Children's understanding of how noise disrupts verbal communication

Aaron Chuey[1,*,+], **Rondeline M. Williams**[1,+], **Catherine Qing**[1], **Michael C. Frank**[1], **and Hyowon Gweon**[1]

[1]Stanford University, Psychology, Stanford, 94305, USA
[*]corresponding author: chuey@stanford.edu
[+]these authors contributed equally to this work

ABSTRACT

Having an abstract understanding of communication means being able to reason not just about its success but also its failure: when and why it fails, what happens as a consequence, and how to fix it. Auditory noise in particular can radically alter the fidelity of verbal communication, yet little is known about how language users reason about its role in communication and when such reasoning emerges in development. The current work investigates the development of children's reasoning about auditory noise. As observers of others' communication, we found that a sample of US children as young as age three understood that auditory noise impedes others' hearing, though the ability to reason about its effect on others' knowledge and behaviors may not develop until slightly later. As communicators, a separate sample of four-year-olds tailored how they themselves communicated by using more gestures when their partner couldn't hear. Taken together, these results suggest that even young children may possess an abstract understanding of communication that is integrated with their ability to reason about others' perception and knowledge. Such an understanding enables children to reason about how noise disrupts communication and how to communicate effectively in the presence of noise.

Introduction

Humans are a communicative species: much of what we know and do is shaped by what others tell us. Yet, communication is an imperfect process. Many factors can prevent the successful transmission of information between a communicator and a listener; some, like ambient noise, affect many people in an environment, while others, like hearing and vision loss, target particular individuals. Thus, an abstract understanding of communication involves the ability to reason about what can disrupt a communicative exchange, the specific consequences of such disruptions, and importantly, how disrupted communication can be remedied or even prevented. The current studies examine children's reasoning about auditory noise, a particularly pervasive factor that affects the fidelity of verbal communication (e.g., 1; 2).

By distorting listeners' auditory perception, noise prevents listeners from comprehending speakers' verbal messages (3; 4) and learning from them (see 5). As adults, we understand how auditory noise disrupts communication and how its effects can be remedied or circumvented. For example, an observer might expect people in a noisy environment to be unable to hear or understand speech, and a speaker may need to repeat what they said. As communicators, we can even anticipate the consequences of noise and proactively choose a means of communication that is not disrupted by noise, such as using gestures instead of speech (6; 7; 8; 9). What are the cognitive processes that underlie our ability to anticipate and respond to the consequences of auditory noise, and how does this ability develop?

One possibility is that humans possess an inferential understanding of the process by which noise affects communication and its impact on listeners. Especially in cooperative or pedagogical contexts, speakers produce messages based on their own mental states and their beliefs about what a listener might want to know; these messages then influence listeners' mental states (10; 11; 12). For instance, if a speaker explains how a toy works to a naive listener, the listener would update their own beliefs about the toy accordingly. However, by obstructing a listener's auditory perception, noise can prevent a speaker's message from influencing a listener's beliefs, affecting the information listeners desire and speakers communicate in kind. Another possibility, however, is that we acquire a series of heuristics about how people react to noise by watching others communicate or by being told what to do in noisy environments. For example, given that we often observe speakers repeat themselves in noisy environments, we might expect that speakers always repeat information masked by noise, or that listeners always desire such information. Although these two accounts make similar predictions about how people will typically respond to noise, they differ in the underlying cognitive mechanisms and their developmental origins. The current studies aim to disambiguate these accounts by examining how young children reason about the communicative consequences of auditory noise, both as observers of communicative exchanges and as communicators themselves.

Understanding how children reason about the consequences of noise may also provide insight into how we reason about others' perception more broadly, including people whose perceptual capabilities differ substantially from our own. In particular, sensory disorders, such as hearing and vision loss, can radically alter people's ability to perceive certain communicative acts. There is some evidence that children are sensitive to such impairments; for instance, hearing children of deaf parents utilize more gestures when communicating with their parents compared to people who are able to hear (see 13). Yet, it remains unclear whether such communicative behaviors are driven by an abstract understanding of others' perceptual access or context-specific expectations about the communicative acts preferred by particular recipients (e.g., one's own parents).

Choosing an appropriate means of communication based on what others can perceive involves a series of inferences about other minds. First, we must represent the idea that other people have multiple sensory channels. While prior work has examined children's ability to independently reason about others' visual and auditory perception, little work has explored their ability to integrate inferences across modalities. Second, we must independently monitor the fidelity of each sensory modality and recognize factors that might disrupt them, such as auditory noise. Third, we must to select and generate an effective communicative act given our partner's perceptual access (e.g., pointing to a menu item instead of naming it to someone who cannot hear well). Nonetheless, it remains unclear whether children's communicative behaviors rely on rich inferences about what others can and cannot perceive versus more narrow, learned heuristics.

Research on children's mental state reasoning has focused primarily on their understanding of others' vision. Many experiments have used occlusion, such as blocking view of an object using an opaque screen, to show that children use what someone can and cannot see to reason about their beliefs (e.g., 14; 15; 16; 17; 18; 19). Yet, far fewer experiments have explicitly examined children's understanding of auditory perception. In the same way that occluded objects have been used to study children's visual perspective taking, auditory noise, which occludes people's ability to comprehend speech, can be used to examine children's auditory perspective taking.

While both modalities—visual and auditory perception—are important for learning, auditory perception is particularly important for learning from others' speech or vocalizations more broadly. Even infants understand that auditory signals can be used to transmit information from one agent to another (20) and that speech in particular systematically influences agents' behaviors (21; 22; 23). By around 2-3 years, children acquire an understanding of how people perceive (24; 25; 26) and remember (27) auditory information more broadly. However, it is unclear when children acquire a more comprehensive understanding of how people influence, or fail to influence, each other's minds, as well as how auditory perception mediates this relationship. For example, do children understand that auditory noise not only makes it more difficult to hear, but can also prevent speakers from transmitting information to listeners more broadly?

The current work aims to fill these gaps by studying not just how children reason about the impact of noise and potential remedies as third-party observers, but also how children themselves communicate to circumvent its impact. In Experiment 1, we asked parents of 2- to 10-year-old children to recall whether and how their children communicate differently in noisy settings. In Experiment 2, we examined how 3- to 5-year-old children reason about the effects of auditory noise on listeners and speakers. Finally, in Experiment 3, we investigated 4-year-old children's ability to flexibly communicate with someone whose auditory access differs from their own. Collectively, these studies suggest that by around age four, children can infer the effects of auditory noise on others' minds and actions, and even tailor how they communicate with people whose auditory perception is inhibited.

Results

Experiment 1

To better understand how children communicate in everyday noisy environments, we surveyed parents about their children's (age 2-10 years) communicative behaviors in noisy environments. A majority of parents (72.5%) said their children communicated differently in noisy environments, and reported first observing these behaviors when their children were roughly three years old. Importantly, relatively fewer parents (45.8%) reported explicitly coaching their children how to behave in noisy environments, suggesting these communicative behaviors do not require explicit instruction. When asked to indicate what their children do to communicate in noisy environments, most parents reported speaking louder (79.3%) and repeating themselves (62.2%), followed by waiting until noise abated to continue speaking (35.1%) and using gestures instead of speech (34.2%).

These results provide some preliminary evidence that children alter their communicative behaviors in environments where their partner's auditory access is masked by noise. However, it is unclear whether these responses are guided by simple heuristics that associate behaviors like speaking louder and gesture use with noisy environments, or a more generative, mentalistic understanding of communication. If children do rely on mental-state reasoning to communicate in noisy environments, they should be able to 1) infer how noise affects what listeners know and predict what other speakers will do to remedy these effects, and 2) as communicators themselves, generate alternative means of communication to circumvent these effects. In subsequent experiments, we therefore tested children's ability to reason about noise as observers (Experiment 2) and communicators (Experiment 3).

Experiment 2

Three preregistered experiments examined whether young children understand that noise inhibits the verbal transmission of information and can reason about its impact on listeners' knowledge. Given parental reports suggesting that children begin to communicate differently in noisy environments around age 3, we recruited three online convenience samples of children from the US between ages 3 and 5.

Experiments 2a and 2b had two within-subject conditions: no-knowledge and partial-knowledge. In each condition, the participant learned the names and functions of two novel toys, and then viewed an animated video of a classroom environment where a teacher described those two novel toys to a student (see Figure 1; condition order counterbalanced). Critically, in both conditions, when the teacher began explaining one of the two toys, the classroom became noisy and masked the teacher's explanation, making it impossible to hear the teacher's explanation of the toy (henceforth the 'noise-masked toy'). The videos in each condition featured two students, whose prior knowledge was manipulated across conditions. In the no-knowledge condition, the student claimed to "know nothing about both toys" whereas in the partial-knowledge condition, the student claimed to "know everything about the toy" that was later masked by noise but "nothing about the other toy". Thus, although the students in both conditions received information about two toys, their final epistemic states differed: in the no-knowledge condition, the student was ultimately ignorant about the noise-masked toy but knowledgeable about the other, whereas the student in the partial-knowledge condition was knowledgeable about both toys.

In Experiment 2a (N=72), children were asked: "which toy does the student want to hear about again?", probing their understanding of the effect of noise on the student's learning. In the no-knowledge condition where the student was still ignorant about the noise-masked toy, children preferentially chose the noise-masked toy ($\beta = 1.58, 95\% CrI = [0.57, 3.43]$); in the partial-knowledge condition, children preferentially chose the toy that was not masked by noise ($\beta = -1.999[-4.45, -0.58]$). Across both conditions, the tendency to choose the noise-masked toy increased with age ($\beta = 0.09[0.01, 0.22]$), but there was no interaction of age and condition ($\beta = 0.00[-0.13, 0.16]$).

Experiment 2b (N=72) used the same procedure, except children were asked: "which toy will the teacher talk about again?", probing their understanding of the effect of noise on the teacher's future communicative behaviors. The results were similar to those of Experiment 2a. Children were more likely to predict that the teacher would repeat the explanation about the noise-masked toy in the no-knowledge condition ($\beta = 1.13, 95\% CrI = [0.24, 2.48]$), and about the other toy in the partial knowledge condition ($\beta = 0.37[0.03, 0.89]$). This effect became more pronounced with age ($\beta = 0.11[0.02, 0.25]$), and again there was no interaction of age and condition ($\beta = -0.09[-0.24, 0.04]$).

These results are not easily explained by inflexible heuristics about how people respond to noise in communicative contexts; instead, young children in these experiments used listeners' epistemic states to infer and predict the consequences of noise. When the student possessed no prior knowledge about the noise-masked toy, they inferred both that the student would want to hear about it again, and that the teacher would repeat information about it. However, when the student already possessed prior knowledge about the noise-masked toy, children were less likely to think that the student would want to hear about it again or that the teacher would repeat information about it. Interestingly, children's inferences about the student and teacher were remarkably similar, perhaps suggesting that inferring how speakers remedy the impact of noise on listeners is not necessarily more challenging than inferring how noise impacts listeners in the first place.

In our samples, children's ability to reason about the consequences of both successful and unsuccessful communication appeared to develop with age. While older children overwhelmingly thought the student who didn't have prior knowledge about the noise-masked toy would want to hear about it again, younger children were more uncertain. Interestingly, young children's responses in the partial knowledge condition suggest that they believed the student who already knew about the noise-masked toy would want to hear about the other toy again; this is consistent with the possibility that children expect people to seek information based on their prior knowledge. However, it remains unclear whether they also expect listeners to update their knowledge based on verbal communication, or understand that auditory noise can corrupt this process.

Experiment 2c (N=24) investigated the possibility that our task may have been too cognitively taxing for younger children, since they had to remember what happened during two communicative exchanges as well as the epistemic states of two students. Three-year-olds viewed a single video where a teacher explained two novel toys to a student, but one of the explanations was masked by auditory noise. After each explanation, the video was paused and children were asked if the student could hear what the teacher said; this made the presence of noise more salient and directly assessed children's understanding of the perceptual, rather than epistemic, consequences of noise. Afterwards, we asked children which toy the student wanted to hear about again.

While 3-year-olds were uncertain whether the student could hear the noise-masked utterance ($\beta = -0.193, 95\% CrI = [-2.394, 1.799]$), they reliably responded that the student could hear the unmasked utterance ($\beta = 2.337[0.220, 5.297]$). In line with these judgments, when asked which toy the student wanted to hear about again, 3-year-olds were more likely to choose the noise-masked toy over the other toy($\beta = 2.794[0.352, 5.801]$). This finding provides some evidence that 3-year-olds understand that auditory noise makes it more difficult for people to hear speech and causes them to desire that information, at least when the presence of noise and listeners' auditory access are made sufficiently salient.

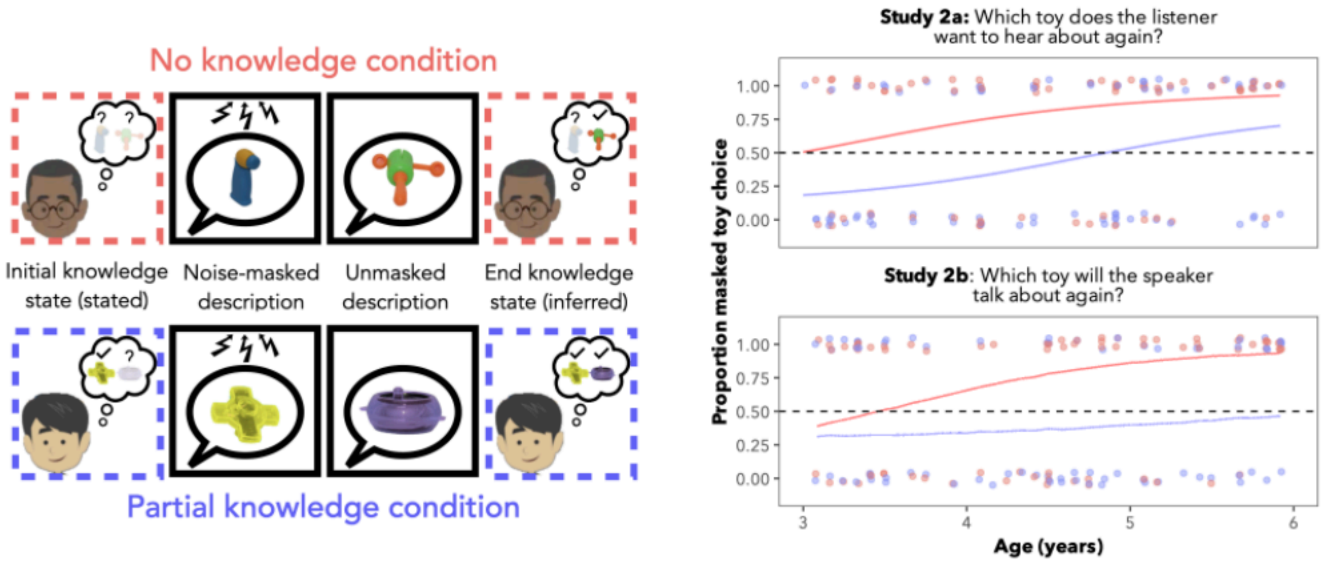

Figure 1. Experiments 2a and 2b procedure and results. Children watched two videos where a teacher describes two toys to a student; during one of the descriptions (noise-masked toy), loud auditory noise masks the teacher's speech. Afterwards, the student expressed a desire to hear about one of the toys again, and children were asked which toy the student would want to hear about again (Experiment 2a) or which toy the teacher would talk about again (Experiment 2b). Crucially, we manipulated the student's prior knowledge: one student (no-knowledge) had no prior knowledge about either toy, while the other student (partial-knowledge) had extensive prior knowledge about the noise-masked toy only. Thus, by the end of the video, the no-knowledge student was still ignorant about the noise-masked toy while the partial-knowledge student was knowledgeable about both toys. Graphs depict Bayesian logistic curves fitted to the data for Experiment 2a (top) and Experiment 2b (bottom), depicting masked-toy choice by participant age and student knowledge.

Experiment 3

Children in Experiment 2 understood that auditory noise prevents listeners from understanding, and ultimately learning from, speakers' utterances. However, an abstract understanding of communication transcends speech alone and should encompass various channels people can use to transmit information to one another. In Experiment 1, for instance, parents reported that their children adapted how they communicate in noisy environments by talking louder, repeating themselves, and even modifying the modality of their communication (e.g., employing gestures instead of speech). While children may simply associate these kinds of behaviors with louder environments, they could be the product of a comprehensive understanding of others' perceptual access.

Compared to waiting or trying to override noise by speaking louder, the ability to switch between speech and gestures is a particularly sophisticated way to circumvent the effects of noise. It requires children to understand that people can acquire information through multiple perceptual channels, such as vision and audition, and tailor their communicative acts accordingly when one channel is compromised. For example, if someone is unable to hear, then a gesture—as opposed to speech—is more effective at transmitting information to them. While the parental reports are anecdotal in nature, the findings from Experiment 2 suggest that children—as young as age 4—may be genuinely capable of such adaptation.

In Experiment 3 (N=48), we experimentally tested whether 4-year-old children can circumvent the effect of noise on communication by using gestures instead of speech. More specifically, given a chance to explain a novel toy to a naïve learner, we predicted that children would be more likely to demonstrate it, and less likely to explain it verbally, when the learner's auditory access was corrupted by noise (see Figure 2). For this experiment, we used noisy headphones as the source of noise instead of ambient noise because this source requires children to track that someone's auditory experience differs substantially from their own, making our paradigm a strong test of children's auditory perspective-taking. This modification also motivated our move to in-person testing.

Our paradigm was as follows. First, we introduced children to a simple toy, and verbally explained that it lights up when five red blocks are stacked on top of a platform; importantly, the blocks were inside a plastic bag next to the toy, and were referred to, but not used, during the explanation. We then introduced children to Gus, a mouse puppet wearing headphones. Crucially, we manipulated whether the headphones were playing loud music (noisy condition) or nothing (noiseless condition) between-subjects. We then asked children whether Gus could hear them and corrected them if they gave the incorrect response. Finally, we told children that Gus wanted to learn how the toy works and asked them to teach Gus however they desired.

In our preregistered main analysis, our key measure was how children taught Gus. Physically demonstrating the toy to Gus (e.g., by stacking the blocks on the toy) or making extensive use of gestures was coded as 'demonstration', while verbally describing the toy to Gus was coded as 'instruction'[1]. Children were significantly more likely to demonstrate in the noisy condition (92%), $\beta = 2.58, 95\% CrI = [1.24, 4.40]$, than in the noiseless condition (54%), $\beta = -2.42[-4.42, -0.81]$. Conversely, children were significantly more likely to verbally instruct in the noiseless condition (62%), $\beta = 1.93[0.70, 3.25]$, relative to the noisy condition (21%), $\beta = -1.40[-2.49, -0.42]$. Thus, children produced more visual demonstrations and less verbal instruction when Gus' auditory perception was inhibited by loud noise. These results demonstrate that young children can tailor how they communicate based on whether their communicative partner can see and hear, even when their own auditory access remains intact and differs from their partner's.

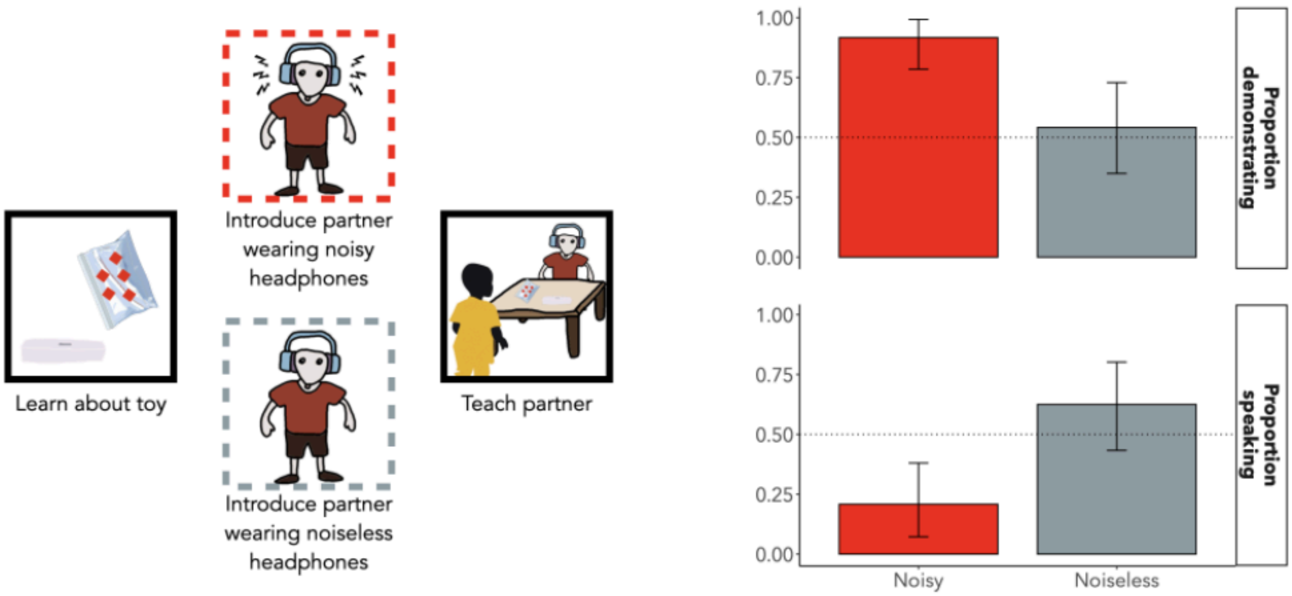

Figure 2. Experiment 3 procedure and results. The experimenter verbally described a novel toy to the child, explaining that stacking five red blocks on a platform causes it to light up. The experimenter then introduced a puppet, Gus, who was wearing headphones that either played loud music or were silent. Finally, the experimenter invited the child to teach Gus how the toy works. We measured whether children 1) physically demonstrated the toy to Gus and 2) verbally instructed Gus. Graphs depict proportion of children who physically demonstrated (top) and verbally instructed (bottom) by condition. Error bars indicate 95% bootstrapped CIs.

Discussion

Successfully reasoning about communication requires not just an understanding of when it succeeds, but also when it fails. Building on parental reports suggesting that children begin to communicate differently in noisy environments in early childhood (Experiment 1), our experiments examined whether young children can predict the effects of auditory noise on communication (Experiment 2) and even adapt their own communicative modalities depending on others' perceptual access (Experiment 3). These studies reveal how children between 3 to 5 years of age develop a mentalistic understanding of auditory noise. In aggregate, children understood that noise prevents listeners from learning from speakers' verbal utterances, causing ignorant listeners to desire the corrupted information (Experiment 2a), and cooperative speakers (i.e., a teacher) to repeat the corrupted information (Experiment 2b). While even 3-year-olds appeared to understand that auditory noise makes it more difficult for people to hear (Experiment 2c), these inferences became increasingly robust with age. As communicators themselves, 4-year-old children readily altered the modality of their communicative acts based on their partner's visual and auditory access. When their partner's auditory access was corrupted by loud music, children used demonstrations instead of verbal instruction to leverage the visual modality (Experiment 3).

Just as physical occlusion has been used in prior work to study children's understanding of others' visual perception, the current work used acoustic noise that occluded speech to study children's understanding of others' auditory perception in communicative contexts. The current results demonstrate that by around 3-4 years, children understand how sound, such

[1]Note that children could perform both teaching behaviors, so demonstration and instruction are not mutually exclusive.

as ambient noise or loud music, can render someone unable to hear and learn from speech. While the loud noise used in Experiment 2 affected children's own ability to hear speech, the headphones in Experiment 3 left children's hearing intact. Thus, young children appear capable of representing auditory perspectives that differ substantially from their own.

More broadly, the current results suggest an early understanding of communication that is integrated with children's ability to reason about others' mental states. To predict the outcomes of noisy exchanges, children in Experiment 2 needed to 1) represent the listeners' initial knowledge, 2) update it based on the speaker's testimony while accounting for the impact of noise, and 3) use the listeners' knowledge to predict what information they would desire (2a) or what information the speaker would repeat (2b). As communicators themselves, children in Experiment 3 needed to 1) represent what their partner could see and hear, 2) account for the selective impact of noise on their partner's auditory (but not visual) perception, and 3) choose a communicative act that would effectively transmit information given their partner's perceptual access.

Overall, successful performance in our studies requires the ability to reason about how others update their beliefs based on information from different perceptual modalities, while accounting for the selective impact of noise on one of these modalities (i.e., audition). Such an ability may be particularly consequential for understanding and communicating with those with sensory disorders. While even toddlers seem to understand the consequence of visual occlusion (e.g., opaque glasses) once they've experienced it themselves (28), children without any sensory disorders might initially assume that everyone can see or hear. Our work suggests that in the domain of auditory perception, by age 4, children are at least capable of understanding that certain individuals may have different auditory access. Future research should investigate how children's communicative adaptations manifest in diverse social contexts with individuals with different sensory capabilities.

Although the current results demonstrate that children can reason about and adapt to auditory noise from a young age, there are also several limitations to consider. First, the current studies used noise that completely prevented someone from hearing other sounds. However, noise often has a more subtle, graded effect on people's hearing, attention, and cognitive resources. As listeners, young children can disambiguate the meaning of degraded speech by relying on their top-down expectations about speakers, suggesting that they understand how noise can interfere with (but not entirely block) communication (29). Nonetheless, it remains unclear whether children—as third-party observers—can infer how noise might alter others' speech comprehension in a similarly nuanced way.

Second, the current studies involved simplified exchanges where a speaker transmitted information to a listener without feedback. Yet, most conversations are more dynamic; for example, listeners can express that they are having difficulty hearing, change their physical position, or ask questions about what they could not hear. In Experiment 2, the students in the videos did not immediately respond to the loud noise, which may have signaled that they could still hear what the teacher said. Future work could examine how children anticipate, interpret, and engage in such strategies.

Third, it is unclear how much children's own experience with noise necessitated or mediated the current results. Acoustic noise exposure varies systematically across environments by socioeconomic status (SES) and population density (30), which can lead to systemic differences in speech processing (31), memory (32; 33), reading ability (34), and word learning (35). In turn, these differences have been shown to drive both language (36; 37) and achievement (38) gaps across childhood and adolescence. Nonetheless, the degree to which early noise exposure affects children's ability to reason about the effects of noise on others remains unknown. Thus, future work should examine the role that experience plays in shaping children's understanding of noise and communication more broadly.

The current studies provide initial evidence about how children think about auditory noise and the conditions under which communication can fail as a whole. Instead of relying on experience-dependent heuristics to predict and adapt to the consequences of auditory noise, children appear to possess a sophisticated understanding of how people acquire information via fallible sensory channels. Such an understanding is crucial for navigating imperfect auditory environments, allowing children to understand when and how communication fails, what happens when it fails, and what they can do to communicate successfully.

Methods

This research was approved by the authors' Institutional Review Board. Parents or legal guardians provided informed consent and children provided verbal assent before participating in each study.

Experiment 1

Participants

Ninety-eight parents of 153 2- to 10-year-old children were recruited using Cloud Research (39) and participated via Qualtrics for $2 payment. An additional 3 parents participated but were excluded for not reporting any children within the requested age range.

Procedure

Caregivers completed a survey about their children's behavior in noisy auditory environments. They first listed the number of children they had between ages 2 and 10 years, and then provided each child's age, gender, race, as well as any language or hearing disorders, delays, or disabilities. In the next set of questions, caregivers were asked: "When it is noisy around your child, have you ever observed them spontaneously do any of the following without being asked when they are trying to communicate with someone?": 1) request listeners to pay closer attention, 2) ask someone if they have any questions about what was communicated, 3) talk louder or yells, 4) repeat themselves, 5) stop talking or wait until noise has stopped or abated, 6) use more gestures (e.g., pointing) or physical demonstrations, or use gestures or physical demonstrations instead of speech. Next, a similar question probed their child's behaviors when someone is trying to communicate with them (e.g., my child has asked someone to talk louder).

The next three questions explored children's acknowledgment of noise ("Have you ever observed your child explicitly acknowledge or refer to auditory noise?"), communicative behaviors in noisy situations ("Have you ever observed your child communicate differently when it is noisy?"), and the extent to which caregivers talked about noise with them ("Have you ever explicitly talked to your child about noise, or how they should act or communicate when it is noisy?"). If they answered "Yes" to any of these questions, children were asked to estimate how old their child was when such events first occurred, and to briefly describe a memorable instance of the behavior if possible. For each story caregivers provided, we coded children's behavior as demonstrating any of the six behaviors described above. Caregivers then completed a demographic questionnaire, including their age, gender, education level, and native language. See OSF repository for a full list of survey questions and answer choices.

Experiment 2a

Our preregistered predictions and analysis plan for Experiment 2a can be found here.

Participants

Seventy-two 3- to 5-year-old children (mean age = 4.56 years; 44.4% Caucasian/White, 29.1% Asian, 2.8% Hispanic/Latino, 1.3% African/Black, 20.8% Biracial/Multiracial, 1.3% Other; 40 female) were recruited through online advertisements. Children had no visual, cognitive, or neurological concerns and heard English at least 75% of the time. An additional three children were ultimately excluded from analysis due to caregiver intervention, experimenter error, or technical difficulties.

Materials

Visual stimuli were designed using Vyond Animation Software in video format. Auditory stimuli were recorded by the first authors and edited in Praat and Audacity. All target speech, including silences and pauses, was equalized to an average sound pressure level of 65dB. The background noise, recorded in a local preschool classroom of 4- and 5-year-old children and adult teachers, was equalized to a default average sound pressure level of 40dB and increased to 85dB during the critical period when the target object label was completely masked by noise. See OSF repository for study stimuli, videos, and PowerPoint presentations.

Procedure

Children completed the experiment with a trained experimenter synchronously via Zoom on either a desktop, laptop, or tablet. Before beginning the experimental session, caregivers were shown an unrelated video and asked to adjust their device's volume to confirm their child could hear the audio clearly. The experimenter then guided children through a slideshow which included both still images and videos.

The experimenter first introduced children to Teacher June, a female adult character, in her classroom. Children then completed two trials (order counterbalanced): the no-knowledge condition and the partial-knowledge condition. In the no-knowledge condition, children saw two novel toys on a table, one on each side of Teacher June; the toy labels were selected from the NOUN Database and matched on complexity (40). The experimenter labeled each object (e.g. "a kern") and described their functions (e.g. "when you squeeze it, it spins around"). To ensure that all children learned the toys' names and functions, the experimenter asked them to repeat this information for both toys. If they were unable to do so, the experimenter reminded them of the correct answers before repeating the question. All children could correctly identify and provide the function of each novel object on the first or second try.

The experimenter then explained that Tim was a new student in Teacher June's class who "had never seen either [of the novel toys] before and knew nothing about them." The experimenter then explained: "Teacher June is going to tell about the toys in her classroom! But the classroom is noisy today and it may be hard for Tim to hear what teacher June is saying sometimes." Light background noise (40dB) played while Teacher June said both the name and function of each toy. However, while she described one of the toys (the noise-masked toy), the volume of the background noise increased substantially (85dB) such that the toy's description was impossible to hear. Then, Tim requested that one of the object labels and functions be repeated. The experimenter then asked children, “which toy do you think Tim wants to hear about again? The [left color] toy or the [right color] toy?” The toy presentation order was counterbalanced across conditions.

The partial-knowledge condition introduced a different pair of novel toys and student, Charles. It was similar to the no-knowledge condition with one key difference: Charles already knew about one of the toys (the noise-masked toy) but was ignorant about the other toy. The experimenter told children that Charles “had the [noise-masked] toy in his old classroom; he played with it a lot and knows everything about it already. But Charles has never seen the [other] toy before and knows nothing about it”. The experimenter repeated this to children to ensure they understood. Critically, noise masked Teacher June’s description of the toy that Charles already knew about, but did not mask the description of the other toy.

Coding and Analysis

Each video and question corresponded to a single trial. For each question, we coded whether children selected the noise-masked or unmasked toy in the partial-knowledge and no-knowledge conditions. Acceptable responses included verbalizing the toy’s name or color, or pointing to the toy on the screen. If it was unclear where the child was pointing, the caregiver could clarify to the experimenter which toy the child was pointing to. Children who did not provide a response for both trials after two verbal prompts from the experimenter were excluded from analysis.

To evaluate whether children monitor someone’s epistemic state to reason about the effect of noise on future behavior, we ran a Bayesian mixed effects model with a random intercept and random slopes using the ‘rstanarm’ package in R. The model centered age to the overall mean (4.56 years) and included an interaction term with the listener’s knowledge state (partial or none). The number of iterations was set to 10,000. From this logistic regression, we extracted the beta coefficients and 95% credible intervals (at 2.5% at the lower bound and 97.5% at the upper bound). For our exploratory analyses, we calculated a posterior predictive distribution based on the previous model using the ‘brms’ R package. We used the default number of chains and draws with 100 iterations. Exact analyses can be found in the OSF repository.

Experiment 2b

Our preregistered predictions and analysis plan for Experiment 2b can be found here.

Participants

Seventy-two 3- to 5-year-old children (mean age = 4.47 years; 43.0% Caucasian/White, 29.1% Asian, 2.8% African/Black, 1.3% Hispanic/Latino, 20.8% Biracial/Multiracial, 2.8% Other; 43 female) were recruited through online advertisements. An additional five children were excluded from analysis due to caregiver intervention, experimenter error, or technical difficulties.

Materials

The materials were identical to those in Experiment 2a.

Procedure

Experiment 2b was similar to Experiment 2a with minor modifications. Children were introduced to Teacher June and completed both the no-knowledge and partial knowledge conditions. In each condition, children were introduced to two novel toys along with their functions. However, the novel object labels were replaced with “blue” and “green” (each toy’s primary color) [2]. Children were then introduced to a new student in Teacher June’s class- Tim or Charles and watched the main video. Unlike in Experiment 2a, the student reported whether he was familiar with one (partial-knowledge) or neither (no-knowledge) of the two toys at the beginning of the video. Teacher June described both toys while noise masked one of the descriptions. The experimenter then explained that “Teacher June is going to tell [Tim/Charles] about one of these toys again,” and asked, “which toy do you think Teacher June is going to tell [Tim/Charles] about again- the [left color] toy or the [right color] toy?” This procedure was repeated for the opposite condition.

Coding and Analysis

The coding and analysis plan were identical to Experiment 2a.

Experiment 2c

Participants

Twenty-four 3-year-old children (mean age = 3.58 years, race = 58.3% Caucasian/White, 20.8% Asian, 4.1% African/Black, 16.7% Biracial/Multiracial; 12 female) were recruited through online advertisements. An additional eight children were excluded from analysis due to not answering or incorrectly answering the inclusion question (n = 4), not responding despite prompting (n = 2), or technical difficulties (n = 2).

[2] We referred to the toys using color for two main reasons. First, color is a salient, stable cue that is easy for children to respond to online and has become standard practice for choice paradigms conducted online (see 41). Second, referring to the color of the toys does not rely on children’s memory (unlike referring to the toys by name) or their ability to correctly identify left and right.

Materials
The materials were identical to those in Experiments 2a and 2b.

Procedure
Experiment 2c was similar to Experiment 2a with several modifications. First, children completed an inclusion check at the beginning of the study to ensure they understood the word “hear”. After viewing a short, unrelated video with sound, children were asked if they could hear the video; if they failed to give an affirmative response, they were asked once more. Then, the experimenter muted the video and asked if children could “hear the video now”. Children who gave an affirmative response were asked once more. If children failed either question after being corrected (e.g., reported hearing the muted video), they were excluded from the final analysis.

Like Experiment 2a, children were first introduced to the two novel toys and their functions, labeled by their primary colors. They were then introduced to Tim, “a new student in Teacher June’s class who wants to learn about these cool toys too.” Importantly, children were not explicitly made aware of Tim’s knowledge about the toys.

Next, the experimenter played a video of Teacher June describing one of the toys to Tim. Then, the video paused and children were asked, “could Tim hear what Teacher June said about the [left color] toy, or could he not hear what Teacher June said?” After providing a response, children viewed the second half of the video where Teacher June described the toy on the right and again asked the previous question.

Afterwards, children were told “Tim wants to hear about one of the toys again,” and the experimenter then asked “which toy do you think Tim wants to hear about again- the [left color] toy or the [right color] toy?”

Coding and Analysis
We ran two Bayesian mixed effects models with random intercepts for each participant using ‘rstanarm’. We evaluated whether noise that masked the toy description would affect whether children reported that the student could or could not hear (model 1) and which toy description should be repeated (model 2). In both models, we ran 10,000 iterations and used the default values for the remaining parameters.

Experiment 3
Our preregistered predictions and analysis plan for Experiment 3 can be found here.

Participants
Forty-eight 4-year-old children [3] (mean age = 4.41 years; 45.8% White, 25.0% Asian American/Pacific Islander, 8.3% Black, 6.2% Hispanic/Latino, 14.6% Mixed-race; 23 female) from a local preschool participated in the study in a separate room. An additional 12 children were excluded for providing incomplete responses (n = 6), technical difficulties (n = 1), or failing an auditory check question (n = 5).

We recruited 4-year-old children for two reasons. First, while parents in Experiment 1 reported that children began displaying adaptive behaviors as early as age 3, these behaviors may not necessarily reflect their understanding of others’ auditory access that differs from their own, which may be more challenging and later emerging. Second, pilot testing suggested that a majority of children under age 4 had difficulty understanding and verbally describing how the toy worked.

Materials
A flat platform with a light in the middle (10 x 6 cm), along with five red blocks in a clear plastic bag, was introduced to participants as a toy. It was described to be out of batteries (in fact, the light was inert) in order to minimize children’s desire to activate the toy for themselves. An anthropomorphic mouse puppet (approximately 64 cm in height) who wore a pair of over-the-ear Bluetooth headphones was introduced as Gus. The headphones either played loud guitar music (“Eruption” by 42) or nothing.

Procedure
Children were seated at a table with the toy and bag of red blocks. The experimenter sat beside the child and said: “I have a really cool toy to show you! It’s out of batteries right now, but I’m going to tell you about it anyways”. The experimenter explained that the toy normally lights up when five red blocks are stacked on top of each other in the center of the toy. Importantly, the experimenter only provided a verbal description and pointed to the bag of red blocks, but did not use the blocks to demonstrate the stacking. The experimenter then asked children to verbally explain how the toy worked, and provided the explanation again if children answered incorrectly.

[3] Note that we had preregistered a sample of 96 children, including 48 4-year-olds and 48 5-year-olds. However, we encountered difficulty in recruiting a sufficient number of 5-year-old children and only report results with the full sample of four-year-olds here. The results of Experiment 3 do not qualitatively differ whether or not 5-year-olds are included; see supplemental materials.

Next, the experimenter retrieved Gus from under the table and introduced him as her friend. Prior work suggests that children readily treat a puppet as an agent, especially when others also treat it as an agent, such as referring to it as their friend (43). In our study, using a puppet instead of a human confederate helped constrain children's communicative behaviors (i.e., verbally describing or demonstrating the toy, rather than asking the puppet to turn the volume down or take off his headphones). Children were assigned to one of two conditions. In the noisy condition, Gus' headphones were playing loud music that was audible from several feet away but did not significantly disrupt the participant's own auditory access. In the noiseless condition, Gus' headphones were silent. The experimenter first asked children what Gus was wearing on their head. If children said "headphones" or something similar (e.g., earphones), the experimenter agreed; if children were not sure or gave an incorrect answer, the experimenter explained that Gus was wearing headphones that can play music.

The experimenter then asked children if they could hear anything coming from Gus' headphones. In the noisy condition, the experimenter remarked that Gus was listening to loud music. In the noiseless condition, the experimenter said that the headphones were not playing anything and that Gus liked to wear them because they "look cool". The experimenter asked children if they thought Gus could hear them talking. If they answered correctly (i.e., could not hear in the noisy condition and could hear in the noiseless condition), the experimenter agreed. If they answered incorrectly, the experimenter corrected them (i.e., "Actually, I think Gus can/can't hear us right now") and repeated the question. If they again answered incorrectly, they were excluded from subsequent analyses.

The experimenter then informed children: "Gus really wants to learn how the toy works; can you teach Gus how the toy works? You can teach Gus however you like, just let me know when you are finished". To minimize the possibility of children directing their explanation to the experimenter instead of Gus, the experimenter then moved to a chair approximately 1.2 meters from the table and appeared to be reading something on a clipboard. Once children indicated they completed their explanation, the experimenter praised their teaching efforts and concluded the session.

Coding and Analysis

Children's communicative behaviors were divided into two primary categories- demonstration and instruction. Children's actions were considered demonstrations if they took the blocks out of the bag and stacked them onto the toy or approximated this behavior. Touching or pointing at the bag alone did not count as a demonstration. Children's actions were considered instructions when they verbalized how to use the toy. For instruction, children must have mentioned stacking the blocks in their instruction, but it was unnecessary that they described each step. Children who demonstrated and instructed were coded as having completed both. Children who directed their explanation to the experimenter or did not produce either behavior were coded as incomplete responses, which were excluded from subsequent analyses.

Given there were two possible distinct but overlapping communicative behaviors, we ran a Bayesian multivariate logistic regression in 'brms' to calculate whether children adjusted both their demonstration and instruction behavior given Gus' auditory access. We used a Bernoulli distribution but retained default values for all other parameters. The beta coefficients and 95% credible intervals (at 2.5% at the lower bound and 97.5% at the upper bound) estimated these effects.

References

1. Buxton, R. T. *et al.* Noise pollution is pervasive in us protected areas. *Science* **356**, 531–533 (2017).

2. Summers, W. V., Pisoni, D. B., Bernacki, R. H., Pedlow, R. I. & Stokes, M. A. Effects of noise on speech production: Acoustic and perceptual analyses. *The J. Acoust. Soc. Am.* **84**, 917–928 (1988).

3. Calandruccio, L., Porter, H. L., Leibold, L. J. & Buss, E. The clear-speech benefit for school-age children: Speech-in-noise and speech-in-speech recognition. *J. Speech, Lang. Hear. Res.* **63**, 4265–4276 (2020).

4. Puglisi, G. E., Warzybok, A., Astolfi, A. & Kollmeier, B. Effect of reverberation and noise type on speech intelligibility in real complex acoustic scenarios. *Build. Environ.* **204**, 108137 (2021).

5. Mealings, K. The effect of classroom acoustic conditions on literacy outcomes for children in primary school: A review. *Build. Acoust.* **29**, 135–156 (2022).

6. Clark, H. Grounding in communication. *Perspectives on socially shared cognition/American Psychol. Assoc.* (1991).

7. Clark, H. H. Depicting as a method of communication. *Psychol. review* **123**, 324 (2016).

8. Sekine, K. & Özyürek, A. Children benefit from gestures to understand degraded speech but to a lesser extent than adults. *Front. Psychol.* **14**, 1305562 (2024).

9. Trujillo, J., Özyürek, A., Holler, J. & Drijvers, L. Speakers exhibit a multimodal lombard effect in noise. *Sci. reports* **11**, 16721 (2021).

10. Goodman, N. D. & Frank, M. C. Pragmatic language interpretation as probabilistic inference. *Trends cognitive sciences* **20**, 818–829 (2016).

11. Gweon, H. Inferential social learning: Cognitive foundations of human social learning and teaching. *Trends cognitive sciences* **25**, 896–910 (2021).

12. Shafto, P., Goodman, N. D. & Griffiths, T. L. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cogn. psychology* **71**, 55–89 (2014).

13. Singleton, J. L. & Tittle, M. D. Deaf parents and their hearing children. *J. Deaf. studies Deaf. education* **5**, 221–236 (2000).

14. Flavell, J. H., Shipstead, S. G. & Croft, K. Young children's knowledge about visual perception: Hiding objects from others. *Child Dev.* 1208–1211 (1978).

15. Lempers, J. D., Flavell, E. R. & Flavell, J. H. The development in very young children of tacit knowledge concerning visual perception. *Genet. Psychol. Monogr.* **95**, 3–53 (1977).

16. Kim, E. Y. & Song, H.-j. Six-month-olds actively predict others' goal-directed actions. *Cogn. Dev.* **33**, 1–13 (2015).

17. Luo, Y. & Baillargeon, R. Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition* **105**, 489–512 (2007).

18. Luo, Y. & Johnson, S. C. Recognizing the role of perception in action at 6 months. *Dev. science* **12**, 142–149 (2009).

19. Surian, L., Caldi, S. & Sperber, D. Attribution of beliefs by 13-month-old infants. *Psychol. science* **18**, 580–586 (2007).

20. Tauzin, T. & Gergely, G. Communicative mind-reading in preverbal infants. *Sci. reports* **8**, 1–9 (2018).

21. Martin, A., Onishi, K. H. & Vouloumanos, A. Understanding the abstract role of speech in communication at 12 months. *Cognition* **123**, 50–60 (2012).

22. Vouloumanos, A., Martin, A. & Onishi, K. H. Do 6-month-olds understand that speech can communicate? *Dev. Sci.* **17**, 872–879 (2014).

23. Vouloumanos, A., Onishi, K. H. & Pogue, A. Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proc. Natl. Acad. Sci.* **109**, 12933–12937 (2012).

24. Melis, A. P., Call, J. & Tomasello, M. 36-month-olds conceal visual and auditory information from others. *Dev. Sci.* **13**, 479–489 (2010).

25. Williamson, R. A., Brooks, R. & Meltzoff, A. N. The sound of social cognition: Toddlers' understanding of how sound influences others. *J. Cogn. Dev.* **16**, 252–260 (2015).

26. Yaniv, I. & Shatz, M. Children's understanding of perceptibility. (1988).

27. Moll, H., Carpenter, M. & Tomasello, M. Two-and 3-year-olds know what others have and have not heard. *J. Cogn. Dev.* **15**, 12–21 (2014).

28. Meltzoff, A. N. & Brooks, R. Self-experience as a mechanism for learning about others: a training study in social cognition. *Dev. psychology* **44**, 1257 (2008).

29. Yurovsky, D., Case, S. & Frank, M. C. Preschoolers flexibly adapt to linguistic input in a noisy channel. *Psychol. science* **28**, 132–140 (2017).

30. Casey, J. A. *et al.* Race/ethnicity, socioeconomic status, residential segregation, and spatial variation in noise exposure in the contiguous united states. *Environ. health perspectives* **125**, 077017 (2017).

31. Fernald, A., Marchman, V. A. & Weisleder, A. Ses differences in language processing skill and vocabulary are evident at 18 months. *Dev. science* **16**, 234–248 (2013).

32. St. John, A. M. & Tarullo, A. R. Neighbourhood chaos moderates the association of socioeconomic status and child executive functioning. *Infant Child Dev.* **29**, e2153 (2020).

33. Waters, N. E., Ahmed, S. F., Tang, S., Morrison, F. J. & Davis-Kean, P. E. Pathways from socioeconomic status to early academic achievement: The role of specific executive functions. *Early Child. Res. Q.* **54**, 321–331 (2021).

34. Cohen, S., Glass, D. C. & Singer, J. E. Apartment noise, auditory discrimination, and reading ability in children. *J. experimental social psychology* **9**, 407–422 (1973).

35. McMillan, B. T. & Saffran, J. R. Learning in complex environments: The effects of background speech on early word learning. *Child Dev.* **87**, 1841–1855 (2016).

36. Dailey, S. & Bergelson, E. Language input to infants of different socioeconomic statuses: A quantitative meta-analysis. *Dev. science* **25**, e13192 (2022).

37. Hart, B., Risley, T. R. & Kirby, J. R. Meaningful differences in the everyday experience of young american children. *Can. J. Educ.* **22**, 323 (1997).

38. Hanushek, E. A., Light, J. D., Peterson, P. E., Talpey, L. M. & Woessmann, L. Long-run trends in the us ses—achievement gap. *Educ. Finance Policy* **17**, 608–640 (2022).

39. Litman, L., Robinson, J. & Abberbock, T. Turkprime. com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behav. research methods* **49**, 433–442 (2017).

40. Horst, J. S. & Hout, M. C. The novel object and unusual name (noun) database: A collection of novel images for use in experimental research. *Behav. research methods* **48**, 1393–1409 (2016).

41. Chuey, A. *et al.* Moderated online data-collection for developmental research: Methods and replications. *Front. psychology* **12**, 734398 (2021).

42. Van Halen. Eruption. online streaming (1978).

43. Asaba, M., Li, X., Yow, W. Q. & Gweon, H. Children selectively demonstrate their competence to a puppet when others depict it as an agent. *Cogn. Dev.* **62**, 101186 (2022).