Development of Children's Sensitivity to Overinformativeness in Learning and Teaching

Hyowon Gweon
Stanford University

Patrick Shafto
Rutgers University – Newark

Laura Schulz
Massachusetts Institute of Technology

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 4- to 7-year-olds' developing sensitivity to overinformative communication and their ability to decide how much information is appropriate depending on the learner's prior knowledge. In Experiment 1 ($N = 184$, age $= 4.09–7.98$ years), 5- to 7-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 (i.e., common ground). However, 4-year-olds did not show a clear preference. In Experiment 2 ($N = 80$, age $= 4.05–6.99$ years), we asked whether children flexibly modulated their own teaching based on learners' knowledge. Five and 6-year-olds, but not 4-year-olds, were more likely to provide exhaustive demonstrations to naïve learners than to knowledgeable learners. These results suggest that by 5 years of age, children are sensitive to overinformativeness 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, common ground

Supplemental materials: http://dx.doi.org/10.1037/dev0000580.supp

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 4- to 8-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 underinformativeness. We then discuss why resisting *overinformativeness* 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 & Topolovec, 2001). Between 14 and 18 months, infants become increasingly adept at using their understanding of shared knowledge to interpret others' ambiguous referents (Liebal, Behne, Carpenter, & Tomasello, 2009; Moll, Richter, Carpenter, & Tomasello, 2008; Saylor & Ganea, 2007; Southgate, Chevallier, & Csibra, 2010); by 2 years of age, 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, like older children (Bonawitz et al., 2011), 2-year-olds treat the omission of information as informative in itself. For instance, when a knowledgeable teacher pedagogically demonstrates a single function of a novel toy, they not only

learn the demonstrated function but also infer that the toy does not have any additional functions (Shneidman, Gweon, Schulz, & Woodward, 2016). By 3 years of age, children distinguish people who provide true versus false information and preferentially learn from informants who were previously accurate (Birch, Vauthier, & Bloom, 2008; Corriveau & Harris, 2009; Koenig, Clément, & Harris, 2004; 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 they can (Matthews, Lieven, Theakston, & Tomasello, 2006).

Between 4 and 8 years, children become increasingly sensitive to teachers who are underinformative; 6-year-olds readily judge underinformative teachers and compensate with additional exploration (Gweon, Pelton, Konopka, & Schulz, 2014), but 4-year-olds do so only if they have first seen an example of a fully informative teacher (Gweon & Asaba, 2018). This finding is consistent with a recent body of work on pragmatic inference; although considerable work suggests that even early school-age children struggle to identify underinformative speakers (e.g., a speaker who says "I ate *some* of the cookies" when he in fact ate *all* of them; Barner, Chow, & Yang, 2009; Huang & Snedeker, 2009; Hurewitz, Papafragou, Gleitman, & Gelman, 2006; Noveck, 2001; Papafragou & Tantalou, 2004), children as young as 4 years 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, although 5-year-olds readily consider a speaker's prior observation of an event to draw appropriate pragmatic inferences, 4-year-olds struggle to do so unless the task is made simpler by reducing the demands for mental-state inference (Papafragou, Friedberg, & Cohen, 2017).

Children also become increasingly proficient at teaching others (Strauss, Ziv, & Stein, 2002; Ziv & Frye, 2004); they selectively transmit conventional behaviors (Clegg & Legare, 2016) or information that they themselves were unable to discover as learners (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 et al., 2017). By 7 years, 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 4 and 8 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, 2018), 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.

Overinformativeness 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 4-year-olds' understanding of informativeness goes beyond mere accuracy. As noted, however, previous work has overwhelmingly focused on children's sensitivity to underinformative communication. Critically, effective communicators should also recognize overinformativeness 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 underinformative, sufficient, or superfluous depending on what the learner knows, and avoid overinformative 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, 1975; relevance theory: Sperber & Wilson, 1986; 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 young children are sensitive to the cost of goal-directed actions (Gergely & Csibra, 2003; Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015; 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 overinformativeness has been scarce. Although a few studies have shown that 5-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; Morisseau, Davies, & Matthews, 2013), prior work also suggests that overinformative utterances are common in school-age children and even adults (e.g., Deutsch & Pechmann, 1982; Ford & Olson, 1975). Importantly, in linguistic communication, speakers may be tolerant of overinformative communication not because they are insensitive to overinformativeness 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 overinformativeness, it makes sense to examine cases in which the cost of transmitting information (for both the speaker and listener) is relatively high. This may be most true in the context of nonverbal demonstrations in which 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 nonlinguistic demonstrations (Baldwin, Loucks, & Sabbagh, 2008), there is little empirical support for children's understanding of overinformativeness per se.

Motivated by the work on the development of children's understanding of underinformativeness through middle childhood (e.g., Barner et al., 2011; Foppolo et al., 2012; Gweon & Asaba, 2018; Gweon et al., 2014; Skordos & Papafragou, 2016), here, we investigate whether 4- to 8-year-olds can determine how much information is appropriate or overinformative 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 et al., 1983), such that the same communicative act should be evaluated differently depending on what the informant *thinks* the learner already knows. Given that 4-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 underinformative communication (e.g., Barner et al., 2011; Gweon & Asaba, 2018; Katsos & Bishop, 2011), but become more sensitive to linguistic pragmatics starting at 5 years of age (e.g., Foppolo et al., 2012; Papafragou et al., 2017; Skordos & Papafragou, 2016), we predict a similar developmental change here. Specifically, we hypothesize that 5-, 6-, and 7-year-olds, but not 4-year-olds, will be sensitive to overinformative 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 will show an 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 underinformative and misleading. However, for a learner who already knows that
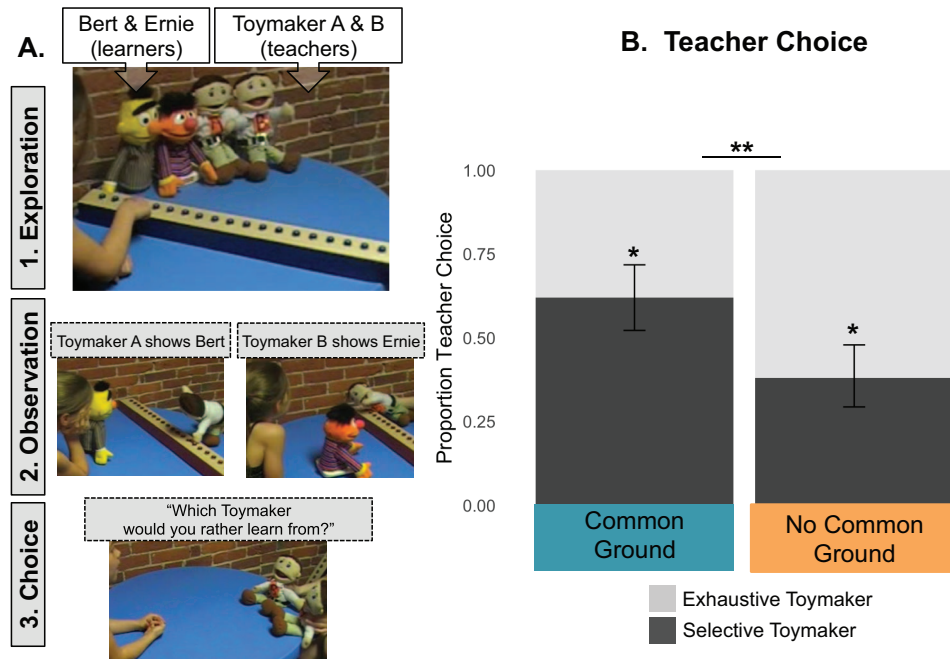
Figure 1. (A) Stimuli and procedures for Experiment 1. During exploration, children first explored the blue and green toys, either in the presence (common-ground) or absence (no-common-ground) of the puppets. During observation, one toymaker pressed three working buttons on yellow and red toys (selective), whereas 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% confidence intervals. * $p < .05$. ** $p < .005$. See the online article for the color version of this figure.

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 4-year-olds may struggle with these inferences, but 5-, 6-, and 7-year-olds should prefer different informants depending on the presence or absence of common ground.

Method

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 pre-schoolers' sensitivity to informant reliability (e.g., Koenig et al., 2004); we recruited a total of 184 children ($M_{age} = 6.10$, $SD = 1.10$, range = 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 because of 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 the Massachusetts Institute of Technology (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

---

[1] The original sample targeted 5-, 6-, and 7-year-olds, and the decision to recruit 4-year-olds was made later to further explore developmental trends. The same procedure was used for 4-year-olds, except that easier memory check questions were used for exclusion (see Procedure). Nevertheless, many 4-year-olds had difficulty answering the exclusion questions (see Method), 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%, $p = 1.0$ (4-year-olds), 67.1% vs. 42.3%, $p = 0.002$ (older children).

[L] × 8[H] × 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, and 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.

Introduction. The common-ground condition started with the experimenter introducing the toymakers and the learners (Bert and Ernie) to the participant. First, 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!" Then participants were introduced to the Toymakers: "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!"

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 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.

Observation. The experimenter pointed to the remaining two toys (red and yellow) and said that, "Toymaker A and Toymaker B want to show Bert and Ernie how these toys work". The experimenter then added, "The Toymakers 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 that many 4-year-olds were unable to explain the difference, 4-year-olds were additionally asked (after each toymaker's demonstration of a toy) whether the toymaker pressed all 20 buttons or just three 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.

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.

By using these stimuli and experimental design, we were able to fix the relative costs of selective and exhaustive evidence across conditions (i.e., three 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 the online supplemental materials 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, $M = 3.2$, $SD = 0.8$; no-common-ground, $M = 3.1$, $SD = 1.1$; $z = 1.27$, $p = .2$, 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 ~ Condition × Age Group). There was a main effect of condition, $\chi^2(1, 182) = 10.62$,

$p = .001$, but no effect of age group or interaction ($p < .25$). 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%, $p = .002$, 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, $p = .03$, two-sided binomial); by contrast, children in the no-common-ground condition chose the exhaustive toymaker above chance (57 of 92, $p = .03$, two-sided binomial; see Figure 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 4-year-olds and older children (% choice selective toymaker in common-ground vs. no-common-ground: 4 years, 45.0% vs. 38.0%; 5 years, 66.6% vs. 38.1%; 6 years, 62.5% vs. 37.0%; 7 years, 70.4% vs. 39.1%; see online supplemental materials, 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, $\chi^2(1, 90) = 4.28$, $p = .04$, 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 4-year-olds (% choice selective teacher in common-ground vs. no-common-ground: 45% vs. 38%; $p = .76$, Fisher's exact test); 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%, $p < .001$, Fisher's exact test). Looking within each condition, children in the common-ground condition preferred the selective teacher (48 of 72, $p = .006$, two-sided binomial) and marginally preferred the exhaustive teacher in the no-common-ground condition (27 of 71, $p = .057$, two-sided binomial). Within the older group, we did not find an effect of age in either condition: common-ground, $\chi^2(1, 70) = 0.82$, $p = .37$; no-common-ground, $\chi^2(1, 69) = 0.0003$, $p = .99$.

Discussion

The results from Experiment 1 suggest that children aged 5 years and older show sensitivity to overinformativeness, 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, 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 4-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 three 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 three 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 4-year-olds, 5- to 8-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); because they themselves knew the ground truth about the toys, they may have failed to recognize that a naïve learner would not (see also Lagattuta et al., 2010, and 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 4-year-olds in either condition. One possibility is that these results reflect a genuine lack of sensitivity to overinformativeness in 4-year-olds. However, prior work on preschool-aged children's limited understanding of underinformativeness 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, 2018; Katsos & Bishop, 2011). Given the high exclusion rate in 4-year-olds, it is possible that 4-year-olds' failure is due to the high task demands rather than a lack of pragmatic competence. A simpler task might reveal a sensitivity to overinformativeness in this age group, just as previous work has demonstrated successful pragmatic inferences in 4-year-olds in simplified tasks (e.g., Barner et al., 2011; Liebal, Carpenter, & Tomasello, 2013; Papafragou et al., 2017; Stiller, Goodman, & Frank, 2015). Thus, in Experiment 2, we used a task in which 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 4-year-olds failed to show this sensitivity in Experiment 1, but children as young as 5 and 6 did, we specifically targeted a group of 4-year-olds and a group of 5- and 6-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 5- and 6-year-olds, but not 4-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.

Method

Subjects. Forty 5- and 6-year-olds ($M_{age} = 5.79$, $SD = 0.55$), range $= 5.00 – 6.99$; 18 girls; "older" group) and 40 4-year-olds ($M_{age} = 4.49$, $SD = 0.27$), 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 six children were recruited but excluded from analyses because of 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 (1 cm 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.

Exploration. In both conditions, children were first given the green toy to explore. In contrast to the preceding experiments in which 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 that 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, the 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.

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.

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." The experimenter then 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 (a) whether children provided exhaustive evidence by pressing all of the buttons rather than just the working buttons, and (b) 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, three 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, $\chi^2(1, 78) = 4.15$, $p = .04$, and a Condition × Age Group interaction, $\chi^2(1, 76) = 4.24$, $p = .04$, but no effect of age group ($p < .4$). Given these results, we ran subsequent analyses separately in each age group.

As predicted, in the older group (5- and 6-year-olds), children were more likely to provide exhaustive demonstrations in the no-common-ground than in the common-ground condition (common-ground vs. no-common-ground: 25% vs. 70%, $p = .01$, Fisher's exact test); in the common-ground condition, only five of 20 provided exhaustive evidence ($p = .04$, two-sided binomial), whereas in the no-common-ground condition, 14 of 20 provided exhaustive evidence ($p = .11$, 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 (common-ground vs. no-common-ground: 40% vs. 40%, $p = 1$, Fisher's exact test), with eight of 20 children providing exhaustive evidence in both conditions ($p = .5$, 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 (common-ground vs. no-common-ground: $M = 9.6$, $SD = 7.4$ vs. $M = 16.55$, $SD = 6.8$; $d = 0.97$, $z = 2.69$, $p = .007$, Mann–Whitney U). However, this difference was not significant in the younger group (common-ground vs. no-common-ground: $M = 10.1$, $SD = 8.3$ vs. $M = 12.25$, $SD = 8.3$; $d = 0.26$, $z = 0.55$, $p = .58$, Mann–Whitney U; see Figure 2C).

_____

[2] In the older group, two 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 ($p = 0.01$).
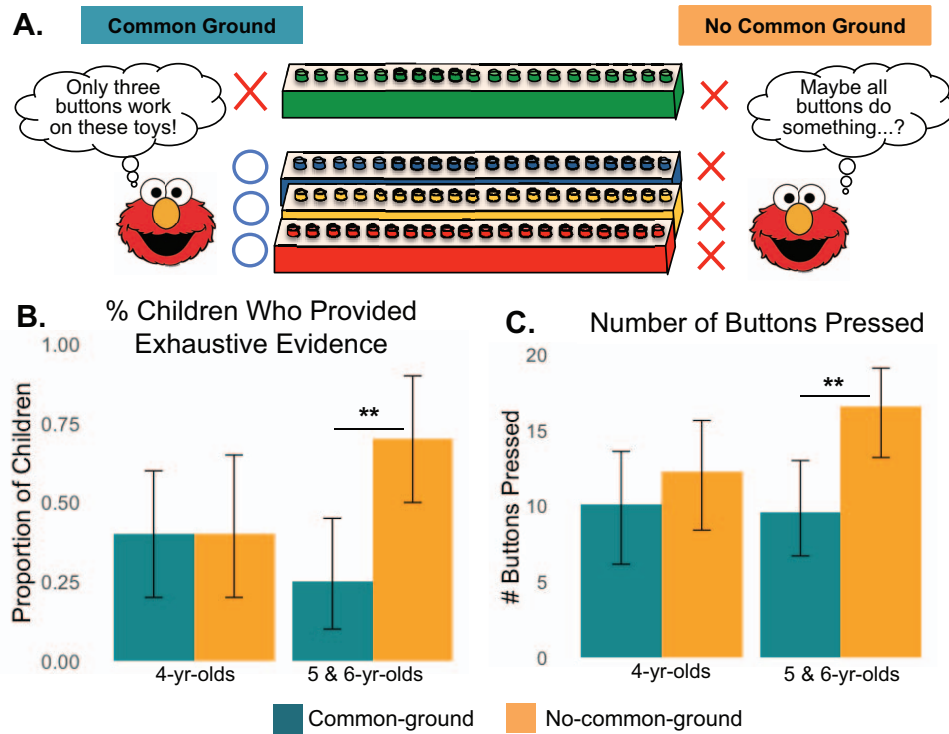
Figure 2. (A) Schematic of experimental design in Experiment 2. In both conditions, children were asked to teach Elmo about the green toy (top), which he had never seen before. In the common-ground condition, Elmo was present while the child explored three other toys; in the no-common-ground condition, Elmo was absent during the exploration phase. (B) Results from Experiment 2. Error bars indicate bootstrapped 95% confidence intervals. * $p < .05.$ ** $p \leq .01.$ See the online article for the color version of this figure.

Discussion

Experiment 2 results suggest that 5- and 6-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 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, 5- and 6-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 4-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 versus exhaustive demonstrations.

Four-year-olds' responses in this task suggest that they were not simply confused or responding randomly; 17 of 40 4-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 4-year-olds provided either selective or exhaustive evidence (82.5%, $p < .001$), as did older children (72.5%, $p = .006$), suggests that 4-year-olds understood the purpose of the task and were motivated to teach the learner (see Figure S2 of the online supplemental materials for a histogram of button presses). Nevertheless, unlike older children, 4-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 the General Discussion.

General Discussion

Across two experiments, we looked at the development of children's sensitivity to overinformative communication and whether they could decide how much information was appropriate for naïve and knowledgeable learners. By 5 years of age, 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 4-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 5- and 6-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 (a) understand how common ground influences others' epistemic states (Clark et al., 1983; Wilson & Sperber, 2004), (b) can use others' epistemic states to infer the "right amount" of information needed for accurate inference, and (c) 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 (Birch et al., 2008; Corriveau & Harris, 2009; Koenig et al., 2004). By evaluating informants' testimony on familiar items given their own knowledge of the world, learners can decide whether to trust the informant in other contexts in which 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 even young 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, 2018).

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., Baldwin et al., 2008; Clark, 1996; Grice, 1975; Wilson & Sperber, 2004), and has recently been formalized in computational models of rational speech act (RSA) and pedagogical communication (e.g., Goodman & Frank, 2016; Shafto, Goodman, & Frank, 2012; Shafto et al., 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 (Degen, Franke, & Jäger, 2013; Levy & Jaeger, 2007; Piantadosi, Tily, & Gibson, 2011; 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., time, physical effort, processing load, opportunity cost). 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 overinfomativeness (e.g., Deutsch & Pechmann, 1982; Davies & Katsos, 2010; Engelhardt, Bailey, & Ferreira, 2006; Ford & Olson, 1975; Morisseau et al., 2013); thus, our results might seem in tension with previous findings. However, because people's tolerance for overinformative communication should depend on how costly that information is, different experimental paradigms may influence the degree to which participants resist overinformativeness. In our study, we limited the channel of communication to time-consuming nonverbal 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 overinformativeness in linguistic tasks (Davies & Katsos, 2010; Morisseau et al., 2013), we found that children as young as 5 years 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 the online supplemental materials 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 *overinformativeness* and *underinformativeness* 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 overinformativeness can be defined as communication that has a net negative utility to the listener, but for different reasons: Underinformative communication has too little value to the listener relative to its costs, whereas overinformative communication has too high a cost relative to its value. This account provides a parsimonious explanation of why we resist over- and underinformativeness in some contexts and tolerate them in other contexts. More, this account explains people's reaction to overinformative and underinformative 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 et al., 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., Clark, 1996; Goodman & Frank, 2016; Grice, 1975; Sperber & Wilson, 1986; Wilson & Sperber, 2004) and utility-based theories of social cognition (Jara-Ettinger et al., 2016).

Our results also help distinguish clearly infelicitous overinformativeness 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 overinformative 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 5- and 6-year-olds distinguish this kind of overinformativeness 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 the current study. 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. Yet, it is again possible that children were also driven by their desire to save the learner's 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, 2018), 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 under-

standing 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 did not 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 nonobvious affordances that were perceptually distinct (thus, children had no reason to generalize anything from the demonstration of one 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, in which 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.

The current work provides some suggestive evidence for a developmental change between 4 and 5 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 5 to 7 (Barner et al., 2011; Davies & Katsos, 2010; Foppolo et al., 2012; Gweon & Schulz, 2018; Morisseau et al., 2013; Nadig & Sedivy, 2002), it remains unclear why 4-year-old children had difficulty with our tasks. Recent studies have found developmental differences in children's sensitivity to underinformativeness (Gweon & Asaba, 2018) and their ability to draw pragmatic inferences given speaker knowledge (Papafragou et al., 2017) with 5- and 6-year-olds outperforming younger children. One possibility is that our results reflect a similar developmental change in children's sensitivity to overinformativeness 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 (Sobel & Letourneau, 2015, 2016; Strauss et al., 2002; Ziv & Frye, 2004) 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 4-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 & Tantalou, 2004;

Stiller et al., 2015). In Experiment 1, there was a relatively high exclusion rate among 4-year-olds, suggesting that they struggled to follow which teacher provided exhaustive or selective demonstration. 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 although 4- to 6-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 overinforamativeness 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 nonverbal 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 therefore ask 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 effect of these 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 for knowledgeable learners; 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, although artificial, the task design allowed us to make precise claims about children's understanding of overinformativeness 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,* 635–645. http://dx.doi.org/10.2307/1131837

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. http://dx.doi.org/10.1093/acprof:oso/9780195188370.003.0006

Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants' ability to draw inferences about nonobvious object properties: Evidence from exploratory play. *Child Development, 64,* 711–728. http://dx.doi.org/10.2307/1131213

Barner, D., Brooks, N., & Bale, A. (2011). Accessing the unsaid: The role of scalar alternatives in children's pragmatic inference. *Cognition, 118,* 84–93. http://dx.doi.org/10.1016/j.cognition.2010.10.010

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,* 195–219. http://dx.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* (pp. 99–104). Austin, TX: Cognitive Science Society.

Birch, S. A., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science, 18,* 382–386. http://dx.doi.org/10.1111/j.1467-9280.2007.01909.x

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,* 1018–1034. http://dx.doi.org/10.1016/j.cognition.2007.12.008

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,* 322–330. http://dx.doi.org/10.1016/j.cognition.2010.10.001

Boyd, R., & Richerson, P. J. (1988). *Culture and the evolutionary process.* Chicago, IL: 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 A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (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, UK: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511620539

Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground at the understanding of demonstrative reference. *Journal of Verbal Learning & Verbal Behavior, 22,* 245–258. http://dx.doi.org/10.1016/S0022-5371(83)90189-5

Clegg, J. M., & Legare, C. H. (2016). Instrumental and conventional interpretations of behavior are associated with distinct outcomes in early childhood. *Child Development, 87,* 527–542. http://dx.doi.org/10.1111/cdev.12472

Corriveau, K., & Harris, P. L. (2009). Preschoolers continue to trust a more accurate informant 1 week after exposure to accuracy information. *Developmental Science, 12,* 188–193. http://dx.doi.org/10.1111/j.1467-7687.2008.00763.x

Davies, C., & Katsos, N. (2010). Over-informative children: Production/comprehension asymmetry or tolerance to pragmatic violations? *Lingua, 120,* 1956–1972. http://dx.doi.org/10.1016/j.lingua.2010.02.005

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* (pp. 376–381). Austin, TX: Cognitive Science Society.

Deutsch, W., & Pechmann, T. (1982). Social interaction and the development of definite descriptions. *Cognition, 11,* 159–184. http://dx.doi.org/10.1016/0010-0277(82)90024-5

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,* 554–573. http://dx.doi.org/10.1016/j.jml.2005.12.009

Foppolo, F., Guasti, M. T., & Chierchia, G. (2012). Scalar implicatures in child language: Give children a chance. *Language Learning and Development, 8,* 365–394. http://dx.doi.org/10.1080/15475441.2011.626386

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. http://dx.doi.org/10.1016/0022-0965(75)90068-5

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences, 7,* 287–292. http://dx.doi.org/10.1016/S1364-6613(03)00128-1

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://dx.doi.org/10.1017/S0140525X15001247

Goodman, N. D., & Frank, M. C. (2016). Pragmatic Language Interpretation as Probabilistic Inference. *Trends in Cognitive Sciences, 20,* 818–829. http://dx.doi.org/10.1016/j.tics.2016.08.005

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics. Volume 3: Speech acts* (pp. 41–58). New York, NY: Academic Press.

Gweon, H., & Asaba, M. (2018). Order matters: Children's evaluation of underinformative teachers depends on context. *Child Development, 89,* e278–e292. http://dx.doi.org/10.1111/cdev.12825

Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition, 132,* 335–341. http://dx.doi.org/10.1016/j.cognition.2014.04.013

Gweon, H., & Schulz, L. (2018). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development.* Advance online publication. http://dx.doi.org/10.1111/cdev.13059

Gweon, H., Shafto, P., & Schulz, L. (2018). *Development of children's sensitivity to over-informativeness in learning and teaching* [Data repository on Open Science Framework]. http://dx.doi.org/10.17605/OSF.IO/6MPR5

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,* 1723–1739. http://dx.doi.org/10.1037/a0016704

Hurewitz, F., Papafragou, A., Gleitman, L., & Gelman, R. (2006). Asymmetries in the acquisition of numbers and quantifiers. *Language Learning and Development, 2,* 77–96. http://dx.doi.org/10.1207/s15473341lld0202_1

Jaeger, T. F., & Levy, R. P. (2007). Speakers optimize information density through syntactic reduction. In B. Schlökopf, J. Platt, & T. Hoffman (Eds.), *Advances in neural information processing systems* (pp. 849–856). Cambridge, MA: MIT Press.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences, 20,* 589–604. http://dx.doi.org/10.1016/j.tics.2016.05.011

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. http://dx.doi.org/10.1016/j.cognition.2015.03.006

Katsos, N., & Bishop, D. V. M. (2011). Pragmatic tolerance: Implications for the acquisition of informativeness and implicature. *Cognition, 120,* 67–81. http://dx.doi.org/10.1016/j.cognition.2011.02.015

Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science, 15,* 694–698. http://dx.doi.org/10.1111/j.0956-7976.2004.00742.x

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,* 1417–1432. http://dx.doi.org/10.1037/a0021062

Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In B. Schlökopf, J. Platt, & T. Hoffman (Eds.), *Advances in neural information processing systems NIPS* (vol. 19, pp. 849–856). Cambridge, MA: MIT Press.

Liebal, K., Behne, T., Carpenter, M., & Tomasello, M. (2009). Infants use shared experience to interpret pointing gestures. *Developmental Science, 12,* 264–271. http://dx.doi.org/10.1111/j.1467-7687.2008.00758.x

Liebal, K., Carpenter, M., & Tomasello, M. (2013). Young children's understanding of cultural common ground. *British Journal of Developmental Psychology, 31,* 88–96. http://dx.doi.org/10.1111/j.2044-835X.2012.02080.x

Liszkowski, U., Carpenter, M., & Tomasello, M. (2008). Twelve-month-olds communicate helpfully and appropriately for knowledgeable and ignorant partners. *Cognition, 108,* 732–739. http://dx.doi.org/10.1016/j.cognition.2008.06.013

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,* 1038–1041. http://dx.doi.org/10.1126/science.aag2132

Magid, R. W., Yan, P., Siegel, M. H., Tenenbaum, J. B., & Schulz, L. E. (2018). Changing minds: Children's inferences about third party belief revision. *Developmental Science, 21*(2). http://dx.doi.org/10.1111/desc.12553

Matthews, D., Butcher, J., Lieven, E., & Tomasello, M. (2012). Two- and four-year-olds learn to adapt referring expressions to context: Effects of distracters and feedback on referential communication. *Topics in Cognitive Science, 4,* 184–210. http://dx.doi.org/10.1111/j.1756-8765.2012.01181.x

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,* 403–422. http://dx.doi.org/10.1017/S0142716406060334

Moll, H., Richter, N., Carpenter, M., & Tomasello, M. (2008). Fourteen-month-olds know what "we" have shared in a special way. *Infancy, 13,* 90–101. http://dx.doi.org/10.1080/15250000701779402

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,* 26–39. http://dx.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,* 329–336. http://dx.doi.org/10.1111/j.0956-7976.2002.00460.x

Noveck, I. A. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition, 78,* 165–188. http://dx.doi.org/10.1016/S0010-0277(00)00114-1

O'Neill, D. K. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development, 67,* 659–677. http://dx.doi.org/10.2307/1131839

O'Neill, D. K., & Topolovec, J. C. (2001). Two-year-old children's sensitivity to the referential (in)efficacy of their own pointing gestures. *Journal of Child Language, 28,* 1–28. http://dx.doi.org/10.1017/S0305000900004566

Papafragou, A., Friedberg, C., & Cohen, M. L. (2017). The role of speaker knowledge in children's pragmatic inferences. *Child Development*. Advance online publication. http://dx.doi.org/10.1111/cdev.12841

Papafragou, A., & Musolino, J. (2003). Scalar implicatures: Experiments at the semantics-pragmatics interface. *Cognition, 86,* 253–282. http://dx.doi.org/10.1016/S0010-0277(02)00179-8

Papafragou, A., & Tantalou, N. (2004). Children's computation of implicatures. *Language Acquisition: A Journal of Developmental Linguistics, 12,* 71–82. http://dx.doi.org/10.1207/s15327817la1201_3

Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *PNAS Proceedings of the National Academy of Sciences of the United States of America, 108,* 3526–3529. http://dx.doi.org/10.1073/pnas.1012551108

Rhodes, M., Gelman, S. A., & Brickman, D. (2010). Children's attention to sample composition in learning, teaching and discovery. *Developmental Science, 13,* 421–429. http://dx.doi.org/10.1111/j.1467-7687.2009.00896.x

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. http://dx.doi.org/10.1016/j.jecp.2015.09.032

Saylor, M. M., & Ganea, P. (2007). Infants interpret ambiguous requests for absent objects. *Developmental Psychology, 43,* 696–704. http://dx.doi.org/10.1037/0012-1649.43.3.696

Shafto, 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. http://dx.doi.org/10.1016/j.cogpsych.2013.12.004

Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning From Others: The Consequences of Psychological Reasoning for Human Learning. *Perspectives on Psychological Science, 7,* 341–351. http://dx.doi.org/10.1177/1745691612448481

Shafto, 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,* 723–735. http://dx.doi.org/10.1111/cdev.12502

Skordos, D., & Papafragou, A. (2016). Children's derivation of scalar implicatures: Alternatives and relevance. *Cognition, 153,* 6–18. http://dx.doi.org/10.1016/j.cognition.2016.04.006

Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference.

Psychological Review, 120, 779–797. http://dx.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://dx.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://dx.doi.org/10.1016/j.jecp.2015.10.009

Sobel, D. M., & Letourneau, S. M. (2018). Preschoolers' understanding of how others learn through action and instruction. *Child Development, 89,* 961–970. http://dx.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,* 907–912. http://dx.doi.org/10.1111/j.1467-7687.2009.00946.x

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition.* Oxford, UK: Blackwell.

Stiller, A., Goodman, N. D., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development, 11,* 176–190. http://dx.doi.org/10.1080/15475411.2014.927328

Strauss, S., Ziv, M., & Stein, A. (2002). Teaching as a natural cognition and its relations to preschoolers' developing theory of mind. *Cognitive Development, 17,* 1473–1487. http://dx.doi.org/10.1016/S0885-2014(02)00128-4

Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology, 28,* 509–529. http://dx.doi.org/10.1146/annurev.anthro.28.1.509

Tomasello, M., & Haberl, K. (2003). Understanding attention: 12- and 18-month-olds know what is new for other persons. *Developmental Psychology, 39,* 906–912. http://dx.doi.org/10.1037/0012-1649.39.5.906

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development, 72,* 655–684. http://dx.doi.org/10.1111/1467-8624.00304

Wilson, D., & Sperber, D. (2004). Relevance theory. In L. Horn & G. Ward (Eds.), *Handbook of pragmatics* (pp. 607–632). Oxford, UK: Blackwell.

Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science, 10,* 288–297. http://dx.doi.org/10.1111/j.1467-7687.2007.00590.x

Ziv, M., & Frye, D. (2004). Children's understanding of teaching: The role of knowledge and belief. *Cognitive Development, 19,* 457–477. http://dx.doi.org/10.1016/j.cogdev.2004.09.002