Sensitivity to over-informativeness

Running head: Sensitivity to over-informativeness

Development of children's sensitivity to over-informativeness in learning and teaching.

Hyowon Gweon¹*, Patrick Shafto², & Laura Schulz³

¹ Department of Psychology, Stanford University

² Department of Mathematics and Computer Science, Rutgers University - Newark

³ Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

KEYWORDS: Social learning, theory of mind, communication, pragmatics,
informativeness, common ground

*This article has been accepted for publication in *Developmental Psychology* (05/2018).

Please use the following to cite this work until full citation is available.

Gweon, H., Shafto, P., & Schulz, L. (in press). Development of children's sensitivity to over-informativeness in learning and teaching. *Developmental Psychology*. DOI: 10.1037/dev0000580

© 2018, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/dev0000580

Corresponding Author: Hyowon Gweon

Department of Psychology, Stanford University

450 Serra Mall, Jordan Hall (Bldg. 420)

Stanford, CA 94305

Email: gweon@stanford.edu

Phone: (650) 498-1194

Abstract

Effective communication requires knowing the “right” amount of information to provide; what is necessary for a naïve learner to arrive at a target hypothesis may be superfluous and inefficient for a knowledgeable learner. The current study examines four- to seven-year-olds’ developing sensitivity to over-informative communication and their ability to decide how much information is appropriate depending on the learner’s prior knowledge. In Experiment 1 (N=184, Age: 4.09 - 7.98), five- to seven-year-old children preferred teachers who gave costly, exhaustive demonstrations when learners were naïve, but preferred teachers who gave efficient, selective demonstrations when learners were already knowledgeable given their prior experience (common ground). However, four-year-olds did not show a clear preference. In Experiment 2 (N=80, Age: 4.05 – 6.99), we asked whether children flexibly modulated their own teaching based on learners’ knowledge. Five and six-year-olds, but not four-year-olds, were more likely to provide exhaustive demonstrations to naïve learners than to knowledgeable learners. These results suggest that by age five, children are sensitive to over-informativeness and understand the trade-off between informativeness and efficiency; they reason about what others know based on the presence or absence of common ground and flexibly decide how much information is appropriate both as learners and as teachers.

Keywords: social learning, theory of mind, communication, pragmatics, informativeness, common ground

Introduction

Humans benefit from a bedrock of accumulated knowledge that would be impossible for individuals to acquire on their own (Boyd & Richerson, 1988; Tomasello, 1999). In informal social interactions and formal educational contexts, humans contribute to the process of cultural transmission both as learners and as teachers. As learners, they evaluate others' informativeness and selectively learn from better teachers; as teachers, they share their knowledge with others by providing information for learners. However, knowing what counts as effective teaching is a nontrivial challenge; the kind of information that is sufficient, relevant, and most beneficial to others varies with context. Thus, both identifying effective teachers and teaching effectively require an abstract understanding of informativeness that goes beyond a simple preference for "more information".

In real-world communicative contexts, the time and effort involved in providing and processing information are inherently limited; thus communicators often have to sacrifice informativeness for efficiency, or vice versa. When informational gain competes with the pressure to be economical, communicators can make better decisions by understanding what the interlocutor already knows, and what information is necessary or superfluous given the interlocutor's epistemic state. Imagine for instance, that you are asked directions to the nearest hardware store. If you are asked by someone you recognize as a local, you might respond with a broad directive ("Take the freeway and get off at exit 24"). However, if you are asked by a neighbor who just moved in next door, you would provide more detailed, turn-by-turn instructions. Providing detailed information to the local or sparse information to the newcomer would both constitute communicative failures. Intuitions like these have been captured by Grice's *Maxim of Quantity* (Grice, 1975) -- we expect a speaker to be only as informative as required, and we conform to these expectations as speakers ourselves.

Figuring out the "right" amount of information for someone can be especially challenging, because others' epistemic states are not fixed; they change with experience. Thus we need to

consider our mutual experiences with others (i.e., *common ground*, Clark, Schreuder, & Buttrick, 1983) and flexibly update our beliefs about what others know given this history of mutual experiences. To return to our earlier example, if you had just given your new neighbor a tour of the area around the hardware store, it would be rather bizarre to provide detailed directions; a simple reminder would suffice. Importantly, this is not because the detailed directions are false or irrelevant; indeed, they might still confer some modest benefit on your neighbor. However, given how much she already knows, the cost of providing the detailed information (e.g., your time, physical effort, and cognitive effort) likely exceeds the added value (e.g., your neighbor's certainty about the directions).

By avoiding “too much information” -- information that is *true, relevant, but unnecessary* -- we can achieve our communicative goals in ways that are both efficient and effective. Here we ask whether four- to eight-year-old children have an abstract understanding of informativeness that goes beyond a simple preference for “more information”. We first review prior work on children's understanding of informativeness with a particular focus on its development from the preschool to early school years, noting that it has overwhelmingly focused on children's evaluation of under-informativeness. We then discuss why resisting *over-informativeness* might be challenging for young children, and how this ability might develop in early childhood. Finally, we introduce two experiments designed to investigate the development of children's sensitivity to common ground and trade-offs between informative and efficient communication as learners (Experiment 1) and as teachers (Experiment 2).

The development of children's sensitivity to informativeness

Starting early in childhood, children show surprisingly sophisticated abilities to use others' knowledge and ignorance both to interpret others' communicative acts and to engage in effective communication themselves. As early as 12-months, infants track what others do and do not know and selectively communicate information (e.g., by pointing to fallen objects) when their conversational partner is ignorant (Liszkowski, Carpenter, & Tomasello, 2008; Tomasello &

Haberl, 2003; see also O'Neill, 1996; O'Neill & Topolevec, 2001). Between 14 and 18-months, infants become increasingly adept at using their understanding of shared knowledge to interpret others' ambiguous referents (Lebal, Behne, Carpenter, & Tomasello, 2009; Moll, Richter, Carpenter, & Tomasello, 2008, Saylor & Ganea, 2007; Southgate, Chevalier, & Csibra, 2010); by two, children can use common ground to infer the meaning of a novel word (Akhtar, Carpenter, & Tomasello, 1996) and use adult feedback to improve the precision of their own referential expressions (Matthews, Butcher, Lieven, & Tomasello, 2012). Moreover, two-year-olds, like older children (Bonawitz et al., 2011), treat the omission of information as informative in itself. For instance, when a knowledgeable teacher pedagogically demonstrates a single function of a novel toy, children not only learn the demonstrated function but also infer that the toy does not have any additional functions (Shneidman, Gweon, Schulz, & Woodward, 2016). By three, children distinguish people who provide true versus false information, and preferentially learn from informants who were previously accurate (Koenig, Clément, & Harris, 2004; Birch, Vautier, & Bloom, 2008, Corriveau & Harris, 2009; see Sobel & Kushnir, 2013, for a review). Three-year-olds can also tailor the information they communicate to others, using more informative nouns rather than less informative pronouns to describe an event if their conversational partner cannot see it than if she can (Matthews, Lieven, Theakston, & Tomasello, 2006).

Between four and eight, children become increasingly sensitive to teachers who are under-informative; six-year-olds readily judge under-informative teachers and compensate with additional exploration (Gweon, Pelton, Konopka, & Schulz, 2014), but four-year-olds do so only if they have first seen an example of a fully informative teacher (Gweon & Asaba, 2017). This finding is consistent with a recent body of work on pragmatic inference; although considerable work suggests that even early school-aged children struggle to identify under-informative speakers (e.g., a speaker who says "I ate *some* of the cookies" when he in fact ate *all* of them) (Noveck, 2001; Barner, Chow & Yang, 2009; Huang & Snedeker, 2009; Hurewitz, Papafragou,

Gleitman, & Gelman, 2006; Papafragou & Tantalou, 2004), children as young as four can succeed when the relevant alternatives (e.g., that a speaker could have said “all” instead of “some”) are made clear in the context (Barner, Brooks, & Bale, 2011; see also Foppolo, Guasti, & Chierchia, 2012; Katsos & Bishop, 2011; Papafragou & Musolino, 2003, Skordos & Papafragou, 2016). Furthermore, while five-year-olds readily consider a speaker’s prior observation of an event to draw appropriate pragmatic inferences, four-year-olds struggle to do so unless the task is made simpler by reducing the demands for mental-state inference (Papafragou, Friedberg, & Cohen, 2017).

Between five and eight, children also become increasingly proficient at teaching others (Strauss, Ziv, & Stein, 2002; Ziv & Frye, 2004); they selectively transmitting conventional behaviors (Clegg & Legare, 2016) or information that others would have a hard time discovering (Ronfard, Was & Harris, 2016). They also tailor both the quantity and quality of their demonstrations depending on the learners’ goals and relative competence (Gweon & Schulz, 2018; see also Bass, Bonawitz, Shafto, Ramarajan, Gopnik, & Wellman, 2017). By seven, children can appropriately sample exemplars that support accurate learning (Rhodes, Gelman, & Brickman, 2010), and teach information that maximizes the learner’s benefits while minimizing the learner’s costs of exploration (Bridgers, Jara-Ettinger, & Gweon, 2016).

These empirical results are in line with recent work directly probing children’s understanding of what it means to learn and teach. For instance, children’s tendency to explain learning as a “process” (e.g., mentioning the source of information or the strategy by which a learner acquires information or skill) improves between four and eight years of age (Sobel & Letourneau, 2015). Between preschool and early school years, children also become increasingly able to distinguish between learning from their own actions and from instructions (Sobel & Letourneau, 2017) and explain teaching as a process that causes a change in one’s knowledge state (Sobel & Letourneau, 2016), suggesting that children’s developing

understanding of the relationship between learning and teaching as a process of knowledge change may support their sensitivity to informativeness.

Over-informativeness and the problem of inferring the “right amount” of information

Collectively, prior work suggests that although there is considerable developmental change between preschool and middle childhood in children’s ability to evaluate others’ communication (as learners) and to communicate effectively (as teachers), even four-year-olds’ understanding of informativeness goes beyond mere accuracy. As noted however, previous work has overwhelmingly focused on children’s sensitivity to under-informative communication. Critically, effective communicators should also recognize over-informativeness and actively resist providing “too much information” with respect to the learner’s prior knowledge. Do young children understand that a given piece of information may be under-informative, sufficient, or superfluous depending on what the learner knows, and avoid over-informative communication as learners and as teachers?

Note that superfluous information does not lead to inaccurate inferences; it may even yield some benefit to the learner to the extent that it increases the learner’s certainty about the target hypothesis. Nonetheless, as adults we have the intuition that speakers should resist providing more than required in a communicative context (the maxim of quantity, Grice, 1978; relevance theory, Sperber & Wilson, 1995; Wilson & Sperber, 2004). On this account, a helpful informant is not someone who always provides exhaustive information; it is someone who does so only when additional information supports accurate learning. When the learner has enough information to infer the correct hypothesis, an effective informant should stop providing additional information even when this information, in principle, could further reduce the learner’s uncertainty. This account rests on the assumption that information transfer is costly (i.e., requires time, physical effort, and cognitive effort for the teacher to generate evidence, and for the learner to process evidence). Although studies suggest that even infants are sensitive to the cost of goal-directed actions (Gergely & Csibra, 2003; Liu, Ullman, Tenenbaum, & Spelke,

2017), relatively little is known about the development of children's understanding of more abstract representations of costs (see Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016, for a discussion). Critically, to *resist* providing more information when the additional benefit of doing so is minimal, one must integrate the costs of communication (with respect to efficiency, for both the informant and the learner) and the rewards of communication (with respect to informativeness, given the learner's prior knowledge) to determine how much information is enough; thus an ability to avoid "too much information" might be relatively challenging for young children.

Research on children's sensitivity to over-informativeness has been scarce. While a few studies have shown that five-year-olds may be sensitive to superfluous descriptions of objects in some contexts (e.g., they penalize speakers who ask for a "fresh" apple when there is only one apple; Davies & Katsos, 2010, Morriseau, Davies, & Matthews, 2013), prior work also suggests that over-informative utterances are common in school-aged children and even adults (e.g., Deutsch & Pechmann, 1982; Ford & Olson, 1975). Importantly, in linguistic communication, speakers may be tolerant of over-informative communication not because they are insensitive to over-informativeness but because the costs to both the speaker and learner of an additional utterance (e.g., a superfluous modifier) are minimal. Given that we are interested in looking at the development of children's understanding of over-informativeness, it makes sense to examine cases where the cost of transmitting information (for both the speaker and listener) is relatively high. This may be most true in the context of non-verbal demonstrations where informants must expend physical effort to generate goal-directed actions, and both informants and learners must invest time to demonstrate and observe these actions. Although prior theoretical work has suggested that Gricean communicative principles extend to non-linguistic demonstrations (Baldwin, Loucks, & Sabbagh, 2008), there is little empirical support for children's understanding of over-informativeness *per se*.

Motivated by the work reviewed above on the development of children's understanding of under-informativeness through middle childhood (e.g., Barner et al., 2011; Foppolo et al., 2012; Gweon et al., 2014; Gweon & Asaba, 2017; Skordos & Papafragou, 2016), here we investigate whether four to eight-year-olds can determine how much information is appropriate or over-informative based on the learner's prior knowledge. To this end, we manipulated the presence of common ground knowledge between the teacher and the learner (Clark, Schreuder, & Buttrick, 1983) such that the same communicative act should be evaluated differently depending on what the informant *thinks* the learner already knows. Given that four-year-olds often fail to consider speakers' knowledge to predict or interpret their utterances (Papafragou et al., 2017; see also Lagattuta, Sayfan, & Blattman, 2010) or penalize under-informative communication (e.g., Barner et al., 2011, Katsos & Bishop, 2011; Gweon & Asaba, 2017) but become more sensitive to linguistic pragmatics starting at age five (e.g., Foppolo et al., 2012; Skordos & Papafragou, 2016; Papafragou et al., 2017), we predict a similar developmental change here. Specifically, we hypothesize that five, six, and seven-year-olds, but not four-year-olds, will be sensitive to over-informative communication both when they are evaluating teachers (Experiment 1) and when they are teaching themselves (Experiment 2).

Experiment 1

In Experiment 1, we examine whether children, between four and eight years of age, will show an increasing appreciation for the *efficiency* of communication. Although exhaustively demonstrating all affordances on a causal mechanism (e.g., pressing all of the buttons on a toy with 20 buttons) may eliminate all uncertainty about their functions, presenting partial evidence may be sufficient, and even desirable, when the learner has relevant prior knowledge. In Experiment 1, children (as third-party observers) watch two teacher-learner interactions and are asked to choose a preferred teacher based on these interactions (See Figure 1).

We show children toys with 20 visually identical causal affordances (i.e., electronic buttons) on each. Given that buttons on toys are usually functional, a naïve learner should

assume that all 20 buttons on the toys are functional as well. In fact, however, only three buttons work on these kinds of toys; given this incorrect assumption, the naïve learner would benefit the most from seeing an exhaustive (and costly) demonstration of all the buttons; providing low-cost, selective information (i.e., a demonstration of just a few functional buttons on the toy) would be under-informative and misleading. However, for a learner who already knows that only three buttons work on these toys, exhaustive demonstration would be unnecessary and superfluous; the teacher only needs to demonstrate the three working buttons on all future toys of this kind. Thus, the critical manipulation is whether the learners have common ground knowledge about the toys based on their interaction with the teachers (who were fully knowledgeable about the toys) and the child (who explored the toys to learn the ground truth). If children understand the trade-off between informativeness and efficiency of communication, they should be more likely to prefer the selective (efficient) informant to the exhaustive (but costly) informant when the learner is naïve, but show the reverse preference when the learner is knowledgeable. Given prior work on children's developing sensitivity to informativeness, we predict that four-year-olds may struggle with these inferences, but five, six, and seven-year-olds should prefer different informants depending on the presence or absence of common ground.

Methods

Participants Children were recruited from an urban children's museum and university preschool. The sample size for this initial experiment was set to be consistent with previous research on preschoolers' sensitivity to informant reliability (e.g., Koenig et al., 2004); we recruited a total of 184 children ($M_{\text{age}}(\text{SD}) = 6.10 (1.10)$, 4.09 – 7.98) evenly distributed across the age range. Post-hoc analysis showed that this sample size yields a power of .87. Participants were randomly assigned to one of two conditions: Common Ground ($N=92$; 20 4-year-olds, 21 5-year-olds, 24 6-year-olds, and 27 7-year-olds; 51 girls) and No Common Ground conditions ($N=92$; 21 4-year-olds, 21 5-year-olds, 27 6-year-olds, and 23 7-year-olds; 53 girls). A total of 13 children were excluded and replaced due to parental interference ($N=6$),

experimental error (N=4), or not completing the procedure (N=3). An additional 30 children were excluded from analyses because they were unable to pass the memory check questions; 16 of these children were 4-year-olds (see Procedure)¹. All experiments were approved by the institutional review board for human subjects research at MIT (Protocol# 0408000894: Causal Reasoning Study) and Stanford University (Protocol# 31350: Research on Social Cognition in Infants, Children, and Adults).

Materials Four custom-built toys were made from foam board, electrical push-button switches, and simple circuits that played musical tunes. The toys were identical except for color (red, green, blue and yellow.) Each toy was a rectangular box with 20 push-button switches (henceforth buttons) along the top panel (80 (L) x 8(H) x 8(W) cm). Three buttons on each toy were connected to electrical circuits so that pressing each button activated a different musical tune. The positions of the active buttons varied across toys, and the remaining buttons were inert but looked identical to the active buttons. Thus there was no way to tell which buttons would play music without pressing the buttons. Two puppets were used as Toymaker A and Toymaker B. The two Toymaker puppets looked identical except that “A” or “B” was written on their ties. Two other puppets (Bert and Ernie) were used as learners.

Procedure Participants were tested individually in a quiet room inside the museum or preschool. The experiment consisted of four phases: Puppet Introduction, Exploration, Observation, & Choice. In the Common Ground Condition, the phases were presented in that order (described in detail below, see Figure 1A); in the No Common Ground condition, the Exploration phase preceded the Puppet Introduction phase.

¹ The original sample targeted five, six, and seven-year-olds, and the decision to recruit four-year-olds was made later to further explore developmental trends. The same procedure was used for four-year-olds except that easier memory check questions were used for exclusion. Nevertheless, many four-year-olds had difficulty answering the exclusion questions (see Methods), resulting in a higher exclusion rate (28%) than older children (8.9%). Including these children did not change the main results; proportion of children who chose the selective teacher in Common Ground vs. No Common Ground conditions: 48% vs. 48.6%, $p = 1$ (4-year-olds), 67.1% vs. 42.3%, $p = 0.002$ (older children).

1. Puppet Introduction: The Common Ground condition started with the experimenter introducing the Toymakers and the learners (Bert and Ernie) to the participant. The participants were told: “Here’s Bert and Ernie! They’re kids, just like you, and they just got to the museum/preschool so they have never seen toys like this before. They don’t know how these toys work! Here’s Toymaker A, and Toymaker B! They come from all the way across the ocean. They make cool toys like these, so they know all about these toys!”

2. Exploration: In the Common Ground condition, the experimenter then pointed to the four toys and said, “When you press the buttons on these toys, they play music. But importantly, not all the buttons work – only some of them play music. Why don’t you go ahead and play with this blue toy first?” Children were allowed to freely explore the toy. After the child tried all the buttons, the experimenter asked the child how many buttons played music on the blue toy (i.e., “three”) and to tell Bert, Ernie, and the Toymakers how many buttons played music on the blue toy. This procedure was repeated with the green toy. Thus in the Common Ground condition, everyone had a strong prior belief that just a few (i.e., “three”) buttons worked on these toys: the Toymakers had made these toys (and thus knew everything about them), children had played with the blue and green toy themselves, and Bert and Ernie (as well as the Toymakers) had watched the child play and were explicitly told how many buttons worked on those toys. Furthermore, the Toymakers were present during the Exploration phase so that they knew what Bert and Ernie had seen. In the No Common Ground condition, children freely explored the toy and answered how many buttons played music in the absence of Bert, Ernie, and the Toymakers.

3. Observation: The experimenter pointed to the remaining two toys (red and yellow) and said, “Toymaker A and Toymaker B want to show Bert and Ernie how these toys work, but they don’t speak English; they only speak Jabberwocky. Bert and Ernie don’t speak Jabberwocky, so the Toymakers will have to show Bert and Ernie how the toys work”. First, children watched as Toymaker A showed Bert the yellow toy; then Toymaker B showed Ernie the same toy. Before

each Toymaker start pressing buttons on a toy, children were reminded that Toymaker A (or B) knows all about the toy, and that he is going to show Bert (or Ernie) how it works. One Toymaker provided selective evidence, pressing just the three active buttons on the toy. The other Toymaker provided exhaustive evidence, pressing every button on the toy sequentially. Half of the children saw Toymaker A demonstrate selective evidence and Toymaker B demonstrate exhaustive evidence; the other half saw the reverse. To ensure that children had encoded not just the location of the active buttons on each toy but the difference in the demonstrations, children were given memory check questions. When both Toymakers finished demonstrating the yellow toy, the experimenter asked: "What was different about how Toymaker A showed how the toy works and how Toymaker B showed how the toy works?" To pass this question, children had to mention that one pressed all the buttons and the other did not. The same procedure was repeated with the red toy. If a child failed to report the difference between the two Toymakers even after watching their demonstrations on the second toy, the child was excluded from the analysis. Because pilot data suggested many four-year-olds were unable to explain the difference, four-year-olds were additionally asked (after each Toymaker's demonstration of a toy) whether the Toymaker pressed all 20 buttons or just 3 buttons. If the child could accurately answer these questions for both Toymakers, they were included in the analyses even though they could not verbally explain the difference.

4. Choice: Children were told, "See the cabinet over there? It's full of toys just like these, and you need to learn about them. Which Toymaker would you rather learn from: Toymaker A, or Toymaker B?" The experimenter looked down at the table and held the two puppets equidistant from the child until the child made a choice.

In the No Common Ground condition, the order of the Introduction and Exploration phases were flipped so that children explored the blue and green toys first and only then were introduced to Bert, Ernie, and the Toymakers. Thus in the No Common Ground condition, Bert and Ernie never saw the child play with the toys; only the child and the Toymakers knew that

just a few buttons worked on these toys. Critically, this manipulation influenced only the learners' (Bert and Ernie) prior knowledge about how many buttons worked on these toys; the Toymakers and the children always knew that only three buttons worked. (See Results.)

By using these stimuli and experimental design, we were able to fix the relative costs of selective and exhaustive evidence across conditions (i.e., 3 button presses (low-cost) and 20 button presses (high-cost), respectively) while varying the subjective value of the evidence for the learner given their prior knowledge. (See Supplementary Online Material for a more detailed discussion of the toys' costs and rewards.)

Results

After exploring the blue and green toys, children were asked how many buttons worked on each one of the toys; children were equally accurate in the two conditions (Common Ground $M(SD) = 3.2(0.8)$ vs. No Common Ground $M(SD)=3.1(1.1)$; $z = 1.27$, $p = 0.2$, Mann-Whitney U).

Our primary measure of interest was children's choice between the two Toymakers. First, we fit a logistic regression model with Condition (Common Ground, No Common Ground) and Age Group (4, 5, 6, & 7) as predictors (Teacher Choice $\sim$ Condition * AgeGroup). There was a main effect of Condition ($\chi^2(1,182) = 10.62$, $p = 0.001$) but no effect of AgeGroup or interaction ($p < 0.25$). Consistent with this, children as a whole were more likely to choose the selective Toymaker (who pressed just the functional buttons) in the Common Ground condition than in the No Common Ground condition (Common Ground vs. No Common Ground: 62.0% vs. 38.0%, $p = 0.002$, Fisher's Exact Test, See Figure 1B). Children in the Common Ground condition chose the selective Toymaker over the exhaustive Toymaker above chance (57 of 92, $p = 0.03$, two-sided binomial); by contrast, children in the No Common Ground condition chose the exhaustive Toymaker above chance (57 of 92, $p = 0.03$, two-sided binomial). See Fig.1B.

Although Age Group did not predict children's choice above and beyond condition, examining children's teacher choice separately for each age group revealed a potential difference between four-year-olds and older children (% choice selective toymaker in Common

Ground vs. No Common Ground: 4yrs: 45.0% vs. 38.0%; 5 yrs: 66.6% vs. 38.1%; 6 yrs: 62.5% vs. 37.0%; 7 yrs: 70.4% vs. 39.1%, see Supplemental Online Material, Figure S1). To further explore possible developmental trends, we fit a logistic regression model with age as a continuous variable separately for each condition. Age predicted children's Toymaker choice in the Common Ground condition ($\chi^2(1, 90) = 4.28, p = 0.04$) but not in the No Common Ground condition ($\chi^2(1, 90) = 0.003, p = .96$). Exploratory analyses suggested no evidence of a condition difference in four-year-olds (% choice Selective Teacher in Common Ground vs. No Common Ground: 45% vs. 38%; $p = 0.76$, Fisher's Exact); however, consistent with our hypothesis, older children as a group were more likely to choose the Selective Teacher in the Common Ground Condition than in the No Common Ground condition (66.7% vs. 38.0%, $p < 0.001$, Fisher's Exact;). Looking within each condition, children in the Common Ground condition preferred the selective teacher (48 of 72, $p = 0.006$, two-sided binomial) and marginally preferred the exhaustive teacher in the No Common Ground condition (27 of 71, $p = 0.057$, two-sided binomial). Within the older group, we did not find an effect of age in either condition (CG: $\chi^2(1, 70) = 0.82, p = 0.37$; NCG: $\chi^2(1, 69) = 0.0003, p = 0.99$).

Discussion

The results from Experiment 1 suggest that children aged five and older show sensitivity to over-informativeness, understanding that the same communicative act can take on different values depending on the learner's prior knowledge. When the learner's prior knowledge allowed accurate learning even from a selective demonstration, five, six, and seven-year-old children were more likely to choose the teacher who provided this demonstration than the teacher who provided exhaustive information about the toy. However, we did not find evidence for such sensitivity in four-year-olds.

Although children's performance differed across the two conditions, children did not perform at ceiling; a substantial minority of children chose wrongly in both conditions. Why did some children choose the exhaustive teacher in the Common Ground condition? In both

conditions, the selective Toymaker's demonstration was *incomplete* (i.e., "sin of omission"; see Gweon et al., 2014), exposing the learner to some epistemic risk. Having seen only 3 of 20 buttons on each toy, the learner was left uncertain about the status of the remaining 17 buttons (e.g., green and red toys could have more than 3 buttons that play music); only the exhaustive demonstration provided absolute certainty about these toys for the learners. This epistemic risk was considerably lower for a knowledgeable learner (Common Ground condition) than for a naïve learner (No Common Ground condition); a learner who had already seen that only three buttons worked on the blue and yellow toys would reasonably expect that the green and red toys also have three working buttons, especially given strong sampling by a knowledgeable informant (see e.g., Shafto, Goodman, & Griffiths, 2014; Xu & Tenenbaum, 2007).

Nevertheless, when learning from the selective informant, even the knowledgeable learner had to settle for a reasonable inductive inference rather than certainty. Thus, some children might have valued exhaustive information despite its costs. Our results suggest that children begin to appreciate not only informativeness but also efficiency in communication by the end of preschool years; unlike four-year-olds, five to eight-year-old children were (appropriately) more likely to accept the small epistemic risk and settle for a reasonable inductive inference rather than absolute certainty when the learners had common ground to support accurate inference.

Perhaps more puzzlingly, although children preferred the exhaustive informant overall in the No Common Ground condition, a substantial minority of children chose the selective informant even though the learner was naïve. One possibility is that at least some of the children might have failed to recognize that the sparse data could mislead the naïve learner (i.e., leading them to conclude from the evidence of three working buttons that all the buttons worked). These children may have been afflicted with a "curse of knowledge" (e.g., Birch & Bloom, 2007); since they themselves knew the ground truth about the toys, they may have failed to recognize that a naïve learner would not (see also Lagattuta, 2010, Papafragou et al., 2017, for a similar failure in pragmatic inference tasks). Additionally, children may have been

impatient with the demonstration themselves and considered their own utilities rather than the learners', thus preferring the quicker, more efficient teacher.

Finally, note that there was no clear preference for either teacher among four-year-olds in either condition. One possibility is that these results reflect a genuine lack of sensitivity to over-informativeness in four-year-olds. However, prior work on preschool-aged children's limited understanding of under-informativeness suggests that young children's pragmatic competence may be masked by task demands or insufficient contextual support (e.g., Barner et al., 2011, Gweon & Asaba, 2017; Katsos & Bishop, 2011). Given the high exclusion rate in four-year-olds, it is possible that four-year-olds failure is due to the high task demands rather than a lack of pragmatic competence. A simpler task might reveal a greater understanding of over-informativeness in this age group, just as previous work has demonstrated successful pragmatic inferences in four-year-olds in simplified tasks (e.g., Barner et al., 2011; Liebal, Carpenter, & Tomasello, 2013; Papafragou et al., 2017; Stiller, Goodman, & Frank, 2014). Thus in Experiment 2, we used a task where the participant did not have to track and remember the demonstrations of multiple informants, but simply had to teach one toy to one learner who either did or did not have common ground with the child; navigating the trade-off between efficient and informative communication may be easier if children themselves are incurring the costs of providing information. Given that four-year-olds failed to show this sensitivity in Experiment 1, but children as young as five and six did, we specifically targeted a group of four-year-olds and a group of five and six-year-olds for a comparison between the two age groups.

Experiment 2

Based on both prior work and results in Experiment 1, we predicted that five- and six-year-olds, but not four-year-olds, would provide exhaustive or selective information depending on the prior knowledge of the learner; more specifically, we predicted that older children would be more likely to provide exhaustive evidence (i.e., press all 20 buttons) when the learner is naïve than when the learner has prior knowledge about the toys.

Methods

Subjects Forty 5 & 6-year-olds (M_{age} (SD) = 5.79 (0.55), range: 5.00 – 6.99, 18 girls; Older Group) and forty 4-year-olds (M_{age} (SD) = 4.49 (0.27), range: 4.05 – 4.97, 17 girls; Younger Group) were recruited from a local children's museum and university preschool and were randomly assigned to one of two conditions: Common Ground and No Common Ground. Sample size was set assuming a relatively large condition difference in the number of buttons pressed ($d = 0.8$) for a power of .75 ($N=20$ in each condition, in each age group, total $N=80$). An additional 6 children were recruited but excluded from analyses due to parental interference ($N=2$), experimental error ($N=2$), or not completing the procedure ($N=2$).

Materials The same 20-button toys in Experiment 1 were used, but we attached magnetic stripes on the side of each toy so that small magnets (1cm diameter) could be placed to indicate which buttons worked. An Elmo puppet was used as the learner.

Procedure The procedure consisted of an Exploration phase, Puppet Introduction phase, and Teaching phase. In the Common Ground condition, the Puppet Introduction occurred during the Exploration phase (see description below); in the No Common Ground condition, the Puppet Introduction occurred after the Exploration phase.

1. Exploration: In both conditions, children were first given the green toy to explore. In contrast to the preceding experiments where the experimenter provided minimal guidance during exploration, in Experiment 2 the experimenter ensured that children pressed all the buttons on the toy, and also provided magnets that could attach to the side of the toy so children could mark and remember the buttons that played music. After pressing all buttons, children were asked how many buttons play music. In the Common Ground condition, Puppet Introduction phase started immediately after the child explored the green toy. In the No Common Ground condition, the child continued to explore other three toys (blue, red, and yellow), and the Puppet Introduction phase came only after the child finished exploring all four toys.

2. Puppet Introduction: In the Puppet Introduction Phase, the experimenter brought out the Elmo puppet and said: “This is Elmo. He has never seen these toys before!” In the Common Ground condition, this introduction came after the child explored the first green toy; the child was then allowed to explore the rest of the toys (blue, red, and yellow) while Elmo sat on the table and “watched”. The child was asked to tell the experimenter and Elmo how many buttons worked on each toy after she finished exploring it. Thus in both conditions, all children had explored all four toys exhaustively, correctly answered that three buttons worked on each of the toys, and put visual markers to remember their locations later. Furthermore, Elmo had never seen the green toy. What critically varied across conditions was whether Elmo had observed the child exploring the three other toys.

3. Teaching: In both conditions, the experimenter then asked the child to show Elmo how the green toy worked, which he had never observed before. The experimenter said: “Elmo doesn’t speak English; he only speaks Jabberwocky. So instead of telling Elmo how the toy works, you will have to show Elmo how the toy works.” Then she placed the green toy between Elmo and the child so that only the child could see where the magnets were and walked out of the child’s line of sight.

Results

Almost all children demonstrated the three working buttons on the toy (one 6-year-old and one 4-year-old showed two working buttons). Thus our primary measures of interest were: (1) whether children provided exhaustive evidence by pressing all of the buttons rather than just the working buttons, and (2) how many buttons children pressed on average. We first collapsed the data across the two age groups and asked whether condition and age predict the proportion of children who provided exhaustive evidence (i.e., pressed 20 buttons, 3 functional and 17 inert). As in Experiment 1, we first fit a logistic regression model with Condition and Age Group as categorical predictors (Exhaustive Evidence ~ Condition * Age Group). We found an effect of Condition ($\chi^2(1,78) = 4.15, p = 0.04$) and a Condition by Age Group interaction ($\chi^2(1, 76) =$

4.24, $p = 0.04$), but no effect of Age Group ($p < 0.4$). Given these results, we ran subsequent analyses separately in each age group.

As predicted, in the Older Group (5 & 6-year-olds), children were more likely to provide exhaustive demonstrations in the No Common Ground than in the Common Ground condition (CG vs. NCG: 25% vs. 70%, $p = 0.01$, Fisher's Exact); in the Common Ground condition, only 5 of 20 provided exhaustive evidence ($p = 0.04$, two-sided binomial), whereas in the No Common Ground condition, 14 of 20 provided exhaustive evidence ($p = 0.11$, two-sided binomial)². However, consistent with results in Experiment 1, we did not find this pattern in the Younger Group; their teaching behaviors did not differ across conditions (CG vs. NCG: 40% vs. 40%, $p = 1$, Fisher's Exact), with 8 of 20 children providing exhaustive evidence in both conditions ($p = 0.5$, two-sided binomial). See Figure 2B.

These results were also reflected in the average number of buttons children demonstrated for the learner. In the Older Group, children in the No Common Ground condition demonstrated more buttons than children in the Common Ground condition (CG vs. NCG: $M(SD) = 9.6(7.4)$ vs. $16.55(6.8)$, $d = 0.97$, $Z = 2.69$, $p = 0.007$, Mann-Whitney U). However, this difference was not significant in the Younger Group (CG vs. NCG: $M(SD) = 10.1(8.3)$ vs. $12.25(8.3)$, $d = 0.26$, $Z = 0.55$, $p = 0.58$, Mann-Whitney U). See Figure 2C.

Discussion

Experiment 2 results suggest that five- and six-year-olds flexibly decide how much information to provide to a learner by considering what he already knows and the cost of transmitting information. As in Experiment 1, exhaustive evidence was beneficial in both conditions; selectively demonstrating just the working buttons would expose the learner to some epistemic risk about the status of the 17 remaining buttons. However, this epistemic risk is

² In the older group, 2 children in the No Common Ground condition accidentally skipped a button during teaching, pressing 19 instead of 20; taking these children into account, 16 of 20 children provided exhaustive evidence ($p = 0.01$).

considerably lower in the Common Ground condition than in the No Common Ground condition because the learner would expect the toy to have three working buttons based on his prior knowledge. As teachers, five- and six-year-olds considered the common ground between themselves and the learner and provided costly, exhaustive evidence when it was necessary for the learner, but resisted doing so when such evidence was unnecessary given his prior experience. By contrast, and consistent with the results in Experiment 1, we did not find evidence for this understanding in four-year-old children.

Note that we failed to find this evidence despite the fact that the task was arguably more engaging and simpler than the task in Experiment 1; it involved a single learner (as opposed to multiple learners and teachers) taught by children themselves who were asked to demonstrate the toy (rather than watching others' demonstrations as third-party observers). Children were also given clear, visual cues to the working buttons on the toy, making it easier to distinguish the difference between selective vs. exhaustive demonstrations.

Four-year-olds' responses in this task suggest that they were not simply confused or responding randomly; 17 of 40 four-year-olds pressed exactly the three working buttons among 20 identical buttons (suggesting that they successfully used the stickers to identify the working buttons), and an additional 16 children exhaustively pressed all the buttons. The fact that most four-year-olds provided either selective or exhaustive evidence (82.5%, $p < 0.001$) as did older children (72.5%, $p = 0.006$) suggests that four-year-olds understood the purpose of the task and were motivated to teach the learner (see Fig.S2 in Supplementary Online Material for a histogram of button presses). Nevertheless, unlike older children, four-year-old children did not modulate the amount of evidence they provided based on the learner's prior knowledge. We discuss potential reasons for this developmental difference in General Discussion (see Limitations and future directions).

General Discussion

Across two experiments, we looked at the development of children's sensitivity to over-informative communication, and whether they could decide how much information was appropriate for naïve and knowledgeable learners. By age five, children preferred teachers who provided costly, exhaustive demonstrations to naïve learners and efficient, partial demonstrations to more knowledgeable learners (Experiment 1). Similarly, children flexibly modulated their own teaching; they provided costly exhaustive demonstration to naïve learners who would otherwise be misled, but resisted doing so for learners who shared common background knowledge with the child, as partial evidence sufficed for those learners to make accurate inferences (Experiment 2). We did not find the same pattern of results in four-year-old children; they neither preferentially chose exhaustive versus selective teachers nor modulated their own teaching with respect to the learner's prior knowledge.

Collectively, these results suggest that five and six-year-olds have an abstract concept of informativeness that goes beyond simple heuristics (e.g., "more information is better"). They understand that the same communicative act (e.g., pressing three working buttons) can be insufficient, superfluous, or "good enough" depending on the learner's prior experience. More specifically, they (1) understand how common ground influences others' epistemic states (Clark et al., 1983; Wilson & Sperber, 2004), (2) are capable of using others' epistemic states to infer the 'right amount' of information needed for accurate inference, and (3) are able to make decisions that reflect this understanding both as learners and as teachers.

In order to evaluate informants, learners need relevant background knowledge themselves; a child can only know that someone who calls a cow a "horse" is unreliable if she already knows what a cow and a horse are herself (Koenig et al., 2004; Birch et al., 2008, Corriveau & Harris, 2009). By evaluating informants' testimony on familiar items given their own knowledge of the world, learners can decide whether to trust the informant in contexts where they lack relevant knowledge (e.g., when the informant calls a novel object a "dax"). The

current study builds on this previous work on epistemic vigilance (see Sobel & Kushnir, 2013, for a review) by showing that by five, children can evaluate the quality of informants not just with respect to their own prior world knowledge but also with respect to their inferences about what *other* learners know (see also Magid, Yan, Siegel, Tenenbaum, & Schulz, 2017).

The idea that effective communication requires informants to consider the mental states of listeners has long been at the center of prominent theories of cooperative communication (e.g., Grice, 1975; Clark, 1996; Wilson & Sperber, 2004; Baldwin, Loucks, & Sabbagh, 2008), and has recently been formalized in computational models of rational speech acts (RSA) and pedagogical communication (e.g., Goodman & Frank, 2016; Shafto, Goodman, & Frank, 2012; Shafto, Goodman, & Griffiths, 2014). The current study is consistent with these accounts, and provides important empirical links between computational accounts of pragmatic inference in linguistic communication (e.g., RSA) and action-based demonstrations (pedagogical reasoning).

However, our results also highlight an important factor that has been emphasized in some (Jaeger & Levy, 2007; Piantadosi, Tily, & Gibson, 2011; Degen, Franke, Jäger, 2013, Shafto, Gweon, Fargen, & Schulz, 2012) but not all models of communication. Many models (e.g., Shafto et al., 2012; 2014) assume that a knowledgeable, helpful teacher selects evidence to maximize the learner's belief in the correct hypothesis (and that the learner draws inferences based on this expectation); this predicts that teachers should always add information if possible, at least until they cannot include any additional information that benefits the learner. The current work was motivated by the idea that this simple assumption fails to capture the fact that information transfer involves costs to both the informant and speaker (e.g., in time, physical effort, processing load, opportunity cost, etc.). Thus a rational teacher should not try to maximize the rewards of information, but maximize the *utility* of information: the difference between the costs and the rewards. When communication is costly, the teacher should limit the amount of information she conveys to whatever is required for a rational learner to infer the

correct hypothesis, even if additional evidence might marginally increase the learners' confidence in that hypothesis.

Prior research on pragmatics has shown that both children and adults are relatively tolerant of over-informativeness (e.g., Deutsch & Pechmann, 1982; Ford & Olson, 1975; Engelhardt et al., 2006; Davies & Katsos, 2010; Morriseau, et al., 2013); thus our results might seem in tension with previous findings. However, because people's tolerance for over-informative communication should depend on how costly that information is, different experimental paradigms may influence the degree to which participants resist over-informativeness. In our study, we limited the channel of communication to time-consuming non-verbal demonstrations (rather than verbal utterances), making exhaustive evidence quite costly both to produce (for the teacher) and observe (for the learner). Unlike prior studies suggesting a limited sensitivity to over-informativeness in linguistic tasks (Davies & Katsos, 2010; Morriseau et al., 2013), we found that children as young as five are sensitive to trade-offs between informativeness and efficiency in communication. If the costs involved in exhaustive communication were lower, children may have preferred maximally informative communication throughout (i.e., if the speaker simply said "None of the other buttons work" after showing the three functional buttons; see Supplementary Online Material for an empirical test of this hypothesis). Additionally, note that in this paradigm having complete certainty about the function of each of the 20 buttons was desirable but not critical. If knowing the status of each button with certainty was essential, children again might have shown a preference for maximal informativeness.

More broadly, although the terms "over-informativeness" and "under-informativeness" are widely used in the literature and treated as infelicitous under Grice's maxim of quantity (Grice, 1975), these terms have not been adequately defined in prior work. We suggest that the "right amount of information" is determined by the information gain in the learners' belief in the correct hypothesis, relative to the cost of information (e.g., the time and effort involved in

communicating and processing the information). From this perspective, both under- and over-informativeness can be defined as communication that has a net negative utility to the listener, but for different reasons; under-informative communication has too little value to the listener relative to its costs, while over-informative communication has too high a cost relative to its value. This account provides a parsimonious explanation of why we resist both over- and under-informativeness and also why people tolerate both in some contexts. Moreover, this account explains people's reaction to over and under-informative communication, in ways that are grounded in precise cognitive terms applicable to the interpretation of intentional, goal-directed behaviors broadly (Jara-Ettinger et al., 2016; see also Goodman & Frank, 2016), rather than by appealing to general principles or linguistic conventions. Importantly, our work suggests that these aspects of pragmatics are supported by an early-developing social cognitive capacity for reasoning about the costs and rewards of others' actions (a "naïve utility calculus"; see Jara-Ettinger et al., 2016). Going beyond prior theoretical proposals for pragmatic inferences in action interpretation (Baldwin, Loucks, & Sabbagh, 2008) and empirical work on the role of gestures in linguistic communication (e.g., Goldin-Meadow & Brentari, 2017), our study provides direct evidence for action-based pragmatics at work, providing an important link between classic and recent theories of communication (e.g., Grice, 1975; Sperber & Wilson, 1995; Wilson & Sperber, 2004; Clark, 1996; Goodman & Frank, 2016) and utility-based theories of social cognition (Jara-Ettinger et al., 2016).

Our results also help distinguish clearly infelicitous over-informativeness from more nuanced failures of effective communication. In the current studies, the exhaustive demonstration allowed for a modest information gain by eliminating all uncertainty about the functions of all 20 buttons. Thus the exhaustive informant was over-informative in the sense of providing information that was not essential for the learner to draw an accurate inference about the toys, but not in the sense of providing information that was completely redundant with what the learner already knew. Future work might look at whether five and six-year-olds distinguish

this kind of over-informativeness from mere repetition of identical information, and whether they also recognize that repetition and redundancy can be useful in some contexts (e.g., when information is difficult for the learner to process).

An open question is *whose* costs children considered. In Experiment 1, we asked children to choose a teacher (for themselves) to learn about similar toys. Thus, children's preference for the efficient teacher likely reflects their consideration of their own costs as learners (e.g., a desire not to waste their own time watching the teacher push 20 buttons when they already knew only three worked). However, it remains possible that children were also considering the cost of these exhaustive demonstrations to the teacher. In Experiment 2, children's tendency to provide selective or exhaustive demonstration depending on the learner's knowledge again is again consistent with their desire to minimize their own costs by not providing more evidence than necessary. However, it is again possible that children were also driven by their desire to save the learner time and effort in observing superfluous evidence. Because helpful teachers generally consider the learners' costs and benefits as well as their own, the exact utility calculus of teachers and learners are difficult to tease apart. Some recent work has begun to address these questions by varying the utility of information to the learner while holding costs to the teacher constant (Bridgers et al., 2016).

More generally, in calculating the utility of communicative acts, there are both agent-independent aspects of costs and rewards (e.g., the time or effort involved in pressing buttons, or the degree to which the demonstration of each button reduces the overall uncertainty about the 20-button toy) and agent-dependent aspects of costs and rewards (e.g., how hard it is for a particular agent to press buttons; how different agents might differently value the same information). In line with recent work on children's developing understanding of learning and teaching (e.g., Sobel & Letourneau, 2015; 2016; 2017), future work might investigate the development of children's sensitivity to different aspects of costs and rewards (including individual differences in learners' competence, or cognitive effort such as concentration or

attention), as well as the situational factors that modulate the impact of communicative utilities (e.g., whether either the teacher or learner is in a rush). An important challenge is to better understand how various kinds of costs and rewards can be mapped onto a “common scale” or “common currency” to support a calculation of the net utility of a communicative act. We hope future work will advance our understanding of how learners compute the utility of information online to evaluate efficiency and informativeness.

We also note that our work may seem in tension with previous work on pedagogical communication suggesting that children constrain their inferences just to the demonstrated function of a causal mechanism when it is demonstrated pedagogically (Bonawitz et al., 2011; Shneidman et al., 2016; see also Xu & Tenenbaum, 2007). Given this, one might wonder why the children in these experiments didn’t expect even the naïve learner to (correctly) infer that only three buttons worked after seeing the selective informant pedagogically demonstrate only three working buttons; if this were the case, we would have observed a preference for the efficient teacher regardless of the conditions in both experiments. The critical distinction is that in previous work (Bonawitz et al., 2011; Gweon et al., 2014; Shneidman et al., 2016) the functions were non-obvious affordances that were perceptually different (thus children had no reason to generalize anything from the demonstration of a novel, unique affordance). By contrast, in the current study, the affordances were familiar, identical buttons that strongly encouraged generalization of functions (Baldwin, Markman, & Melartin, 1993). In particular, prior knowledge about electronic buttons implies that they almost always *make something happen*: Having seen three buttons that make music, the most plausible inference for a naïve learner is that all buttons make music. In the No Common Ground condition where the learner had no way of knowing that active buttons were in fact rare, children in Experiment 2 pressed all buttons to counteract this generalization. Because success in our experiment requires some prior knowledge about electronic buttons as causal affordances, it is important to consider alternative ways of designing stimuli in order to study populations that might lack this knowledge.

Limitations and future directions

The current work provides some suggestive evidence for a developmental change between four and five years of age. Although this is generally consistent with prior work on children's pragmatic competence in linguistic communication that suggests an underlying pragmatic competence and context-dependent performance between ages five to seven (Barner et al., 2011; Foppolo et al., 2012; Gweon & Schulz, 2018; Nadig & Sedivy, 2002; Davies & Katsos, 2010; Morriseau et al., 2013), it remains unclear why 4-year-old children had difficulty with our tasks. Recent studies have found developmental differences in children's sensitivity to under-informativeness (Gweon & Asaba, 2017) and their ability to draw pragmatic inferences given speaker knowledge (Papafragou et al., 2017) with five and six-year-olds out-performing younger children. One possibility is that our results reflect a similar developmental change in children's sensitivity to over-informativeness and their ability to consider learner's prior knowledge. Children's developing theory of mind (Wellman, Cross, & Watson, 2001) and their understanding of teaching and learning broadly (Ziv & Frye, 2004; Strauss et al., 2002; Sobel & Letourneau, 2015; 2016) might contribute to the developmental change shown in our task. Additionally, formal schooling might also facilitate children's consideration of the time and effort involved in learning and teaching.

However, we would be cautious about interpreting the four-year-olds' failures too strongly, especially given prior work that reveals pragmatic competence in younger children by reducing task demands, using more sensitive measures, or making relevant variables more salient in context (Barner et al., 2011; Katsos & Bishop, 2011; Liebal et al., 2013; Papafragou and Tantalou, 2004; Stiller et al., 2014). In Experiment 1, there was a relatively high exclusion rate among four-year-olds, suggesting that they struggled to follow which teacher provided exhaustive or selective demonstration, and although Experiment 2 had lower task demands than Experiment 1, it also required sophisticated mental-state inferences about the learner's prior knowledge based on the presence or the absence of common ground. Recent work

provides suggestive evidence that while four- to six-year-old children can tailor their demonstrations appropriately given explicit information about the learner's goals and competence, their ability to sequence information to best support others' learning increases with age (Gweon & Schulz, 2018). Future work might look at whether younger children's sensitivity to over-informativeness can be identified when the demand for sophisticated mental-state inference is reduced (see Papafragou et al., 2017) and whether children's ability to integrate others' prior knowledge and the cost of information manifests in other aspects of pragmatic reasoning.

Our task opens possibilities for further computational and empirical work for studying the role of information costs in communication. By using demonstrations of causal affordances (e.g., buttons on a toy), our task grounds the notion of costs in concrete, physical terms (e.g., time and effort required for each button press) and allows researchers to manipulate them in a quantifiable way, rather than appealing to cognitive or mental costs that are arguably more difficult to measure or manipulate experimentally. However, we also note that the context for these tasks is, in many respects, artificial. To fix the modality of communication, we constrained informants to non-verbal demonstrations; to ensure that learner's prior knowledge could render the same set of demonstrations effective in one condition but misleading in another, we designed stimuli that violated prior knowledge (i.e., only some of many identical buttons worked). One might ask therefore how far these findings extend to real world contexts.

We believe the idea that people consider the inferential value and the cost of communication is quite general. At the same time, because spoken language provides a remarkably efficient communication channel, the role of communication costs may not always be obvious in verbal communication. Nonetheless, there are a number of real world contexts that demonstrate trade-offs analogous to those in this experiment. For instance, effective classroom teachers typically explain all the classroom rules in detail on the first day of school (when learners are naïve) but thereafter may use only a single word to make their point (e.g.,

trusting that “hands up!” suffices to communicate that students should raise their hands and close their mouths). Similarly, effective teachers are more likely to provide multiple examples of a problem to naïve learners and a single example to more knowledgeable ones. This is not necessarily because the information contained in additional words or examples is irrelevant or redundant; after all, there may still be some information to be gained, just as exhaustive evidence contributes information even in the Common Ground condition of our studies). Rather, these behaviors emerge from the understanding that less information may suffice for accurate learning for knowledgeable learners. Thus while artificial, the task design allowed us to make precise claims about children’s understanding of over-informativeness and its dependence on the learner’s epistemic state. Even in this laboratory setup, children had to consider the learners’ prior knowledge, the evidence the teacher sampled, the cost of the agents’ actions and how object properties can be generalized from data. Understanding the complex interplay of these cognitive capacities and how they develop in early childhood remains an exciting direction for future research.

References

- Akhtar, N., Carpenter, M., & Tomasello, M. (1996). The role of discourse novelty in early word learning. *Child Development*, 67(2), 635-645.
- Baldwin, D., Loucks, J., & Sabbagh, M. (2008). Pragmatics of human action. In T. F. Shipley & J.M. Zacks (Eds.), *Understanding events: From perception to action*. (pp. 96-129). New York, NY: Oxford University Press.
- Baldwin, D., Markman, E., & Melartin, R. (1993). Infants' ability to draw inferences about nonobvious object properties: evidence from exploratory play. *Child Development*, 64(3), 711–728.
- Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition*, 118(1), 84–93.
- Barner, D., Chow, K., & Yang, S.-J. (2009). Finding one's meaning: a test of the relation between quantifiers and integers in language development. *Cognitive Psychology*, 58(2), 195–219. <http://doi.org/10.1016/j.cogpsych.2008.07.001>
- Bass, I., Bonawitz, E., Shafto, P., Ramarajan, D., Gopnik, A., & Wellman, H. (2017). I know what you need to know: Children's developing theory of mind and pedagogical evidence selection. *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Birch, S., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, 18(5), 382-386.
- Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, 107(3), 1018–1034.
- 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), 322–330.

- Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process*. University of Chicago Press.
- Bridgers, S., Jara-Ettinger, J. & Gweon, H. (2016). Children consider others' expected costs and rewards when deciding what to teach. In Papafragou, A., Grodner, D., Mirman, D., & Trueswell, J.C. (Eds.) *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 710-715). Austin, TX: Cognitive Science Society.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground at the understanding of demonstrative reference. *Journal of verbal learning and verbal behavior*, 22(2), 245–258.
- Clegg, J. M., & Legare, C. H. (2016). Instrumental and conventional interpretations of behavior are associated with distinct outcomes in early childhood. *Child Development*, 87(2), 527-542.
- Corriveau, K.H. & Harris, P.L. (2009). Preschoolers continue to trust a more accurate informant 1 week after exposure to accuracy information. *Developmental Science*, 12, 188-193.
- Davies, C., & Katsos, N. (2010). Over-informative children: Production/comprehension asymmetry or tolerance to pragmatic violations? *Lingua*, 120(8), 1956–1972.
- Degen, J., Franke, M., & Jäger, G. (2013). Cost-based pragmatic inference about referential expressions. In M. Knau, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*, 376–381.
- Deutsch, W., & Pechmann, T. (1982). Social interaction and the development of definite descriptions. *Cognition*, 11(2), 159-184.
- Engelhardt, P. E., Bailey, K. G., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity?. *Journal of Memory and Language*, 54(4), 554-573.
- 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.

- Ford, W., & Olson, D. (1975). The elaboration of the noun phrase in children's description of objects. *Journal of Experimental Child Psychology*, 19, 371–382.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: the naive theory of rational action. *Trends in Cognitive Sciences*, 7(7), 287–292.
- Goldin-Meadow, S., & Brentari, D. (2017). Gesture, sign, and language: The coming of age of sign language and gesture studies. *Behavioral and Brain Sciences*, 40, e46.
<http://doi.org/10.1017/S0140525X15001247>
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and Semantics: Vol. 3, Speech Acts* (pp. 41-58). New York: Academic Press.
- Gweon, H. & Asaba, M. (2017). Order matters: Children's evaluation of under-informative teachers depends on context. *Child Development*. DOI: 10.1111/cdev.12825
- Gweon, H., Pelton, H., Konopka, J.A., & Schulz, L.E. (2014). Sins of omission: Children selectively explore when agents fail to tell the whole truth. *Cognition*, 132, 335-341.
- Gweon, H., Shafto, P., & Schulz, L. (2018). Development of children's sensitivity to over-informativeness in learning and teaching. *Data repository on OSF*:
<http://doi.org/10.17605/OSF.IO/6MPR5>.
- Gweon, H., & Schulz, L.E. (2018). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development*. DOI: 10.1111/cdev.13059
- 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.
- Hurewitz, F., Papafragou, A., Gleitman, L., Gelman, R. (2006). Asymmetries in the acquisition of numbers and quantifiers. *Language Learning and Development*, 2(2), 77–96.
- Jaeger, T. F., & Levy, R. P. (2007). Speakers optimize information density through syntactic reduction. In *Advances in Neural Information Processing Systems* (pp. 849-856).

- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20(8), 589-604.
- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition*, 140, 14–23.
- Katsos, N., & Bishop, D. V. M. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition*, 120(1), 67–81.
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33), 12002-12007.
- Koenig, M., Clément, F., & Harris, P. (2004). Trust in Testimony. *Psychological Science*, 15(10), 694-698.
- Liebal, K., Behne, T., Carpenter, M., & Tomasello, M. (2009). Infants use shared experience to interpret pointing gestures. *Developmental science*, 12(2), 264-271.
- Lagattuta, K. H., Sayfan, L., & Blattman, A. J. (2010). Forgetting common ground: six- to seven-year-olds have an overinterpretive theory of mind. *Developmental Psychology*, 46(6), 1417–1432. <http://doi.org/10.1037/a0021062>
- Liebal, K., Carpenter, M., & Tomasello, M. (2013). Young children's understanding of cultural common ground. *British Journal of Developmental Psychology*, 31(1), 88-96.
- Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition*, 108(3), 732–739.
- Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, 358(6366), 1038–1041. <http://doi.org/10.1126/science.aag2132>

- Magid, R. W., Yan, P., Siegel, M. H., Tenenbaum, J. B., & Schulz, L. E. (2017). Changing minds: Children's inferences about third party belief revision. *Developmental Science*. <http://doi.org/10.1111/desc.12553>
- 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.
- Morriseau, T., Davies, C., & Matthews, D. (2013). How do 3- and 5-year-olds respond to under- and over-informative utterances? *Journal of Pragmatics*, 59, 26–39.
- Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, 13(4), 329-336.
- Noveck, I. A. (2001). When children are more logical than adults: experimental investigations of scalar implicature. *Cognition*, 78(2), 165–188.
- O'Neill, D. K. (1996). Two-Year-Old Children's Sensitivity to a Parent's Knowledge State When Making Requests. *Child Development*, 67(2), 659–677.
- O'Neill, D. K., & Topolovec, J. C. (2001). Two-year-old children's sensitivity to the referential (in) efficacy of their own pointing gestures. *Journal of Child Language*, 28(01), 1-28.
- Papafragou, A., Friedberg, C., & Cohen, M. L. (2017). The Role of Speaker Knowledge in Children's Pragmatic Inferences. *Child Development*. <http://doi.org/10.1111/cdev.12841>
- Papafragou, A., & Musolino, J. (2003). Scalar implicatures: experiments at the semantics-pragmatics interface. *Cognition*, 86(3), 253–282.
- Papafragou, A., & Tantalou, N. (2004). Children's computation of implicatures. *Language Acquisition*, 12(1), 71-82.
- Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, 108(9), 3526-3529.
- Ronfard, S., Was, A. M., & Harris, P. L. (2016). Children teach methods they could not discover for themselves. *Journal of Experimental Child Psychology*, 142, 107–117.

- Shafro, P., Goodman, N.D., & Griffiths, T.L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55-89.
- Shafro, P., Gweon, H., Fargen, C., & Schulz, L. (2012). Enough is enough: Inductive sufficiency guides learners' ratings of informant helpfulness. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 977-982). Austin, TX: Cognitive Science Society
- Shneidman, L., Gweon, H., Schulz, L. E., & Woodward, A. L. (2016). Learning From Others and Spontaneous Exploration: A Cross-Cultural Investigation. *Child Development*, 87(3), 723–735.
- Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicatures: Alternatives and relevance. *Cognition*, 153, 6–18.
- Smith, N. J., Goodman, N., & Frank, M. (2013). Learning and using language via recursive pragmatic reasoning about other agents. *Advances in Neural Information Processing Systems*.
- Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: how children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, 120(4), 779–797.
<http://doi.org/10.1037/a0034191>
- Sobel, D. M., & Letourneau, S. M. (2015). Children's developing understanding of what and how they learn. *Journal of Experimental Child Psychology*, 132, 221–229.
<http://doi.org/10.1016/j.jecp.2015.01.004>
- Sobel, D. M., & Letourneau, S. M. (2016). Children's developing knowledge of and reflection about teaching. *Journal of Experimental Child Psychology*, 143, 111–122.
<http://doi.org/10.1016/j.jecp.2015.10.009>
- Sobel, D. M., & Letourneau, S. M. (2017). Preschoolers' Understanding of How Others Learn Through Action and Instruction. *Child Development*. <http://doi.org/10.1111/cdev.12773>

- Southgate, V., Chevallier, C., & Csibra, G. (2010). Seventeen - month - olds appeal to false beliefs to interpret others' referential communication. *Developmental Science*, 13(6), 907-912. <http://doi.org/10.1111/j.1467-7687.2009.00946.x>
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*, Oxford: Blackwell.
- Stiller, A., Goodman, N. D., & Frank, M. C. (2014). Ad-hoc implicature in preschool children. *Language, Learning, & Development*. DOI: 10.1080/15475441.2014.927328
- Strauss, S., Ziv, M., & Stein, A. (2002). Teaching as a natural cognition and its relations to preschoolers' developing theory of mind. *Cognitive Development*, 17(3), 1473–1487.
- Tomasello, M. (1999). The human adaptation for culture. *Annual review of anthropology*, 509-529.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development*, 72(3), 655–684.
- Wilson, D. & Sperber, D., (2004). Relevance Theory. In L. Horn and G. Ward (Eds.). *Handbook of Pragmatics*. Oxford: Blackwell.
- Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental science*, 10(3), 288-297.
- Ziv, M., & Frye, D. (2004). Children's understanding of teaching: The role of knowledge and belief. *Cognitive development*, 19(4), 457-477.

Figure 1. A: Stimuli and procedures for Experiment 1. During Exploration, children first explored the blue and green toys, either in the presence (Common Ground) or absence (No Common Ground) of the puppets. During Observation, one Toymaker pressed three working buttons on yellow and red toys (selective), while the other Toymaker pressed all buttons on these toys (exhaustive). During choice, children were asked: "which toymaker would you rather learn from?" B: Results from Experiment 1. Error bars are bootstrapped 95% CI. * $p < 0.05$; ** $p < 0.005$

Figure 2. A: Schematic of experimental design in Experiment 2. In both conditions, children were asked to teach Elmo about the green toy, which he had never seen before. In the Common Ground condition, Elmo was present while the child explored three other toys; in the No Common Ground condition Elmo was absent during the Exploration phase. B: Results from Experiment 2. Error bars indicate bootstrapped 95% CI. * $p < 0.05$; ** $p \leq 0.01$.

Acknowledgments

Thanks to Grace Bennett-Pierre, Hannah Pelton, Nathaniel Kim, and Veronica Chu for help with data collection. We also thank Josh Tenenbaum, Rebecca Saxe, Mike Frank, and Sidney Strauss for helpful discussions about this work, and three anonymous reviewers for their thoughtful comments. A part of this work has been presented in the Proceedings of the 36th Annual Conference of the Cognitive Science Society (pp. 565-570). All data and analysis scripts are available at: <https://osf.io/6mpr5/> (Gweon et al., 2018; DOI: 10.17605/OSF.IO/6MPR5). This work was supported by the Varieties of Understanding Project at Fordham University and The John Templeton Foundation to H.G., NSF grant DRL-1149116 to P.S., and the Center for Brains, Minds, and Machines (CBMM), funded by NSF STC awards CCF-1231216.

**Supplementary Online Material: Development of children's sensitivity to
"too much information" in learning and teaching.**

Hyowon Gweon, Patrick Shafto, & Laura Schulz

1. Supplementary Figures

Figure S1. Experiment 1. Exploratory plot showing children's choice of teachers by age group (4, 5, 6, and 7-year-olds) and condition.

Figure S2. Experiment 2. Histogram of children's button presses (out of 17 inert buttons).

2. Supplementary Experiment

The experiments in the main manuscript provide the critical tests for our hypothesis that children avoid over-informativeness (i.e., “too much information” in learning and teaching. The experiment reported in this supplementary material presents an additional experiment that confirms the basic preference for informativeness. Recent prior work has shown children’s sensitivity to informativeness; even when two teachers provided the same demonstration of an interesting function of a causal mechanism, children provide lower ratings for the teacher who demonstrated a toy that had additional, undemonstrated functions (Gweon et al., 2014; Gweon & Asaba, in press). This suggests that children evaluate teachers based whether or not the demonstration provided complete information about the causal mechanism, rather than sheer quantity of information, suggesting an abstract understanding of informativeness. However, our experimental design had features that might weaken children’s preference for informative teachers. Previous studies have used causal mechanisms that had non-obvious affordances that were difficult to identify visually (e.g., novel toys with levers or mechanisms that produced unfamiliar effects), such that the informant’s testimony alone determined the learner’s belief about the toy. In our study, we used simple push-buttons that were perceptually identical, and similar musical tunes as effects. Furthermore, the similarity in appearance between the toys also strongly supported learning from prior knowledge and inductive generalization.

Because all children had experience with two of the four toys in Experiment 1, it was important to show that children would still show a preference for more informative informant, and that there is nothing in our experimental setup that would interfere with this preference. Thus using stimuli and procedures almost identical to those in Experiment 1, we verified this basic preference for informativeness in a context where there was little reason to consider the cost of information. In order to establish that children show a clear baseline preference for the fully informative teacher, in this supplementary experiment both teachers demonstrated just three working buttons but one added a short sentence that eliminated all uncertainty about the buttons: “None of the other buttons play music.” after the demonstration. Because this teacher provided full certainty about the toy without incurring a significant cost, there was little reason to reject this fully informative teaching.

Methods

Subjects Given the relatively large effect size observed in Experiment 1 and the stronger manipulation planned for this experiment we assumed a large effect (25% or fewer choosing the selective Toymaker), yielding a power of .7 at $N=20$. Thus forty children were recruited from an urban children's museum (M_{age} (SD) = 6.7 (0.8), range: 5.03 – 7.86, 24 girls) and were randomly assigned to one of two conditions: Common Ground ($N=20$, 12 girls) and No Common Ground conditions ($N=20$, 12 girls). Children were dropped and replaced due to parental interference ($N=2$), experimental error ($N=5$), not completing the procedure ($N=1$), because the parent reported that the child had a developmental disorder ($N=1$) or did not speak English ($N=1$), or for failing to report the difference between the two Toymakers ($N=26$; 15 in Common Ground condition, 11 in No Common Ground condition).

Materials The same toys and puppets used in Experiment 1 were used.

Procedure The procedures were identical to those used in Experiment 1 except as follows. We omitted the part of the Puppet Introduction in which we said that the Toymakers spoke Jabberwocky rather than English. During the Observation phase, both toymakers started their demonstrations by saying "Hi, Bert!" or "Hi, Ernie!". The toymaker who provided selective evidence said "Bye!" after pressing the three buttons; the toymaker who provided exhaustive evidence waved his hand across the toy and said, "None of the other buttons make music! Bye!". To ensure that the children encoded the difference in the Toymakers' communication as well as the location of the active buttons, the experimenter asked the children to identify what was different about the two demonstrations (as in Experiment 1). To pass this question, children had to mention that one of the Toymakers said that none of the other buttons work. The same procedure was repeated with the other toy. If a child failed to report the difference between the two Toymakers even after watching their demonstrations on the second toy, the child was excluded from analysis. The final Choice phase was the same as in Experiment 1.

Results

After exploring the blue and green toys, children were asked how many buttons worked on each one of the toys. Collapsing across both toys, children's mean responses

were not different across conditions (Common Ground $M(SD) = 3.2(1.4)$ vs. No Common Ground $M(SD)=2.7(1.2)$, $t(74) = 1.75$, $p = ns$).

As predicted, in this experiment, children in both conditions showed a strong preference for the exhaustive Toymaker. Only 2 of 20 children in the Common Ground condition (10%; 95% CI: 1 – 31%, $p < 0.001$, two-sided binomial) and 3 of 20 in the No Common Ground condition (15%; 95% CI: 3 - 38%), $p < 0.001$, two-sided binomial; Figure 2) chose the selective Toymaker. A logistic regression on all children (collapsed across conditions, with age as the main predictor) showed no effect of age on children's choice of Toymaker (Nagelkerke $R^2 < 0.01$).

As in Experiment 1 in the main manuscript, we had made an a priori decision to exclude children who failed to report the difference between the Toymakers' demonstrations. However, given the relatively large number of children excluded (presumably due to the quite subtle difference in their behaviors in this experiment), we also looked at the data for children overall. The exclusion criteria did not affect the findings: 75% of children in Common Ground condition and 81% in No Common Ground condition chose the exhaustive informant.

Discussion

As expected, children showed a robust preference for the fully informative teacher, regardless of the learner's prior knowledge. This ensured that the current stimuli and procedure provide results that are consistent with prior work.

The Exhaustive Toymaker's demonstration was only a few seconds longer than the Selective Toymaker's demonstration, adding a few words that were minimally effortful to say. Thus it is unlikely that costs for information, either in time or effort, was a prominent factor for children's decisions. The pattern of results suggest that when costs is not under consideration, children exhibit a robust preference for the teacher who provides more informative evidence, even when the learner has some prior knowledge about the toys.

Note that the Exhaustive Toymaker was not only more informative about the status of the buttons, but also more informative about the reasons for his behaviors; that is, by saying "none of the other buttons work", he effectively explained why he is not pressing the rest of the buttons. Thus it is possible that children's preference for the Exhaustive Toymaker was driven by both kinds of informativeness. Nevertheless, these results show that children are sensitive to the relative informativeness of the teachers,

and that they prefer to learn from the more informative of the two. Thus, stimuli and the dependent measure used in our main experiment produce results that conceptually replicate prior work. These results provide empirical grounds for testing our main hypotheses: rather than preferring exhaustive evidence across the board (as in this supplementary experiment), if children are sensitive to the fact that information transfer

is costly, they should prefer selective communication when the learners' prior knowledge makes it inductively sufficient.

Figure S3. Proportion of children who chose the Selective Toymaker (shown in black) and those who chose the Exhaustive Toymaker (shown in gray).

References

Gweon, H. & Asaba, M. (in press). Order matters: Children's evaluation of under-informative teachers depends on context. *Child Development*. DOI: 10.1111/cdev.12825

Gweon, H., Pelton, H., Konopka, J.A., & Schulz, L.E. (2014). Sins of omission: Children selectively explore when agents fail to tell the whole truth. *Cognition*, 132, 335-341.

Acknowledgments

Thanks to Grace Bennett-Pierre, Hannah Pelton, Nathaniel Kim, and Veronica Chu for help with data collection. We also thank Josh Tenenbaum, Rebecca Saxe, Mike Frank, and Sidney Strauss for helpful discussions about this work, and three anonymous reviewers for their thoughtful comments. A part of this work has been presented in the Proceedings of the 36th Annual Conference of the Cognitive Science Society (pp. 565-570). All data and analysis scripts are available at: <https://osf.io/6mpr5/> (Gweon et al., 2018; DOI: 10.17605/OSF.IO/6MPR5). This work was supported by the Varieties of Understanding Project at Fordham University and The John Templeton Foundation to H.G., NSF grant DRL-1149116 to P.S., and the Center for Brains, Minds, and Machines (CBMM), funded by NSF STC awards CCF-1231216.