

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Does Amy Know Ben Knows You Know Your Cards? A Computational Model of Higher-Order Epistemic Reasoning

Permalink

<https://escholarship.org/uc/item/2kk1h4b2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Cedegao, Zhang
Ham, Huang
Holliday, Wesley H.

Publication Date

2021

Peer reviewed

Does Amy Know Ben Knows You Know Your Cards? A Computational Model of Higher-Order Epistemic Reasoning

Zhang Cedegao*

cedzhang@berkeley.edu
University of California, Berkeley

Huang Ham*

hamhuang@sas.upenn.edu
University of Pennsylvania

Wesley H. Holliday

wesholliday@berkeley.edu
University of California, Berkeley

*These authors contributed equally to this work.

Abstract

Reasoning about what other people know is an important cognitive ability, known as *epistemic reasoning*, which has fascinated psychologists, economists, and logicians. In this paper, we propose a computational model of humans' epistemic reasoning, including *higher-order* epistemic reasoning—reasoning about what one person knows about another person's knowledge—that we test in an experiment using a deductive card game called “Aces and Eights”. Our starting point is the model of perfect higher-order epistemic reasoners given by the framework of *dynamic epistemic logic*. We modify this idealized model with bounds on the level of feasible epistemic reasoning and stochastic update of a player's space of possibilities in response to new information. These modifications are crucial for explaining the variation in human performance across different participants and different games in the experiment. Our results demonstrate how research on epistemic logic and cognitive models can inform each other.

Keywords: higher-order theory of mind; deductive reasoning; epistemic logic; computational modeling

Introduction

During social interactions, people constantly need to reason about other people's mental states, an ability known as Theory of Mind (ToM) (Premack & Woodruff, 1978; Frith & Frith, 2005). ToM encompasses a wide range of cognitive abilities, since people possess many kinds of mental states, such as beliefs, intentions, and emotions (Apperly, 2010; Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017; Ong, Zaki, & Goodman, 2019). As we reason about other people's mental states, others may be doing the same. Therefore, the ability to engage in *higher-order* reasoning, i.e., to reason about other people's reasoning, is also essential and prevalent in everyday interactions. One type of higher-order reasoning is higher-order *epistemic* reasoning, which is reasoning about what others know about what others know (and so on). For example, in email exchanges, higher-order epistemic reasoning guides the choice of CC vs. BCC: it depends on whether I want my recipients to *know* that the other recipients of the message also *know* the content of my message.

Much of the prior research on higher-order epistemic reasoning focuses on developmental aspects—when and how do children acquire the ability to think about what others know (Gopnik & Wellman, 1992)? Through the influential empirical paradigm of the “false-belief task”, where children read stories about people having asymmetric knowledge about the location of some object and answer questions about them, researchers found that children around 4 years old can carry

out first-order epistemic reasoning, and they begin to possess second-order reasoning abilities 2 or 3 years later (Perner & Wimmer, 1985; Wellman, Cross, & Watson, 2001).

A major area of research studying higher-order epistemic reasoning in adults is behavioral game theory, where economists debate the interpretations of people's seemingly irrational reasoning strategies in games that require such reasoning (Ho, Camerer, & Weigelt, 1998; Camerer, 2006). A potential limitation of this area of research is that most games are humans playing against humans, so the reasoning ability of their opponents is not controlled. Hence it is hard to assess to what extent people deviate from rational behaviors due to their accommodating the cognitive limitations of others.

Closest to the present study is a line of work on formal modeling and epistemic reasoning that combines insights from logic, game theory, and psychology (e.g., Verbrugge, 2009; van Maanen & Verbrugge, 2010; de Weerd, Verbrugge, & Verheij, 2015). One notable finding is that an appropriate definition of logical complexity predicts cognitive difficulty in reasoning games (Szymanik, Meijering, & Verbrugge, 2013; Zhao, van de Pol, Raijmakers, & Szymanik, 2018). Our work contributes to this tradition by introducing new cognitive models directly based on logic and studying reasoning behaviors that are more complicated than those in previous studies and yet manageable by some people in a highly controlled experimental setting. Our work is also related to the “cognitive logics” research program (Ragni, Kern-Isberner, Beierle, & Sauerwald, 2020), with which we share the pursuit of bridging logic and the psychology of reasoning. However, so far they have mainly focused on non-monotonic logics as the formalism and classical tasks such as syllogism and conditional reasoning as the problems of interest (e.g., Ragni, Eichhorn, & Kern-Isberner, 2016; da Costa, Saldanha, Hölldobler, & Ragni, 2017), which differ from the focus of the present study.


Here we present a computational framework based on *epistemic logic* to model the *deductive* component of higher-order epistemic reasoning. We show that our proposed model is able to explain people's performance in a card game that involves sophisticated higher-order epistemic reasoning, known as “Aces and Eights” (Fagin, Halpern, Moses, & Vardi, 2003). To our knowledge, our model is the first to account for people's performance on third- and fourth-order epistemic reasoning (cf. van Maanen & Verbrugge, 2010).

More generally, our research aims to further facilitate interactions between logic and cognitive science: logics beyond Boolean and first-order logic can also be helpful for modeling cognition, and human cognition can inspire new logics.


Epistemic Logic

Our starting point for modeling higher-order epistemic reasoning in humans is the model of perfect reasoners provided by epistemic logic, the main formal model of reasoning about knowledge in epistemology since the 1960s (Hintikka, 1962). Epistemic logic provides semantic models of ideal agents' knowledge about the world and each others' knowledge, as well as syntactic formalizations of their reasoning about knowledge, with numerous applications in computer science (Fagin et al., 2003), game theory (Aumann, 1999), and philosophy (Holliday, 2018). For contemporary surveys on epistemic logic, see van Ditmarsch, Halpern, van der Hoek, and Kooi (2015) and Rendsvig and Symons (2021).

The basic semantic models, which we call *epistemic structures*, are directed graphs with edges labeled by agents. Nodes of the graph represent *possible states of the world*, and an edge from state w to state v labeled for an agent A indicates that *if the state of the world is w , then agent A cannot rule out the possibility that the state of the world is v* . Thus, edges between nodes labeled for A represent A 's *uncertainty*. For example, the following epistemic structure represents a scenario in which Amy just flipped a coin and saw which side landed face up, while Ben didn't see which side landed face up, but he knows Amy did (throughout this section and in Figure 2, we assume there is a loop edge between each state and itself labeled for each agent, but we do not draw them):



The two possible states of the world are that *Amy sees heads* (H) or *Amy sees tails* (T); Ben's uncertainty is represented by the edges between H and T labeled for B ; and Amy's lack of uncertainty is represented by there being no edges labeled for A . By contrast, the following structure represents a scenario in which Amy just flipped a coin and observed which side landed face up, and Ben didn't see which side landed face up, but also *Ben doesn't know whether Amy saw it*:




Now there are four possible states: Amy sees heads (H), Amy sees tails (T), the coin lands heads but Amy doesn't see it (H'), or the coin lands tails but Amy doesn't see it (T'). Ben is fully uncertain between all states, as represented by the edges labeled for B between all states, but Amy is only uncertain in the cases H' and T' in which she does not see which side lands face up, as represented by the bottom edges labeled for A . Since Ben does not know whether the actual state of the

world is among the top two or the bottom two states, Ben has *higher-order uncertainty* about Amy's uncertainty.

Having seen two concrete epistemic structures, one can understand the formal definition of knowledge in a structure: where a *proposition* is a set of states, *in state w agent A knows proposition P* if and only if w and all nodes reachable from w by a directed edge labeled for A belong to P . For example, in state H' above, Amy does *not* know the proposition $\text{Heads} = \{H, H'\}$, because there is an edge from H' to T' labeled for A , and T' is not in Heads . Also in T and T' , Amy does not know Heads , because T and T' are not in Heads . By contrast, in state H , Amy *does* know Heads , because there is no edge for A from H to a state outside of Heads . Thus, Amy knows Heads only in state H . We may then consider "Amy knows Heads " as a proposition of its own, denoted $K_A(\text{Heads})$, with $K_A(\text{Heads}) = \{H\}$. Similarly, where $\text{Tails} = \{T, T'\}$, we have $K_A(\text{Tails}) = \{T\}$. We can then consider propositions such as "Amy knows which side landed face up," i.e., she knows Heads or knows Tails , which is simply the *union* $K_A(\text{Heads}) \cup K_A(\text{Tails}) = \{H, T\}$. Amy's *not* knowing which side landed up is then the *complement* $\sim(K_A(\text{Heads}) \cup K_A(\text{Tails})) = \{H', T'\}$. Ben's ignorance, in every state, about whether Amy knows which side landed face up can be formalized by saying that in every state, Ben knows neither $K_A(\text{Heads}) \cup K_A(\text{Tails})$ nor its complement. That is, $K_B(K_A(\text{Heads}) \cup K_A(\text{Tails})) = \emptyset$ and $K_B \sim(K_A(\text{Heads}) \cup K_A(\text{Tails})) = \emptyset$.

While epistemic logic provides a way of representing the knowledge of multiple agents at a single time, *dynamic epistemic logic* (DEL) provides a way of modeling how their knowledge changes in response to new information (van Ditmarsch, van der Hoek, & Kooi, 2008; van Benthem, 2011). In the case of a *public announcement* to all agents of some proposition P , we model their knowledge acquisition by simply deleting from the initial epistemic structure all states outside of P . For example, if Amy announces to Ben, "I don't know which side landed face up," represented by the proposition $\sim(K_A(\text{Heads}) \cup K_A(\text{Tails})) = \{H', T'\}$, then we delete states H and T from the previous structure, resulting in:



In both states of this *updated* epistemic structure, Ben knows that Amy does not know which side landed face up: $K_B \sim(K_A(\text{Heads}) \cup K_A(\text{Tails})) = \{H', T'\}$. In this case, Amy's announcement of her ignorance did not give Ben further information about the non-epistemic or "ontic" fact of whether the coin landed heads or tails. But sometimes an announcement of ignorance or knowledge on the part of one agent *can* give another agent knowledge of ontic facts, such as how a coin landed or how cards were distributed in a game. This is precisely the phenomenon that occurs in the Aces and Eights game used in the main task of our experiment.

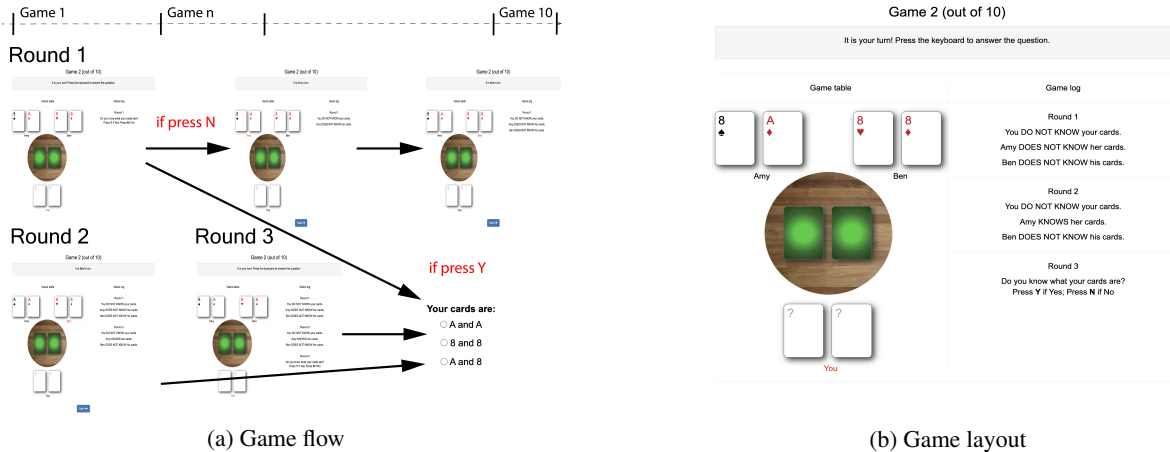


Figure 1: Example game. Each game has at most 3 rounds. If the participant announces “I know my cards” (by pressing Y) or makes an incorrect announcement, the game ends immediately. The correct announcements for this game are “I don’t know”.

Experiment

Participants

We recruited 306 participants on the Prolific platform. Only participants between 18–60 years old who resided in the US, had an approval rate of 95%+, and completed at least 10 studies were eligible to sign up. Participants received \$5 of base payment (aimed at \$10/hr) and could earn substantial bonuses based on their performance (see details below). Three participants were excluded due to data recording errors (two missing and one incomplete). Ninety-two participants were excluded for failing more than 20% of our attention checks, giving impossible responses according to the rules, or spending more than 87 minutes (maximum time set by Prolific). This set of exclusion criteria left us with 211 participants ($Age = 30.1$, 95 females) for analysis and modeling. The experimental protocol was approved by the local IRB.

Procedure

After accepting our ad on Prolific, participants proceeded to a new browser window. They were told they were about to play a card game called “Aces and Eights” with two computer agents, Amy and Ben. They would play 10 such games, and their main goal was to win as many games as possible.

An Aces and Eights game is a 3-player card game using 4 Aces and 4 Eights. The cards are shuffled and each player is dealt 2 cards. The remaining 2 cards are put facing down on the “table”. Each player cannot see her own cards but the other 2 players can. Then, following a predetermined order, the players take turns announcing whether they know their cards. Because information can be obtained from knowing whether other players know, at least one player should eventually know her own cards no matter how the cards were dealt. In the context of our task, Amy, Ben, and the participant were the players. The participant won a game if all her announcements in the game were correct, i.e., she announced “I don’t know” when there was not enough information to know her

cards and announced “I know” while correctly reporting her cards when there was enough information. Since Amy and Ben were computer agents, they would always announce correctly, and participants knew this.

The task had *three phases*: practice, main games, and post-games survey. First, the participants read detailed instructions about how to play the game and then completed 3 practice games similar but different from the ones in the main-game phase. After each practice game, explanations were provided to make sure participants understood the objectives. The 10 main games followed the practice phase, and following each game was a feedback page where the participants learned whether they had won the previous game and responded to an attention check question. Participants earned \$0.5 bonus for each game they won. After finishing all 10 games, the participants were asked to type a whole number between 1 and 100 and told that whoever’s number is closest to half of the average number across all participants would receive an additional \$5. This is the classic “p-beauty contest” task (Camerer, 2006), included to check the construct validity of our task. The experiment ended with a demographic survey.

In what follows, we call one Aces and Eights game a **game**, each rotation of announcement-making by all 3 players a **round**, the order of player announcements the **order** of a game, the assignment of cards to players the **state** of a game, and the state-order pair a **configuration** of a game. We denote each state by a string of 6 letters where the first two letters encode the participant’s cards, the second two encode Amy’s, and the last two encode Ben’s. Within the two letters, “A” always precedes “8” by our convention. For example, AAA888 is the state where the participant holds AA, Amy holds A8, and Ben holds 88. When designing the experiment, we considered some orders and states to be *equivalent* to each other. For example, the order “Amy, Participant, Ben” is equivalent to “Ben, Participant, Amy,” assuming no effect of the computer agents’ names. We stipulated that Amy always announces after the participant, resulting in 3 distinct




Figure 2: Initial epistemic structure. Nodes represent game states and edges indicate a player’s uncertainty. States with the same color are states considered equivalent in our task.

orders denoted by the acronyms “ABP”, “BPA”, and “PAB”. Also note that “A” and “8” are dummy card symbols, so swapping them does not change what announcements are correct. For example, AAA888 is equivalent to 88A8AA. For each pair of equivalent states (having the same color in Figure 2), we randomly selected one for a participant to play.¹

Modeling

The Epistemic Structure

The knowledge states of ideal players in an Aces and Eights game can be determined by the epistemic structure in Figure 2. The nodes represent game states. The edges labeled for a player indicate the player’s uncertainty between states, as previously explained. Suppose the actual game state is AAAA88, and the announcement order is ABP. For Round 1, Amy should announce “I don’t know my cards” because the AAAA88 node has outgoing A edges connecting to AA8888 and AAAA88, indicating that Amy cannot rule them out. After Amy announces “I don’t know”, the graph is updated by eliminating all nodes that do not have any outgoing A edges, namely 88AA88 and AA88AA, because Amy would have known her cards had those states been the actual state. Then Ben’s correct announcement is determined in the updated structure in the same fashion as for Amy; the structure is further reduced by eliminating all nodes incompatible with Ben’s announcement; and finally the participant’s correct announcement is computed in the same way.

Though updating the epistemic structure tells us how ideal players play, it is not a plausible model for how humans reason in this game, because they do not always announce correctly. We therefore designed and compared four computational models to better account for human performance.

¹Data and code for this work are available at <https://github.com/HuangHam/CogSci2021-Epistemic-Reasoning>.

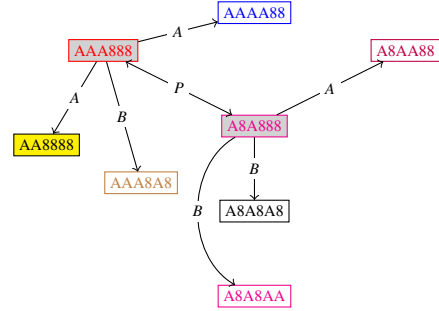


Figure 3: The bounded epistemic structure for the game state AAA888 (red) with $\ell = 1$. Interior states are shaded in gray.

Computational Models

Noisy DEL As our baseline, this model is based on DEL, introduced above, plus a *noise* $\in [0, 1]$ parameter. For each response, the model has a *noise* probability of guessing and a $1 - \text{noise}$ chance of answering correctly according to DEL.

Stochastic update with epistemic bound (SUWEB) This is the target model we propose. The three parameters are *epistemic level* $\in \{0, 1, 2, 3, 4\}$, *update prob* $\in [0, 1]$, and *noise*.

The *epistemic level* (ℓ) determines a bounded version of the epistemic structure in Figure 2 for a given game state w . This is used to compute the SUWEB model’s predictions of how the participant will respond to questions. To construct the structure, given w , let $R(w)$ be the set containing w and all states reachable from w in 1 step along P edges. Thus, $R(w)$ is the set of the participant’s initial epistemic possibilities, i.e., the states she initially considers possible based on what she sees at the beginning of the game. Because the participant may also reason about what other players know (if $\ell > 0$), next we add the states that the other players consider possible starting from states in $R(w)$: for $k \geq 0$, let $R_k(w)$ be the set of all states reachable from $R(w)$ in at most k steps from states in $R(w)$ along edges labeled for *any agent* (we stipulate that $R_{-1}(w) = R_0(w) = R(w)$). The *bounded epistemic structure* is obtained from the graph in Figure 2 by restricting the set of nodes to $R_\ell(w)$ and retaining only those directed edges whose source belongs to $R_{\ell-1}(w)$. We call $R_{\ell-1}(w)$ the set of *interior states* and $R_\ell(w) \setminus R_{\ell-1}(w)$ the set of *peripheral states*. Peripheral states have no outgoing edges in the bounded epistemic structure. We think of these states as being silent about what players know or don’t know, providing information only about the assignment of the player’s cards. Yet the presence of peripheral states in the structure helps determine the extent of the players’ knowledge in interior states. Figure 3 shows an example of a bounded epistemic structure.

When a player announces “I know my cards” (resp. “I don’t know my cards”) we formalize the announced proposition as *the set of interior states in which the player knows* (resp. *does not know*) *their cards*. For each interior state w that is not in that set, we delete w (and thus all edges connecting to w) with

probability *update_prob*. The result of this stochastic deletion of states is the *updated epistemic structure*. Define $R'(w)$ in the same way as $R(w)$ above but in the updated structure, so $R'(w)$ is the participant's new set of epistemic possibilities after the announcement. The participant announces she knows her cards if $R'(w)$ is a singleton set, which means she has no uncertainty. We repeat this updating process if there is another announcement before the participant's turn.

The *noise* parameter aims to capture participants' guessing behavior. If the participant's current set of epistemic possibilities is a singleton, she never guesses; if it is empty, she guesses "know" or "don't know" with equal probability, and when asked what her cards are, she randomly chooses a state from the initial $R(w)$; otherwise she guesses "know" with probability *noise* and reports her cards by randomly choosing a state from her current set of epistemic possibilities.

As an example of the SUWEB model in action, suppose we start with the game state AAA888 (red in Figure 3), the announcement order BPA, $\ell = 1$, *update_prob* = 1, and *noise* = 0. First, Ben will announce "I don't know," due to his uncertainty between AAA888 and AAA8A8 (brown). If the participant were a perfect logician, she would delete AAAA88 (blue) in response to Ben's announcement, since in the full epistemic structure in Figure 2, Ben knows his cards in AAAA88. However, in the bounded epistemic structure in Figure 3, AAAA88 is a peripheral state, so it cannot be deleted. Thus, while the perfect logician gains information from Ben announcing "I don't know," the level 1 player does not. By contrast, in a bounded epistemic structure with $\ell = 2$, AAAA88 is an *interior* state, so it is deleted in response to Ben's announcement. Thus, a level 2 player also gains information from Ben's announcement. Of course, if *update_prob* < 1, then even a level 2 player may fail to delete AAAA88.

After Ben's announcement, the participant will announce "I don't know." At this point, Amy knows her cards (since she operates as a perfect DEL agent) and announces "I know." This contradicts what the participant believed about Amy, since in both interior states of the bounded epistemic structure, Amy does not know her cards. This inconsistency is modelled by both of the participant's interior states being deleted (assuming *update_prob* = 1). When the participant is next asked to make an announcement, she has to guess. By contrast, a level 2 player can solve this game correctly.

Stochastic intake with epistemic bound (SIWEB) This model differs from SUWEB only in the mechanism of updating the bounded epistemic structure. After each announcement, SIWEB either ignores the announcement or reacts to it. The probability of reaction is a parameter *intake_prob*. If it does react, it updates like SUWEB with *update_prob* = 1.

Stochastic update with no epistemic bound (SUWNB) This model is obtained by removing the *epistemic_level* parameter from the SUWEB model. In other words, it assumes that everyone starts a game with the full epistemic structure.

Results

On average out of 10 games, the participants won 4.16 games and spent 28.19 minutes completing the experiment. The average of numbers guessed by the participants in the p-beauty contest was 32.39. Moreover, the performance of the Aces and Eights task was negatively correlated with the number guessed in the p-beauty contest, where a lower number suggests greater capability for higher-order reasoning (Camerer, 2006), supporting the construct validity of our task ($r = -.27, p < 10^{-4}$). All computational models were fitted using maximum likelihood estimation, by minimizing the negative log-likelihood using the SciPy function *minimize* with 20 random starting points from the parameter space. Due to the size of the combinatorial space, the likelihood of SUWEB and SUWNB models was estimated through samples obtained by simulating the model 200 times. We were able to construct Bayesian networks to calculate the exact likelihood of Noisy DEL and SIWEB.

To evaluate how well our models capture human performance, we plotted the predicted accuracy against participant accuracy for all game configurations except those with states filled by yellow in Figure 2. Because participants who learned the rules would know their cards immediately in those states, our exclusion criteria dropped everyone who did not. The baseline Noisy DEL model does not correlate well with human performance; among the other three models, SUWEB fits the data best, as suggested by RMSE, and best explains the variance in performance across games, as suggested by the R^2 adjusted for model complexity (Figure 4).

We further compared the models using average AIC and BIC (Akaike and Bayesian information criteria). While AIC clearly finds SUWEB the best fitting model, BIC does not find SUWEB significantly better than SUWNB (Figure 5a). For model validation, we divided all 30 game configurations into 5 conditions according to the epistemic level required. The epistemic level required by a game configuration is the smallest ℓ_r such that for all $\ell \geq \ell_r$, the SUWEB model with *epistemic_level* = ℓ , *update_prob* = 1, and *noise* = 0 can guarantee winning. We then plotted the average accuracy adjusted for the difference in expected accuracy of pure guessing for each condition (Figure 5b). We observe that for humans, the accuracy of games that require level 0 or level 1 is significantly higher than games that require higher levels ($ps < 10^{-14}$). However, the accuracy of games that require level 2, 3, and 4 are not significantly different from each other ($ps > 0.08$). By this behavioral measure, our SUWEB model captures human performance reasonably well. In contrast, the SUWNB model fails to validate despite its low BIC.

Finally, we examined how our modeling paradigm compares to the Cognitive Hierarchy Model, which predicts that the distribution of higher-order reasoning levels in a population follows roughly a Poisson distribution with $\lambda = 1.5$ (Camerer, 2006). We plotted the empirical distribution of SUWEB's fitted *epistemic_level* parameters and compared it with a Poisson distribution whose λ is estimated from these




Figure 4: Model performance. Each state in the legend represents its equivalent states. Error bars represent s.e.m.

parameters using maximum likelihood estimation with the R function *optim* (Figure 5c). We observe that the two distributions are shaped similarly and the estimated $\lambda = 1.87$ is close to 1.5, where the difference may be due to a higher quality of our sample. Nonetheless, the SUWEB model predicts many more level 1 reasoners, suggesting a Poisson distribution may only be an approximation of the true population distribution which concentrates much more at level 1.

Discussion

In this paper, we proposed the SUWEB model of higher-order deductive epistemic reasoning and showed that in capturing human behavior it outperforms alternative models that either lack a restriction on epistemic bound (SUWNB) or lack stochasticity in eliminating possibilities (SIWEB). Moreover, SUWEB is capable of validating the trend that people significantly underperform in games that require more than first-order reasoning. These logic-inspired models suggest that the constraint on the epistemic level at which one reasons is a dissociable factor from the mere stochasticity of cognition in explaining why and how much people underperform. In addition to using ideas from logic to inspire our cognitive models, we believe the cognitive models may inspire new directions in logic, such as stochastic versions of DEL for formalizing the probability that a bounded agent will know some proposition after a stochastic update with an announcement.

We identify three limitations of our work. First, the estimate of our models' likelihoods may not have come from a sufficiently large sample, potentially biasing AIC and BIC. In the future, we will explore other model comparison schemes such as Bayesian fitting and cross-validation. Second, SUWEB did not validate human performance on games that require level 1 reasoning as well as SIWEB. We will explore the possibility of improving SUWEB by combining features in SIWEB or incorporating psychological biases such as the conformity effect. The third limitation is the lack of ecological validity. Our experimental paradigm can only examine *deductive* aspects of higher-order reasoning, as the game used had unique correct answers. It would be an inter-

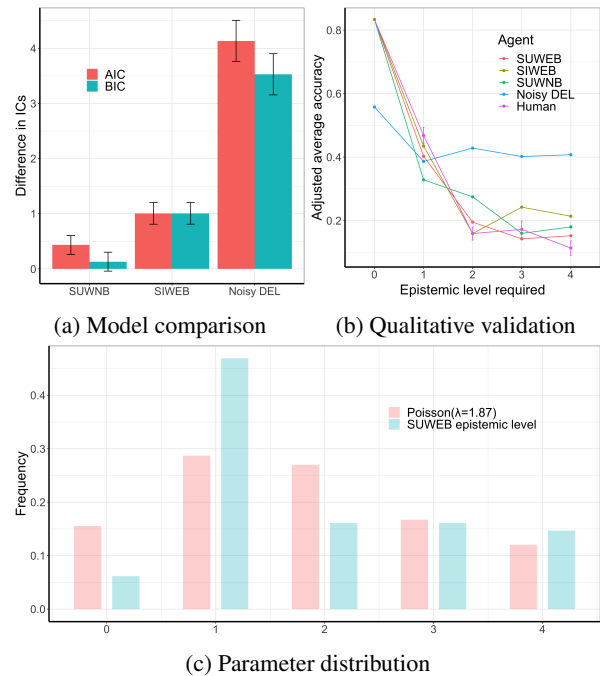


Figure 5: Error bars represent s.e.m. (a) The AIC/BIC score of each model is subtracted by that of SUWEB. (b) Accuracy is adjusted by subtracting the average accuracy of guessing. (c) Poisson samples with values ≥ 4 were revalued as 4.

esting future direction to generalize our framework to cases where participants cannot rely on others' announcements and to connect with other models built for game-theoretic tasks.

More broadly, we hope to explore how our work can expand to other domains of ToM and shed light on how much of social cognition can be explained by non-social cognitive faculties (Apperly, 2010). Also, building on classic work in pragmatics (Clark & Marshall, 1981; Smith, 1982), we hope to explore how higher-order epistemic reasoning, by which mutual and common knowledge is formed, is related to norms of linguistic behaviors such as writing and conversation.

Acknowledgments

We thank Adam Bjorndahl for introducing us to the Aces and Eights game, and we thank Johan van Benthem, Yifeng Ding, Thomas Icard, Steven Piantadosi, the Computation and Language Lab at UC Berkeley, the Jenkins Lab at UPenn, and the anonymous reviewers for helpful comments.

References

- Apperly, I. (2010). *Mindreaders: The cognitive basis of “theory of mind”*. Psychology Press.
- Aumann, R. J. (1999). Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28, 263–300.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 1–10.
- van Benthem, J. (2011). *Logical dynamics of information and interaction*. Cambridge University Press.
- Camerer, C. F. (2006). When does “economic man” dominate social behavior? *Science*, 311(5757), 47–52.
- Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 10–63). Cambridge University Press.
- da Costa, A. O., Saldanha, E.-A. D., Hölldobler, S., & Ragni, M. (2017). A computational logic approach to human syllogistic reasoning. In *Proceedings of the 39th annual meeting of the cognitive science society* (pp. 883–888).
- van Ditmarsch, H., Halpern, J., van der Hoek, W., & Kooi, B. (Eds.). (2015). *Handbook of epistemic logic*. College Publications.
- van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2008). *Dynamic epistemic logic*. Springer.
- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (2003). *Reasoning about knowledge*. MIT Press.
- Frith, C., & Frith, U. (2005). Theory of mind. *Current biology*, 15(17), R644–R645.
- Gopnik, A., & Wellman, H. M. (1992). Why the child’s theory of mind really is a theory. *Mind & Language*, 7(1–2), 145–171.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Cornell University Press.
- Ho, T.-H., Camerer, C., & Weigelt, K. (1998). Iterated dominance and iterated best response in experimental “p-beauty contests”. *The American Economic Review*, 88, 947–969.
- Holliday, W. H. (2018). Epistemic logic and epistemology. In S. O. Hansson & V. F. Hendricks (Eds.), *Introduction to formal philosophy* (pp. 351–369). Springer.
- van Maanen, L., & Verbrugge, R. (2010). A computational model of second-order social reasoning. In *Proceedings of the 10th international conference on cognitive modeling* (pp. 259–264).
- Ong, D. C., Zaki, J., & Goodman, N. D. (2019). Computational models of emotion inference in theory of mind: A review and roadmap. *Topics in Cognitive Science*, 11(2), 338–357.
- Perner, J., & Wimmer, H. (1985). “John thinks that Mary thinks that...” Attribution of second-order beliefs by 5- to 10-year-old children. *Journal of Experimental Child Psychology*, 39(3), 437–471.
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4), 515–526.
- Ragni, M., Eichhorn, C., & Kern-Isberner, G. (2016). Simulating human inferences in the light of new information: A formal analysis. In *Proceedings of the 25th international joint conference on artificial intelligence* (pp. 2604–2610).
- Ragni, M., Kern-Isberner, G., Beierle, C., & Sauerwald, K. (2020). Cognitive logics – features, formalisms, and challenges. In *Proceedings of the 24th european conference on artificial intelligence* (pp. 2931–2932).
- Rendsvig, R., & Symons, J. (2021). Epistemic Logic. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2021 ed.). Metaphysics Research Lab, Stanford University.
- Smith, N. V. (1982). *Mutual knowledge*. Academic Press.
- Szymanik, J., Meijering, B., & Verbrugge, R. (2013). Using intrinsic complexity of turn-taking games to predict participants’ reaction times. In *Proceedings of the 35th annual meeting of the cognitive science society* (pp. 1426–1431).
- Verbrugge, R. (2009). Logic and social cognition: The facts matter, and so do computational models. *Journal of Philosophical Logic*, 38(6), 649–680.
- de Weerd, H., Verbrugge, R., & Verheij, B. (2015). Higher-order theory of mind in the Tacit Communication Game. *Biologically Inspired Cognitive Architectures*, 11, 10–21.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655–684.
- Zhao, B., van de Pol, I., Raijmakers, M., & Szymanik, J. (2018). Predicting cognitive difficulty of the Deductive Mastermind Game with dynamic epistemic logic models. In *Proceedings of the 40th annual meeting of the cognitive science society* (pp. 2789–2794).