

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The better part of not knowing: Virtuous ignorance

Permalink

<https://escholarship.org/uc/item/8232v65n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Kominsky, Jonathan F

Langthorne, Philip

Keil, Frank C

Publication Date

2015

Peer reviewed

The better part of not knowing: Virtuous ignorance

Jonathan F. Kominsky (jonathan.kominsky@yale.edu), Philip Langthorne (p.langthorne@gmail.com) & Frank C. Keil (frank.keil@yale.edu)

Yale University Department of Psychology, 2 Hillhouse Ave.
New Haven, CT 06520 USA

Abstract

For cases in which precise information is practically or actually unknowable, certainty and precision can indicate a lack of competence, while expressions of ignorance may indicate greater expertise. In two experiments, we investigated whether children and adults are able to use this “virtuous ignorance” as a cue to expertise. Experiment 1 found that adults and children older than 9 years selected confident informants for knowable information and ignorant informants for unknowable information. However, 5-7-year-olds overwhelmingly favored a confident informant, even when such precision was completely implausible. In Experiment 2, we demonstrated that 5-8-year-olds and adults are both able to distinguish between knowable and unknowable items when asked how difficult the information would be to acquire, but those same children still failed to reject the precise and confident informant for unknowable items. We suggest that children have difficulty integrating information about the knowability of particular facts into their evaluations of expertise.

Keywords: cognitive development; credibility; informants; confidence; epistemological beliefs

Introduction

Sometimes the most impressive intellectual achievement can be recognizing the boundaries of one’s own knowledge, or knowledge more generally. Indeed, this idea is a classic philosophical theme across a wide range of cultures, whether it be Socrates (“...I am wiser than he is to this small extent, that I do not think I know what I do not know” [Plato, *Apology* 21d, tr. Tredennick, 1954]), or Confucius (“Real knowledge is to know the extent of one’s ignorance.” [from Dunning, Johnson, Erlinger, & Kruger, 2003]).

Knowing what one does not know may require considerable sophistication. More than just knowing what portion of available knowledge one possesses, one must have some sense of the full extent of what is available and potentially knowable as well as what may never be knowable but still relevant. More knowledgeable individuals usually have more accurate senses of their abilities and limitations, including of their knowledge and explanatory understandings, while less knowledgeable individuals tend to be miscalibrated and overconfident (Dunning, 2012; Dunning et al. 2003). Children are particularly severely miscalibrated about their own understanding of various phenomena (Mills & Keil, 2005).

Here, we turn to a different, but closely related, problem: Identifying when others are overconfident about their own knowledge. In particular, we ask whether children and

adults are capable of recognizing when an informant who says that they do not know the answer to a question is actually more knowledgeable than one who provides a confident and precise answer. The literature on self-assessment suggests that claims of detailed knowledge are not invariably signs of expertise, and in some cases may be signs of incompetence or ignorance. In fact, under some circumstances claims of ignorance may indicate expertise and knowledge, and in these cases we can distinguish mere ignorance from “virtuous ignorance,” or admitting that certain knowledge cannot be possessed. For some questions or problems, the more expert individual may better understand how certain forms of information are unknowable. For example, a novice might feel quite confident that one can know both the exact position and exact velocity vector of a particle at a given point in time, but an expert familiar with Heisenberg’s Uncertainty Principle would freely admit that they could not know both precisely, and recognize any claims that one did as indicating ignorance rather than knowledge.

We expect that adults can identify many of these cases in which virtuous ignorance is a marker of greater expertise. In many instances, to know that one does not know (or *cannot* know) a specific piece of information arises from a sophisticated understanding of the physical world and how uncertainty, chaotic systems, or causal complexity make certain forms of precision or predictions highly implausible, if not impossible. Thus, one can evaluate the credibility of another’s testimony in terms of how it meshes with one’s own understanding of the plausibility of actually being able to attain the attested knowledge. In such cases, an informant might seem to have especially low credibility if that individual expresses great certainty about information for which certainty is highly inappropriate. This facet of assessing confidence has been largely neglected in empirical studies with adults, but in extreme cases, it seems intuitively clear that adults will reject confident declarations about logically unknowable things. Thus, most adults would discount the confident individual who claims to know something with implausible precision and prefer a virtuously ignorant informant. For example, it makes no sense for someone to claim that they know exactly how many leaves there are in all the trees in a large national park, and anyone who makes such claim with confidence should be regarded with great skepticism.

Children, however, may have great difficulty rejecting a confident and precise answer, even when that answer is highly implausible. There are two key challenges that

children must master. The first is an epistemological challenge: Children must recognize that the information is implausible or impossible to possess. To know that it is not feasible to have a precise number for all the leaves in Yellowstone Park, one needs to have a sense of the immensity of the number, of the challenges of getting a snapshot of all leaves at a moment in time, and even of the ambiguities of when a budding or decaying leaf becomes or is no longer a leaf. Therefore, one way in which children might fail to recognize virtuous ignorance as a cue to expertise is by not realizing that the ignorance is in fact virtuous. Rather, they might favor whatever information they are given, no matter how absurd it would sound to an adult. However, there are presumably at least a few domains in which children would recognize that some knowledge is impossible to possess.

The second challenge is that even in domains where children could recognize that possessing some knowledge is implausible or impossible, children may not be able to reject a confident informant. That is not to suggest that children simply accept everything they are told. On the contrary, the extensive literature on testimony reveals that even preschoolers take evaluative stances towards claims made by others and will take into account many source characteristics. These attributes can include: a source's past record of accuracy, a source's departure from consensus view, a source's current mental states and access to information, and a source's apparent dependency on other sources (for a review of this literature, see Robinson & Einav, 2014). In short, well before the start of formal schooling children appreciate that different sources should be trusted to different extents (Buschbaum et al., 2014; Harris, 2012).

That said, young children might over-value confidence early on as a marker of source information quality, perhaps at the expense of other factors. A sensitivity to confidence in demonstrations of object use emerges early in development. Young speakers are sensitive to linguistic indicators of certainty and confidence (Matsui, 2014), and by the time they enter the early school years, they are quite sophisticated evaluators of certainty-related expressions (Moore, Bryant, & Furrow, 1989).

In particular, five and six year olds are heavily influenced by a person's confidence but have difficulty, relative to older children and adults, in calibrating informants, i.e. taking into account the relationship between an informant's accuracy and confidence (e.g., is a person's confidence diagnostic of their credibility?). Two recent studies have provided evidence that children's difficulties with such calibrations are due to executive processing limitations. Tenney, Small, Kondrad, Jaswal, & Spellman (2011) found that children would not reject poorly calibrated informants while adults would, but when adults were placed under significant cognitive load, they performed as children did. In addition, Jaswal et al. (2014) demonstrated that some children will often accept obviously false adult testimony about an event that the child themselves witnessed, but the

likelihood of rejecting the testimony was positively correlated with performance on a spatial inhibitory control task.

Predictions

Based on these observations, we can make an overarching prediction: Young children will favor an implausibly confident informant over a virtuously ignorant one.

Further predictions emerge from the two challenges that could lead children to fail at this task: epistemological challenges and challenges integrating epistemology with expressions of confidence. These two accounts are not contradictory. In some domains children may fail to understand the implausibility of certainty; but we suggest that even when they do, the draw of confidence may be too strong to overcome. We make three specific predictions:

First, a simple certainty/confidence heuristic will be the first available strategy used by a child because they tend to focus on individual-centered markers of credibility more than the information itself, and because doubting certainty often requires detailed understandings of the world.

Second, older children and adults will favor virtuous ignorance over certainty, based on their understanding of the plausibility of precision in particular contexts.

Finally, the relationship between limited executive processing and the certainty bias in young children (Tenney et al., 2011; Jaswal et al., 2014) may be so strong that even when young children know that it is implausible to attain a certain piece of knowledge, they will still favor an implausibly certain informant. Thus even if children of a certain age can grasp that it might be extremely difficult or impossible to have a certain level of numerical precision about a given topic, they cannot inhibit their tendency to believe a confident informant. Indeed, given that inhibitory control improves radically between the ages of 5 and 10 (Williams, Ponesse, Schachar, Logan, & Tannock, 1999), these challenges may extend well into middle childhood.

We test these predictions with two experiments. In Experiment 1, we demonstrate that young children do indeed favor an implausibly confident informant, while older children and adults will favor a "virtuously ignorant" informant. In Experiment 2, we ask children and adults to evaluate how difficult it would be to possess certain types of knowledge, and demonstrate that even when children recognize the difficulty of knowing of specific pieces of information, they favor certainty over virtuous ignorance.

Experiment 1

Experiment 1 examined how children and adults evaluate implausible claims of precise numerical certainty. Two sets of questions were constructed: those for which it was implausible or impossible to make claims of numerical precision (e.g., the number of leaves in Yellowstone National Park) and those for which claims of certainty about precise numbers were plausible (e.g., the number of keys on a piano). We selected participants between ages five and ten years, expecting significant developmental shifts during that

period, including in inhibitory control (Williams et al., 1999).

Methods

Participants 105 children aged 5 to 10 years were divided into 4 grade clusters. Specifically, 26 children in Kindergarten ($M = 67.1$ mos., $SD = 5.57$ mos., 14 females), 27 children in first grade ($M = 77.67$ mos., $SD = 4.38$ mos., 14 females), 20 children in grades 2 and 3 ($M = 98.5$ mos., $SD = 8.9$ mos., 9 females), and 41 children in grades 4, 5 and 6 ($M = 126.5$ mos., $SD = 8.9$ mos., 17 females) participated in Experiment 1. Children were recruited in three ways: from regional schools, from regional science museums, and through visits into the experimenters' laboratory. All age groups were recruited by all three methods. In addition, 53 adults were recruited on Amazon's Mechanical Turk system ($M = 38.6$ years, $SD = 14$ years, 41 females).

Materials We created sixteen stimulus items, eight knowable and eight unknowable. The stimuli were presented to children on iPads, one at a time, with images to help them track the claims, along with two silhouettes representing the informants. Items were presented in a random order. The overall effect was to create a storybook appearance for each stimulus page. At the bottom of each page was a small button to advance to the next item.

For each item, one of the two silhouetted respondents answered the question at the top of the page with the word "exactly" and a precise integer, and the other responded, "I don't know because it is not possible to answer that question precisely." Participants were asked to identify which of the two respondents, who had both claimed to be "experts" when we questioned them, was actually a better expert than the other. An example display and complete list of items can be found at <http://goo.gl/NoyqSs>

Procedure Adult participants viewed the stimuli through a web browser from their home computers and responded by clicking on the silhouette preferred. All child participants viewed the stimuli on Apple iPads and touched the screen to endorse a specific response. Children in grades K-3 had the text read aloud to them by the experimenter across all of the trials, whereas children in fourth grade and above were read aloud the instructions and the first item by the experimenter, who then allowed the children to read and advance through the remaining items themselves.

The experimenter explained to the child participants that they were about to play a detective game involving experts, and asked if they knew what the word "expert" meant. The experimenter then defined the word for children who did not know it, or redirected the definition provided by children who were able to generate one, to be someone who "understood something really, really well," and made clear that this new definition was the one to be used for the game. Children were then asked if they understood something

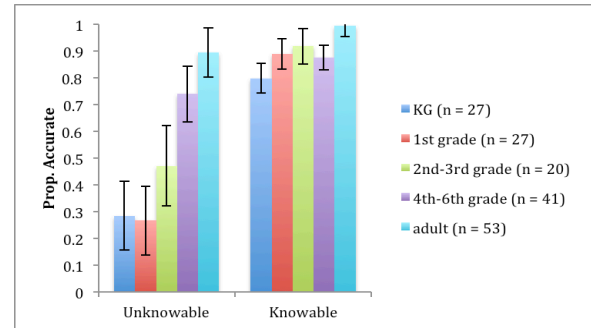


Figure 1: Results of Experiment 1. Error bars represent 95% CIs. Performance improved with age for both item types, but to a greater degree for unknowable items.

really well and so could be an expert in that thing, to ensure comprehension. Next, the rules of the game were explained: the experimenter claimed to have had a list of questions and found people who said that they were experts about the topics, but that the experimenter suspected that some of the "experts" might really be better than others. The participants were told, "You're going to see a bunch of questions and answers from different expert, and I want you to help me decide who was really the better expert for each question."

Adults, participating online, read the experiment instructions on a page after answering age and sex demographics questions and before the first item, and did not complete the interactive comprehension check of identifying something in which they could be experts and receiving feedback.

Invalid trials (missing data or stray clicks outside of the target regions covering the figures and their responses) were recorded as null responses for each item. (This was infrequent and only occurred on 3.5% of all trials.) Finally, any participant who had no responses recorded for either category was excluded from analysis.

Results

We calculated an "accuracy" score for all the implausible and plausible numerical items, based on the "correct" choice (more accurately, the expected adult response) of either the expert who gave a precise response for the knowable items or an "I don't know" response for the unknowable items. These scores were computed as the proportion of valid responses for each item type that were "correct". Fig. 2 shows the results of Experiment 1 by item type and age group. There was a main effect of item type, with much higher accuracy overall for knowable ($M = .90$, $SD = .16$) than unknowable items ($M = .60$, $SD = .42$), $F(1,163) = 108.67$, $p < .001$, $\eta_p^2 = .400$. This main effect was qualified by a significant interaction between age group and item type, $F(4,163) = 10.592$, $p < .001$, $\eta_p^2 = .206$.

Further analyses revealed a main effect of age group for unknowable items, $F(4, 163) = 25.063$, $p < .001$, $\eta_p^2 = .381$. Adults ($M = .90$, $SD = .27$) were significantly more likely to choose the "I don't know" expert than kindergarteners ($M = .28$, $SD = .30$), first graders ($M = .27$, $SD = .37$), and second

and third graders ($M = .47$, $SD = .43$), all $ps < .001$. (All p -values reported for pairwise comparisons use Bonferroni correction for multiple comparisons.) Adults did not differ from children in grades four through six ($M = .74$, $SD = .36$), $p = .28$. Among the younger age groups, there were significant differences between children in grades four through six and kindergarteners ($p < .001$), children in first grade ($p < .001$), and children in 2nd-3rd grade ($p = .039$).

There was also an effect of age group for knowable items, $F(4, 163) = 8.80$, $p < .001$, $\eta_p^2 = .178$. Follow-up analyses indicated a slight developmental improvement in accuracy on knowable items, but a much smaller and less consistent improvement than was found for unknowable items (notably, 2nd-3rd graders were as accurate as adults). Furthermore, all age groups were well above chance accuracy for knowable items, $ps < .001$.

Discussion

As predicted, all ages favored confident informants over ignorant ones when the information provided was plausibly knowable, but when the information provided was implausibly precise, older children and adults selected the virtuously ignorant informant, while younger children continued to favor the confident informant.

The question remains as to the extent to which younger children are completely unaware of such patterns of “knowability” or whether they simply weigh them differently and let poor knowability be trumped by the trappings of precision and confidence. In other words, could children who show a strong bias to pick the confident expert in all conditions nonetheless be aware of the (im)plausibility of such answers? Further examination of the results of Experiment 1 provides some indication of this: One can re-analyze the responses in terms of how frequently the certain expert was chosen, and ask whether there is a difference in this frequency based on item type. These post-hoc analyses found that kindergarteners in Experiment 1 were marginally less likely to select the certain expert for unknowable items than knowable ($p = .06$) and first graders were significantly less likely to choose the certain expert for unknowable items ($p = .01$), despite their overall poor accuracy on unknowable items. This suggests some awareness of knowability, but an inability to inhibit a bias to believe a confident informant until around age 10, which aligns well with developments in inhibitory control (Williams et al., 1999).

Experiment 2

To test for a trumping influence of confidence even when information is seen as difficult or impossible to possess by younger children, in Experiment 2, we constructed the strongest possible contrast, in which we first asked children how difficult it would be to know particular things, and then immediately afterwards had them do the same task as Experiment 1 with items closely related to the ones that they

had just rated. We changed the unknowable items in a very minor way by expanding the scope (e.g., counting the number of blades of grass in Central Park vs. in New York State) so that they were not pure repetitions, but still close enough that a judgment of implausibility for one would entail implausibility for the other, and making sure that the initial judgment was always a narrower scope of entities. We validated these stimuli with adult piloting in MTurk, to verify that participants felt the rating items were comparable but of narrower scope. We predicted that young children would identify some types of knowledge as unknowable, but still be drawn to pick a confident expert.

Methods

Participants 31 children age 5-6 ($M = 69.7$ mos., $SD = 14.0$ mos., 17 females), 24 children age 7-8 ($M = 95.8$ mos., $SD = 7.5$ mos., 12 females) and 40 adult participants from MTurk ($M = 31.8$ years, $SD = 10.1$ years, 16 females) participated in Experiment 2. Children were recruited by the same methods as Experiment 1, and had a similar demographic profile, but none had participated in Experiment 1.

Materials & Procedure The procedure for Experiment 2 was similar to the procedure used in Experiment 1, but with fewer forced choice items and preceded by a difficulty rating task. Participants first completed four training items in which they learned to use the rating scale by rating the size of four animals (squirrel, cat, cow, and horse) on a scale that went from “SMALL” to “BIG.” All participants gave higher ratings for the cow and horse than squirrel and cat, indicating proficiency with the response method. Participants then saw a series of six items which depicted a person climbing a set of very low steps on the left side with the word “EASY” and another climbing up a very steep cliff on the right side with the word “HARD”. The six screens presented training items that had been adapted from the results of Experiment 1, as the items on which participants in all age groups performed the most consistently. However, instead of presenting the exact item from Experiment 1 (e.g., 'How many windows are on the White House'), the item's scope was narrowed to represent a subset of the original (e.g., 'How many windows are on the President's office in the White House'). An image was presented along with each question, as in Experiment 1.

For these rating items, participants were told to touch the screen on the blue line, just as before, to indicate whether it would be easy, hard, or somewhere in between to find out the answer to the question presented. The six rating screens eliciting a hard or easy judgment were presented in a random order to every participant. The second part of the experiment was identical to the procedure from Experiment 1 (including training), but only using the six items from Experiment 1 that corresponded with the rating items. The order of presentation was independently randomized and not tied to the presentation of the rating items.

Results

We analyzed the difficulty ratings in terms of the absolute x-coordinate of the recorded mouse-click or touch (within a constrained y-coordinate range), yielding a scale from 1 to 900 (pixels) with lower numbers representing “easier” responses. For difficulty ratings, there was a main effect of item type, $F(1, 92) = 322.7, p < .001, \eta_p^2 = .778$, as well as a significant interaction between item type and age group, $F(2, 92) = 37.83, p < .001, \eta_p^2 = .451$. There was no significant main effect of age group, $F(2, 92) = 1.927, p = .151$. The average ratings by age group and item type can be found in Table 1. All age groups gave higher ratings for unknowable items (paired-sample *t*-tests, $ps < .001$), indicating that even the youngest children were able to distinguish “knowable” from “unknowable” information.

Further analyses showed main effects of age group for both knowable and unknowable items, one-way ANOVAs, $ps < .001$. For knowable items, adults gave lower ratings than both 5-6-year-olds and 7-8-year-olds, Bonferroni-corrected pairwise comparisons $ps < .001$, but the younger age groups did not differ from each other, $p > .9$. For unknowable items, there was significant developmental improvements between every age group, $ps \leq .045$, indicating some developmental improvement in the ability to recognize epistemological challenges. However, given that younger children were able to distinguish knowable from unknowable items, the question remains whether they can use this information to reject overconfident experts.

For the expert evaluation task, there were also main effects of item type, $F(1, 92) = 50.30, p < .001, \eta_p^2 = .353$, and age group, $F(2, 92) = 52.72, p < .001, \eta_p^2 = .534$, and a significant interaction, $F(2, 92) = 6.71, p < .001, \eta_p^2 = .127$. As in Experiment 1 there were main effects of age group for both knowable, $F(2,92) = 5.33, p < .001, \eta_p^2 = .104$, and unknowable items, $F(2, 92) = 26.62, p < .001, \eta_p^2 = .367$. As can be seen in Fig. 2, for unknowable items, adults ($M = .825, SD = .292$) differed significantly from both 5-6-year-olds ($M = .237, SD = .346$) and 7-8-year-olds ($M = .431, SD = .423$), $ps < .001$, but the younger age groups did not differ from each other, $p = .127$. For knowable items, there was again a much smaller small developmental improvement, and all age groups performed significantly better than chance ($ps < .001$). Thus, even though children distinguish knowable and unknowable items in their difficulty ratings, they seem unable to use this information to reject an overconfident expert.

An alternative interpretation might be that, even though 5-6-year-olds gave higher difficulty ratings for unknowable items, their ratings are still below some threshold of “implausibility,” but if any individual children recognized this implausibility, they would be able to reject the confident expert. This account predicts that, if difficulty ratings for unknowable items and age group were pitted against each other as predictors of performance on the expert evaluation task, age group should be a weaker predictor than difficulty rating. However, this is not the case. In a stepwise regression of performance on

Table 1: Difficulty ratings for unknowable and knowable items in Experiment 2.

Age Group	Unknowable	Knowable
5-6-year-olds	556.5 (163.7)	371.4 (129.9)
7-8-year-olds	648.2 (137.7)	342.6 (125.2)
Adults	733.5 (100.6)	172.2 (102.5)

NOTE: numbers in parentheses are standard deviations.

unknowable items in the expert evaluation task with age group, difficulty ratings of unknowable items, and an interaction term as predictors, only age group emerges as a significant predictor ($\beta = .601, p < .001$). Thus, developmental improvements in endorsing virtuous ignorance are independent of the ability to assess the unknowability of information.

Discussion

Experiment 2 demonstrated that even the youngest children easily distinguished between knowable and unknowable items in their difficulty ratings. However, despite this, children were unable to reject the confident expert. While there are developmental improvements in the ability to recognize the difficulty of certain types of knowledge, they were unrelated to the ability to reject the confident expert. This suggests that the key developmental shift is not just in assessing epistemic challenges, but also in a growing ability to integrate information about knowability and confidence when evaluating experts.

General Discussion

Children have difficulty using their epistemological knowledge to recognize when a person who speaks with confidence might not know what he or she is talking about. Children have an early-developing sensitivity to expressions of confidence and certainty, and can use these expressions to evaluate statements. Here we have shown that younger children are so swayed by certainty that they do not take into account those cases where professed ignorance is actually the stronger indicator of expertise or, equivalently, confidence and precision are indicators of incompetence.

Experiment 1 showed that young school children choose




Figure 2: Results of the forced-choice task in Experiment 2. Error bars represent 95% CIs.

experts who claimed to have implausibly precise numerical knowledge whereas older children and adults clearly recognized the absurdity of such claims and chose the expert who professed ignorance and/or uncertainty. Experiment 2 showed that, even when children recognize the implausibility of certain types of knowledge, they are unable to integrate it with expressions of confidence to recognize that an informant is miscalibrated.

This inability may be due to a failure of executive processing, such that cues to confidence cannot be integrated with an understanding of the implausibility of possessing such information. This account is similar to accounts of how children fail to appropriately calibrate the confidence and accuracy of an expert over repeated trials (Tenney et al., 2011; Jaswal et al., 2014). Indeed, the underlying process may be nearly identical, merely differing in what information is being integrated with confidence. Rather than observed accuracy, in this case children must integrate *plausible* accuracy with confidence to determine if an expert is appropriately calibrated.

Despite these failures of integration, it is impressive that even five-year-olds in Experiment 2 distinguished between knowable and unknowable information. This early ability suggests that, with some assistance integrating the information they already possess, these children might learn the value of virtuously ignorant informants. However, there may still be epistemological challenges, even for adults. Adults could also easily be seduced by an inappropriately confident informant if they did not understand that such confidence was implausible. For example, recalling the example from the introduction, an adult who did not know that it was impossible to know both the position and velocity vector of a particle might favor a confident and precise expert. In the reported experiments we validated our stimuli with adults to ensure that the selected items were recognized as implausible to be certain about, but in day-to-day life it is unclear how well adults can actually identify the plausibility of knowing something.

This problem is compounded if we consider how such plausibility information might be learned. Given the intricate web of deference needed to successfully navigate a complex world (Keil, Stein, Webb, Billings, & Rozenblit, 2008), a lay sense of knowability may often come from the very experts that we are trying to evaluate. For example, of the readers of this paper who knew of Heisenberg's Uncertainty Principle, it is unlikely that any of them have direct evidence for it or proved it themselves. Indeed, the authors themselves only know it through deference to physics experts. If there were an equal population of experts that claimed that such information was plausibly knowable with precision, how would we be able to evaluate whether a confident, precise response was appropriate?

It is also worth noting that our "virtuously ignorant" informants did not simply say "I don't know." They provided a specific reason for their ignorance, i.e., that the information could not be known. It is not necessarily a cue of expertise to merely express ignorance, even when

something is in fact unknowable. One could claim ignorance because one is not an expert, as easily as one could claim ignorance because one *is* an expert. It would be somewhat surprising if someone who merely expressed ignorance without providing further information would ever be seen as an expert. While not specifically tested in these experiments, we would expect that the additional statement that specific information is unknowable is important for identifying a virtuously ignorant expert over someone who simply knows nothing.

References

- Buchsbaum, D., Bridgers, S., Whalen, A., Seiver, E., Griffiths, T. L., & Gopnik, A. (2012). Do I know that you know what you know? Modeling testimony in causal inference. In *Proceedings of the 34th annual conference of the cognitive science society*.
- Dunning, D. (2012). *Self-insight: Roadblocks and detours on the path to knowing thyself*. Psychology Press.
- Dunning, D., Johnson, K., Ehrlinger, J., & Kruger, J. (2003). Why people fail to recognize their own incompetence. *Current Directions in Psychological Science*, 12(3), 83-87.
- Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge; Harvard University Press
- Jaswal, V. K., Pérez-Edgar, K., Kondrad, R. L., Palmquist, C. M., Cole, C. A., & Cole, C. E. (2014). Can't stop believing: Inhibitory control and resistance to misleading testimony. *Developmental Science*, Online First.
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science*, 32(2), 259-300.
- Matsui, T. (2014). Children's understanding of linguistic expressions of certainty and evidentiality. *Pragmatic Development in First Language Acquisition*, 10, 295.
- Mills, C. M., & Keil, F. C. (2005). The development of cynicism. *Psychological Science*, 16, 385-390
- Moore, C., Bryant, D., & Furrow, D. (1989). Mental terms and the development of certainty. *Child Development*, 60, 167-171.
- Plato, & Tredennick, H. (1954). *The Last Days of Socrates. The Apology, Crito, Phaedo. Translated and with an Introduction by Hugh Tredennick*. Penguin Books.
- Robinson, E. J., & Einav, S. (Eds.). (2014). *Trust and skepticism: Children's selective learning from testimony*. Psychology Press.
- 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, 1065-1077.
- Williams, B. R., Ponesse, J. S., Schachar, R. J., Logan, G. D., & Tannock, R. (1999). Development of inhibitory control across the life span. *Developmental Psychology*, 35(1), 205-213.