You must know something I don't: risky behavior implies privileged information

Anonymous Cog Sci Submission

Abstract

People infer each other's mental states by assuming that agents act rationally to fulfill their desires, given their beliefs. These inferences often reveal that we are pluralistic with regard to preference—we accept that others like things we dislike and vice versa. Here, we propose that people are much more conservative when it comes to risk: we doubt others' risk assessments can differ radically from our own, searching instead for other mental-state explanations. We present two Studies demonstrating that when adults see agents violate utility maximizing behavior, they infer that the agent has privileged knowledge justifying their choices. In Study 1, participants were quicker to imitate an agent who immediately made extremely risky bets than one who started out making low-risk bets that became progressively riskier. In Study 2 we show that participants' choices depend on mental state inferences rather than 'contagious' but mind-blind risk-seeking behavior.

Keywords: theory of mind, risk aversion, rational agent models, social cognition

Introduction

Imagine that you're at a cocktail party and you overhear a hushed conversation where someone is urging their friend to immediately buy a large set of shares from a company whose stock is at an all-time low with no signs of recovery. By just hearing a few words in a brief interaction you can infer that the speaker knows more about this company than you do, that the company's stock value is probably about to soar, and that these two people are engaging in insider trading.

While at first sight these conclusions may feel effortless and even obvious, the reasoning behind them is far from trivial. In principle, we could make sense of the speaker's exhortations by appealing to a simpler explanation: the speaker was simply irrationally risk-seeking. After all, buying stock from failing company will clearly lead to a regretful outcome. Yet, if we allow the speaker a minimal level of competence, explaining their actions as motivated by a preference for extreme risk is unsatisfactory: is there really anyone so unlike ourselves as to deliberately pay enormous costs for no gains? Rather than writing off the speaker as irrational, we wonder if they know something we don't. Why?

When we interpret other people's actions, we do so by appealing to unobserved mental states – beliefs, desires, and intentions. But because we cannot see these mental states, we must infer them by watching how people behave. Though behavioral economists have challenged the long tradition in Western thought that considered humans to be "rational animals" endowed with a "faculty of reason" (Aristotle, *trans.* 2011; Kant, 1787/1929; Tversky & Kahneman, 1981; Kahneman & Tversky 1979; Camerer et al., 1989; Birch & Bloom, 2007), research in social cognition suggests that laypeople do tacitly assume that other agents are rational. In other words, when we watch other people, we assume that they are acting in ways that they believe will help them achieve their goals; to deliberately thwart one's own goals would be "irrational".

Computational work suggests that this expectation for rational action can be formalized as an expectation that agents act to maximize their subjective utilities — the difference between the costs they incur and the rewards they obtain (Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016; Jern, Lucas, & Kemp, 2017; Lucas, et. al., 2014). If, for instance, we see an agent walk past a nearby coffee shop to get coffee at a shop located a few blocks farther away, we can infer that they think the quality of the coffee shops is different. Otherwise, going to the closer coffee shop would have yielded a higher utility. Empirical work suggests that this expectation that agents maximize utilities drives how we infer other people's beliefs (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017), desires (Velez-Ginorio, Siegel, Tenenbaum, & Jara-Ettinger, 2017), preferences (Jern, Lukas & Kemp, 2017; Pesowski, Denison, & Friedman, 2016), and goals (Baker, Saxe, & Tenenbaum, 2009), and that these inferences are already at work from a very young age (Liu, Ullman, Tenenbaum, & Spelke, 2017; Gergely & Csibra, 2003; Bridgers, Jara-Ettinger, & Gweon, 2016).

Critically, to our knowledge, these inferences have been largely studied in contexts where we infer mental states by attending to the costs that agents incur. Yet, in cases like the insider trading example, costs are as yet uncertain. What prompted our suspicion about their knowledge was not the cost associated with buying stock from a failing company, but the risk. Here we propose that, beyond expecting agents to rationally trade off the costs they incur with the rewards they obtain, we also believe that agents' tolerance for risk is far from infinite – and far more limited than we believe the range of qualitative reward preferences to be.

We propose that although risk falls within the same computational framework, it also operates under different principles than costs and rewards. Costs and rewards are subjective. Risk is less so. When we watch other people act, we recognize that their costs (Jara-Ettinger, Schulz, & Tenenbaum, 2015) and rewards (Repacholi & Gopnik, 1994) may be different from our own. Some people enjoy going to the movies while others find the experience to be tedious. Some love durian while others hate it. When we learn that someone else's preferences differ from our own, we simply take this into account when making sense of their actions. When reasoning about risk, however, it is difficult to accept that others see it differently than we do. If someone else appears to think something isn't risky when we think it is, we feel an urge to correct their beliefs. Yet, the more flagrantly they violate our expectations about acceptable levels of risk, the more difficult it becomes to suppose that they're unaware of the risk. If this is the case, we may be more likely to assume that they have privileged knowledge which would make their actions rational (as in the case of insider trading). Crucially, if one agent has inside information that will benefit them, then observers can also benefit by copying their behavior. An observer who makes a strong inference that the agent is knowledgeable may therefore be more likely to take an otherwise unacceptable risk.

In this paper we test this prediction in two studies. Participants completed a gambling task (Figure 1a-b) in which they had an opportunity to take a gamble of varying levels of risk after observing a stranger's risky bet. In Study 1 we manipulate the order in which the risk changes (Figure 1b). In one condition we progressively increase the risk to extreme levels, gradually making it less likely that the agent merely assesses risk differently that the participants. In the second condition we decrease the risk from extreme to more acceptable levels, making it implausible at the outset that the agent merely assesses risk differently that the participants. In Study 2 we test whether seeing the stranger switch to taking the safer bet prompts participants to return to their normal risk-averse behavior.

Study 1

In Study 1, we offered participants a series of choices between two prizes: a 3-token "Safe Bet" and a 4-token "Risky Bet", hidden in opaque boxes (Figure 1a). The absolute value ratio was always 4:3, but the expected value ratio varied from 1:2 to 1:20 in favor of the Safe Bet (Table 1) as a result of the number of boxes a Risky Bettor had to choose between. In two test conditions, the number of boxes increased (Increasing Risk (IR) condition) or decreased (Decreasing Risk (DR) condition) across 14 trials (Figure 1b). Crucially in these conditions, participants saw a that 3rd party bettor (the "Agent") had chosen one of the risky boxes prior to making their own choice. This Agent was described as being in one of two conditions from an earlier study: either they knew the contents of each box, or they were simply betting. If more risky behavior from a 3rd party bettor elicits stronger inferences that the bettor has "inside information", then participants in the IR condition should be increasingly likely to copy the Agent's bet across trials as risk increases. Since participants in the DR condition see the strongest violations of expected value maximization right away, they should begin copying the Agent's bets immediately and continue copying through the Study. Hence, we predict stronger correlations of bet copying with trial number in the IR condition than in the DR condition.

Table 1. Expected values were manipulated over 14 trials by varying the number of boxes to choose from. Expected values always favored the Safe Bet.

	# of Boxes	P(Gain)	E(Bet)	SafeBet Advantage: E(SafeBet) - E(RiskyBet)
SafeBet	1	1.00	$3.00	
RiskyBet	2	0.50	$2.00	$1.00
RiskyBet	3	0.33	$1.33	$1.67
RiskyBet	4	0.25	$1.00	$2.00
RiskyBet	5	0.20	$0.80	$2.20
RiskyBet	10	0.10	$0.40	$2.60
RiskyBet	15	0.07	$0.27	$2.73
RiskyBet	20	0.05	$0.20	$2.80
RiskyBet	25	0.04	$0.16	$2.84
RiskyBet	30	0.03	$0.13	$2.87
RiskyBet	35	0.03	$0.11	$2.89
RiskyBet	40	0.03	$0.10	$2.90
RiskyBet	50	0.02	$0.08	$2.92
RiskyBet	60	0.02	$0.07	$2.93
RiskyBet	80	0.01	$0.05	$2.95

Method

Participants. 120 participants were recruited through Amazon's Mechanical Turk platform (n=40 per condition). One participant was not included in the study

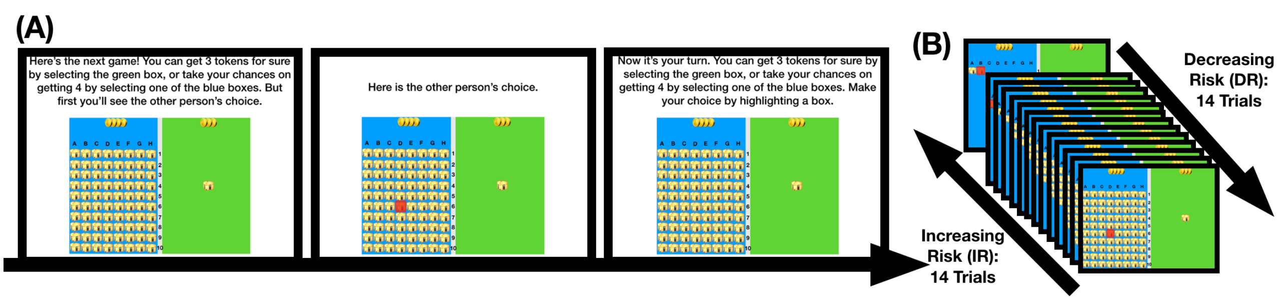

Figure 1a-b: (a) Example of a trial (1:80 odds). Participants first saw the two options: a guaranteed 3 tokens in the box in green, a choice between 1 of 80 boxes, one of which contained 4 tokens. Next, they saw the Agent's choice. Finally, they made their own choice. (b) Participants each saw 14 trials in one of two orders.

due to a server error. Ten additional respondents were recruited but not allowed to participate because they failed to pass comprehension checks three times in a row.

Procedure. Participants were told they would be presented with a series of choices between two bets: a risky bet, and a safe bet. Participants saw an array of boxes (Figure 1a) on blue fields and green fields (called the "Risky Bet" and the "Safe Bet", respectively). The Safe Bet was always a single box said to contain a guaranteed bonus payment. The number of boxes in the Risky Bet field varied across trials from 2 to 80 boxes (Figure 1b), and participants were told that one of the boxes contained a bonus payment, but the others were empty. The absolute values of each bet were announced to participants. Participants were told that their final bonus would be the sum of all of their winnings, but that they would not see the outcomes of their bets. Next, participants were told that they would be paired with a person from an earlier experiment for the duration of the study, and would see which box this person had chosen on each trial. This person was described as either *knowing the contents of each box*, or *just betting*. After answering 6 questions to confirm they had read and understood the instructions, participants proceeded to the task.

Participants completed 14 trials. In the test conditions, they first saw the bets they were being offered, then saw the Agent's choice highlighted in red, and finally made their choice. The Agent always chose a box from the Risky Bet array, following the higher absolute value. In the IR condition, each trial offered a Risky Bet of increasingly poor odds (from 1 in 2 to 1 in 80). In the DR condition, the reverse was true. In a control condition, participants saw the bets they were being offered and made their choice *before* seeing the Agent's choice; hence, they could not take advantage of the Agent's knowledge even if they inferred them to be knowledgeable. Participants in all conditions made their choice by clicking directly on the box.

Finally, participants rated the likelihood that the Agent *knew* what was in the box on a scale of 1 to 9. Participants indicating 5 or higher were also asked to indicate the trial during which they first began to suspect that the Agent was "*knew*" which box to choose rather than just guessing".

We predicted that participants in the two test conditions would take significantly more Risky Bets than those in the control condition. More specifically, we predicted a stronger correlation of trial number with over-betting for the IR condition than DR condition. Over-betting was defined as the difference between the observed rate participants choosing the same bet as the Agent and the rate expected by chance on each trial. Additionally, we predicted that participants in all conditions would think it more likely that the Agent knew which box choose, and that participants in the DR condition would report believing this earlier than participants in the IR condition and the Control condition (in which risk also increased over trials).

Results and Discussion

Figure 2 shows the results from Study 1. In total, 34.4% of bets in the IR condition were the same as the Agent's, and only 3.6% of bets were *New* Risky bets; while in the DR condition, 31.4% of bets were the same as the Agent's, and only 7.3% were *New* Risky bets. To test our main predictions, we calculated an Over-Betting Index (OBI) for each trial in each condition). OBI was defined as the difference between the proportion of participants copying the Agent's bet on each trial and the proportion that would be expected on that trial due to chance if everyone took the Risky Bet. We then calculated the correlation between OBI and trial number. In the IR condition, the Agent's choices grew increasingly irrational across trials from a naive perspective (Spearman's ρ =.947, p<.001), while in the DR condition, they started out maximally irrational and remained so (Spearman's ρ =.176, p=.547). A permutation test (n=10,000) confirmed that these correlations were significantly different, p=.003.

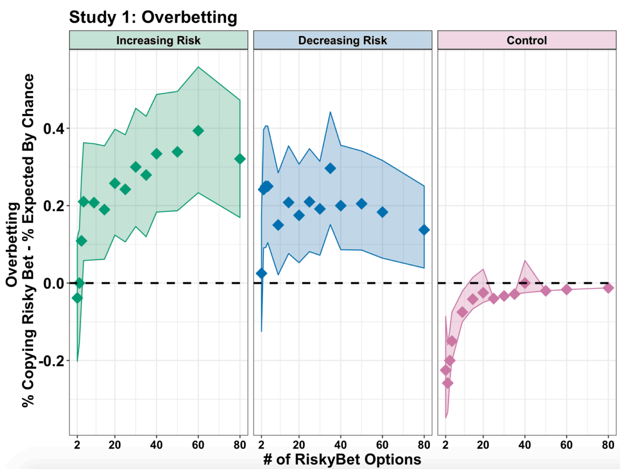

Figure 2: Over-betting in Study 1. Participants did not immediately copy the Agent's relatively low-risk bets in the first trial, but grew increasingly likely to do so (green). In contrast, when the Agent's first bets were already extremely risky (blue), participants immediately began to copy, but did not grow more likely to do so. When the Agent bet after the participant, participants rarely bet at all (purple).

As predicted, participants in all conditions thought it more likely that the Agent knew the location of the bonus than that the Agent was guessing (M_{IR} = 5.87, M_{DR} = 5.38, $M_{Control}$ = 6.02). This belief did not differ between conditions (F(2,116)=0.817, 77 iterations, p=.714, using the aovp function from the 'lmPerm' package in R), but did differ significantly from chance, p<.001 (n=10,000 permutations using the one sample one-tailed permutation test from the 'jmuOutlier' package in R).

Contrary to our prediction, the participants in the DR condition did not report suspecting that the Agent knew the location of the rewards earlier than in the IR or control conditions ($F(2,86)=.518$, 56 iterations, $p=.7321$, using the aovp function from the 'lmPerm' package in R).

In the control condition, participants were unable to copy the Agent's choices. Unlike in the test conditions, these participants were hesitant to make any bets at all. 91.8% of bets in the control condition were Safe, 3.9% were Same, and 4.3% were New Risky Bets. This pattern of betting differed significantly from the patterns in the test conditions (χ^2 (4) = 191.54, $p<.001$); of particular interest is that very few risky bets were made in the control condition, while almost all risky bets in the test conditions were copies of the Agent's bet rather than new risks.

The rate at which participants in the test conditions over-bet is remarkable, considering that they were given no outcome information about their bets, uncertain information about the Agent, and the expected values of copying an ignorant Agent's bet would predict a net loss. The overwhelming preference for the Safe Bet in the control condition underlines the extent to which the participants in the test conditions diverged from normal betting behavior.

Study 2

To further confirm that the participants were inferring Agent knowledge on the basis of otherwise irrational behavior, we conducted a second study in which the Agent *stopped* taking the Risky Bet part way through the task. If participants in Study 1 were indeed betting merely on the strength of their inference that the Agent *knew* the locations of the Risky Bet rewards, then seeing the Agent make choices that suggest that they are *not* knowledgeable should have a dramatic effect on participants' willingness to take a Risky Bet: as soon as the Agent stops taking risks, so should the participants.

Method

Participants. We recruited 80 participants through MTurk (n=40 per condition). 9 additional respondents who failed to the comprehension screening were not given access to the Study.

Procedure. The methods for Study 2 were identical to the Increasing Risk condition in Study 1, with the following exception: the Agent "switched" from taking the Risky Bet partway through the 14 trials. In the Early Switch condition, the Agent started taking the Safe Bet on trial 7; in the Late Switch condition, they switched on trial 11. We therefore predicted that participants in the Late Switch condition would stop copying the Agent's bets later than the participants in the Early Switch condition. Additionally, after being asked whether they believed the Agent knew the location of the rewards, any participant who had made at least one risky bet was asked whether they "wished they had taken the risky bet less often". They responded on a scale of 1 ("No, I'm happy with my bets") to 9 ("Yes, I wish I'd risked less"). If participants copied the Agent's bet only because they believed the Agent to be knowledgeable, then they should regret taking the risky bet more often; moreover, participants in the Late Switch condition should regret taking the risky bet more than those in the Early Switch condition, who had taken fewer overall risky bets.

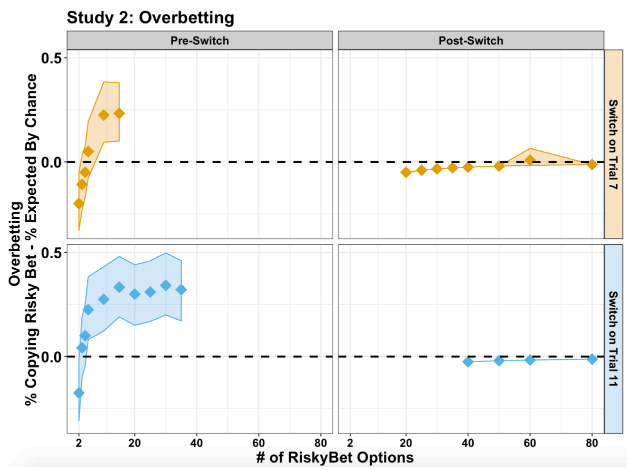

Figure 3: Over-betting in Study 2. Ribbons represent 95% CIs. Following the Agent, participants took increasingly risky bets until the Agent abruptly switched to Safe Bets on trial 7 (yellow) and 11 (blue), respectively.

Results and Discussion

Figure 3 shows the results from Study 2. In all, 62.5% of participants took at least one Risky Bet (either the same as the Agent, or a new Risky Bet). Before the Agent switched to taking the Safe Bet, 26.7% of bets in the Early Switch condition were the same as the Agent's and only 5.0% of bets were *new* Risky Bets; while in the Late Switch condition, 36.8% of bets were the same as the Agent's, and only 3.8% were *new* Risky Bets. After the Agent switched to taking Safe Bets, nearly all further choices were Safe Bets: 95.3% of bets in the Early Switch condition, and 96.3% of bets in the Late Switch condition. Moreover, the OBI of the Pre-Switch trials in Study 2 strongly correlated with the OBI of the same trials in the IR condition of Study 1, Spearman's ρ =. 939, $p<.001$).

To test our main prediction for Study 2, we conducted a Mann-Whitney U test comparing the median switch-point of participants who took at least one Risky Bet (n=50). The switch-point was defined as the last trial on which they chose the Same Bet or a new Risky Bet. The median switch-point in the Early Switch condition was trial 6, while the median switch-point in the Late Switch condition was trial 10 (W= 122, p<.001), confirming our prediction that participants would interpret the Agent's switch as a sign of ignorance, and switch to taking the Safe Bet as soon as the Agent did. Indeed, trials 6 and 10 were the participants' last opportunities to copy, in their respective conditions; they reverted to risk-averse behavior instantly and nearly without exception.

As predicted, participants in Study 2 considered it unlikely that the Agent knew the location of the rewards, M_{Early} = 2.82, M_{Late}=3.6, p<.001 (n=10,000 permutations using the one sample one-tailed permutation test from the 'jmuOutlier' package in R).

However, contrary to predictions, participants in the Late Switch condition did *not* report wishing they had taken the Risky Bet less often, in the Late Switch condition than in the Early Switch condition (M_{Late} = 4.11, M_{Early} = 4.04, (F(2,54)=.008, 51 iterations, p=.941, using the aovp function from the 'lmPerm' package in R), and actually appeared to regret their decisions *less* than would be expected by chance on a scale of 1 to 9, p = .0237 (n=10,000 permutations using the one sample two-sided permutation test from the 'jmuOutlier' package in R).

Study 2 suggests that people were indeed inferring at first that the Agent was knowledgeable, and were willing to base their own choices on this conjecture alone. When provided with evidence that they were mistaken, participants quickly changed their minds. Why they did not regret their choices is an interesting question for further research.

General Discussion

Human behavior can often seem irrational; despite some broadly consistent patterns, differences in specific and general abilities, access to information, preferences, and resources result in widely varying behavioral responses to the same stimulus. To navigate the social world, agents must be able to make sense of behaviors that seem irrational at first glance. One way to account for irrational behaviors is to assume that other agents have a different utility function; they may prefer risk to certainty, small rewards to large, or have qualitative preferences that radically differ from our own. Another is to assume other agents are entirely irrational — that their choices are not goal driven, or are intentionally chaotic. We find that people reject both of these possibilities when reasoning about risk. Instead, they appear to infer that an "irrational" agent shares their preferences for high-reward/low-risk choices, but is simply more knowledgeable than they are. Moreover, our participants in both studies accepted considerable risk by adjusting their betting strategies to benefit from this conjecture, though they had no evidence that the new strategy was more profitable, beyond their belief in the Agent's rationality. Importantly, it appears that the more irrational the Agent's behavior, the more participants believe the Agent must in fact be knowledgeable. However, participants' trust is not blind; when in Study 2 the potentially knowledgeable Agent's behavior suddenly began to suggest that they were merely exceptionally risk-seeking rather than knowledgeable, participants immediately reinterpreted the Agent's earlier behavior, and reverted to their risk-averse betting.

Our studies have several limitations. Firstly, the inference that the Agent was knowledgeable was not spontaneous: we explicitly stated that the Agent could know the location of the Risky Bet reward. While this likely promoted the inference that the Agent was knowledgeable, our results cannot be attributed to task demands: if participants had merely assumed the 3rd party was knowledgeable on the basis of that instruction, no difference in correlation strength between the DR and IR conditions would have been found. We are currently conducting studies to show that that similar tasks can prompt spontaneous inferences of 3rd party knowledge. Secondly, the Agent's betting strategy may have been implausible for reasons other than its unusual riskiness; expected values for all 14 trials favored the Safe Bet, and in Study 1, the Agent's choice for all 14 trials was the Risky Bet, as well as 6 or 10 of 14 in Study 2. This one-dimensional betting strategy may have itself served as a cue that they came from the "knowledgeable" group mentioned in the instructions. Anecdotally, consider the how suspicious people are of card counters or the game hackers who repeatedly win against long odds; while in our study outcome information was not available, the willingness to repeatedly take enormous risk may be suspicious enough. Further work will examine people's ability to infer knowledge when the agent adopts a more complex betting strategy. Finally, we tested people's intuitions in a betting game where risk was apparent and the context may have made salient the possibility of 'cheaters', regardless of their strategy, but specific content-based cues may overcome the effects documented here in other contexts. For example, we are unlikely to infer that reckless teenagers or haggard gambling addicts are in fact more knowledgeable than we are simply because they are taking outlandish risks. Mental state inferences incorporate context to account for the enormous variety of potential mental states, and can just as easily strengthen a basic bias as weaken it. However, given no outcome information and little context, our results suggest that people may be hard pressed to motivate exceptionally risky behavior beyond positing privileged knowledge.

Our interpretation of these results is also consistent with recent work suggesting that people infer speaker knowledge from over-specification (Rubio-Fernandez, 2017; Keysar et al., 2003). The Maxim of Quantity (Grice, 1975), dictates that a speaker should be as informative as necessary, and *no more* informative. Grice's maxims can be understood as a utility function dictating costs (utterance length) and rewards (being understood), and violating these maxims can be semantically significant. In Rubio-Fernandez's (2017) task, two naive participants played a referential communication game in which players' visual access to objects differed. When the speaker violated the maxim of quantity (e.g., by saying "blue fish" when "fish" would have been sufficient), 76% of listeners reported suspecting that the speaker could see the orange fish that was allegedly visible only to the listener. While there is no risk involved in Rubio-Fernandez's (2017) task, the results illustrate another way in which people may infer privileged knowledge when at loss to rationally motivate other's behavior.

In two studies, we showed that people reinterpret "irrational" risk-seeking to preserve agents' rationality, at potentially high cost to themselves. The more irrational the behavior appears, the more people strive to rationalize it by assuming that the agent has access to information that they themselves do not have. However, they do not simply abandon their prior rational preferences; given evidence that the agent is does *not* have inside information, our participants quickly reverted to risk-averse betting. These results suggest that while our intuitive psychology is flexible enough to accommodate pluralistic preferences, strong deviations from standard norms are taken to be indicative of differences in information access rather than reflective of true preferences.

Acknowledgments

Removed for blind review.

References

Aristotle, (2011) *Nicomachean ethics*. (J. Sachs, Trans.). Hackett Publishing. (Original work published 350 B.C.E).

Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 0064.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3),329-349.

Birch, S. A., & Bloom, P. (2004). Understanding children's and adults' limitations in mental state reasoning. *Trends in cognitive sciences*, *8*(6), 255-260.

Birch, S., & Bloom, P. (2007). The Curse of Knowledge in Reasoning about False Beliefs. *Psychological Science*, *18*(5), 382–386.

Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2016). Children consider others' expected costs and rewards when deciding what to teach. *In Proceedings of the 38th annual conference of the cognitive science society* (pp. 559-564). Cognitive Science Society.

Camerer, C., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An Studyal analysis. *Journal of political Economy*, *97*(5), 1232-1254.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998-998.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naıve theory of rational action. *Trends in cognitive sciences*, 7(7), 287-292.

Grice, H. P. (1975). *Logic and conversation.* 1975, 41-58.

Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335-341.

Kahneman, D., & Tversky, A. (1979). "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47, 263-291.

Keysar, B., Barr, D. J., Balin, J. A., & Brauner, J. S. (2000). Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science*, 11(1), 32-38.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The Naïve Utility Calculus: Computational Principles Underlying Commonsense Psychology. *Trends in Cognitive Sciences*, *20*(8), 589–604.

Jern, A., Lucas, C. G., & Kemp, C. (2017). People learn other people's preferences through inverse decision-making. *Cognition*, 168, 46-64.

Kant, I. (1934) *Critique of Pure Reason.* (N. Kemp-Smith, Trans.). London Macmillan. (Original work published 1787).

Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, 358(6366), 1038-1041.

Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., ... & Hu, J. (2014). The child as econometrician: A rational model of preference understanding in children*. PloS one*, 9(3), e92160.

Pesowski, M. L., Denison, S., & Friedman, O. (2016). Young children infer preferences from a single action, but not if it is constrained. *Cognition*, 155, 168-175.

Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: evidence from 14-and 18-month-olds. *Developmental psychology*, 33(1), 12.

Rubio-Fernández, P. (2017). The director task: A test of Theory-of-Mind use or selective attention?. *Psychonomic bulletin & review*, 24(4), 1121-1128.

Tversky, A., & Kahneman, D. (1981). *Evidential impact of base rates* (No. TR-4). Stanford Univ Ca Dept of psychology.

Velez-Ginorio, J., Siegel, M., Tenenbaum, J. B., & Jara-Ettinger, J. (2017). Interpreting actions by attributing compositional desires.(2017).