Young children choose informative referring expressions to describe the agents and patients of transitive events.

Melissa Kline¹, Ted Gibson & Laura Schulz

Massachusetts Institute of Technology

Abstract

A key component of most models of pragmatics is that speakers consider more than one way of conveying a message, and how informative each version is in context. Theories of pragmatics, and particularly pragmatic development, are hampered by the fact that while we can often observe what a participant does (either as a speaker or a listener), we can rarely observe the choices they consider. Here, we focus on the problem that remains once two alternative utterances are available: selecting the most informative one. By providing 3-6 year old children with a choice between transitive sentences that drop either an uninformative argument (i.e. one that can be directly recovered from the visual context) or an informative one, we can directly test their ability to select between alternate ways of conveying the same message, abstracting away from the many other processes required to produce a multi-word expression. While 3 year olds select between these alternatives at chance (a pattern which, based on control version, may be partially due to task demands), we find clear evidence that children aged 4 and older preferentially select the more informative option. This result shows that (1) children's awareness of informativity contrasts extends to sentence-level messages, not just object referent selection and (2) directly testing the development of specific cognitive abilities theorized to be involved in pragmatic communication can both clarify and constrain our theories of language development.

¹ Corresponding author. Inquiries may be sent to mekline@mit.edu

Introduction

Despite the humor, confusion, and social consequences of pragmatic breakdowns, by and large we succeed at getting our messages across to each other - to the point that indirectness and ambiguities in language often go unnoticed. A central feature of language production that supports this success is the variability in how we refer to particular entities, from full vs. pronoun descriptions of objects and people to the ways that conversational partners quickly converge on shared terms (Brennan & Clark, 1996). Our referential abilities have provided a well-studied microcosm of pragmatic abilities: in a context where we might choose to talk about one of several things, we attempt to refer to the one we want, in a way that ensures the listener gets the message we intend. However, the status of these abilities in language acquisition has remained relatively unclear, in part because the process of designing and comprehending referentially informative messages requires coordinating several distinct cognitive abilities. Here, we target a specific part of this process in a novel experimental linguistic context, whereby the participant selects the most informative utterance from two ways of conveying the same message about a simple transitive event. A key component of most models of pragmatics is the idea that speakers conceive of more than one possible way of communicating an intended message given a particular non-linguistic context, and select the one best suited to communicate the message in that context. By directly probing this step, we can begin to understand how and when children solve the problem of referential communication, and ultimately come to be helpful, adult-like communicators.

Domain-general and language-specific awareness of others' knowledge

Young children have a number of surprisingly sophisticated abilities for understanding the social world around them, including detecting other people's intentions and knowledge states from the social context of language (e.g. Bonawitz et al., 2011). In nonlinguistic cognition, infants and toddlers show evidence of tracking other people's goals, knowing something about their knowledge states, and using communicative gestures like pointing to interpret meaning (Brandone, Horwitz, Aslin, & Wellman, 2014; Carpenter,

Akhtar, & Tomasello, 1998; Csibra, 2003; Phillips & Wellman, 2005; Tomasello, Carpenter, Call, Behne, & Moll, 2005). These early social and pragmatic abilities stand in stark contrast to the extended difficulties that young children are reported to have with specifically linguistic abilities, in referential communication tasks and related language processes.

Referential communication tasks, in which participants must either describe or find one object or person out of an array, have been a key method for testing the predictions of pragmatic theories in both adults and young children (Sedivy, Tanenhaus, Chambers, & Carlson, 1999). While the complexity of stimuli and measures used vary across populations, the basic setup of referential communication tasks is largely the same for adult and child studies. For a given array of objects, a particular way of referring to one of them can either be sufficiently informative, uniquely narrowing the array down to a single alternative, or under-informative (*that dog* when there is both a pug and a terrier in the scene). (Keysar, Barr, Balin, & Brauner, 2000; Keysar, Barr, & Horton, 1998, though see Heller, Grodner & Tanenhaus 2008). During language production, adults take in a great deal of information – including information that is not contextually available to the listener – and then converge on referential expressions which appear to not just ensure that listener's comprehension, but even speed it up, for instance by selectively producing overinformative color adjectives when the color is atypical for the object or when doing so can speed the listener's visual search of the array (Rubio-Fernández, 2016; Sedivy, 2003). One context where adults fail to be informative is when they do not successfully *notice* a possible ambiguity, producing underinformative, ambiguous expressions like merely referring to 'a bat' when both the animal and the sports equipment are in the shared context (Ferreira, Slevc, & Rogers, 2005).

The general pattern of findings for children in this paradigm is decidedly murky. Implicit measurements such as eye movements and response speeds suggest that children as young as four notice and are affected by the processing of an ambiguous referring expression during comprehension, but like adults, they will spend time looking towards a possible distractor even if the speaker can't see it (Morisseau, Davies, & Matthews, 2013; Nadig & Sedivy, 2002; Nilsen, Graham, Smith, & Chambers, 2008). When asked to

describe referents for adults who either can or cannot see a possible distractor, children perform inconsistently – like adults, they are more *likely* to provide a disambiguation when it's needed ('a blue cup' when there are both red and blue cups that the listener can see) than when the information isn't necessary, but they are less likely overall to produce the adjective (Nadig & Sedivy, 2002). When asked to predict whether utterances will be enough to help a puppet find a hidden prize, children are sensitive to some kinds of Gricean maxim violations (understanding that an irrelevant statement like 'I like all these cups' is unhelpful for finding a sticker under one of those cups,) but not others (e.g. accepting an underinformative instruction such as finding a sticker 'under the duck' when there is a big duck and a small one) (Eskritt, Whalen, & Lee, 2008; Grice, 1975; Nilsen et al., 2008). These findings parallel evidence from corpora and everyday experience: among other phenomena, children under age six use definite noun phrases (*that one*) and pronouns (*SHE did it*) even when the listener cannot tell who or what is being referred to (Brown, 1973; Sonnenschein & Whitehurst, 1984), and they appear to particularly struggle with other pragmatic phenomena such as scalar implicatures (e.g. that 'some apples' implies some but not all of the apples (Davies & Katsos, 2010; Huang & Snedeker, 2009, i.a.)

How can we understand the disconnect between toddlers' early social abilities and the failure of preschoolers to perform like adults in communication tasks like these? Critically, the process of designing or comprehending a sentence that is referentially informative requires coordinating several cognitive abilities, both domain general and specific to language. Rabagliati & Robertson (2017) show that children's tendency to produce ambiguous messages is related to a combination of weak inhibitory skills (leading them to produce 'default' descriptions of objects) and changes in how children monitor the nonlinguistic context for possible ambiguity both before and after their own speech. In areas where children are known to struggle, small task variations that compensate for specific deficits can lead to strikingly more adult-like performance (Srivastavan et al 2017, Foppolo, Guasti, & Chierchia, 2012). Even relatively incidental differences such as a greater willingness to credit a speaker with a correct answer can obscure children's knowledge – children report some kinds of pragmatic violations only when

given the option to give that speaker a 'middle sized' reward rather than just no reward versus the big reward (Katsos & Bishop 2011). While these findings enrich our understanding of the cognitive processes required to succeed at referential communication, they also make it clear that detecting children's underlying computations about what kind of information will help a listener is very challenging. Understanding the development of pragmatic abilities will require us to de-convolve these various cognitive processes: on the one hand, tracking how these basic cognitive skills emerge, and on the other, establishing how and whether children's underlying model or use of informativity changes over development.

Comparisons between adults and children that are confounded with variations in non-pragmatic areas like the above are at a risk of either over- or under-estimating children's abilities. In many referential communication experiments, the critical choice a participant faces is whether or not to mention a distinguishing adjective or trait, and in general, longer sentences will be more informative than shorter sentences. Even if we knew children were considering only two utterance options, 'cup' or 'red cup', a child who says 'cup' may have known that 'red cup' would be more informative in context, but ended up producing the easier, shorter sentence. On the other hand, a child who says 'red cup' may simply be following a general strategy of preferring longer utterances, not because they know it conveys more information in a particular context. One experimental study, by controlling the specific distribution of utterances and possible referents seen by five-year-old children, showed that when they heard a speaker who gave descriptions that were either more informative, longer, or both, only length mattered for children's performance (Whitehurst, Sonnenschein, & Ianfolla, 1981). But this is a necessarily indirect inference from children's speech: because de-confounding more informative from longer utterances in natural speech is so challenging, this argument rests on the (in)effectiveness of a particular contrast in the input that children hear in the study.

Because production tasks recruit the entire processing pathway (not just pragmatics, but lexical retrieval, sentence construction, articulation, and so on), an important complementary strategy is to test children's

pragmatic abilities with tasks that remove or alter processing limitations. For instance, Stiller, Goodman & Frank (2015) argued that three- and four-year-olds are able to *comprehend* scalar relationships pragmatically (a key skill for correctly using terms like *some*) by showing that they succeed in a game that requires this information. In particular, given a choice of three smiley-faces, these children seem to know that when referents form an ad-hoc scale (e.g., a bare smiley-face, one with glasses, one with glasses and a hat), “My friend wears glasses” probably refers to the second smiley face, not the third. This is a natural pragmatic inference because if the speaker had meant to refer to the third smiley, they could have said "My friend has glasses and a hat". Stiller et al. propose that children succeed on this task because they are making exactly this inference: implicitly evaluating the possible things a speaker could have said to describe *each* smiley face and using this information to determine the most rational target. If this explanation is correct, then children are accomplishing a central pragmatic task – comparing the information carried by alternative statements – even if they struggle to deploy such computations in their own productions.

Figure 1: Possible single-feature utterances generated by two children in the Stiller et al. (2015) paradigm. Bolded descriptions will be consistent with the prompt "My friend has glasses...can you find my friend?"

However, a key limitation of this study is that we cannot directly observe which alternate utterances the child was actually considering. If they focused on the utterances the experimenters expected (My friend wears glasses/My friend wears glasses and a hat), then the pragmatic enrichment from 'glasses' to 'glasses but no hat' is clear. But the pattern of choices could also be explained in other ways. Suppose the children had taken the following strategy: look at each of the smiley faces, generate one possible description, and then check if it matches what the experimenter asks for. Imagine that two children participate in this experiment and happen to generate the sets in Figure 1. When the experimenter says 'my friend has glasses', Child 1 will succeed, as expected, but Child 2 will be faced with an ambiguous choice, potentially answering at random between faces two and three.² In fact, we know that children around this age will happily choose a correct but ambiguous referent when asked to select 'the man with the hat' from a larger set that includes more than one, and will sometimes fail to access alternative ways of describing an object, for instance when scalars like 'some' are involved (Barner, Brooks, & Bale, 2011). What this thought experiment shows is that the same patterns of behavior – even removing the more difficult, naturalistic referential production task – may arise in more than one way. If the child fails to generate the alternatives we expect, their behavior cannot necessarily be taken as evidence of a particular pragmatic strategy. Solving this quandary will require converging evidence from multiple tests of how children evaluate what sorts of utterances are informative or uninformative.

The current experiments

One way to put theories of early pragmatic development to the test is to alter the *types of messages* being conveyed in experimental tasks. Many models of referential communication have focused principally on utterances describing physical objects or people, meaning that any findings about how children produce and comprehend referential messages may not generalize beyond the particulars of combining nouns and adjectives and comparing them to arrays of objects. Here, we focus instead on mentioning participants (i.e. arguments) in event descriptions. While noun reference requires a speaker and listener to converge on

² Thanks to Jesse Snedeker for this insight

the same object, a sentence expresses a proposition about an event or state in the world, focusing us in on some elements of the ongoing stream of experience. A simple sentence like *Jane fed the dog* expresses a compositional meaning with several parts, while leaving out other possible arguments and adjuncts (compare e.g. *Jane fed the dog broccoli*, *Jane fed the dog last night*, *Jane fed the dog for an hour*.) In casual conversation these omissions can go farther. Even languages like English that require the presence of a lexical subject in a neutral context can easily accommodate exchanges like:

(1) What's Jane doing?

(___) Feeding the dog.

(2) What's the dog doing?

(___) Eating (___).

In recent work, we have shown that argument dropping can provide a powerful test case for understanding how adults evaluate the informativity of a particular way of describing an event in context (Kline, Schulz, & Gibson, In press). In this task, participants are shown simple events that take place in varying non-linguistic contrasts and, critically, are asked to describe them in a limited communication channel: simple transitive sentences like *Alexa ate the apple* must be conveyed in just two words. In that study, we show that adults' choices about what to say are highly sensitive to informativity, and consistent with 'rational speaker' computational models of pragmatics, which have been successful in quantitatively capturing many aspects of pragmatic behavior (Frank & Goodman, 2012; Goodman & Frank, 2016; Kao, Wu, Bergen, & Goodman, 2014).

Here, we extend this approach to children: do they know when a sentence fragment like MONKEY EAT is helpful and when it is not? Complementing the comprehension approach used by Stiller et al. (2015), we focus on the question of how children evaluate possible utterances in light of a particular non-linguistic context. We make the task easier for children and more straightforward for analysis by asking

Figure 2: Agent-Ambiguous (a) and Patient-Ambiguous (b) arrays, paired with the target sentence “The monkey eats the orange” (c).

children to explicitly identify which utterances out of a set are helpful, rather than to predict other downstream effects like whether a listener will or won't be able to solve a problem (Nilsen et al., 2008). To accomplish this, we present three- to six-year-old English speaking children with puppets who use periphrastic, incomplete language to tell a story to the child's caregivers about a hidden event they had seen, one puppet's sentence including the agent (MONKEY EAT) and one including the patient (EAT ORANGE). The key manipulation of the experiment is the common-ground context set, illustrated in Figure 2: for the same event, children could see either a context group with a monkey, a banana, and an orange (Patient-Ambiguous, PA), or a monkey, a duck, and an orange (Agent-Ambiguous, AA).³ Although the two kinds of shortened sentences mentioned above are equal in length and express the same amount of information in a neutral context, rational speaker models of early pragmatics

predict that children will recognize that the informativity of these utterances depends on the set of

³ We assume that children will be able to identify which of these are more likely to be agents or patients; we are not testing whether participants know that a monkey is more likely to be the agent of EAT than an orange is. Rather, we test whether they take into account the *possible events* (a monkey eating an orange, a monkey eating a banana etc.) that correspond to a contrast set when deciding which speaker is most helpful.

possible events. If the possible events are a monkey eating an orange or a monkey eating a banana (PA), mentioning that the agent was the monkey is true but superfluous, because there is no other salient choice. However, mentioning the orange reduces the set of possible events from two to one, and successfully gets the message across. In contrast the reverse is true when the possible events are a monkey eating an orange or a duck eating an orange (AA): now MONKEY EAT identifies the actual event, while EAT ORANGE is unhelpful. If young children's pragmatic failures are due to a true inability to identify informative utterances in context (rather than to generating those relevant utterances, or to suppressing easier but less helpful utterances), their opinion of the puppets' statements should be insensitive to the non-linguistic context: both puppets leave out a word from the basic description of the event, so both are equally bad (or good).

Given the diversity of claims about children's referential abilities, we tested a broad range of children – aged three through six. To make sure that any age-related trends were not due to general challenges with the task, we also included a control condition for the youngest children that was identical to the test condition except that one puppet gave a false (rather than merely unhelpful) answer. Identifying the age at which children can make these comparisons can provide a converging timeline with object-reference studies for the development of pragmatic abilities and can help us to understand how the interrelated processes involved in language production come together in development, eventually allowing children and adults to successfully produce informative expressions of many kinds.

Methods

Participants 163 monolingual (no more than 25% input in another language) children between the ages of 3;0 and 6;11 participated at the Boston Children's Museum. 109 children participated in the main version of the experiment (helpful/unhelpful), and 54 3-4-year-olds participated in a control task (true/false). Data from 27 additional children is not included because of: (a) unexpected experimental

errors (8 children); the child did not complete the session (10); or because a sibling or other family member either saw and reported the hidden scenes or answered for the participating child (9)⁴.

Materials Digital versions of all materials created for this study are available in the online repository of Supplemental Materials for this paper, at <http://osf.io/a2ih3/>. The critical scenarios and test sentences were presented to children in two trial booklets presented in a fixed order (*Eating* and *Feeding*). Each booklet presented four event participants (e.g. boy, girl, dog, cat), followed by images of the four possible events that could occur (e.g. the boy/girl feeds the dog/cat). The position of the target event in this series was varied across verbs. The next page showed the critical common-ground array, consisting of the agent of the target event that would follow, the patient, and a distractor. The arrangement of the agent, patient, and distractor was fixed for each verb, but different between booklets. The distractor was either the other possible agent (Agent-Ambiguous) or the other possible patient (Patient-Ambiguous). Finally, the last page of the trial showed the array again, next to a flap of paper covering a picture of the target event. The Agent-Ambiguous and Patient-Ambiguous arrays along with the target event for *Eating* is shown in Figure 2.

In addition to the storybooks, the experiment used sets of paper character cutouts presented on popsicle sticks mounted on a base. For each trial, two similar characters dressed in distinctive ways (e.g., a pig in a striped shirt and a pig polka-dot shirt) served as the two possible communicators.

Procedure

All sessions were recorded for later off-line coding using a digital video camera. Each participant was assigned to either the Agent-Ambiguous or Patient-Ambiguous condition just described; the same

⁴ Errors we made in training research assistants resulted in two kinds of additional data exclusion. First, due to a misunderstanding of the paradigm, a number of children were run in a version that did not provide critical elements (i.e. the puppets having seen the hidden event they were supposed to describe). All sessions from the period of data collection where this error occurred were excluded. Second, a pilot version of the task booklet containing additional trials (which the piloting showed younger children to fatigue on), was inadvertently used with a subset of the subsequent children. Because this was unintended, we did not analyze these additional trials. The full datasets for all participants, including those excluded from the present analyses, are available in supplemental materials.

booklets were used for both the main helpful/unhelpful experiment and the control true/false experiment. All five- and six-year-old participants completed the helpful/unhelpful task, while three- and four-year-olds completed either one or the other (the control version of the experiment was introduced midway through data collection, so children in this age group were not assigned randomly to main vs. control, but *were* randomly assigned to condition *within* each task). Below we describe the main experiment while noting alterations made for the true/false control task.

Participants were told that they would be reading a set of storybooks, with help from a series of puppets:

Here are my two puppets who are going to help us tell some stories. These puppets are from another place, so they don't speak English that well. They can only say one or two words at a time, so we have to listen to them very closely. One of them says "Me puppet!" and the other one says, "Me puppet too!"

So there's two of them! One of them is going to do a really good job telling the story, and one is just going to do an okay job, all right? I want you to help me decide which one does a better job.

The puppets in the helpful/unhelpful study were described as 'really good' and 'just okay' to support the eventual judgment children would make (since the speakers only ever made true statements about the events and children are often tolerant to pragmatic violations, Davies & Katsos, 2010). For the true/false version, children were told that the puppets would *get the story right* or *get the story wrong*. This explicit description of the speakers as 'right' and 'wrong' was included to make the problem to be solved as clear as possible for the youngest children, and to prepare them for the fact that half the puppets would say a clearly false thing about the event under the flap.

The experimenter then proceeded to lead the child through the first storybook. The experimenter labeled each potential participant in the event (e.g., monkey, duck, banana, and orange), and then labeled the pictures of the possible events (*Look, here the duck is eating the orange.*) The experimenter used neutral prosody for each event, and was careful to use full nouns when they repeated between scenes (e.g., *Now, the duck is eating the banana*, rather than *Now, he's eating the banana*). Sometimes participants started to

label the images they saw; the experimenter said “Yeah!” or corrected the child if they mis-identified the scene.

After the four possible events had been presented, the experimenter showed the array (either Agent-Ambiguous or Patient-Ambiguous) to the child and pointed out the ambiguity:

Ohh, look, now what do we have here? Now we have three things! Look, there’s (a duck)/(some bananas), some oranges, and a monkey. Do you want to find out what happens next?

Note that in this script, the experimenter does not describe the specific ambiguity (e.g., *there’s two fruits*) or describe the possible events corresponding to the list of entities. The experimenter then proceeded to the target page, which showed the array again, and a flap of paper covering the target event. The back of the storybook was lifted and oriented so that the experimenter, the child, and the two puppets could see this page, but the child’s parent/guardian could not (*Oh, see this flap here? Your mom doesn’t know what’s under there. You’ll get to see what’s under there in a second, but your mom won’t know yet!*)

The child was invited to lift the flap and see the picture of the actual event (e.g., *The monkey eats the orange*). The experimenter confirmed with the child that they knew what happened, and brought each puppet over separately to look at the picture as well. They then instructed the child to close the flap, and laid the storybook back on the table.

The experimenter reminded the child that both puppets would ‘try and tell the story’, and then spoke for each of the two puppets:

Helpful/Unhelpful: *This puppet over here in the polkadots says MONKEY EAT. And this puppet over here in the stripes says EAT ORANGE. Okay? So this one says MONKEY EAT, and this one says EAT ORANGE. So, which puppet did a better job telling your mom the story?*

True/False: *This puppet over here in the polkadots says MONKEY EAT. And this puppet over here in the stripes says DUCK EAT. Okay? So this one says MONKEY EAT, and this one says DUCK EAT. So, which puppet got the story right?*

Thus, in both tasks, one puppet said something both true and helpful: they named the verb, plus the ambiguous argument (e.g., the Agent in the Agent-Ambiguous condition, MONKEY EAT). However, the alternative answer differs across the tasks: in the main task, for instance in the Agent-Ambiguous condition, the other puppet said the correct but unhelpful “EAT ORANGE”; in the true/false control, the other puppet incorrectly reported “DUCK EAT”. The full set of utterances for each condition is presented in Table 1.

We note here a limitation of the study: the first utterance heard was always the Agent-Verb utterance (MONKEY EAT): we fixed the order of these utterances to reduce task demands following initial piloting. For the True/false version, the position of the correct answer was matched to the main task, with the false utterances replacing the unhelpful ones. The consequence of this is that it not possible to interpret children’s overall preference for speakers who mention agents or patients: any baseline preference for A-V or V-P utterances would be conflated with order; instead, we determine whether children’s preferences were significantly affected by the correspondence between particular utterances and the changing context scene.

After hearing both utterances, the child pointed to one of the two puppets (if they did not do so spontaneously, the experimenter asked them, *Can you point to one of the puppets?*). Participants

	Agent-Ambiguous	Patient-Ambiguous
Helpful/unhelpful	MONKEY EAT EAT ORANGE	MONKEY EAT EAT ORANGE
Control (True/false)	MONKEY EAT DUCK EAT	EAT BANANA EAT ORANGE

Table 2: Utterances made by the puppet communicators in each version of the task. The true, helpful utterance is bolded in each cell. The left-hand utterance of each pair was always produced first (see discussion of task demands above.)

sometimes spontaneously described the event instead of choosing a puppet, either repeating one of the puppets (*Monkey eat!*) or describing the whole event (*The monkey ate the orange!*) In either case, the experimenter responded by asking the child to touch the one who did a better job telling the story. In all cases, the experimenter did not provide explicit positive or negative feedback, but said “okay!” in a positive tone and moved on. After this, the experimenter introduced a new set of puppets (so that opinions about which puppet was most informative would not be carried between trials). On both trials, the puppet who mentioned the agent spoke first, but the position of the puppet speaking first (left or right) was balanced across the two trials.

Children’s points were coded offline by a research assistant with the child’s first clear point directed at a puppet taken as their choice. 30% of experimental sessions were re-coded by the first author; agreement was at 95% and disagreements were resolved by discussion

Results

All data and code to produce analyses and graphs presented below are available in the project repository, at <http://osf.io/a2ih3/>. Because the Agent-Verb utterance was always produced first (see above), we collapse across conditions and code and plot responses as correct (a point to the puppet who made the helpful or true utterance) or incorrect (a point to the unhelpful or incorrect puppet). The results for both versions are shown in Figure 3, divided by age group. We discuss first the results of the main experiment and then the control version. For the main experiment, there was a trend toward an effect of ambiguity type in the helpful/unhelpful scenario: this effect of condition on whether children chose the puppet mentioning the patient was marginally significant by a mixed-effects logistic regression⁵ ($X^2=2.82$, $df=6$, $p=0.093$).

When children saw arrays with an ambiguous agent, they were numerically more likely to choose the puppet

⁵ The full model specification for this regression is given in R code by `glmer(chosePatientDrop ~ Condition + (1|trial) + (1|Participant, data=sub.long, family="binomial")`. Random slopes for trial were attempted but the model did not converge.

who had mentioned the agent (MONKEY EAT) as the informative speaker (52%); and when they saw arrays with ambiguous patients, they were more likely to choose puppets who mentioned the patient (EAT ORANGE) (39%). Age significantly moderated this effect: there was an interaction between age and condition effects ($X^2 = 6.35$, $df=4$, $p<0.05$), indicating that younger children were performing differently on this task than their older counterparts. Binning by year, children ages 4, 5, and 6 tended to choose informative speakers, while 3 year olds did not (Exact multinomial test: three-year-olds, $p = 0.57$; four-year-olds, $p < 0.01$; five-year-olds, $p < 0.05$; six-year-olds, $p < 0.01$).

Did these age trends result from developing pragmatic abilities, or from difficulties with the task among younger children? The control experiment was designed to determine whether children were overall able to choose between speakers despite both speaking ungrammatically. The results of the true/false task are shown on the right hand side of Figure 2. A linear mixed model with maximal random effect structure showed a non-significant effect of age on selecting the true speaker ($X^2 = 2.66$, $df = 4$, $p = 0.1$). Both three and four year olds selected the true puppet at rates greater than chance (Exact multinomial tests,

Figure 2: Children's choices of informative speakers, collapsing across ambiguity conditions. Plot shows number of children in each age group selecting the informative/true speaker on 0, 1, or 2 trials.

three-year-olds, $p < 0.05$; four-year-olds, $p < 0.001$); there was a marginal but not significant difference between the three and four year olds (Fisher's exact test, $p = 0.08$). While three-year-olds in particular struggled with the pragmatic questions (main task), they were above chance on the control task, indicating that the task itself was not entirely too challenging for them - but it *was* possibly more challenging for them than for the older children.

Further evidence that performance on the task improves during this developmental period is that by treating pragmatic and truthful as 'correct' answers, we do not detect a significant interaction between the child's age and the experimental task ($X^2 = 1.07$, $df = 6$, $p = 0.3$). In other words, it may be the case that both task performance and pragmatic abilities are increasing during this window. Thus, this experiment shows strong evidence that pragmatic abilities do improve between age 3 to age 6, but we would hesitate to specify a lower bound to children's ability to make the required pragmatic inference. We do not rule out the possibility that four year olds may be the youngest age at which children can compare alternative utterances for informativity, but also note that we will clearly need other, simpler tasks to further explore pragmatic abilities at this age range.

Discussion

4- to 6-year olds judged two-word descriptions of transitive events according to their relative informativity, preferring sentence fragments that disambiguated the ambiguity present in context and communicated what had happened to a naïve listener. When there were two possible agents in an array and just one possible patient, children preferred speakers who mentioned the agent of the sentence, disambiguating the target event. In contrast, when the arrays contained a single agent and two possible patients, they preferred the speaker who mentioned the patient. We also show clear evidence for developing referential communication abilities: selection of more informative puppets increased with age, and the youngest children (3-year-olds) chose between helpful and unhelpful puppets at chance. Our findings suggest 4-6-year-olds use the relative informativity of event descriptions to determine which of

two speakers is the most helpful. These results are consistent with a ‘rational speaker’ model in which participants choose what to say and what to omit in a way that maximally reduces uncertainty for a listener (Frank & Goodman 2012). We have shown that a rational-speaker style computational model of argument dropping in more complex contexts (e.g. 2 agents and 5 patients) successfully captures the dynamics of adults’ performance (Kline, Schulz & Gibson In press). The older children’s performance here – on a simpler but analogous task - suggests that they are using similar strategies.

How should we interpret the at-change performance of the three-year-olds in this study? One possibility is that they struggle with the intended computational challenge of selecting the most informative from among multiple ways of communicating a message. This study does not address whether children can *generate* the relevant alternatives themselves, but even here, when the alternatives are made explicit, three-year-olds’ performance is not reliable. This explicit task is unnatural in the sense that children do not normally have to choose between two speakers’ alternate ways of encoding a message, but has the advantage that it ‘solves’ one part of the pragmatic problem for children, and avoids the problem of having to infer from our own adult intuitions what alternate utterances the child may be generating themselves. However, 3-year-olds do succeed (if not perfectly) at this kind of two-alternative choice task for other linguistic questions, including the control true/false version of this study and other linguistic studies that ask children as young as 2 years old to make judgments between alternatives (Franklin, Clifford, Williamson, & Davies, 2005; Noble, Rowland, & Pine, 2011). This tentative evidence for success on the central question of which utterance informativity lends support to theories that propose that online executive function and limited processing capacity – as much or more than awareness of what constitutes an informative utterance – are what produce the striking and characteristic failures of pragmatic speech by young children (Brown-Schmidt, 2009; Nilsen & Graham, 2009; Rabagliati & Robertson, 2017).

On the other hand, while three-year-olds’ performance on the control task is above chance, they are nowhere near ceiling performance, and we see possible improvement at the task itself in four-year-olds.

Thus, another interpretation is that 3-year-olds are in fact able to make these informativity comparisons, but struggle to apply them in this setting. Children in this study heard puppets use non-standard utterances like “Monkey eat” neither was saying what the child themselves would say to describe the event - namely, the full sentence “The monkey eats the orange”. A subtler challenge is that success in this task requires not only remembering what was under the flap and what was said, but also the knowledge states of at least three people other than themselves: their caregiver, who hasn't seen what's under the flap, and the two puppets, who have. This might present challenges for children whose explicit theory-of-mind abilities are still developing (Wimmer & Perner, 1983), in contrast to a task like Stiller et al. (2015) (see Figure 1) that doesn't require children to explicitly consider a conversation partner who knows less than they do. In everyday pragmatic production, keeping track of what someone else knows is a central part of evaluating how helpful a sentence will be to them. Understanding how the rest of the pragmatic 'skill set' (monitoring alternative utterances and alternative messages, calculating informativity) interacts with developing theory of mind abilities will be an important question for future research.

Referential communication at the sentence level

A key difference between this study and most previous work is the nature of the messages being transmitted (and thus the utterances used to describe them). In the majority of referential communication studies, children are dealing with the identification of an object referent, typically labeled with a single noun phrase or novel label, out of a set of referents. Here, they must do something slightly different: given a set of people and objects, they must represent the possible set of *events* that might have occurred and possible ways of referring to those events. In the context of the current studies, this transformation is potentially very simple: children could ignore the unambiguous participant (e.g. the single agent) and simply remember what could fill the other role (e.g. either the cat or the dog as patient). However, it is possible that in general children treat sentence-level communication differently from simple reference because the set of alternative messages is both broader and more vague. A major simplifying feature of the particular events that we consider is that they are unambiguous with respect to the role of participants

in the verb argument structure. The overwhelming likelihood is that, when there is a human and a fruit in the context, the human would be the agent and the fruit would be the patient of an eating event, not the other way around. In the real world, entities, especially humans, are often candidates to be either the agent or the patient of an event (e.g., *hugging*). Furthermore, while we present scenes that are given a clear interpretation and focus (the two speakers agree that we should talk about 'eating'), the true 'context set' for simple sentences includes everything about an event that you might want to describe (even assuming you want to talk about an event that is currently going on). This includes not just a choice of agents and patients, but what perspective you want to take on each possible event: 'A monkey is eating an orange', and 'A Monkey is sitting under a tree' might be equally good descriptions of a particular situation.

Somewhat analogous problems arise in referential communication tasks for objects (should I refer to the big dog in the display thinking of it as a Dalmatian, a mammal, a plastic toy? cf. Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Xu & Tenenbaum, 2007.) However, it is a truly open question how far we should expect what we know about referential communication – getting your conversation partner to agree with you about what object you are focusing on – to extend to the larger problem of getting your conversational partner to understand you at the sentence level (i.e. agree on the event and event perspective you are taking).

Studying children's understanding of referential communication at the sentence level also has implications for how we think about children's early multi-word speech, and in particular for theories of subject-dropping. Argument dropping is a common phenomenon in many languages, and represents the broadest reference selection choice at the sentence level: whether to mention a particular participant in the event at all (beyond whether it is more appropriate to say *it*, or *the horse*, or *my favorite big green horse* in a given context). Toddlers go through a stage in language acquisition where, even in a language like English that doesn't allow this, they tend to drop the subjects from sentences, often producing utterances like "eat apple." The distribution of argument dropping in child speech has been explicitly or implicitly attributed to pragmatic influences across many theories (Allen, 1997, 2008; Bloom, 1990; Greenfield & Zukow,

1978; Valian, 1991). That is, the particular times that children produce subjects when they are in the subject-drop stage are those cases when the subject is informative, and likewise drops subjects when they can be recovered from context. A major impediment to understanding the pragmatic status of dropped arguments is that most of the data on subject dropping comes from corpora of natural speech between children and their caregivers: these children are too young to participate reliably in many experimental paradigms that can elicit complex speech. In corpus studies, toddlers appear more likely to mention arguments of sentences when those arguments are new to the dialog or otherwise inaccessible to their conversational partner (cf. Allen, 2008). Unfortunately, these corpus studies suffer from the necessary limitation that (especially when only audio recordings or transcripts are available) nonlinguistic context is not only uncontrolled but may be unrecoverable, leaving researchers to rely on linguistic context alone to understand just how informative these young children's utterances actually were in context. An important direction for future language acquisition research is to identify how the surprisingly sophisticated social abilities in early toddlerhood (Brandone et al., 2014; Carpenter et al., 1998; Carpenter, Call, & Tomasello, 2005; Csibra, 2003; Heyes, 2014) might be leveraged to create controlled experimental contexts that can test for toddlers' sensitivity to the information conveyed in language in a way that clearly connects to their own natural speech during the same developmental period.

We hope that the present study takes a step toward uniting three literatures: studies of the pragmatic abilities of preschoolers who are using relatively complex language, the early non-linguistic social awareness of what others know and want, and the specific role of information content and what other people know in children's earliest multi-word speech. Findings in each of these domains should constrain theory-making in the others: if preschoolers do not show evidence of being able to design referentially informative expressions at the sentence level, it becomes more difficult to ascribe pragmatic 'motives' to the elliptical speech that these much younger toddlers are producing. On the other hand, if even young infants are quite pragmatically aware in their nonlinguistic cognition but don't fully extend these abilities to language, we need to understand why. We show in this study that it is possible to make progress on the

emergence of a specific pragmatic skill – comparing and selecting between a more and less informative way of referring – in the context of a new domain of referential communication, namely simple sentence-level meanings. Fully understanding this process will require answering many of the questions being addressed in the noun/object referential communication literature, including as how children's developing pragmatic skill reflects initial patterns of attention, processing capacity, the ability to monitor whether your own utterances are ambiguous, and the ability to generate the relevant possible alternatives of both what you could say and what you could mean (Brown-Schmidt, 2009; Nilsen & Graham, 2009; Rabagliati & Robertson, 2017, Stiller et al. 2015). Critically, as this study shows, ensuring that tasks are both appropriate for an age group and being interpreted by children in the way we expect is just as important as the key results as well: experimental results can both over- and under-estimate children's abilities or the source of their difficulties. Fleshing out these interacting influences into a coherent theory of referential communication and its development is a central goal for the study of pragmatics, and a diversity of both experimental tasks and types of referential communication will be necessary to reach the end goal of modeling and predicting children's actual performance in settings outside the lab and understanding the development of pragmatic abilities during the early years of life.

References

- Allen, S. (1997). A discourse-pragmatics explanation for the subject-object asymmetry in early null arguments: The principle of informativeness revisited. In A. Sorace, C. Heycock, & R. Shillcock (Eds.), *Proceedings of the GALA '97 Conference on Language Acquisition*. Edinburgh: University of Edinburgh Press.
- Allen, S. (2008). A discourse-pragmatic explanation for argument representation in child Inuktitut. *Linguistics*, 483–521. <https://doi.org/10.1515/ling.38.3.483>
- 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>

- Bloom, P. (1990). Subjectless Sentences in Child Language. *Linguistic Inquiry*, 21(4), 491–504.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322–30. <https://doi.org/10.1016/j.cognition.2010.10.001>
- Brandone, A. C., Horwitz, S. R., Aslin, R. N., & Wellman, H. M. (2014). Infants' goal anticipation during failed and successful reaching actions. *Developmental Science*, 17(1), 23–34. <https://doi.org/10.1111/desc.12095>
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482–1493.
- Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.
- Brown-Schmidt, S. (2009). The role of executive function in perspective taking during online language comprehension. *Psychonomic Bulletin & Review*, 16(5), 893–900. <https://doi.org/10.3758/PBR.16.5.893>
- Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, 21(2), 315–330. [https://doi.org/10.1016/S0163-6383\(98\)90009-1](https://doi.org/10.1016/S0163-6383(98)90009-1)
- Carpenter, M., Call, J., & Tomasello, M. (2005). Twelve- and 18-month-olds copy actions in terms of goals. *Developmental Science*, 8(1), F13–F20.
- Csibra, G. (2003). Teleological and referential understanding of action in infancy. *Phil. Transcripts of the Royal Society of London*, 358, 447–458.
- Davies, C., & Katsos, N. (2010). Over-informative children: Production/comprehension asymmetry or tolerance to pragmatic violations? *Lingua*, 120(8), 1956–1972. <https://doi.org/10.1016/j.lingua.2010.02.005>
- Eskritt, M., Whalen, J., & Lee, K. (2008). Preschoolers can recognize violations of the Gricean maxims. *British Journal of Developmental Psychology*, 26(3), 435–443. <https://doi.org/10.1348/026151007X253260>

- Ferreira, V. S., Slevc, L. R., & Rogers, E. S. (2005). How do speakers avoid ambiguous linguistic expressions? *Cognition*, 96(3), 263–284. <https://doi.org/10.1016/j.cognition.2004.09.002>
- Foppolo, F., Guasti, M. T., & Chierchia, G. (2012). Scalar Implicatures in Child Language: Give Children a Chance. *Language Learning and Development*, 8(4), 365–394. <https://doi.org/10.1080/15475441.2011.626386>
- Franklin, A., Clifford, A., Williamson, E., & Davies, I. (2005). Color term knowledge does not affect categorical perception of color in toddlers. *Journal of Experimental Child Psychology*, 90(2), 114–141. <https://doi.org/10.1016/j.jecp.2004.10.001>
- 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>
- 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>
- Greenfield, P. M., & Zukow, P. G. (1978). Why do children say what they say when they say it? An experimental approach to the psychogenesis of presupposition. In K. Nelson (Ed.), *Children's Language* (Vol. 1, pp. 287–336). New York, NY: Gardner Press.
- Grice, H. (1975). Logic and Conversation. *Syntax and Semantics*, 3, 41–58.
- Heller, D., Grodner, D., & Tanenhaus, M. K. (2008). The role of perspective in identifying domains of reference. *Cognition*.
- Heyes, C. (2014). False belief in infancy: a fresh look. *Developmental Science*, 17(5), 647–659. <https://doi.org/10.1111/desc.12148>
- Huang, Y. T., & Snedeker, J. (2009). Semantic meaning and pragmatic interpretation in 5-year-olds: Evidence from real-time spoken language comprehension. *Developmental Psychology*, 45(6), 1723–1739. <https://doi.org/10.1037/a0016704>
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33), 12002–12007. <https://doi.org/10.1073/pnas.1407479111>

- Keysar, B., Barr, D. J., Balin, J. A., & Brauner, J. S. (2000). Taking Perspective in Conversation: The Role of Mutual Knowledge in Comprehension. *Psychological Science*, 11(1), 32–38.
<https://doi.org/10.1111/1467-9280.00211>
- Keysar, B., Barr, D. J., & Horton, W. S. (1998). The Egocentric Basis of Language Use: Insights From a Processing Approach. *Current Directions in Psychological Science*, 7(2), 46–50.
<https://doi.org/10.1111/1467-8721.ep13175613>
- Kline, M., Schulz, L., & Gibson, E. (In press). Partial truths: Adults choose to mention agents and aatients in proportion to informativity, even if it doesn't fully disambiguate the message. *Open Mind*, 1–13. https://doi.org/10.1162/OPMI_a_00013
- Morisseau, T., Davies, C., & Matthews, D. (2013). How do 3- and 5-year-olds respond to under- and over-informative utterances? *Journal of Pragmatics*, 59, Part A, 26–39.
<https://doi.org/10.1016/j.pragma.2013.03.007>
- 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.
- Nilsen, E. S., & Graham, S. A. (2009). The relations between children's communicative perspective-taking and executive functioning. *Cognitive Psychology*, 58(2), 220–249.
<https://doi.org/10.1016/j.cogpsych.2008.07.002>
- Nilsen, E. S., Graham, S. A., Smith, S., & Chambers, C. G. (2008). Preschoolers' sensitivity to referential ambiguity: evidence for a dissociation between implicit understanding and explicit behavior. *Developmental Science*, 11(4), 556–562. <https://doi.org/10.1111/j.1467-7687.2008.00701.x>
- Noble, C. H., Rowland, C. F., & Pine, J. M. (2011). Comprehension of Argument Structure and Semantic Roles: Evidence from English-Learning Children and the Forced-Choice Pointing Paradigm. *Cognitive Science*. <https://doi.org/10.1111/j.1551-6709.2011.01175.x>
- Phillips, A. T., & Wellman, H. M. (2005). Infants' understanding of object-directed action. *Cognition*, 98, 137–155.

- Rabagliati, H., & Robertson, A. (2017). How do children learn to avoid referential ambiguity? Insights from eye-tracking. *Journal of Memory and Language*, 94, 15–27.
<https://doi.org/10.1016/j.jml.2016.09.007>
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439. [https://doi.org/10.1016/0010-0285\(76\)90013-X](https://doi.org/10.1016/0010-0285(76)90013-X)
- Rubio-Fernández, P. (2016). How Redundant Are Redundant Color Adjectives? An Efficiency-Based Analysis of Color Overspecification. *Frontiers in Psychology*, 7.
<https://doi.org/10.3389/fpsyg.2016.00153>
- Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: evidence for effects of informativity expectations. *Journal of Psycholinguistic Research*, 32(1), 3–23.
- Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71(2), 109–147.
- Sonnenschein, S., & Whitehurst, G. J. (1984). Developing Referential Communication: A Hierarchy of Skills. *Child Development*, 55(5), 1936. <https://doi.org/10.2307/1129940>
- Stiller, A. J., Goodman, N., & 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>
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675–735.
- Valian, V. (1991). Syntactic subjects in the early speech of American and Italian children. *Cognition*, 40(1-2), 21–81.
- Whitehurst, G. J., Sonnenschein, S., & Ianfolla, B. J. (1981). Learning to Communicate from Models: Children Confuse Length with Information. *Child Development*, 52(2), 507–513.
<https://doi.org/10.2307/1129168>
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114, 245–272. <https://doi.org/10.1037/0033-295X.114.2.245>