

Is it OK if Alexa doesn't know?: Children and adults' beliefs about smart speaker virtuous ignorance

Lauren N. Girouard-Hallam (girouala@umich.edu)

Department of Psychology, University of Michigan
East Hall, 1004, 530 Church St, Ann Arbor, MI 48109, USA

Allison J. Williams-Gant (awillia7@bu.edu)

Wheelock College of Education and Human Development, Boston University
621 Commonwealth Ave, #319, Boston, MA 02215, USA

Judith H. Danovitch (j.danovitch@louisville.edu)

Department of Psychological and Brain Sciences, University of Louisville
355 Life Sciences, Louisville, KY 40208, USA

Kathleen H. Corriveau (kcorriv@bu.edu)

Wheelock College of Education and Human Development, Boston University
2 Silber Way, Room 23, Boston, MA 02215, USA

Abstract

This study examines 124 5- to 10-year-old children's and their parent's beliefs about smart speakers that provide ignorant or exact responses to knowable and unknowable questions. More specifically we explored whether participants value virtuous ignorance (i.e., the admission of not knowing the answer to a question because the information is unknowable) when considering explanations provided by smart speakers. With age, children and adults increasingly indicate that smart speakers that provide virtuously ignorant responses are more credible than smart speakers that provide exact responses to unknowable questions. Interestingly, children, but not adults, indicated that speakers with virtuously ignorant responses are better for unknowable future event related questions than unknowable number related questions. Children's evaluations of technology's virtuous ignorance changed with age and question type, perhaps reflecting changes in children's expectations about the kinds of responses devices can reasonably provide. Implications for children's and adult's learning and understanding from smart speakers are discussed.

Keywords: children; cognitive development; virtuous ignorance; smart speaker; internet; technology

Introduction

Humans are social learners and gain information from interacting with other people. One type of informant children and adults rely on for new information is smart speakers (e.g., Alexa, Siri, and Google). Smart speakers are becoming increasingly common in households and both parents and children report interacting with these speakers (Rideout & Robb, 2020). Smart speakers are programmed to search for answers to questions using internet sources but are limited in their scope and ability to answer questions. Smart speakers sometimes respond with "I don't know" or "I'm not sure" when they are unable to find a response to the question. Young children can become frustrated when smart speakers give these sorts of responses (Oranç & Ruggeri, 2021; Yarosh

et al., 2018) despite the reality of technological limitations when answering questions. Children trust smart speakers for general knowledge information and children may perceive these devices to have more access to information than humans do (Girouard-Hallam, Streble, & Danovitch, 2022). It is therefore important to understand whether children and adults believe that smart speakers can know things where humans would need to provide an ignorant response. This study therefore examines how children and adults evaluate the answers that smart speakers provide when the questions push the boundaries of reasonable knowledge.

Children's Evaluation of Virtuous Ignorance in Human Informants

Kominsky, Langthorne, & Keil (2016) introduced the term "virtuous ignorance" to describe the admission of not knowing the answer to a question because the information is impossible or implausible to obtain. Kominsky and colleagues were interested in examining children's and adults' ability to use virtuous ignorance as a characteristic of expertise. This research stemmed from a history of research examining the various epistemic characteristics children use when evaluating another person's testimony. Children as young as 4-years-old prefer to learn from informants who demonstrate epistemic cues (e.g., accuracy and expertise; Tong, Wang, & Danovitch, 2020). Importantly, 3-year-olds show an understanding of expertise when seeking out information (Aguiar, Stoess & Taylor, 2012; Lutz & Keil, 2002) and recognize that certain experts are especially knowledgeable about specific domains of expertise (Koenig & Jaswal, 2011; Kushnir, Vredenburg, & Schneider, 2013; Nguyen, 2012; Sobel & Corriveau, 2010).

However, informants can also display other characteristics that people use to gauge the accuracy of their responses, such

as confidence (Bridgers, et al., 2016). Previous research suggests that although adults can correctly evaluate an informant's confidence and accuracy to determine if they are a reliable source, children as young as 5- and 6-years-old do not have this ability (Tenney et al., 2011). In this study, children and adults were provided conflicting information from confident and cautious informants. Initially, both children and adults trusted the confident informant more than the cautious one. However, once participants were provided with additional information that the confident informant was previously inaccurate, adults, but not children, then trusted the cautious informant more than the confident one. This finding suggests that confidence is a strong cue of credibility for children, even if it is unfounded.

Additional research has suggested that cues of uncertainty and ignorance (e.g., saying "I don't know") provoke distrust in children (Sabbagh & Baldwin, 2001; Mills et al., 2011). This is especially important for children's information seeking, such that children will avoid directing questions to an ignorant informant when a more knowledgeable informant is available (Corriveau, Meints, & Harris, 2009; Mills et al., 2011). In these studies, children often treat an ignorant informant as similar to an incorrect informant, presumably because the information is available, but the informant is not aware of it. However, they are more willing to forgive an ignorant response in cases when the answer to a question is impossible to obtain.

To expand on these findings, Kominsky et al. (2016) measured children's and adults' ability to use virtuous ignorance as a characteristic of expertise in human informants. In the first two experiments, children and adults heard pairs of experts provide confident exact responses or ignorant responses to questions concerning knowable (e.g., If you count the number of windows on the White House, how many will you get?) and unknowable information (e.g., If you count all the leaves on all trees in the entire world, how many will you get?). In Experiment 1, all unknowable questions involved numerical responses and, in Experiment 2, all unknowable questions involved predictions about the future. Kominsky et al. (2016) demonstrated in Experiment 1 and 2 that it is not until around 9-years-old that children began to recognize that human experts who demonstrate virtuous ignorance are better experts than experts who provide confident and exact answers to unanswerable questions. Importantly, human informants are not the only sources children seek out information from and it is important to understand if children make similar judgments for non-human informants.

Children's Trust in Internet-Based Devices

Children frequently rely on the internet and internet-based devices to obtain answers to questions that they have about the world around them (e.g., Girouard-Hallam & Danovitch,

2022; Lovato, Piper, & Wartella, 2019; Wang, Tong, & Danovitch, , 2019). Children use epistemic cues like prior accuracy to decide when to trust technological informants like computers (Danovitch & Alzahabi, 2013). Judgments pertaining the accuracy of the informant extend from computers to the internet (Wang et al., 2019) and humanoid robots (e.g., Breazeal et al., 2016) as well. Just like with human sources, there is evidence to suggest that children monitor internet-based sources for their accuracy (Tong et al., 2022). Previous research examining children's trust and learning from smart speakers has established that children are willing to trust smart speakers (Girouard-Hallam & Danovitch, 2022), yet it has not yet extended to include instances in which smart devices provide incorrect answers or are unable to answer the question.

There are obvious limits to the ability of smart speakers to provide answers to questions, and "I don't know" is a plausible response from technological tools like smart speakers. A practical use case of this is Amazon's Alexa smart speaker line, which has beta-tested a feature where if the smart speaker does not "know" the answer to the question in real time, it will send a text to the owner of the speaker later with an answer should it find one (Kinsella, 2018). Yet child-computer interaction research suggests that young children find it particularly frustrating when smart speakers are unable to answer their questions and will attempt to rephrase their questions to prompt a response from the device (Yarosh et al., 2018). This frustration also appears to be related to a child's age, such that older children are more willing to change topic or ask another question when a smart speaker does not know the answer (Oranç & Ruggeri, 2021). Given the frustration that young children may experience, and the reality that they are likely to receive these responses from smart speakers, it is important to evaluate whether children recognize that "I don't know" is a more reasonable response from these devices than an impossible answer.

Children's Understanding of Smart Speaker

Ontology

Children's understanding of smart speaker ontology is a potential mechanism by which children calibrate their trust in these kinds of digital informants. Previous research has shown that children view devices like smart speakers as having both artifact-like and human-like qualities (Festerling & Siraj, 2020; Flanagan et al., 2023; Girouard-Hallam et al., 2021); thus, children may view smart speakers as having a unique ontological status. This evidence supports the New Ontological Category (NOC) hypothesis (e.g., Kahn et al., 2012), which states that children evaluate technology as having both human-like and object-like traits. The NOC hypothesis has also been supported by evidence that children form parasocial relationships with smart speakers (Hoffman Owen, & Calvert, 2021) and that they ascribe human-like and artifact-like properties to technological devices

simultaneously (Xu & Warschauer, 2020). Beyond qualities that artifacts and people might share, some children also ascribe near omniscience to some technological sources, such as search engines (Girouard-Hallam & Danovitch, 2024), further suggesting that children treat technological informants differently from their human or machine-like counterparts.

Given the evidence that children may in fact view technological devices as belonging to their own unique ontological category, questions then arise about children's understanding of the capacities and limitations of smart speakers. It may be that children ascribe human-like fallibility to smart speakers and understand that sometimes they cannot reasonably provide answers to certain kinds of questions. However, it may also be that children view technological devices like smart speakers as possessing special abilities or access to information that people do not have, leading even older children to believe that smart speakers should be able to provide exact answers or predictions about events in the distant future reliably even in cases where people could not. Therefore, this study tests children's beliefs about the validity of virtuously ignorant smart speakers (e.g., smart speakers that acknowledge that they "do not know" the answers) when answering questions with unknowable answers.

Current Study

The current study uses similar methods as Kominsky et al.'s (2016) study of virtuous ignorance in humans, but instead of two human informants providing information, the informants are two smart speakers. In addition, the age range was expanded from 5- to 10-year-olds to 5- to 12-year-olds to include children who have a more complex theoretical understanding of the internet and related devices (Yan 2005, 2006, 2009).

We hypothesize that, as participant age increases, children will prefer the smart speaker that provides admissions of ignorance in response to questions about unknowable information over a smart speaker that provides precise answers to those questions. However, for knowable questions, regardless of age, children will prefer the smart speaker that provides informative answers over one that says "I don't know." We predict that adults will act similarly to older children, preferring the smart speaker that provides admissions of ignorance in response to questions about unknowable information.

Method

Participants

Participants included 124 5- to 12-year-old children ($M_{\text{age}} = 9.18$, $SD = 2.33$, 65 girls, 59 boys) and 115 of their parents ($M_{\text{age}} = 40.63$, $SD = 4.87$, 97 mothers, 14 fathers, 4 unspecified parents). Two additional children and their caregivers were excluded (one because a sibling interfered

with the interview and another because the child was unable to complete the interview). Fifty-three percent of children were identified by their parents as White, 20% were Asian, 15% were Mixed Race/Ethnicity, 4% were Black/African American, 2% were Hispanic/Latinx, and 6% did not provide an answer. The majority of parents (56%) reported having a graduate or professional degree and around a third of families (33%) reported an average household income of more than \$150,000 a year. Therefore, this sample can be best described as middle to upper class.

To determine our sample size, a one-sample t-test proxy for group size (see Murayama et al., 2022) was used, which determined that a sample of 100 participants would be sufficient for a mixed-effects analysis. However, the best practices for determining sample size with binary dependent variable mixed-effects models are still developing. Given that Murayama et al.'s (2022) suggested power analysis method is best for mixed-effects models with continuous dependent variables, we planned to add an additional 20 participants to our planned sample size to ensure a robust sample for determining the effect size within our binary mixed-effects models.

Procedure

Study data were collected between 2023 and 2024. This study was conducted with approval from the [omitted for blind review] Institutional Review Board. This study was pre-registered on AsPredicted.org. Data and materials are available upon request. Around half of the data was collected using Children Helping Science for participant management and data collection. In addition to the automated recruitment of families in the Children Helping Science database, the other half of the participants were recruited through lab databases from [omitted for blind review] and participated over Zoom.

Smart Speaker Introduction

Participants were first given a description of smart speakers: "Smart speakers are a device that you can use to get answers to questions. When you have a question, you can ask a smart speaker your question, and it will search the internet for the answer. Then, it can say the answer out loud. Smart speakers come in all different shapes, sizes, and colors. These are examples of smart speakers."

Images of speakers varying in shape and color were shown on the screen. Participants were told that the experimenter had a list of questions that they wanted to find the answers to, so they asked different smart speakers their questions. Participants would see the questions asked and the pairs of answers the smart speakers gave. Participants were then instructed: "For each pair of smart speakers, please pick the smart speaker that, in your opinion, is actually the better smart speaker that provides answers about the topic."

Preference Trials

Using a similar method from Kominsky et al. (2016), participants then completed 12 trials including four question types (see Table 1): four questions about numbers that cannot be exactly expressed (i.e., Unknowable Number), four questions about the future with answers that cannot be predicted (i.e., Unknowable Future), two questions about numbers that can be exactly expressed (i.e., Knowable Number), and two questions about current events with answers that can be reasonably predicted or estimated (i.e., Knowable Present). Unknowable Number, Unknowable Future, and Knowable Number were modified based on questions used in Kominsky et al. (2016). Knowable Present questions were created specifically for the current study and were comprised of items about current events where the answer could change, but were relatively predictable (e.g., the most popular pet this year or the tallest building in the world this year). Questions were presented in one of two fixed randomized orders.

For each trial, participants saw the question printed at the top of the screen and two smart speakers in the middle of the screen. Each trial contained two unique speakers varying in color and shape that did not repeat across any trials, resulting in 12 different pairs of speakers. The experimenter read the question aloud (e.g., “The question is: If you count all the leaves on all the trees in the entire world, how many will you get?”) and then each smart speaker provided one either an exact response (e.g., “There are exactly 809,343,573,353,235 leaves on all trees in the world”) or a “virtuously ignorant” response (i.e., “I do not know because it is not possible to answer that question precisely”). The virtuously ignorant response was the same for all trials. Responses by the smart speaker were presented through audio and text such that an automated female voice provided the response while a speech bubble appeared above the image of the speaker on the screen. Participants were then asked to pick “which one do you think is the better smart speaker?”

Familiarity with Smart Speakers (Parent Only)

Parent participants were then asked if they use smart speakers and, if so, which ones. They were then asked if their child uses smart speakers and, if so, which ones.

Results

Preliminary analyses did not reveal significant effects of order, so this variable was excluded from further analyses.

To examine children’s beliefs about the virtuous ignorance of smart speakers, we developed a generalized linear mixed model (GLMM) using the `glmer` function with the `bobyqa` optimizer in the `lme4` package (Bates et al., 2015) in R version 4.2 (R Core team, 2022). The fixed effects in the model were question category (Unknowable Future and Knowable Present, Unknowable Number and Knowable Number) and child age (mean-centered) was included as a continuous predictor. The model also included 2-way interactions between smart speaker response type and age,

and random intercepts for participant and item. A logit link function was used because the dependent variable was the binary selection of the Ignorant response (coded as 1) or the Exact response (coded as 0).

Table 1: Example of Questions and Exact Answers for each Category of Question

| Category | Example Question | Example Exact Answer |
|--------------------|--|--|
| Knowable Present | What is the most popular household pet this year? | “The most popular household pet this year is definitely a dog.” |
| Knowable Numbers | If you count the number of keys that are normally on a piano, how many will you get? | “There are exactly 88 keys on a piano.” |
| Unknowable Future | What will be the most popular boys’ name three years from now? | “The most popular boys’ name three years from now will definitely be Blaise.” |
| Unknowable Numbers | If you count the number of blades of grass that sprouted in New York state last year, how many will you get? | “In the last year, exactly 537,454,265,729, 986,534 blades of grass sprouted in New York state.” |

There was a significant main effect of age ($p < .001$, $\beta = 1.45$, 95% CI = [1.04, 1.86]), such that younger children selected the smart speaker that gave the exact response more frequently than older children did. There was also a significant simple main effect of category ($ps < .021$ across all reference categories) such that children more frequently selected the smart speaker that gave an exact response most in cases involving knowable numbers and more in cases involving knowable present than in cases of *unknowable* numbers or future events.

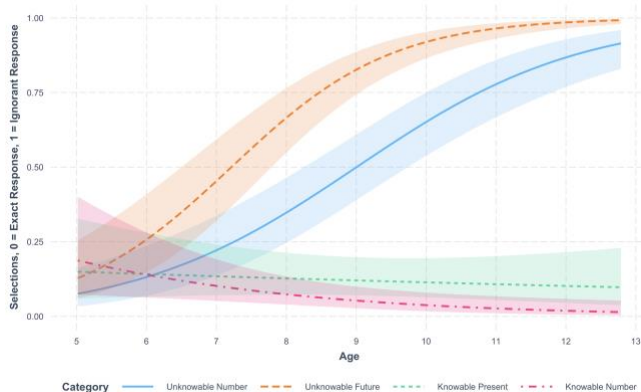
These simple main effects were subsumed by a significant two-way interaction between Age and Category. As children’s age increased, they more frequently chose the ignorant smart speaker for the unknowable number and unknowable future categories, but not for the knowable present or knowable number categories (see Table 2). There was a significant difference between responses to unknowable number and unknowable future questions by age ($p = .009$, see Table 2) such that children chose the ignorant response more frequently from a younger age for unknowable future items than for unknowable number items. Notably, by age 8, children began selecting the ignorant informant the majority of the time for unknowable future items, but not

until age 10 did they consistently select it for unknowable number items. Children treated knowable present items as more difficult to provide an exact answer for than knowable numbers, though children selected the ignorant response less than 25% of the time in both categories (see Figure 1).

Table 2: Two-Way Interaction Results for Age by Condition Across Reference Categories

| Significant Effect | β | CIs | p |
|--------------------|---------|--------------|-------------|
| Unknowable Number | | | |
| Unknowable Future | 0.57 | 0.14, 1.00 | $p = 0.009$ |
| Knowable Present | -1.60 | -2.09, -1.10 | $p < .001$ |
| Knowable Number | -2.28 | -2.89, -1.66 | $p < .001$ |
| Unknowable Future | | | |
| Knowable Present | -2.17 | -2.72, -1.62 | $p < .001$ |
| Knowable Number | -2.84 | -3.50, -2.19 | $p < .001$ |
| Knowable Present | | | |
| Knowable Number | -0.68 | -1.32, -0.03 | $p = .039$ |

Figure 1: Children’s Endorsement of the Ignorant Response by Age and Question Condition

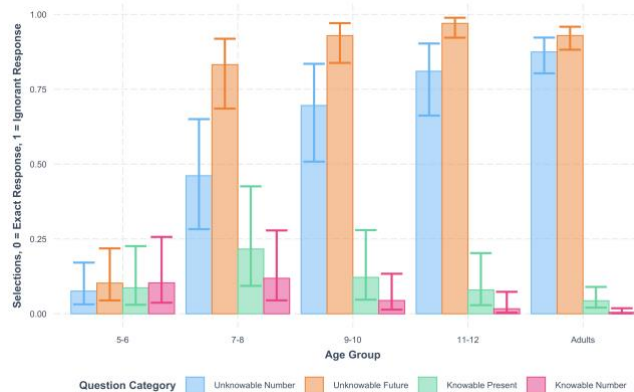


To compare adult’s perceptions of the virtuous ignorance of smart speakers to children’s perceptions, we ran a second GLMM. Instead of treating age continuously, age was treated as a categorical fixed effect where participants were binned into age groups (5-6, 7-8, 9-10, 11-12) and adults were treated as a single age group. Otherwise, the setup of the analysis remained the same.

There was a significant simple main effect of age group (all $ps < .001$) subsumed by a significant two-way interaction between age group and category. As participant age group increased, they more frequently chose the ignorant smart speaker for the unknowable number and unknowable future categories, but not for the knowable present or knowable number categories (all $ps < .001$). With age, children’s responses became more like adult responses; however, only adults did not significantly differentiate between unknowable

future and unknowable number items when deciding which smart speaker had the better answer ($\beta = 0.30, p = 0.503, 95\% \text{ CI} = [-0.59, 1.20]$).

Figure 2: Children and Adult’s Endorsement of the Ignorant Response by Age and Question Condition



Familiarity with Smart Speakers

Of the 115 parents surveyed, 53% said that they used smart speakers themselves and 59% said that their children used smart speakers. The majority of parents who said that they had smart speakers in their home (67%) stated that they had at least one Amazon speaker (Alexa or Echo) at home. Less frequent responses included Google Home and Siri. One parent could not remember the kind of smart speaker they had.

In addition, a third GLMM was run that included the binary independent variable of whether the children were familiar with smart speakers. This analysis revealed a significant three-way interaction between Familiarity, Category, and Age, specifically related to the Unknowable Number ($\beta = 1.17, p = 0.022, 95\% \text{ CI} = [-1.04, -0.08]$) and Unknowable Future ($\beta = -1.83, p = .003, 95\% \text{ CI} = [-1.45, -0.29]$) items. Children who were familiar with voice assistants were *more likely* to embrace exact responses for questions with unknowable answers, particularly at older ages.

Discussion

The current study explored children’s and adults’ beliefs about smart speakers who provide ignorant or exact responses to knowable and unknowable questions. Similarly to Kominsky et al.’s (2016) findings, it was not until children were 8- to 10-years-old that they recognized that an ignorant smart speaker gave a better answer than one that gave a confident but impossible (exact) answer to unknowable questions (i.e., a smart speaker that was virtuously ignorant).

Although the overall age at which this shift occurred was similar to Kominsky et al. (2016)’s findings, children responded differently in their considerations of virtuous ignorance with smart speakers than they did in Kominsky et al.’s studies with human informants. In Kominsky et al., the

developmental shift to acceptance of virtuous ignorance occurred *later* for future prediction items than in the current study. In contrast, children in Kominsky et al.'s studies accepted "I don't know" responses *earlier* when the items were about numerical estimates. Children view technology as limited in its ability to predict future information (Girouard-Hallam & Danovitch, 2025) but they believe that it has near-omniscience for providing answers to general knowledge questions, which could include exact numerical estimates (Girouard-Hallam & Danovitch, 2024; 2025). One possibility is that because children evaluate technology as having both human-like and object-like traits (e.g., Kahn et al., 2012), they assess technology as limited and fallible in forecasting the future, just as they would expect from a person, but extremely capable of computing complex algorithms, just as they would expect from a computer. Future research should examine these results in the context of the NOC hypothesis by examining whether children believe that a smart speaker's object-like traits (e.g., inexhaustibility, algorithmic ability) contribute to its ability to exactly answer unknowable number questions.

Moreover, just as in Kominsky et al. (2016), children's acceptance of ignorant responses generally increased with age, and adults in the current study did not treat unknown future predictions and exact number estimates differently, suggesting that with development, people become better able to conceptualize virtuous ignorance in smart speakers just as they do with humans. According to the "computers as social actors" paradigm (Nass & Moon, 2000), people will sometimes interpret information provided by a computer as if they were learning from a social human partner. It may be that, as children develop and understand more about the nature of information obtained from technology (i.e., that *people* still have to put information online; Girouard-Hallam & Danovitch, 2024), they increasingly treat smart speakers as social actors in terms of limitations in answering questions about unknowable phenomena. The current findings extend evidence for children's extension of virtuous ignorance to technological informants, although they suggest a somewhat different developmental trajectory for this understanding.

One explanation for the differences observed between children and adults could be explained by the dual processing theory of selective trust (Hermes, Behne, & Rakoczy, 2018; Kahneman & Anderson, 2003). Specifically, younger children in the current study may be primarily relying on System 1 thinking and quickly make the decision that an exact answer is better than no answer at all. Older children and adults may be utilizing System 2 thinking and evaluating each response provided by the speakers carefully and considering whether the response is knowable or not.

Similarly, Kominsky et al. (2016) suggests that the effect of age on children's appreciation for virtuously ignorant responses is linked to executive processing and, specifically,

integrative capacity and inhibitory control. Integrative capacity improves throughout middle childhood and into adolescence (Dempster, 1981). In the current study, young children might have struggled to integrate that the exact responses given by the smart speaker signaled confidence but were also impossible to know. Given that young children have a bias to believe and trust verbal testimony (Jaswal et al., 2010), the younger children may also believe any response that provides additional information is better than an ignorant response. However, older children with better inhibitory control might be able to overcome this bias and evaluate the virtues of an ignorant response depending on the question being asked.

Familiarity with smart speakers also appears to play a role in children's judgments about smart speaker virtuous ignorance. The children who were more familiar with smart speakers may be more likely to have used smart speakers as calculators or homework aids, bolstering their perception that smart speakers are more capable of answering questions than children who have not experienced these functions before. However, future work is needed to explore this possibility.

There are several potential limitations to this work that should be addressed in future research. First, this study did not consider individual differences in children's numerical reasoning or future projections. Numerical reasoning may be more challenging for children under 10 than reasoning about far future events, leading to the discrepancy in the observed developmental patterns. Future work should explore numerical and future reasoning abilities as potential levers for children's understanding of virtuous ignorance. Second, this study largely included middle- and upper-class families from a single country (the United States). Culture may impact children's understanding of whether or not responses of "I don't know" are acceptable, and future research should include cross-cultural populations for comparison. Finally, future work should consider how often and for what kinds of tasks families use smart speakers. For example, children who use smart speakers for math questions more frequently may be more likely to believe that these devices can answer unknowable number questions with exact responses.

Taken together, our results demonstrate that children's evaluations of virtuous ignorance in smart speakers change between 5- and 12-years-old. As participants' age increased, children's preference for smart speakers that provided ignorant responses to unknowable questions over smart speakers that provided exact answers to those questions increased. Children preferred the ignorant response over the exact response for questions regarding the unknowable future events than unknowable numbers. Only adults preferred the virtuously ignorant response for both types of unknowable questions. These findings highlight the importance of teaching children about the capacity of internet connected devices as well as the value of virtuous ignorance.

References

- Aguiar, N. R., Stoess, C. J., & Taylor, M. (2012). The development of children's ability to fill the gaps in their knowledge by consulting experts. *Child Development, 83*, 1368–1381. <https://doi.org/10.1111/j.1467-8624.2012.01782.x>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Breazeal, C., Harris, P. L., DeSteno, D., Kory Westlund, J. M., Dickens, L., & Jeong, S. (2016). Young children treat robots as informants. *Topics in Cognitive Science, 8*(2), 481–491. <https://doi.org/10.1111/tops.12192>
- Bridgers, S., Buchsbaum, D., Seiver, E., Griffiths, T. L., & Gopnik, A. (2016). Children's causal inferences from conflicting testimony and observations. *Developmental Psychology, 52*(1), 9–18. <https://doi.org/10.1037/a0039830>
- Corriveau, K. H., Meints, K., & Harris, P. L. (2009). Early tracking of informant accuracy and inaccuracy. *British Journal of Developmental Psychology, 27*(2), 331–342. <https://doi.org/10.1348/026151008X310229>
- Danovitch, J. H., & Alzahabi, R. (2013). Children show selective trust in technological informants. *Journal of Cognition and Development, 14*(3), 499–513. <https://doi.org/10.1080/15248372.2012.689391>
- Dempster, F. N. (1981). Memory span: Sources of individual and developmental differences. *Psychological Bulletin, 89*, 63–100. <http://dx.doi.org/10.1037/0033-2909.89.1.63>
- Festerling, J., & Siraj, I. (2020). Alexa, what are you? Exploring primary school children's ontological perceptions of digital voice assistants in open interactions. *Human Development, 64*(1), 26–43. <https://doi.org/10.1159/000508499>
- Flanagan, T., Wong, G., Kushnir, T., & Pérez-Edgar, K. (2023). The minds of machines: Children's beliefs about the experiences, thoughts, and morals of familiar interactive technologies. *Developmental Psychology, 59*(6), 1017–1031. <https://doi.org/10.1037/dev0001524>
- Girouard-Hallam, L. N., & Danovitch, J. H. (2022). Children's trust in and learning from voice assistants. *Developmental Psychology, 58*(4), 646–661. <https://doi.org/10.1037/dev0001318>
- Girouard-Hallam, L. N., & Danovitch, J. H. (2024). How does Google get its information?: Children's judgements about Google search. *British Journal of Developmental Psychology, 42*(3), 334–347. <https://doi.org/10.1111/bjdp.12487>
- Girouard-Hallam, L. N., & Danovitch, J. H. (2025). Children's trust in Google's ability to answer questions about the past, present, and future. *Computers in Human Behavior, 165*, 108496–. <https://doi.org/10.1016/j.chb.2024.108496>
- Girouard-Hallam, L. N., Strebler, H. M., & Danovitch, J. H. (2021). Children's mental, social, and moral attributions toward a familiar digital voice assistant. *Human Behavior and Emerging Technologies, 3*(5), 1118–1131. <https://doi.org/10.1002/hbe2.321>
- Hermes, J., Behne, T., & Rakoczy, H. (2018). The development of selective trust: Prospects for a dual-process account. *Child Development Perspectives, 12*(2), 134–138. <https://doi.org/10.1111/cdep.12274>
- Hoffman, A., Owen, D., & Calvert, S. L. (2021). Parent reports of children's parasocial relationships with conversational agents: Trusted voices in children's lives. *Human Behavior and Emerging Technologies, 3*(4), 606–617. <https://doi.org/10.1002/hbe2.271>
- Jaswal, V. K., Croft, A. C., Setia, A. R., & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science, 21*(10), 1541–1547. <https://doi.org/10.1177/0956797610383438>
- Kahn, P. H., Kanda, T., Ishiguro, H., Gill, B. T., Ruckert, J. H., Shen, S., Gary, H. E., Reichert, A. L., Freier, N. G., & Severson, R. L. (2012). Do people hold a humanoid robot morally accountable for the harm it causes? *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, 33–40*. <https://doi.org/10.1145/2157689.2157696>
- Kahneman, D., & Anderson, N. B. (2003). A Perspective on Judgment and Choice: Mapping Bounded Rationality. *The American Psychologist, 58*(9), 697–720. <https://doi.org/10.1037/0003-066X.58.9.697>
- Koenig, M. A., & Jaswal, V. K. (2011). Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Development, 82*, 1634–1647. <https://doi.org/10.1111/j.1467-8624.2011.01618.x>
- Kominsky, J. F., Langthorne, P., & Keil, F. C. (2016). The better part of not knowing: Virtuous ignorance. *Developmental Psychology, 52*(1), 31–45. <https://doi.org/10.1037/dev0000065>
- Kushnir, T., Vredenburgh, C., & Schneider, L. A. (2013). “Who can help me fix this toy?” The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology, 49*, 446–453. <https://doi.org/10.1037/a0031649>
- Kinsella, 2018. *Amazon Alexa is testing answer updates feature that will notify users when it learns a new answer*. Voicebot.ai. Retrieved from: <https://voicebot.ai/2018/08/06/amazon-alexa-testing-answer-updates-feature-that-will-notify-users-when-it-learns-a-new-answer/>
- Lovato, S. B., Piper, A. M., & Wartella, E. A. (2019). Hey Google, do unicorns exist? Conversational agents as a path to answers to children's questions. *Proceedings of the 18th ACM International Conference on Interaction Design and Children, 301–313*. <https://doi.org/10.1145/3311927.3323150>
- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development, 73*, 1073–1084. <https://doi.org/10.1111/1467-8624.00458>
- Mills, C. M., Legare, C. H., Grant, M. G., & Landrum, A. R. (2011). Determining whom to question, what to ask, and

- how much information to ask for: The development of inquiry in young children. *Journal of Experimental Child Psychology*, 110, 539-560. <https://doi.org/10.1016/j.jecp.2011.06.003>
- Murayama, K., Usami, S., & Sakaki, M. (2022). Summary-statistics-based power analysis: A new and practical method to determine sample size for mixed-effects modelling. *Psychological Methods*, 27(6), 1014–1038. <https://doi.org/10.1037/met0000330>
- Nass, C. and Moon, Y. (2000), Machines and Mindlessness: Social Responses to Computers. *Journal of Social Issues*, 56: 81-103. <https://doi.org/10.1111/0022-4537.00153>
- Nguyen, S. P. (2012). The role of external sources of information in children's evaluative food categories. *Infant and Child Development*, 21, 216–235. <https://doi.org/10.1002/icd.745>
- Oranç, C., & Ruggeri, A. (2021). “Alexa, let me ask you something different” Children’s adaptive information search with voice assistants. *Human Behavior and Emerging Technologies*, 3(4), 595–605. <https://doi.org/10.1002/hbe2.270>
- R Core Team. (2022). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from: <https://www.R-project.org/>.
- Rideout, V., & Robb, M. (2020). *The Common Sense Census: Media Use by Kids Age Zero to Eight, 2020*. Common Sense Media. Retrieved January 22, 2025, from <https://www.commonsensemedia.org/research/the-common-sense-census-media-use-by-kids-age-zero-to-eight-2020>
- Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers’ theory of mind and semantic development. *Child Development*, 72(4), 1054–1070. <https://doi.org/10.1111/1467-8624.003344>
- Scott K., & Schulz L. (2017). Lookit (part 1): A new online platform for developmental research. *Open Mind (Cambridge, Mass.)*, 1(1). https://doi.org/10.1162/opmi_a_00002
- Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development*, 81, 669–679. <https://doi.org/10.1111/j.1467-8624.2009.01422.x>
- Tenney, E. R., Small, J. E., Kondrad, R. L., Jaswal, V. K., & Spellman, B. A. (2011). Accuracy, confidence, and calibration: How young children and adults assess credibility. *Developmental Psychology*, 47(4), 1065–1076. <https://doi.org/10.1037/a0023273>
- Tong, Y., Wang, F., & Danovitch, J. (2020). The role of epistemic and social characteristics in children’s selective trust: Three meta-analyses. *Developmental Science*, 23(2), 1-18. <https://doi.org/10.1111/desc.12895>
- Tong, Y., Wang, F., Danovitch, J., & Wang, W. (2022). When the internet is wrong: Children’s trust in an inaccurate internet or human source. *British Journal of Developmental Psychology*, 40(2), 320–333. <https://doi.org/10.1111/bjdp.12405>
- Wang, F., Tong, Y., & Danovitch, J. (2019). Who do I believe? Children’s epistemic trust in internet, teacher, and peer informants. *Cognitive Development*, 50, 248–260. <https://doi.org/10.1016/j.cogdev.2019.05.006>
- Xu, Y., & Warschauer, M. (2020). What are you talking to?: Understanding children’s perceptions of conversational agents. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. <https://doi.org/10.1145/3313831.3376416>
- Yarosh, S., Thompson, S., Watson, K., Chase, A., Senthilkumar, A., Yuan, Y., & Brush, A. J. B. (2018). Children asking questions: speech interface reformulations and personification preferences. *Proceedings of the 17th ACM Conference on Interaction Design and Children*, 300–312. <https://doi.org/10.1145/3202185.3202207>