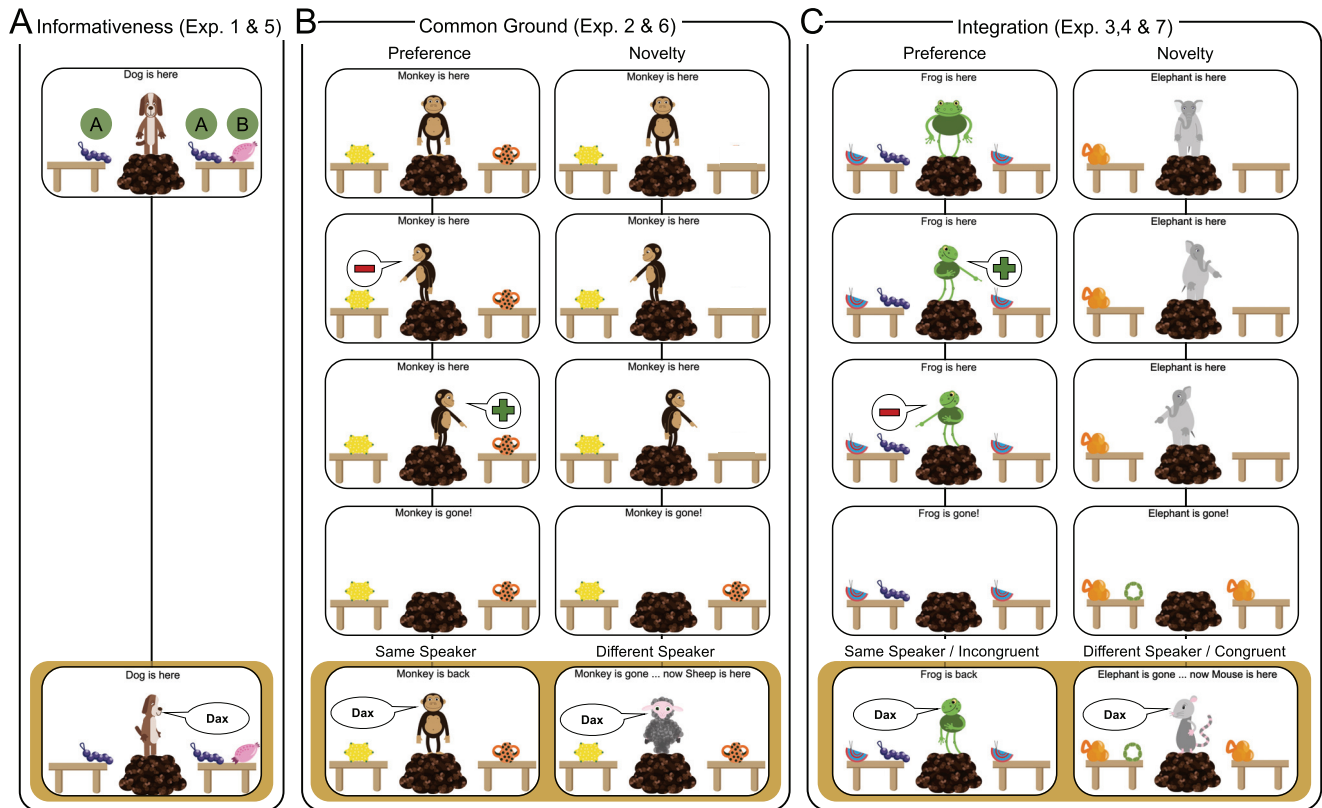
Experiment 1

Method

In Experiment 1, participants could learn which object a novel word referred to by assuming that the speaker communicated in an informative way (Frank & Goodman, 2014). The speaker was located between two tables, one with two novel objects, A and B, and the other with only object A (Figure 1A; side counterbalanced). At test, the speaker turned and pointed to the table with the two objects (A and B) and used a novel word to request one of them. The same utterance was used to make a request in all adult studies (“Oh cool, there is a [nonword] on the table, how neat, can you give me the [nonword]?”). Participants could infer that the word referred to object B via the counterfactual inferences that, if the (informative) speaker had wanted to refer to object A, they would have pointed to the table with the single object (this being the least ambiguous way to refer to that object). This inference rests on the assumption that the speaker is communicating about an object category or type (object A or B) and not a particular object token (e.g., object A on the left table). In the control condition, both tables contained both objects and no inference could be made based on the speaker's behavior. Participants received six trials, three per condition.

Results

Participants selected object B above chance in the test condition ($M = .74$, 95% confidence interval, CI of M [.65, .83], $t(39) = 5.51$, $p < .001$, $d = .87$) and more often compared with the control condition ($\beta = 1.28$, $SE = .29$, $p < .001$; see Figure 2). This finding replicates earlier work showing that adult listeners expect speakers to communicate in an informative way.

Figure 1*Schematic Experimental Procedure With Screenshots From the Adult Experiments*

Note. In all conditions, at test (bottom, golden shade), the speaker ambiguously requested an object using a nonword (e.g., “dax”). Participants clicked on the object they thought the speaker referred to. Speech bubbles represent prerecorded utterances. Informativeness (A) translated to making one object (object B) less frequent in context. Common ground (B) was manipulated by making one type of object preferred by or new to the speaker. Green plus signs represent utterances that expressed preference, and red minus signs represent utterances that expressed dis/preference. When expressing dis/preference, the object was temporarily enlarged (see main text for details). Integration (C) combined informativeness with the same/different speaker conditions for both common ground manipulations. Here we only show two (out of eight) integration conditions: preference - same speaker - incongruent (left) and novelty - different speaker - congruent (right). See the online article for the color version of this figure.

Experiment 2

Method

In Experiments 2A and 2B, we tested if participants use common ground information that is specific to a speaker to identify the referent of a novel word (Akhtar et al., 1996; Diesendruck et al., 2004; Saylor et al., 2009). In Experiment 2A, the speaker expressed a preference for one of two objects (Figure 1B, left). There was an object on each table. The animal introduced themselves, then turned to one of the tables (left or right: counterbalanced) and expressed either that they liked (“Oh wow, I really like that one”) or disliked (“Oh bleh, I really do not like that one”) the object before turning to the other side and expressing the respective other attitude. Next the animal disappeared and, after a short pause, either the same or a different animal returned and requested an object while facing straight ahead. Participants could use the speakers preference to identify the referent when the same speaker returned but not when a different speaker appeared whose preferences were unknown.

In Experiment 2B, common ground information came in the form of novelty (Figure 1B, right). There was an object on one of the tables,

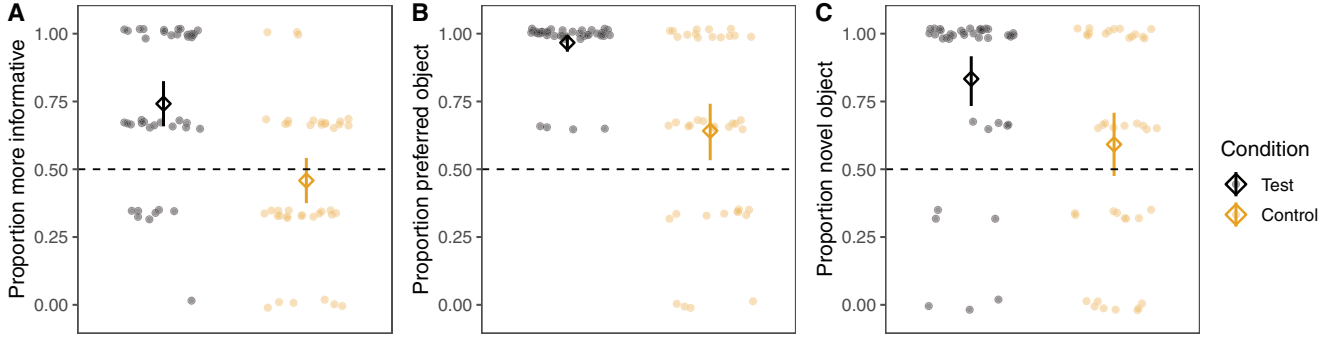
while the other was initially empty (side counterbalanced). The animal turned to one of the tables (left or right: counterbalanced) and commented either on the presence (“Aha, look at that”) or the absence (“Hm . . . , nothing there”) of an object before turning to the other side and commenting in a complementary way. Later, a second object appeared on the previously empty table. Then the speaker used a novel word to request one of the objects. The referent of the novel word could be identified by assuming that the speaker uses it to refer to the object that is new to them. This inference was not licensed when a different speaker returned to whom both objects were equally new. For both novelty and preference, participants received six trials, three with the same and three with the different speaker.

Results

In Experiment 2A, participants selected the preferred object above chance when the same speaker returned ($M = .97[.93; 1]$, $t(39) = 29.14$, $p < .001$, $d = 4.61$) and more so compared with when a different speaker returned ($\beta = 2.92$, $SE = .57$, $p < .001$).

In Experiment 2B, participants selected the novel object above chance when the same speaker made the request ($M = .83[.73;$

Figure 2
Results From Experiments 1 (A), 2A (B), and 2B (C) for Adults



Note. For preference and novelty, control refers to a different speaker (see Figure 1B). Transparent dots show data from individual participants (slightly jittered to avoid overplotting), diamonds represent condition means, error bars are 95% confidence intervals (CIs). Dashed line indicates performance expected by chance. See the online article for the color version of this figure.

.93], $t(39) = 6.77$, $p < .001$, $d = 1.07$) and more often compared with when a different speaker made the request ($\beta = 6.27$, $SE = 1.96$, $p = .001$, see Figure 2).

Modeling Information Integration

Experiments 1 and 2 confirmed that adults make pragmatic inferences based on information provided by the utterance as well as by common ground and provided quantitative estimates of the strength of these inferences for use in our model. We modeled the integration of utterance informativity and common ground as a process of socially guided probabilistic inference, using the results of Experiments 1 and 2 to inform key parameters of a computational model. The Rational Speech Act (RSA) model architecture introduced by Frank and Goodman (2012) encodes conversational reasoning through the perspective of a listener (“he” pronoun) who is trying to decide on the intended meaning of the utterance he heard from the speaker (“she” pronoun). The basic idea is that the listener combines his uncertainty about the speaker’s intended meaning—a prior distribution over referents $P(r)$ —with his generative model of how the utterance was produced: a speaker trying to convey information to him. To adapt this model to the word learning context, we enrich this basic architecture with a mechanism for expressing uncertainty about the meanings of words (lexical uncertainty)—a prior distribution over lexica $P(\mathcal{L})$ (Bergen et al., 2016).

$$P_L(r, \mathcal{L} | u) \propto P_S(u | r, \mathcal{L}) \cdot P(\mathcal{L}) \cdot P(r)$$

In the above equation, the listener is trying to jointly resolve the speaker’s intended referent r and the meaning of words (learning the lexicon \mathcal{L}). He does this by imagining what a rational speaker would say, given the referent she is trying to communicate and a lexicon. The speaker is an approximately rational Bayesian actor (with degree of rationality α), who produces utterances as a function of their informativity. The space of utterances the speaker could produce depends upon the lexicon $P(u | \mathcal{L})$; simply put, the speaker labels objects with the true labels under a given lexicon \mathcal{L} (see online supplemental materials available online for details):

$$P_S(u | r, \mathcal{L}) \propto \text{Informativity}(u; r)^\alpha \cdot P(u | \mathcal{L})$$

The informativity of an utterance for a referent is taken to be the probability with which a naive listener, who only interprets utterances according to their literal semantics, would select a particular referent given an utterance.

$$\text{Informativity}(u; r) = P(r | u) \propto P(r) \cdot \mathcal{L}_{\text{point}}$$

The speaker’s possible utterances are pairs of linguistic and nonlinguistic signals, namely labels, points, and gaze. Because the listener does not know the lexicon, the informativity of an utterance comes from the speaker’s point and gaze, the meaning of which is encoded in $\mathcal{L}_{\text{point}}$ and is simply a truth-function checking whether or not the referent is at the location picked out by the speaker’s point/gaze. Though the speaker makes her communicative decision assuming the listener does not know the meaning of the labels, we assume that in addition to pointing and/or gazing at the location, the speaker produces a label consistent with her own lexicon \mathcal{L} , described by $P(u | \mathcal{L})$. We assume that each label in the lexicon refers to an object type (e.g., object A) and not an object token (e.g., object A on the left table; Csibra & Gergely, 2009; see online supplemental materials for modeling details).

This computational model provides a natural avenue to formalize quantitatively how the informativeness of an utterance and conversational priors trade-off during word learning. As mentioned above, we treat common ground as a conversational prior over meanings, or types of referents, that the speaker might be referring to. That is, we assume that the interactions around the referents in the present context (i.e., preference or novelty; Experiment 2A and B) result in a speaker-specific prior distribution over referents. As mentioned in the introduction, we do not model the processes that give rise to this distribution explicitly, we simply measure them via the results from Experiment 2. For example, in Experiment 2, for the preference/same speaker participants chose the object the speaker liked (e.g., object B) with a proportion of .97 and the object the speaker disliked (object A) with .03. In Experiments 3 and 4, this measurement determined the prior distribution over objects in cases whenever the same manipulation was

used (preference/same speaker). Note that Experiment 3 involved three objects while Experiment 2 only involved two. We nevertheless used the exact proportions measured in Experiment 2 for each object as unnormalized probabilities in the prior. This approach conserved the relative relation between object types. Thus, when utterance and common ground information were aligned (i.e., object B was the more informative referent) the unnormalized distribution over objects was [$P(A_1) = 0.03$, $P(B) = 0.97$, $P(A_2) = 0.03$] and after normalizing it was [.03, .94, .03]. When information sources were dis-aligned (i.e., object A was the more informative referent), the object distribution was [.97, .03, .97] or [.49, .02, .49] after normalizing.

The in-the-moment, contextual informativeness of the utterance is captured in the likelihood term, whose value depends on the rationality parameter α . Assumptions about rationality may change depending on context and we used the data from Experiment 1 to specify α . We performed a Bayesian analysis in which we used the *integration model* (assuming equal prior probability over referents) with an unknown a priori value of α , and conditioned on the data from Experiment 1 to compute a posterior distribution over α ; in turn, the model generates posterior predictions for the proportion of correct responses in Experiment 1. We computed the maximum a posteriori (MAP) estimate and used this value for α to generate model predictions for Experiment 3 and 4. For additional information on parameter estimation we ask the reader to consult the online supplemental materials.

Based on these parameters, the model generates predictions for situations in which utterance and common ground expectations are jointly manipulated (Figure 1C). In the online supplemental materials, we include a worked example in which we walk the reader through the steps of computing model predictions from the parameters and the model equations. We recommend going through this example to get a better understanding of the model.

In addition to the parameters fit to the data from previous experiments, we include a noise parameter, which can be thought of as reflecting the cost that comes with handling and integrating multiple information sources. Technically, the noise parameter represents the proportion of responses better explained by a process of random guessing than by pragmatics; we estimate this parameter from the observed data (Experiment 3). Including the noise parameter greatly improved the model fit to the data (see online supplemental materials for details). We did not preregister the inclusion of a noise parameter for Experiment 3 but did so for all subsequent experiments.

Experiment 3

Method

In Experiment 3, we combined the procedures of Experiment 1 and 2A or 2B. The test setup was identical to Experiment 1, before making a request, the speaker interacted with the objects so that some of them were preferred by or new to them (Figure 1C). This combination resulted in two ways in which the two information sources could be aligned with one another. In the congruent condition, the object that was the more informative referent in the present context was also the one that was preferred by or new to the speaker. In the incongruent condition, the object that was the less informative referent in the present context was the one that was preferred by or new to the speaker.

In the preference condition, the speaker turned to one table, pointed to the object and expressed either liking or disliking using the same utterances as in Experiment 2A. To make it clear which object the speaker was referring to while pointing to the table with two objects, the referred-to object was temporarily enlarged. Whether the speaker first turned to the table with a single object or to the one with the two objects was counterbalanced. In the congruent condition, the preferred object was also the one that was unique to the table with the two objects. In the incongruent condition, the preferred object was also present on the other table.

In the novelty condition, the scene began with only one object on one of the tables. After commenting on the presence and absence of objects in the same way as in Experiment 2B, the speaker disappeared and two additional objects appeared, one on the previously empty table and one on the other table. Whether the speaker first turned to the empty table or to the one with an object was counterbalanced. In the congruent condition, two different objects appeared so that the object that was unique to the table with the two objects was also new in context. In the incongruent condition, two identical objects appeared so that the object that was unique to the table was the one that was old in context. The test event was the same for preference and novelty: the speaker turned to the table with the two objects and used the same request as in Experiment 1.

Taken together, there were 2 (novelty or preference) \times 2 (same or different speaker) \times 2 (congruent or incongruent) = 8 conditions in Experiment 3. For each of these eight conditions, we generated model predictions using the modeling framework introduced above. To arbitrate between hypotheses about how information is integrated, we compared the three models introduced in the introduction: The *integration model* in which both information sources are flexibly combined, the *no conversational prior model* that focused only on utterance-level information and the *no informativeness model* that focused only on common ground information.

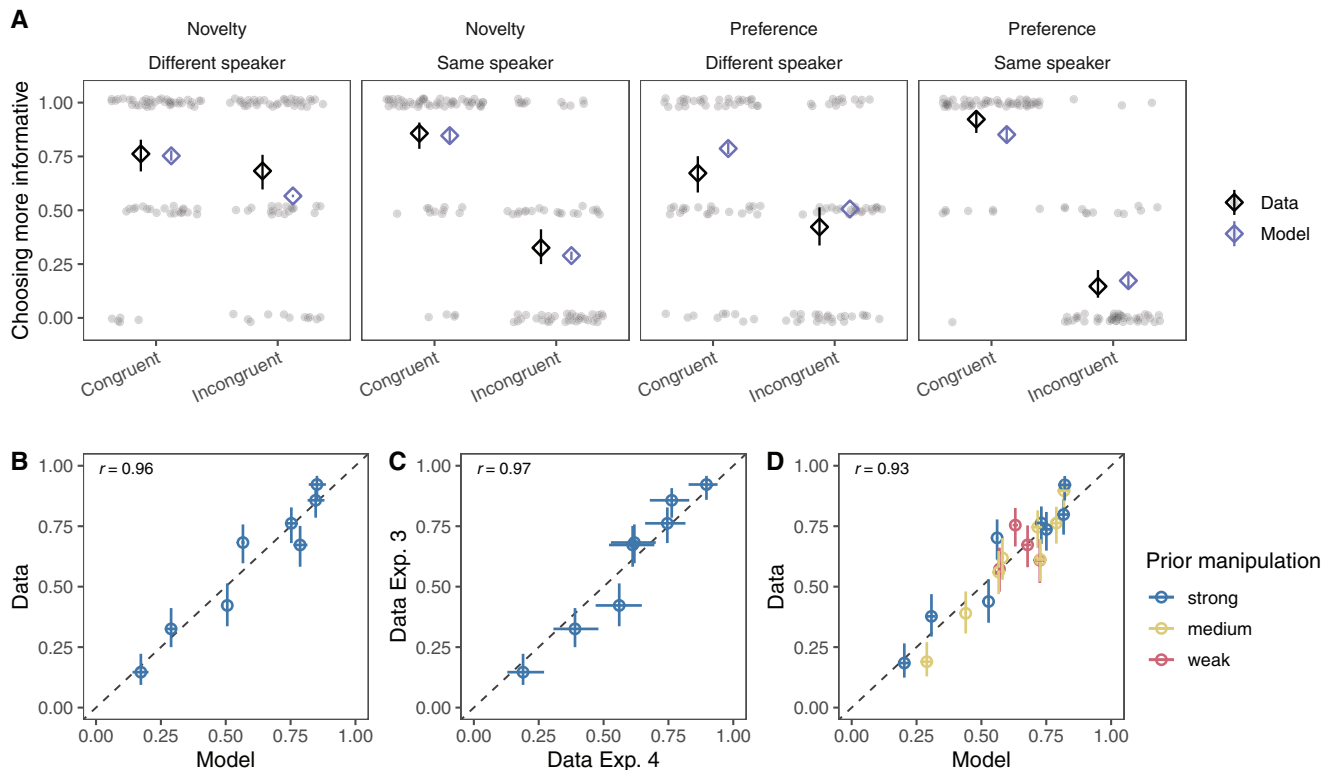
Participants completed eight trials for one of the common ground manipulations with two trials per condition (same/different speaker \times congruent/incongruent). Conditions were presented in a random order. We discuss and visualize the results as the proportion with which participants chose the more informative object (i.e., the object that would be the more informative referent when only utterance information is considered).

Results

As a first step, we used a GLMM to test whether participants were sensitive to the different ways in which information could be aligned. We found that participants distinguished between congruent and incongruent trials when the speaker remained the same (model term: alignment \times speaker; $\beta = -2.64$, $SE = .48$, $p < .001$). Thus, participants were sensitive to the different combinations of manipulations.

As a second step, we compared the cognitive model predictions to the data. Participants' average responses were highly correlated with the predictions from the *integration model* in each condition (Figure 3B). When comparing models, we found that model fit was unambiguously better for the *integration model* compared with the *no conversational prior model* (Bayes Factor (BF) = 4.2e + 53) or the *no informativeness model* (BF = 2.5e + 34), suggesting that participants considered and integrated both sources of information.

Figure 3
Results From Experiment 3 and 4 for Adults



Note. Data and model predictions by condition for Experiment 3 (a). Transparent dots show data from individual participants (slightly jittered to avoid overplotting), diamonds represent condition means. Correlation between model predictions and data in Experiment 3 (b), between data in Experiment 3 and the direct replication in Experiment 4 (c) and between model predictions and data in Experiment 4 (d). Coefficients and p -values are based on Pearson correlation statistics. Error bars represent 95% highest density intervals (HDIs). See the online article for the color version of this figure.

Finally, we examined the noise parameter for each model. The estimated proportion of random responses according to the *integration model* was .30 (95% highest density interval (HDI) = [.23–.36]). This parameter was substantially lower for the *integration model* compared with the alternative models (*no conversational prior model*: .60 [.46–.72]; *no informativeness model*: .41 [.33–.51]), lending additional support to the conclusion that the *integration model* better captured the behavioral data. Rather than explaining systematic structure in the data, the alternative models achieved their best fit only by assuming a very high level of noise.

Experiment 4

Method

To test the scope of the *integration model*, we first replicated and then extended the results of Experiment 3 to a broader range of experimental conditions. Specifically, we manipulated the strength of the common ground information (three levels—strong, medium, and weak—for preference and two levels—strong and medium—for novelty) by modifying the way the speaker interacted with the objects before the request. The procedural details and statistical analysis for these manipulations are described in the online supplemental materials. For Experiment 4, we paired each level of

prior strength manipulation with the informativeness inference in the same way as in Experiment 3. This resulted in a total of 20 conditions, for which we generated a priori model predictions in the same way as in Experiment 3. That is, we conducted a separate experiment for each level of prior strength and common ground manipulation to estimate the prior probability of each object following this particular manipulation (analogous to Experiment 2). This prior distribution was then passed through the model for the congruent and incongruent conditions, resulting in a unique prediction for each of the 20 conditions. Given the graded nature of the prior manipulations, Experiment 4 basically tests how well the model performs with different types of prior distributions.

The strong prior manipulation in Experiment 4 was a direct replication of Experiment 3 (see Figure 3C). Each participant was randomly assigned to a common ground manipulation and a level of prior strength and completed eight trials in total, two in each unique condition in that combination.

Results

The direct replication of Experiment 3 within Experiment 4 showed a very close correspondence between the two rounds of data collection (see Figure 3C). GLMM results for Experiment 4 can be found in the online supplemental materials available online. Here we focus on the analysis based on the cognitive models.

Model predictions from the *integration model* were again highly correlated with the average response in each condition (see Figure 3D). We evaluated model fit for the same models as in Experiment 3 and found again that the *integration model* fit the data much better compared with the *no conversational prior* ($BF = 4.7e + 71$) or the *no informativeness model* ($BF = 8.9e + 82$). The inferred level of noise based on the data for the *integration model* was .36 [.31–.41], which was similar to Experiment 3 and again lower compared with the alternative models (*no conversational prior model*: .53 [.46–.62]; *no informativeness model*: .67 [.59–.74]).

Study 2: Children

The previous section showed that competent language users flexibly integrate information during pragmatic word learning. Do children make use of multiple information sources during word learning as well? How does this integration emerge developmentally? While many verbal theories of language learning imply that such integration does occur, the actual process of integration has rarely been described nor tested in detail. Here we provide an explanation in the form of our *integration model* and test if it is able to capture children's word learning. Embedded in the assumptions of the model is the idea that developmental change occurs via changes in the strengths of the individual inferences, which leads to a change in the strength of the integrated inference. As a starting point, our model assumes developmental continuity in the integration process itself (Bohn & Frank, 2019), though this assumption could be called into question by a poor model fit. The study for children followed the same general pattern as the one for adults. We generated model predictions for how information should be integrated by first measuring children's ability to use utterance-level and common ground information in isolation when making pragmatic inferences. We then adapted our model to study developmental change: We sampled children continuously between 3.0 and 5.0 years of age—a time in which children have been found to make the kind of pragmatic inferences we studied here (Bohn & Frank, 2019; Frank & Goodman, 2014)—and generated model predictions for the average developmental trajectory in each condition.

Participants

Children were recruited from the floor of the Children's Discovery Museum in San Jose, CA. Parents gave informed consent and provided demographic information. Each child contributed data to only one experiment. We collected data from a total of 243 children between 3.0 and 5.0 years of age. We excluded 15 children due to less than 75% of reported exposure to English, five because they responded incorrectly on 2/2 training trials, three because of equipment malfunction, and two because they quit before half of the test trials were completed. The final sample size in each experiment was as follows: $N = 62$ (41 girls, $M_{\text{age}} = 4$) in Experiment 5, $N = 61$ (28 girls, $M_{\text{age}} = 3.99$) in Experiment 6 and $N = 96$ (54 girls, $M_{\text{age}} = 3.96$) in Experiment 7. For Experiment 5 and 6, we also tested 2-year-olds but did not find sufficient evidence that they use utterance and/or common ground information in the tasks we used to justify investigating their ability to integrate the two. Sample sizes in all experiments were chosen to yield at least 80 data points in each cell for each age group.

Materials

All procedures, sample sizes and data analyses were again pre-registered (see online supplemental materials for the specific links for each experiment); materials, data, and analysis code can be found in the associated repository (see Study1). Experiments were implemented in the same general way as for adults. Children were guided through the games by an experimenter and responded by touching objects on the screen of an iPad tablet (Frank et al., 2016).

Experiment 5

Method

Experiment 5 for children was modeled after Frank and Goodman (2014). Instead of appearing on tables, objects were presented as hanging in trees, which facilitated the depiction of a speaker pointing to distinct locations. After introducing themselves, the animal turned to the tree with two objects and said: "This is a tree with a [nonword], how neat, a tree with a [nonword]". Next, the trees and the objects in them disappeared and new trees replaced them. The two objects from the tree the animal turned to previously were now spread across the two trees (one object per tree, position counterbalanced). While facing straight, the animal first said "Here are some more trees" and then asked the child to pick the tree with the object that corresponded to the novel word ("Which of these trees has a [nonword]?"). Children received six trials in a single test condition.

Results

To compare children's performance to chance level, we binned age by year. Four-year-olds selected the more informative object (i.e., the object that was unique to the location the speaker turned to) above chance ($M = .62$ [.53; .71], $t(29) = 2.80$, $p = .009$, $d = .51$). Three-year-olds, on the other hand, did not ($M = .46$ [.41; .52], $t(31) = -1.31$, $p = .198$, $d = .23$). Consequently, when we fit a GLMM to the data with age as a continuous predictor, performance increased with age ($\beta = .38$, $SE = .11$, $p < .001$; see Figure 4). Thus, children's ability to use utterance information in a word learning context increased with age.

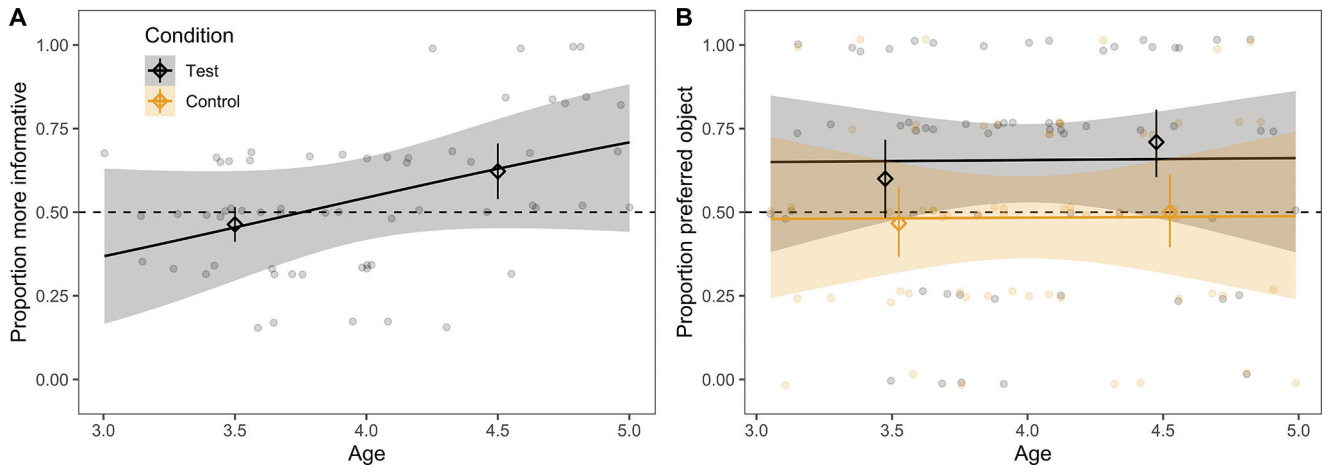
Experiment 6

Method

In Experiment 6, we assessed whether children use common ground information to identify the referent of a novel word. We tested children only with the preference manipulation.³ The procedure for children was identical to the preference manipulation for

³ We initially tested children with the novelty as well as the preference manipulation. We found that children made the basic inference in that they selected the object that was preferred by or new to the speaker, but found little evidence that children distinguished between requests made by the same speaker or a different speaker in the case of novelty. This finding contrasts with earlier work (Diesendruck et al., 2004). Because our focus was on how children integrate informativeness and conversational priors resulting from common ground, we did not follow up on this finding but dropped the novelty manipulation and focused on preference for the remainder of the study. We studied information integration in children using the novelty manipulation in a different study (Bohn, Tessler, Merrick et al., 2021).

Figure 4
Results From Experiment 5 (A) and 6 (B) for Children



Note. For preference, control refers to the different speaker condition (see Figure 1B). Transparent dots show data from individual participants (slightly jittered to avoid overplotting), regression lines show fitted generalized linear models with 95% confidence intervals (CIs). Dashed line indicates performance expected by chance. See the online article for the color version of this figure.

adults. Children received eight trials, four with the same and four with a different speaker.

Results

Four-year-olds selected the preferred object above chance when the same speaker made the request ($M = .71[.61, .81]$, $t(30) = 4.14$, $p < .001$, $d = .74$), whereas 3-year-olds did not ($M = .60[.47, .73]$, $t(29) = 1.62$, $p = .117$, $d = .30$). When the different speaker made the request, performance was at chance level in both age groups (3-year-olds: $M = .47[.36, .57]$; 4-year-olds: $M = .50[.39, .61]$). When we fit a GLMM to the data with age as a continuous predictor, we found an effect of speaker identity ($\beta = .89$, $SE = .24$, $p < .001$) but no effect of age ($\beta = .02$, $SE = .16$, $p = .92$) or interaction between speaker identity and age ($\beta = -.01$, $SE = .23$, $p = .97$; see Figure 4). Thus, children across the age range used common ground information to infer the referent of a novel word.

Modeling Information Integration in Children

Model predictions for children were generated using the same model described above for adults. To incorporate developmental change in the model, we allowed the rationality parameter α (that controls the degree of speaker informativeness) and the prior distribution over objects (a proxy for common ground) to change with age.

We defined α for a given age via a simple linear regression. Thus, instead of inferring a single value across age, we used the data from Experiment 5 to find the intercept (β_0^α) and slope (β_1^α) that best described the developmental trajectory in those data. As for adults, we inferred α via the *integration model* with equal prior probabilities for each object. We computed a posterior distribution for the intercept and the slope of this regression function. In Experiment 7, the speaker optimality parameter for a child of a given age was computed by taking the MAP for the intercept and adding the MAP for the slope times the child's age i : $\alpha_i = \beta_0^\alpha + i \cdot \beta_1^\alpha$.

To estimate the prior distribution over objects, we used the data from Experiment 6 to model the intercepts ($\beta_{0,j}^p$) and slopes ($\beta_{1,j}^p$)

that best described the developmental trajectories in the data for each of the two (j) conditions. This allowed us to generate prior distributions over objects in the cognitive model that were sensitive to the child's age. We used a simple logistic regression to find the intercept and slope (MAP of posterior distribution) that best described children's performance in the two conditions of Experiment 6. In Experiment 7, the prior probability for an object was computed by taking the intercept for the respective condition j , adding the slope times the child's age i and then using a logistic transformation to convert the outcome into proportions: $\rho_{i,j} = \text{logistic}(\beta_{0,j}^p + i \cdot \beta_{1,j}^p)$. Because these proportions corresponded to a two-object scenario, they were then converted to the three-object scenario by assuming equal probabilities for objects of the same type and normalizing. The overall distribution depended on the alignment of information sources in the same way as it did for adults. The online supplemental materials provides additional information on the parameter estimation.

These parameter settings were then used to generate age-sensitive model predictions in 2 (same or different speaker) \times 2 (congruent or incongruent) = 4 conditions. As for adults, all models included a noise parameter, which was estimated based on the data of Experiment 7.

Experiment 7

Method

In Experiment 7, we combined the procedures of Experiments 5 and 6 and collected new data from children between 3.0 and 5.0 years of age in each of the four conditions (Figure 1C). We again inserted the preference manipulation into the setup of Experiment 5. After greeting the child, the animal turned to one of the trees, pointed to an object—that was temporarily enlarged and moved closer to the animal—and expressed either liking or disliking. Then, the animal turned to the other tree and expressed the opposite attitude (disliking or liking) for the other kind of object. Next, the animal disappeared and either the same or a different animal

returned. We counterbalanced whether the speaker first turned to the tree with the two objects or the tree with a single object. The remainder of the trial was identical to the request phase of Experiment 5. Children received eight trials, two per condition (same/different speaker \times congruent/incongruent) in a randomized order.

Results

As a first step, we used a GLMM to test whether children were sensitive to the different ways in which information could be aligned. Children's propensity to differentiate between congruent and incongruent trials for the same or a different speaker increased with age (model term: age \times alignment \times speaker; $\beta = -.89$, $SE = .36$, $p = .013$).

Analyses comparing the model predictions from the probabilistic models to the data suggest that children flexibly integrate conversational priors and informativity information. Furthermore, this integration process is accurately captured by the *integration model* at least for 4-year-olds. For the correlational analysis, we binned model predictions and data by year. There was a substantial correlation between the predicted and measured average response for 4-year-olds, but less so for 3-year-olds (Figure 5B). One of the reasons for the latter was the low variation between conditions. For the model comparison, we treated age continuously. As with adults, we found a much better model fit for the *integration model* compared with the *no conversational prior* (BF = 551) or the *no informativeness model* (BF = 8042).

The inferred level of noise based on the data for the *integration model* was .51 (.26–.77), which was lower compared with the alternative models considered (*no conversational prior model*: .81 [.44–1.00]; *no informativeness model*: .99 [.88–1.00]) but numerically higher than that of adults (see Figure 5C).

The high level of inferred noise moved the model predictions for children in all conditions close to chance level. We compared two additional sets of models with different parameterizations of the noise parameter that emphasized differences between conditions in the model predictions more (see online supplemental materials and Figure 5A). This analysis was not preregistered. Parameter free models did not include a noise parameter and developmental noise models allowed the noise parameter to change with age.

In each case, the *integration model* provided a better fit compared with the alternative models (parameter-free: *integration* vs. *no conversational prior* BF = 334, *integration* vs. *no informativeness* BF = 6.4e + 29; developmental noise: *integration* vs. *no conversational prior* BF = 1926; *integration* vs. *no informativeness* BF = 1.8e + 07). The developmental noise parameter for the *integration model* decreased with age, suggesting that for younger children, the model explained the data by assuming a high rate of random guessing, whereas, for older children, the model explained the data by virtue of the processes that are implemented in its structure (see Figure 5D).

General Discussion

Integrating multiple sources of information is an integral part of human communication. To infer the intended meaning of an utterance, listeners must combine their knowledge of communicative conventions (semantics and syntax) with social expectations about their interlocutor. This integration is especially vital in early

language learning when the different varieties of pragmatic information are among the most important sources of information for learners who may not yet have mastered syntax and semantics (Bohn & Frank, 2019). But how are pragmatic cues integrated during word learning? Here we used a Bayesian cognitive model to formalize this integration process. We studied how utterance-level (Gricean) expectations about informative communication are integrated with conversational priors (resulting from common ground). Adults' and children's learning was best predicted by a model in which both sources of information traded-off flexibly. Alternative models that considered only one source of information made substantially worse predictions.

We begin our discussion by contextualizing our modeling findings and then turn to the developmental implications of the modeling results. We end with some discussion of the limitations of our experimental tasks and our computational model.

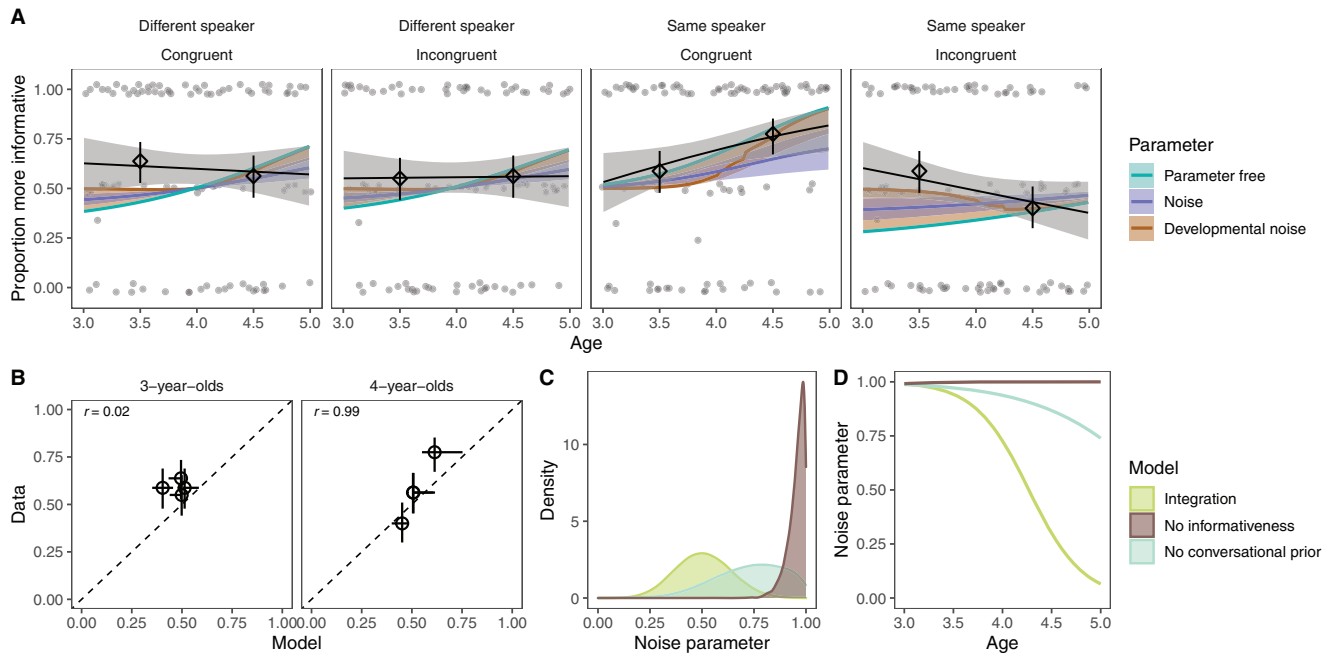
Modeling Contributions

Cue integration in language processing has been extensively studied in recent decades, but the focus of this work has usually been on how adults combine perceptual, semantic or syntactic information (Hagoort et al., 2004; Kamide et al., 2003; Özyürek et al., 2007; Tanenhaus et al., 1995). We extend the study of linguistic cue integration to pragmatics. However, we present a substantive theory of the integration process itself. Real world language comprehension and learning happens in socially dynamic and complex situations that inevitably require integrating multiple pragmatic information sources. The *integration model* provides a formal description of the process of information integration, at least at the computational level of analysis (Marr, 1982). As such, our work complements theorizing about information integration in other domains of language comprehension (e.g., Fourtassi & Frank, 2020; McClelland et al., 2006; Smith et al., 2017).

All of the models we compared here integrated some explicit structure, rather than (e.g., simply weighing information sources by some ratio. Predictions result from models of the task rather than simply models of the data. That is, inferences are not computed separately by the modeler and specified as inputs to a regression model, but instead are the results of an integrated process that operates over a (schematic) representation of the experimental stimuli. Further, our models are variants derived from the broader RSA framework, which has been integrated into larger systems for language learning in context (Cohn-Gordon et al., 2018; Monroe et al., 2017; Wang et al., 2016).

What does the *integration model* tell us about the way in which information is integrated? In a broad sense, it explores the idea that information integration follows a rational Bayesian procedure in that the listener's a posteriori belief is optimal given their a priori beliefs and all available information sources. That is, the listener's inferences about the speaker's intended referent are optimal given the implications of both the preceding social interaction and the speaker's utterance (Chater & Oaksford, 1999; Griffiths & Tenenbaum, 2006; Yuille & Kersten, 2006). The results seem to support this conception of the integration process: First, for adults and older children, the model predicted the data with high accuracy, including seemingly counterintuitive patterns (see below). Second, none of the lesioned models—that assumed that

Figure 5
Results From Experiment 7 for Children



Note. (A) Model predictions (with 95% highest density intervals [HDIs]) and data across age in the four conditions. Transparent black dots show data from individual participants and black lines show conditional means of the data with 95% confidence interval (CI). Black diamonds show the mean of the data for age bins by year and error bars show 95% CIs. (B) Correlation between model predictions (with noise parameter) and condition means binned by year (with 95% HDIs). For 4-year-olds, two conditions yielded the same data means and model predictions and are plotted on top of each other. (C) Posterior distribution of the noise parameter for the different models. (D) Developmental trajectory of the noise parameter for the three developmental noise models; trajectories are based on MAPs of the posterior distribution for the intercept and slope. See the online article for the color version of this figure.

participants disregard one of the information sources—produced equally good predictions.

With respect to the specific information sources we manipulated, the model assumes that the informativeness of an utterance depends on the person-specific conversational priors (resulting from common ground). Broadly speaking, it formalizes the view that common ground is the starting point for the listener to determine how informative a given utterance is. As such, the model gives an explicit and formal description of how common ground may constrain the processing of utterances—something that was unspecified in earlier experimental work on information integration (e.g., Khu et al., 2020; Nadig & Sedivy, 2002).

This conception of information integration also explains the seemingly counterintuitive predictions of the model. For example, one might expect the model to predict a chance level performance in the same speaker—incongruent conditions because the two cues “pull” the listener in opposite directions. Instead, the model predicts a performance below chance, favoring the object implicated by the prior—that also matches adults’ responses. This subtle prediction emerges because the listener assumes that the speaker takes the conversational prior shared between the speaker and the (naïve) listener as a starting point when computing the effect of each utterance. As a consequence, when prior interactions strongly implicate one object as the more likely referent, the speaker reasons that this object will be the inferred referent of any semantically

plausible utterance, even when the same utterance would point to a different object in the absence of a conversational prior.

Taken together, our model advances classic theories on pragmatic language comprehension (Grice, 1991; Sperber & Wilson, 2001) and learning (Bruner, 1983; Tomasello, 2009) by providing an explicit description of the integration process. The model offers a computational description of how information may be integrated during pragmatic word learning. Future work will be required to understand if and how the RSA framework, which typically makes aggregate predictions at the group level, can be used to understand the moment-by-moment and trial-by-trial behavior of individuals. Individuals’ behavior could well result from a heuristic approximation to full RSA-type inference. New methods will likely need to be developed to evaluate this conjecture.

Developmental Findings

The correlational analysis showed that the *integration model* accurately predicted information integration in 4-year-olds. However, the model did not successfully describe 3-year-olds’ inferences; thus, it is possible that they were not able to integrate information sources. Our findings are also consistent with a simpler explanation, namely that the overall weaker responses we observed in the independent measurement experiments (Experiments 5 and 6), combined with some noise in responding, led the

younger age group to appear relatively random in their responses. As a consequence, there was not much variation in the group-level performance of 3-year-olds for the model to explain. The results of the model comparison also support this interpretation. Here, we treated age continuously and found that the *integration model* provided the best fit across the entire age range. Taken together, we may say that as soon as children are sufficiently sensitive to the individual information sources, the *integration model* accurately describes the way that information is integrated. To strengthen this interpretation, future work should use tasks (or age groups) that show a clear and strong response for each information source.

Our model presents a substantive theory of the development of information integration during word learning. The primary source of developmental change in our model is age-related changes in the propensity to make individual inferences. As they get older, children expect speakers to be more informative and more likely to observe common ground. Still, the process by which the two information sources are integrated at any given age is assumed to be the same. The alternative models we considered are plausible accounts of other ways in which information could be integrated, but they also share the assumption of developmental continuity with respect to the integration process. Thus, in future work, it would be important to explore alternative models for the development of the integration process; one possible candidate would be a model in which the integration process itself changes with age.

Bohn and colleagues (Bohn, Tessler, Merrick et al., 2021) explored such an alternative integration model. They used a similar modeling framework but studied different information sources. In addition to an integration model that is structurally comparable to the one described here, they formulated a biased integration model that assumed that children are biased toward some information sources over others. In a developmental version of this biased model, they assumed that the strength of this bias changes with age, which represents an alternative view on development. However, when directly compared, the integration model explained their data better.

The developmental noise model reported for Experiment 7 offers yet another way to address the question of developmental change. This model estimates a developmental trajectory for the proportion of responses that are better explained by random guessing than by the model structure (see Figure 5D). If such a data analytic model would find that model fit is comparable for younger and older children but that the noise parameter through which this fit is achieved decreases with age, we might conclude that cognitive abilities that pertain to task demands are the major locus of change rather than abilities that have to do with integrating information. In the developmental noise model in Experiment 7, we found that noise decreased with age but, at the same time, that the resulting model fit was substantially worse for 3-year-old children. As mentioned in the previous paragraph, we think that a lack of sensitivity to the individual information sources, rather than a failure to integrate them, is the reason for this poor model fit in the younger age group. The strongest evidence for developmental changes in integration would come in a case where younger children showed evidence of above/below-chance judgment in the combined task (Experiment 7) that was distinct from that predicted by the two above/below-chance component tasks (Experiment 5 and 6). Such a comparison would require more precision (either

via more trials or more participants) than our current experiment affords, however.

Limitations

An important limitation of our experimental work is that we studied a single population of American-English speaking children and adults. It is unclear how our results would transfer to different populations. We expect substantial variation across cultures and languages in how sensitive children are to the different information sources. The few studies that investigated pragmatic inferences similar to the ones we observed in our study found substantial variation (Fortier et al., 2018; Su & Su, 2015; Zhao et al., 2021). Extending this work to study information integration will be a very valuable avenue for future research. Nevertheless, we think that our modeling framework provides an excellent tool to study universalities and differences in information integration.

Our modeling work is limited in that we did not model the social-cognitive processes that underlie common ground. Instead, we assumed that the interactions that preceded the utterance (and presumably constitute common ground) result in a person-specific conversational prior. From a modeling perspective, this approach treats common ground as equivalent to more basic manipulations of contextual salience (e.g., in Frank & Goodman, 2012). Thus, our model would not differentiate between a situation in which an object would be salient because it has been the focus of an interaction and one in which it would be more salient because it was big or colorful. Thus, evoking common ground in this context is largely backed-up by the experimental tasks: the fact that participants (children and adults) were sensitive to the identity of the speaker tells us that the contextual salience of the referents resulted from a process of social reasoning. Thus, we feel confident in saying that our results speak to how participants integrated different sources of pragmatic information. Based on a process model of common ground, one could further specify how common ground information (i.e., social context) interacts with other contextual information (Degen et al., 2015; Heller et al., 2016; Mozuraitis et al., 2018; Tessler et al., 2017).

Finally, the studies we presented here are experimental in nature and lack ecological validity. The experimental setting and the fact that speakers were animated characters in a computerized storybook made it impossible for participants to ask questions and seek clarification. Such mechanisms are integral parts of real-world communication (see, e.g., Arkel et al., 2020; H. H. Clark & Brennan, 1991) and it will be important for future modeling work to account for them. A promising step in this direction is the work by Anderson (2021) who proposed a system for conversational updating within the RSA framework. To increase ecological validity, our model of information integration could be used to model real-world interactions. However, this would require a sufficiently large dataset in which certain information sources are presented alone as well as in combination with others (see, e.g., Bohn, Liebal, & Tessler, 2021 for such an approach in the field of great ape communication; see also Scontras et al., 2021; Shukla et al., 2019, for how to connect RSA-style models with real-world data). Taken together, despite important limitations, we think that the approach we have taken here provides a powerful framework for testing theories about information integration during language learning.

Conclusion

Studying how multiple types of pragmatic cues are balanced contributes to a more comprehensive understanding of word learning. In the current study, participants inferred the referent by integrating nonlinguistic cues (gaze and pointing gestures) with assumptions about speaker informativeness and common ground information, going beyond previous experimental work in measuring how these information sources were combined. The real learning environment is far richer than what we captured in our experimental design, however. For example, in addition to multiple layers of social information, children can rely on semantic and syntactic features of the utterances as cues to meaning (E. V. Clark, 1973; Gleitman, 1990). Across development, children learn to recruit these different sources of information and integrate them. RSA models allow for the inclusion of semantic information as part of the utterance (Bergen et al., 2016) and it will be a fruitful avenue for future research to model the integration of linguistic and pragmatic information across development. To conclude, our work here shows how computational models of language comprehension can be used as powerful tools to explicate and test hypotheses about information integration across development.

References

- Akhtar, N., Carpenter, M., & Tomasello, M. (1996). The role of discourse novelty in early word learning. *Child Development*, 67(2), 635–645. <https://doi.org/10.2307/1131837>
- Anderson, C. J. (2021). Tell me everything you know: A conversation update system for the rational speech acts framework. *Proceedings of the Society for Computation in Linguistics*, 4, 244–253. <https://doi.org/10.7275/j6z3-9k74>
- Arkel, J. V., Woensdregt, M., Dingemanse, M., & Blokpoel, M. (2020). A simple repair mechanism can alleviate computational demands of pragmatic reasoning: Simulations and complexity analysis. In R. Fernández & T. Linzen (Eds.), *Proceedings of the 24th Conference on Computational Natural Language Learning* (pp. 177–194). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.conll-1.14>
- 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. <https://doi.org/10.1016/j.cognition.2010.10.010>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Behne, T., Carpenter, M., & Tomasello, M. (2005). One-year-olds comprehend the communicative intentions behind gestures in a hiding game. *Developmental Science*, 8(6), 492–499. <https://doi.org/10.1111/j.1467-7687.2005.00440.x>
- Bergen, L., Levy, R., & Goodman, N. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*. Advance online publication. <https://doi.org/10.3765/sp.9.20>
- Bohn, M., & Frank, M. C. (2018, December 6). *Modeling common ground*. osf.io/u7kxe
- Bohn, M., & Frank, M. C. (2019). The pervasive role of pragmatics in early language. *Annual Review of Developmental Psychology*, 1(1), 223–249. <https://doi.org/10.1146/annurev-devpsych-121318-085037>
- Bohn, M., & Köymen, B. (2018). Common ground and development. *Child Development Perspectives*, 12(2), 104–108. <https://doi.org/10.1111/cdep.12269>
- Bohn, M., Call, J., & Tomasello, M. (2019). Natural reference: A phylo- and ontogenetic perspective on the comprehension of iconic gestures and vocalizations. *Developmental Science*, 22(2), e12757. <https://doi.org/10.1111/desc.12757>
- Bohn, M., Le, K. N., Peloquin, B., Köymen, B., & Frank, M. C. (2021). Children's interpretation of ambiguous pronouns based on prior discourse. *Developmental Science*, 24(3), e13049. <https://doi.org/10.1111/desc.13049>
- Bohn, M., Liebal, K., & Tessler, M. H. (2021). Great ape communication as contextual social inference: A computational modeling perspective. *PsyArXiv*. <https://doi.org/10.31234/osf.io/yd9k4>
- Bohn, M., Tessler, M. H., Merrick, M., & Frank, M. C. (2021). How young children integrate information sources to infer the meaning of words. *Nature Human Behaviour*, 5(8), 1046–1054. <https://doi.org/10.1038/s41562-021-01145-1>
- Bohn, M., Zimmermann, L., Call, J., & Tomasello, M. (2018). The social-cognitive basis of infants' reference to absent entities. *Cognition*, 177, 41–48. <https://doi.org/10.1016/j.cognition.2018.03.024>
- Braginsky, M., Tessler, M. H., & Hawkins, R. (n.d.). *Rwebppl: R interface to WebPPL*. Retrieved from <https://github.com/mhtess/rwebppl>
- Brown-Schmidt, S. (2009). Partner-specific interpretation of maintained referential precedents during interactive dialog. *Journal of Memory and Language*, 61(2), 171–190. <https://doi.org/10.1016/j.jml.2009.04.003>
- Bruner, J. (1983). *Child's talk: Learning to use language*. Norton.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3(2), 57–65. [https://doi.org/10.1016/S1364-6613\(98\)01273-X](https://doi.org/10.1016/S1364-6613(98)01273-X)
- Clark, E. V. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T. Moore (Ed.), *Cognitive development and acquisition of language* (pp. 65–110). Academic Press. <https://doi.org/10.1016/B978-0-12-505850-6.50009-8>
- Clark, E. V. (2009). *First language acquisition*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511806698>
- Clark, E. V. (2015). Common ground. In B. MacWhinney & W. O'Grady (Eds.), *The handbook of language emergence* (Vol. 87, pp. 328–353). Wiley.
- Clark, H. H. (1996). *Using language*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511620539>
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). American Psychological Association. <https://doi.org/10.1037/10096-006>
- Cohn-Gordon, R., Goodman, N., & Potts, C. (2018). Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2* (pp. 439–443). Association for Computational Linguistics.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153. <https://doi.org/10.1016/j.tics.2009.01.005>
- Degen, J., Tessler, M. H., & Goodman, N. D. (2015). Wonky worlds: Listeners revise world knowledge when utterances are odd. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual conference of the cognitive science society* (pp. 548–553). Cognitive Science Society.
- Diesendruck, G., Markson, L., Akhtar, N., & Reudor, A. (2004). Two-year-olds' sensitivity to speakers' intent: An alternative account of Samuelson and Smith. *Developmental Science*, 7(1), 33–41. <https://doi.org/10.1111/j.1467-7687.2004.00320.x>
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429–433. <https://doi.org/10.1038/415429a>
- Fazly, A., Alishahi, A., & Stevenson, S. (2010). A probabilistic computational model of cross-situational word learning. *Cognitive Science*, 34(6), 1017–1063. <https://doi.org/10.1111/j.1551-6709.2010.01104.x>

- Fortier, M., Kellier, D., Flecha, M. F., & Frank, M. C. (2018). *Ad-hoc pragmatic implicatures among Shipibo-Konibo children in the Peruvian Amazon*. <https://doi.org/10.17910/B7.691>
- Fourtassi, A., & Frank, M. C. (2020). How optimal is word recognition under multimodal uncertainty? *Cognition*, 199, 104092. <https://doi.org/10.1016/j.cognition.2019.104092>
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998. <https://doi.org/10.1126/science.1218633>
- Frank, M. C., & Goodman, N. D. (2014). Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology*, 75, 80–96. <https://doi.org/10.1016/j.cogpsych.2014.08.002>
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, 20(5), 578–585. <https://doi.org/10.1111/j.1467-9280.2009.02335.x>
- Frank, M. C., Sugarman, E., Horowitz, A. C., Lewis, M. L., & Yurovsky, D. (2016). Using tablets to collect data from young children. *Journal of Cognition and Development*, 17(1), 1–17. <https://doi.org/10.1080/15248372.2015.1061528>
- Gagliardi, A., Feldman, N. H., & Lidz, J. (2017). Modeling statistical insensitivity: Sources of suboptimal behavior. *Cognitive Science*, 41(1), 188–217. <https://doi.org/10.1111/cogs.12373>
- Galati, A., & Brennan, S. E. (2021). What is retained about common ground? Distinct effects of linguistic and visual co-presence. *Cognition*, 215, 104809. <https://doi.org/10.1016/j.cognition.2021.104809>
- Ganea, P. A., & Saylor, M. M. (2007). Infants' use of shared linguistic information to clarify ambiguous requests. *Child Development*, 78(2), 493–502. <https://doi.org/10.1111/j.1467-8624.2007.01011.x>
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1(1), 3–55. https://doi.org/10.1207/s15327817la0101_2
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829. <https://doi.org/10.1016/j.tics.2016.08.005>
- Goodman, N. D., & Stuhlmüller, A. (2014). *The design and implementation of probabilistic programming languages*. Retrieved from <http://dippl.org>
- Graham, S. A., San Juan, V., & Khu, M. (2017). Words are not enough: How preschoolers' integration of perspective and emotion informs their referential understanding. *Journal of Child Language*, 44(3), 500–526. <https://doi.org/10.1017/S0305000916000519>
- Graham, S. A., Sedivy, J., & Khu, M. (2014). That's not what you said earlier: Preschoolers expect partners to be referentially consistent. *Journal of Child Language*, 41(1), 34–50. <https://doi.org/10.1017/S0305000912000530>
- Grice, H. P. (1991). *Studies in the way of words*. Harvard University Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767–773. <https://doi.org/10.1111/j.1467-9280.2006.01780.x>
- Grosse, G., Moll, H., & Tomasello, M. (2010). 21-month-olds understand the cooperative logic of requests. *Journal of Pragmatics*, 42(12), 3377–3383. <https://doi.org/10.1016/j.pragma.2010.05.005>
- Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science*, 304(5669), 438–441. <https://doi.org/10.1126/science.1095455>
- Hanna, J. E., Tanenhaus, M. K., & Trueswell, J. C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language*, 49(1), 43–61. [https://doi.org/10.1016/S0749-596X\(03\)00022-6](https://doi.org/10.1016/S0749-596X(03)00022-6)
- Heller, D., Parisien, C., & Stevenson, S. (2016). Perspective-taking behavior as the probabilistic weighing of multiple domains. *Cognition*, 149, 104–120. <https://doi.org/10.1016/j.cognition.2015.12.008>
- Horowitz, A. C., Schneider, R. M., & Frank, M. C. (2018). The trouble with quantifiers: Exploring children's deficits in scalar implicature. *Child Development*, 89(6), e572–e593. <https://doi.org/10.1111/cdev.13014>
- Horton, W. S., & Gerrig, R. J. (2016). Revisiting the memory-based processing approach to common ground. *Topics in Cognitive Science*, 8(4), 780–795. <https://doi.org/10.1111/tops.12216>
- Kamide, Y., Scheepers, C., & Altmann, G. T. (2003). Integration of syntactic and semantic information in predictive processing: Cross-linguistic evidence from German and English. *Journal of Psycholinguistic Research*, 32(1), 37–55. <https://doi.org/10.1023/A:1021933015362>
- Khu, M., Chambers, C. G., & Graham, S. A. (2020). Preschoolers flexibly shift between speakers' perspectives during real-time language comprehension. *Child Development*, 91(3), e619–e634. <https://doi.org/10.1111/cdev.13270>
- Knutsen, D., & Le Bigot, L. (2014). Capturing egocentric biases in reference reuse during collaborative dialogue. *Psychonomic Bulletin & Review*, 21(6), 1590–1599. <https://doi.org/10.3758/s13423-014-0620-7>
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT press. <https://doi.org/10.7551/mitpress/5526.001.0001>
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. Freeman.
- Matthews, D., Lieven, E., & Tomasello, M. (2010). What's in a manner of speaking? Children's sensitivity to partner-specific referential precedents. *Developmental Psychology*, 46(4), 749–760. <https://doi.org/10.1037/a0019657>
- Matthews, D., Lieven, E., Theakston, A., & Tomasello, M. (2006). The effect of perceptual availability and prior discourse on young children's use of referring expressions. *Applied Psycholinguistics*, 27(3), 403–422. <https://doi.org/10.1017/S0142716406060334>
- McClelland, J. L., Mirman, D., & Holt, L. L. (2006). Are there interactive processes in speech perception? *Trends in Cognitive Sciences*, 10(8), 363–369. <https://doi.org/10.1016/j.tics.2006.06.007>
- Monroe, W., Hawkins, R. X., Goodman, N. D., & Potts, C. (2017). Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5, 325–338. https://doi.org/10.1162/tacl_a_00064
- Moore, R., Mueller, B., Kaminski, J., & Tomasello, M. (2015). Two-year-old children but not domestic dogs understand communicative intentions without language, gestures, or gaze. *Developmental Science*, 18(2), 232–242. <https://doi.org/10.1111/desc.12206>
- Mozuraitis, M., Chambers, C. G., & Daneman, M. (2015). Privileged versus shared knowledge about object identity in real-time referential processing. *Cognition*, 142, 148–165. <https://doi.org/10.1016/j.cognition.2015.05.001>
- Mozuraitis, M., Stevenson, S., & Heller, D. (2018). Modeling reference production as the probabilistic combination of multiple perspectives. *Cognitive Science*, 42(4), 974–1008. <https://doi.org/10.1111/cogs.12582>
- 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. <https://doi.org/10.1111/j.0956-7976.2002.00460.x>
- Nilsen, E. S., Graham, S. A., & Pettigrew, T. (2009). Preschoolers' word mapping: The interplay between labelling context and specificity of speaker information. *Journal of Child Language*, 36(3), 673–684. <https://doi.org/10.1017/S0305000908009021>
- Noveck, I. A. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition*, 78(2), 165–188. [https://doi.org/10.1016/S0010-0277\(00\)00114-1](https://doi.org/10.1016/S0010-0277(00)00114-1)
- 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. <https://doi.org/10.1017/S03050009000004566>

- Ozyürek, A., Willems, R. M., Kita, S., & Hagoort, P. (2007). On-line integration of semantic information from speech and gesture: Insights from event-related brain potentials. *Journal of Cognitive Neuroscience*, 19(4), 605–616. <https://doi.org/10.1162/jocn.2007.19.4.605>
- R Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Ryskin, R., Stevenson, S., & Heller, D. (2020). Probabilistic weighting of perspectives in dyadic communication. *Proceedings of the 42nd annual conference of the cognitive science society* (pp. 57–63). Cognitive Science Society.
- Saylor, M. M., Ganea, P. A., & Vázquez, M. D. (2011). What's mine is mine: Twelve-month-olds use possessive pronouns to identify referents. *Developmental Science*, 14(4), 859–864. <https://doi.org/10.1111/j.1467-7687.2010.01034.x>
- Saylor, M. M., Sabbagh, M. A., Fortuna, A., & Troseth, G. (2009). Preschoolers use speakers' preferences to learn words. *Cognitive Development*, 24(2), 125–132. <https://doi.org/10.1016/j.cogdev.2008.12.003>
- Scontras, G., Tessler, M. H., & Franke, M. (2021). A practical introduction to the rational speech act modeling framework. *arXiv*. <https://doi.org/10.48550/arXiv.2105.09867>
- Shukla, P., Elmadjian, C., Sharan, R., Kulkarni, V., Turk, M., & Wang, W. Y. (2019). What should I ask? Using conversationally informative rewards for goal-oriented visual dialog. *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 6442–6451). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1646>
- Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicatures: Alternatives and relevance. *Cognition*, 153, 6–18. <https://doi.org/10.1016/j.cognition.2016.04.006>
- Smith, A. C., Monaghan, P., & Huettig, F. (2017). The multimodal nature of spoken word processing in the visual world: Testing the predictions of alternative models of multimodal integration. *Journal of Memory and Language*, 93, 276–303. <https://doi.org/10.1016/j.jml.2016.08.005>
- Sperber, D., & Wilson, D. (2001). *Relevance: Communication and cognition* (2nd ed.). Blackwell Publishers.
- Stiller, A. J., Goodman, N. D., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development*, 11(2), 176–190. <https://doi.org/10.1080/15475441.2014.927328>
- Su, Y. E., & Su, L.-Y. (2015). Interpretation of logical words in mandarin-speaking children with autism spectrum disorders: Uncovering knowledge of semantics and pragmatics. *Journal of Autism and Developmental Disorders*, 45(7), 1938–1950. <https://doi.org/10.1007/s10803-014-2350-0>
- Sullivan, J., Boucher, J., Kiefer, R. J., Williams, K., & Barner, D. (2019). Discourse coherence as a cue to reference in word learning: Evidence for discourse bootstrapping. *Cognitive Science*, 43(1), e12702. <https://doi.org/10.1111/cogs.12702>
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632–1634. <https://doi.org/10.1126/science.7777863>
- Tessler, M. H., Lopez-Brau, M., & Goodman, N. D. (2017). Warm (for winter): Comparison class understanding in vague language. *Proceedings of the 39th annual conference of the cognitive science society*. Cognitive Science Society.
- Tomasello, M. (2008). *Origins of human communication*. MIT press. <https://doi.org/10.7551/mitpress/7551.001.0001>
- Tomasello, M. (2009). *Constructing a language*. Harvard University Press.
- Vouloumanos, A., Onishi, K. H., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences of the United States of America*, 109(32), 12933–12937. <https://doi.org/10.1073/pnas.1121057109>
- Wang, S., Liang, P., & Manning, C. D. (2016). Learning language games through interaction. *54th annual meeting of the association for computational linguistics, ACL 2016* (pp. 2368–2378). Association for Computational Linguistics.
- Wilkes-Gibbs, D., & Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language*, 31(2), 183–194. [https://doi.org/10.1016/0749-596X\(92\)90010-U](https://doi.org/10.1016/0749-596X(92)90010-U)
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272. <https://doi.org/10.1037/0033-295X.114.2.245>
- Yoon, E. J., & Frank, M. C. (2019). The role of salience in young children's processing of ad hoc implicatures. *Journal of Experimental Child Psychology*, 186, 99–116. <https://doi.org/10.1016/j.jecp.2019.04.008>
- Yoon, S. O., & Brown-Schmidt, S. (2014). Adjusting conceptual pacts in three-party conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(4), 919–937. <https://doi.org/10.1037/a0036161>
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, 10(7), 301–308. <https://doi.org/10.1016/j.tics.2006.05.002>
- Zhao, S., Ren, J., Frank, M. C., & Zhou, P. (2021). The development of quantity implicatures in mandarin-speaking children. *Language Learning and Development*, 17(4), 1–23. <https://doi.org/10.1080/15475441.2021.1886935>

Received January 6, 2021

Revision received February 6, 2022

Accepted February 11, 2022 ■