Name of applicant and Yale affiliation:
Rosie Aboody, Graduate Student, Department of Psychology

Title of proposed project:
How do I know what you know? Building a theoretical framework of epistemic inference

Budget with description:

Expense	Cost
Paying adult research participants for Experiments 1 & 2	$3 per participant x 500 participants[1]
Gift cards for child research participants for Experiments 3 & 4	$5 per participant x 300 participants[2]

[1]Cost calculated assuming participants will receive $3 for completing a 15-minute survey, in accordance with Connecticut minimum wage as of September 1st, 2020.
[2]Cost based on the current rate of the Yale Infant and Child Development Group.

Total cost: $3,000.00 USD

Timeline for project:

September & October 2020	Finalize study design of Experiments 1 and 3. Pre-register, collect, and analyze data.
November & December 2020	Finalize study design of Experiments 2 and 4. Pre-register, collect, and analyze data.
January & February 2021	Write draft of manuscript.
March & April 2021	Submit draft of manuscript to collaborators for review; implement edits and conduct any additional analyses requested.
May 2021	Submit manuscript to journal.

Letter of recommendation in support of project work:
My advisor, Dr. Julian Jara-Ettinger, has e-mailed the letter directly to you.

Description of project:

Introduction

In our everyday lives we ascribe rich mental states to others, judging from their actions what they intended, wanted, knew, or believed. These inferences allow us to navigate the social world: to communicate effectively, decide whether to do or teach, punish or forgive. The question of how we understand other minds has received substantial attention in diverse fields, including psychology, cognitive science, philosophy, sociology, and more recently, computer science. While past work has well characterized how we infer others' intentions, goals and desires from their actions, little research has tested how we infer what others know or believe. But by uncovering the principles that enable us to make epistemic inferences — and revealing how these abilities develop — we can better understand the capacities that make humans exceptional teachers, learners, helpers, and communicators.

In this project, we combine approaches from Psychology, Cognitive Science and Computer Science. In Aim 1, we characterize mature epistemic-state inferences by building and validating a computational model of adult epistemic inference. In Aim 2, we investigate how these abilities develop.

Aim 1: Characterizing mature epistemic-state inferences

Past work suggests that our ability to infer others' mental states relies upon an expectation that others will act efficiently, choosing actions that produce the greatest rewards while incurring the fewest costs (see Jara-Ettinger et al., 2016). So, from the effort someone expends pursuing an outcome, we can infer how much they want it. Formally:

$$Utility(action, outcome) = Reward(a, o) - Cost(a, o) \quad (1)$$

We expand this framework to capture epistemic inferences. First, we propose that sometimes, the cost of achieving an outcome is determined by what we know, K:

$$Utility(action, outcome) = Reward(a, o) - Cost_K(a, o) \quad (2)$$

If so, then seeing others' actions and their costs should reveal what they know. For example, if one gold coin is hidden on a small, easy-to-search island, and one is hidden on a huge, hard-to-search one, it would be sensible for an ignorant agent to search the small island. So if you see them search the larger (and seemingly costlier) island, you can infer they knew a lot about it. In this way, if adults expect others to act efficiently based on their epistemic states, they should be able to use the costs of others' actions to infer how much they know.

Now sometimes, we can seek added knowledge, $+K$. This knowledge may come at a cost (like a 5-year PhD!), but it might also allow us to accomplish our goals more easily, decreasing their cost:

$$Utility(action, outcome) = Reward(a, o) - Cost_{+K}(a, o) - Cost(+K) \quad (3)$$

Here, seeing whether agents choose to seek added knowledge — and at what cost — should reveal both how much they know, and how much they believe they can learn. For instance, it would only be efficient for someone to consult a treasure map before searching a tiny

island if they both knew little about the coin's location, and thought that the treasure map would be very helpful.

We will compare participant judgments to those of our model, with Experiments 1 and 2 testing the predictions of Eq. 2 and 3 respectively. Ensuring this approach will be fruitful, we have collected pilot data (n = 30) for our first experiment. We find that participants' judgments correlate precisely with the predictions of our model, $r(28) = .798$, $p < .001$, providing initial support for our account.

Aim 2: Characterizing the development of our epistemic-state inferences

To infer what others know from their actions (as tested in *Aim 1*), we need to understand how knowledge is likely to affect the outcome of those actions — understanding, for example, that an ignorant agent is unlikely to find a treasure buried on a huge island, but might have better luck on a tiny one. The relation between epistemic states and action outcomes is not always straightforward: for example, if only one of the four light switches in your entryway actually works, you'd probably be convinced someone was knowledgeable if they flipped the right switch on the first try. But if all four switches turned the lights on, a success might be less diagnostic.

We introduce four- and five-year-olds to a similar situation, contrasting a "probable" task with a 4/4 chance of success against an "improbable" task with a 1/4 chance of success. We simply ask participants which task would better reveal an agent's knowledge state. If children consider the probability that this agent would succeed by chance, they should say that the improbable task is more diagnostic, because even an ignorant agent would always succeed on the probable one. In a pilot of 18 subjects, we find that 72% of children (13/18) wanted to see the agent complete the improbable task.

Experiment 4 will simply compare a case with a 3/4 chance of success (75%) to the 25% case. It's possible that children only choose the improbable task when the contrast is extreme. A more graded comparison case will reveal just how nuanced children's judgments are.

Relevance to current research program

The proposed experiments embody the core aim of my research program: to characterize how we infer others' epistemic states from their actions. A multidisciplinary approach is central to my work, and offers two key advantages. First, by formalizing a theory of how people might make epistemic inferences, we can precisely predict how adults should respond if our theory is correct. If people respond as the model predicts across many situations, this suggests that participants use a similar process to make epistemic inferences. Second, these predictions can then guide our developmental approach. In our formal account, epistemic inferences depend upon an understanding of how knowledge affects the outcomes of others' actions. In prior work we show that preschoolers already make sophisticated epistemic inferences (Aboody, Zhou & Jara-Ettinger, under review; Aboody, Zhou, Flowers & Jara-Ettinger, 2019; Aboody, Huey & Jara-Ettinger, 2018). If children make these inferences as our account predicts, then they should understand how action outcomes and epistemic states relate. By examining the emergence and development of abilities our account predicts are pre-requisites to epistemic reasoning, we can systematically test the predictions of our theoretical approach.

Taken together, this work will help us better understand the epistemic inferences underpinning many critical and uniquely human capacities, like our abilities to teach, learn, and communicate.

References

Aboody, R., Zhou, C., & Jara-Ettinger, J. (under review). In pursuit of knowledge: preschoolers expect agents to weigh information gain and information's cost when deciding whether to explore.

Aboody, R., Flowers, M., Zhou, C., & Jara-Ettinger, J. (2019). Ignorance = doing what is reasonable: Children expect ignorant agents to act based on prior knowledge. Proceedings of the 41st Annual Conference of the Cognitive Science Society.

Aboody, R., Huey, H., & Jara-Ettinger, J. (2018). Success does not imply knowledge: Preschoolers believe that accurate predictions reveal prior knowledge, but accurate observations do not. Proceedings of the 40th Annual Conference of the Cognitive Science Society.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J.B. (2016). The naive utility calculus: Computational principles underlying commonsense psychology. Trends in Cognitive Sciences.