UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
How Long Should That Take? Reading Minds in Real Time from Decision Speed

Permalink
https://escholarship.org/uc/item/2hh0s5fs

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors
Morris, Ben
Jara-Ettinger, Julian

Publication Date
2026

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

How Long *Should* That Take? Reading Minds in Real Time from Decision Speed

Ben Morris (benjamin.morris@yale.edu)[1] & Julian Jara-Ettinger (julian.jara-ettinger@yale.edu)[1]

[1]Department of Psychology, Yale University

Abstract

Human decision making is richly structured: choices reflect not only what option is selected, but how efficiently evidence is evaluated and conflict is resolved. Do observers represent this structure in others' minds? We investigate how adults use others' decision time to infer both properties of the decision itself (e.g., difficulty) and properties of the decision-maker (e.g., preferences). Across two experiments (n = 200), participants observed agents making simple two-alternative reward decisions that varied in expected difficulty, decision time, and chosen option. We find that observers interpret timing relative to structured expectations about how long a decision should take, using deviations from these expectations to draw nuanced inferences about the state of the decision problem under uncertainty (Exp 1) and about an agent's preferences (Exp 2). Crucially, identical choices supported different conclusions depending on how long they took. These results suggest that people invert expectations about how decisions unfold to read others' minds in real time, treating decision time as a cue to conflict, uncertainty, and value.

Keywords: theory of mind, decision-making, social cognition

Introduction

Imagine someone on a game show choosing between two gift cards: $100 for Store A, or $3 for Store B. No-brainer: A. But what if they pause, glancing back and forth, muttering "hmm…", before finally picking Store A? That same choice feels surprising with hesitation, even suspicious. Maybe they dislike Store A *that* much? Now imagine the choice is $52 vs. $48. A pause here feels natural, the tradeoff is tighter, and the decision genuinely harder.

These intuitions reflect a core feature of human decision-making: people are motivated to make good decisions. Even infants readily approach a higher value option in simple choice contexts (e.g., Feigenson et al., 2002). As social creatures, this is something that people naturally understand about each other. Using this expectation, people can reconstruct other people's preferences based on what they choose (e.g., Jern et al., 2017; Jara-Ettinger et al., 2016). Even young children possess intuitive theories of choice, with expectations about what rational agents should do given the features of a decision problem, such as relative rewards and costs (e.g., Jara-Ettinger et al., 2016; Liu et al., 2017). This basic capacity to infer other people's unobservable desires based on their observed choices is foundational for people's ability to quickly understand and predict each other's behavior.

However, the opening example shows more than that. It highlights another core feature of decision-making: good decisions take time. Moreover, the time it takes us to make the right choice increases as the problem gets more difficult (e.g., Ratcliff & McKoon, 2008). As a result, brief pauses or hesitations can be a powerful signal about what is happening inside someone's mind. A snap decision can signal an easy decision or strong preference; hesitation can suggest uncertainty, ambivalence, or conflict. Crucially, however, the meaning of a delay depends on how long the decision was expected to take in the first place.

Classical theories of social cognition largely emphasized how observers interpret overt actions and outcomes (e.g., Gergely & Csibra, 2003). But a growing body of recent work suggests that people also reason about how behavior unfolds over time. Observers use pauses, hesitation, response delays, and other temporal dynamics to infer hidden mental states including preferences, confidence, knowledge, and uncertainty (Bavard et al., 2024; Berke et al., 2023; Frydman & Krajbich, 2021; Morris et al., 2026; Richardson & Keil, 2022; Ziano & Wang, 2021). Particularly relevant, Gates et al. (2021) developed a rational model showing how observers can use response times to infer others' preferences.

Yet decisions themselves are structured: some choices are computationally easy, whereas others involve closer competition, greater conflict, or noisier evidence. As a result, the same hesitation can support very different inferences depending on how difficult a decision was expected to be. Existing work has often treated slower decisions as a relatively direct cue to weaker preferences, uncertainty, or conflict. Less clear is whether observers interpret decision time relative to structured expectations about how a decision should unfold.

Interpreting decision time requires solving an inferential puzzle. A hesitation could reflect something about the decision: perhaps the options were close, or the stakes unclear. Alternatively, it could reflect properties of the

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society* ©2026 the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

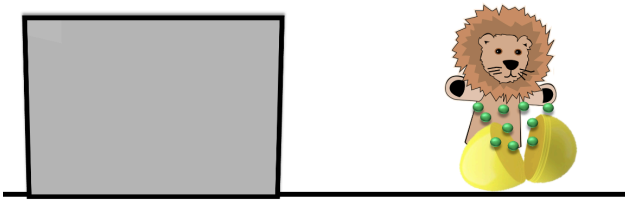

Figure 1. A still image from Experiment 1 showing the basic trial setup where "Lion" is selecting between two eggs, one hidden behind the wall (left) and one visible (right) in this case with 10 candies. In this example, Lion has selected the visible option.

agent: perhaps they had weak preferences, inner conflict, or low confidence. To use decision time effectively, observers must reason jointly about the structure of the world and the mind navigating it.

In the present studies, we test how nuanced this intuitive theory of choice dynamics is. We ask whether observers use timing cues relative to expectations about decision difficulty to infer both (a) the state of the world (i.e. the structure of a given choice), and (b) the states of an agent (i.e. the preferences of the decision-maker). By holding choices constant while varying deliberation time and expected difficulty, we test whether observers attribute hesitation to features of the problem, the agent, or their interaction. Together, these experiments test the nuance of people's representations of others' decision-making minds, revealing how expectations about decision difficulty and deliberation dynamics shape social inference.

Importantly, expectations about how long a decision should take are not arbitrary. In numerical comparison tasks, decisions become slower and more error-prone as quantities become harder to discriminate, consistent with Weber's law (Feigenson et al., 2004; Halberda, 2008). We leverage this well-characterized structure to manipulate and formalize decision difficulty in our tasks. Our claim is not that observers explicitly represent ANS computations, but that numerical comparison provides a principled domain in which observers can form systematic expectations about how difficult a choice ought to be under neutral preferences.

Experiment 1: Learning about the World

Experiment 1 tests how fluidly participants can recover the value of an unobserved option by combining multiple sources of partial information about an agent's choice. Participants watched short videos of an agent choosing between two alternatives with different amounts of rewards (see Figure 1). Critically, one alternative was always occluded from participants' own view. Participants were tasked with inferring the magnitude of the hidden reward. We systematically manipulated (i) the visible option's value, (ii) which option the agent ultimately chose, and (iii) the speed of the agent's decision.

We expected that participants would be able to integrate these cues to form graded inferences about the hidden value. In particular, we predicted that observers would infer higher hidden rewards when the agent chose the occluded option despite a strong visible alternative, and that slower decision speeds would be interpreted as evidence of closer competition between options. The hypotheses, design, exclusion criteria, and primary analyses for Experiment 1 were preregistered prior to data collection.

Methods

Participants 100 participants were recruited online through Prolific. Following preregistered criteria, participants for whom more than 20% of trials violated minimal rationality constraints (see Analysis) were excluded and resampled (n = 10). All reported effects are robust to retaining these participants.

Stimuli Each trial consisted of a short animated video depicting an agent ("Lion") choosing between two options (see Figure 1). In each trial, two toy eggs would fall and break open, revealing their contents. Any given egg could contain a reward from 0-20 candies. One egg would always fall behind a wall such that the participant could not observe its value, and the other egg was always visible to participants.

Trials varied across three factors: visible option size, decision speed, and agent's choice. For visible option size, the visible egg could contain 1, 3, 5, 10, 15, or 19 candies. Decision speed was operationalized as the number of apparent "looks" the agent made toward each of the two options prior to choosing, reflecting different degrees of consideration (looking at each option once, twice, or three times). Each look lasted for 1.5 seconds. Although decision speed can in principle be dissociated from looking behavior, in naturalistic decision-making the two are tightly coupled; thus, the number of looks provided a clean, transparent, and ecologically-interpretable proxy for decision time. We return to this distinction in the General Discussion. Lastly, the agent either selected the visible or occluded option. Note that the occluded option was never revealed, and if it was selected the agent would go behind the occluder.

The full stimuli design thus yielded 6 (value levels) x 3 (speed levels) x 2 (choice levels) trial types for a total of 36 possible trials. We did not include six possible trials in which extended deliberation occurred despite trivially small visible quantities (e.g., 1 or 3 candies), as such cases would violate observers' expectations about rational numerical comparison and risk reframing the agent as incompetent.

Procedure Participants were introduced to the game "Egg Drop" and told that they would be watching "Lion" play the game. Participants were told that each egg would have

Figure 2. Predicted value of the hidden option as a function of the value of the visible option for each speed condition (number of looks), separately for the trials where the agent chose the hidden option (left) and the visible option (right).

anywhere from 0 - 20 candies, that Lion liked candies and was trying to get as many candies as possible, and that all the candies would be the same kind (color and flavor).

On each trial, participants watched a short video of the agent's decision. Participants were then asked to estimate the value of the occluded reward ("*How many candies do you think are in the egg behind the wall*?"). Responses were made using a continuous slider ranging from 0 - 20 candies. Participants received no feedback on their decision. Each participant completed 18 trials, with equal numbers of trials for each choice type and each speed level. For each trial type, stimulus values were randomly sampled. Trials were presented in a randomized order.

Results

Exclusions. Following our preregistered analysis plan, we excluded trials that violated the minimal numerical constraint implied by the agent's choice. In the task, all candies were identical, and the agent was explicitly described as attempting to obtain as many candies as possible. Thus, when the agent selected the hidden option, valid responses should indicate that the hidden option contained more candies than the visible option; conversely, when the agent selected the visible option, valid responses should indicate that the hidden option contained fewer candies.

To allow for uncertainty in numerical discrimination, we used a Weber-scaled tolerance derived from a normative ANS model (w = 0.265; Halberda et al., 2008). Using this model, we computed a one-sided tolerance interval corresponding to a 95% confidence bound around the visible quantity. Responses were excluded only if they contradicted the agent's choice beyond this tolerance (e.g., if the agent chose the visible option with 10 candies, participant responses inferring that the hidden option contained substantially more than 10 candies (i.e. beyond the ANS-derived tolerance) were excluded). To avoid over-penalizing small numerosities, a minimum tolerance of 1 candy difference was allowed.

At the participant level, individuals for whom more than 20% of trials violated the above constraint were excluded from all analyses. Using these criteria, 10 participants were excluded and resampled, and data from a further 2.5% of trials were excluded. Critically, the same qualitative pattern was observed in analyses including all participants and trials.

Overall Analyses. To test whether observers use decision time to make inferences about the decision, we fit a linear mixed-effects model predicting participants' inferred value of the hidden option from the agent's decision time (number of looks, treated as a continuous predictor), the agent's choice (visible vs. hidden), and the visible option's value. The model included interactions between decision time and choice, and between decision time and visible value, with a random intercept for participant. Note that for analyses involving speed we removed trials where the visible option was 1 or 3 because speed did not vary for these trials by design.

We observed our key interaction effect between decision speed and choice ($\beta = 1.99$, $p < .001$), such that when agent chose the hidden option, longer decision times were associated with *lower* inferred values of the hidden option, whereas when the agent chose the visible option, longer decision times were associated with *higher* inferred values of the hidden option. In other words, identical increases in decision time led to opposite inferences depending on which option the agent ultimately selected.

In addition, inferred hidden values increased with the value of the visible option ($\beta = 3.40$, $p < .001$), and this effect also interacted with decision speed ($\beta = 0.75$, $p < .001$), such that the influence of visible value on

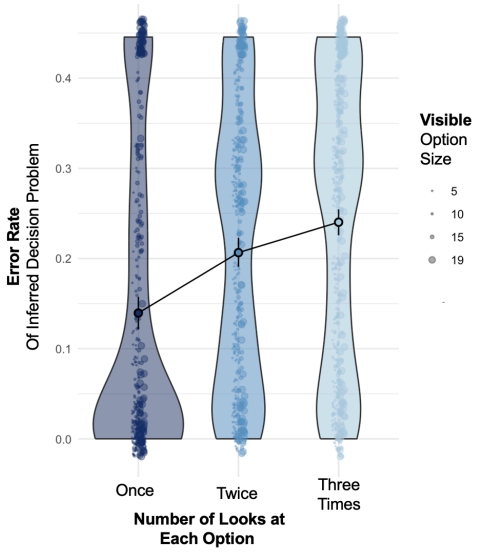

Figure 3. A violin plot showing the expected error rate based on a normative ANS model for a decision of the visible quantity and the participants inferred hidden option, demonstrating that as the number of looks increases participants infer that the underlying choice must have been more difficult. Black point ranges show averages with 95% CIs and colored dots show raw data with jitter for readability.

inferences was stronger for slower decisions. Together, these results indicate that observers integrate decision time with both choice outcomes and decision structure to draw graded inferences about the underlying difficulty of the choice.

ANS results. Are people's inferences well captured by the assumption that increased looking means an increasingly difficult decision problem? As a complementary analysis, we used a normative ANS function described above from prior work to estimate expected difficulty quantitatively (Halberda et al., 2008). The ANS model provides a way to formalize the intuition that numerical comparisons become harder as quantities become more similar (e.g., distinguishing 10 from 11 is harder than distinguishing 10 from 18). For each participant's prediction on a given trial, we generated an expected error value that reflected how hard that comparison should be for an idealized ANS observer. For example, if the visible option for a given trial contained 5 candies, and the participant inferred that the hidden option contained 7 candies, we can compute the expected error rate for that hypothetical comparison problem (about 0.19 in this case).

With this measure, we then calculated the average inferred difficulty as a function of the agent's speed of decision (see Figure 3). A mixed effects model predicting difficulty as function of decision speed (with a random effect of subject) showed the expected main effect of decision speed ($\beta = 0.1$, $p < .001$) indicating that longer decision times led observers to infer comparisons that would be increasingly difficult for an approximate number system. This pattern further supports the idea that observers interpret extended deliberation in terms of structured expectations about numerical comparison difficulty.

Discussion

Experiment 1 demonstrates that observers can infer latent reward values from sparse behavioral evidence, integrating choice outcomes with temporal cues from decision dynamics. Even when one option was fully occluded, participants systematically adjusted their estimates in response to how quickly an agent decided and which option they selected. Consistent with this interpretation, an analysis grounded in a normative Approximate Number System (ANS) model showed that slower decisions led participants to infer comparisons with higher expected difficulty, suggesting that observers interpret extended deliberation relative to structured expectations about how hard a decision should be. These findings provide initial evidence that people treat decision speed as informative about underlying value structure, not merely as noise or idiosyncratic variation. Importantly, these inferences reflect reasoning about the structure of the decision problem; in Experiment 2, we extend this approach to ask whether similar timing cues are used to infer properties of the agent themselves.

Experiment 2: Learning about an Agent

In Experiment 2, we ask whether observers use decision difficulty, choice, and decision speed to reason about agent-internal properties. Specifically, we make both options visible and test how observers attribute choice dynamics to infer an agent's relative preference for the two options. Thus, this task reverses the challenge of Experiment 1: rather than estimating the world from a known agent, participants now estimate an agent from a known world. Observers can use response time as a cue to infer others' preferences, consistent with process-level models of decision-making such as drift–diffusion frameworks (Gates et al., 2021; Morris et al., 2026; Roberts et al., 2011; Roberts & Francis, 2013).

Methods

Experiment 2 used a variation of the 'Egg Drop' game. Each trial showed an agent choosing between two toy eggs, but unlike Experiment 1, both eggs were visible. The two eggs always contained different amounts of candy, which differed in color and flavor. The ratio between the two options, the agent's choice, and the agent's decision time were manipulated. After each decision, participants judged which candy that agent preferred more. The hypotheses, design, exclusion criteria, and primary analyses for Experiment 2 were preregistered prior to data collection.

Participants. Our final sample consisted of 100 participants recruited online through Prolific. Trial- and participant-level exclusions followed the preregistered

criteria. Trials in which participants indicated a strong preference for the *unchosen* option (<= 3 on the recoded scale) were excluded as likely reflecting inattention. Participants for whom more than 20% of trials were excluded were removed from all analyses and resampled. These criteria resulted in excluding and resampling data from 1 additional participant, and a further 1% of trial data. All reported effects are robust to retaining these participants and trials.

Stimuli. Experiment 2 used 8 different trial types of the 'Egg Drop' game. On each trial, two toy eggs would fall and break open, but unlike Experiment 1, both were visible on screen. Additionally, unlike Experiment 1, the two eggs always contained different colored candies (e.g., one egg with blue candies and the other with green). Trials varied on the ratio of the number of candies in each option as an operationalization of decision difficulty (8 vs. 10 candies, so-called 'hard trials' because the two options were close; and 3 vs. 13 candies, so-called 'easy trials' because the two options were not close). Trials also varied on which option the agent chose, either the larger 'value-maximizing' option or the smaller 'non-value-maximizing' option (value- maximizing defined relative to rewards under neutral preferences). And lastly, trials varied the agent's decision time. Experiment 2 used just two speed conditions for simplicity, such that the agent either decided quickly (after looking at each option once) or slowly (after looking at each option three times).

Procedure. Participants were introduced to the "Egg Drop" game as in Experiment 1, but told that both eggs would be visible. Unlike Experiment 1, in Experiment 2 participants were told that each egg would have different kinds of candy: different colors with different flavors. Participants were told that agents were trying to get as many candies as possible, but also might like some candies more than others. On each trial, participants watched a short video of an agent deciding between two options of different amounts. Participants were then asked which candy the agent liked more and asked to respond on a 1-9 scale from 1 - *Definitely the [purple] candy on the left*" to 9 - "*Definitely the [green] candy on the right*" (with 5 indicating "*about the same*"). Participants received no feedback on their decision. Each participant completed a total of 8 such trials, each with a different agent selecting between two different colored candy options. Trial order was fully randomized across participants.

Results

First the scale was recoded such that 1 indicated an extreme preference for the unchosen candy and 9 indicated an extreme preference for the chosen candy. To do so, we first reverse coded scales for any trials where the larger value maximizing option was presented on the left, so that all scales reflected inferences about the value-maximizing vs. non-value-maximizing distinction. We then also reverse scored scales for trials where the agent picked irrationally, yielding comparable scales where the values would reflect increasing preference for the chosen candy. This recoding allowed for simpler interpretation and comparison.

To test the preregistered hypotheses about how observers use decision dynamics to infer agent preferences, we fit a linear mixed-effects model predicting inferred preference strength from agent choice (value-maximizing vs. non–value-maximizing), decision speed (fast vs. slow), and decision difficulty (easy vs. hard), including all two-way interactions and the three-way interaction. The model included a random intercept for participant.

As predicted, there was a significant negative effect of speed, such that slower decisions were interpreted as reflecting weaker preferences overall for the chosen option than faster decisions ($\beta = -0.75$, $p < .001$). Additionally, there was a significant effect of choice, such that choosing the option with less (i.e. the non-value maximizing choice) was interpreted as reflecting *stronger* preferences than value-maximizing choices ($\beta = 0.59$, $p < .001$).

As expected, there was also a significant interaction between choice and decision speed which indicated that decision speed had a larger impact on inferred preference strength when the agent chose the option with more than when the agent chose the option with less ($\beta = -0.47$, $p = .029$). That is, when agents chose the larger, value-maximizing option, pausing may have been unexpected and thus caused a particularly large hit to perceived preference.

In contrast, and unexpectedly, decision difficulty did not reliably moderate the effects of either choice or decision speed. This pattern suggests that observers primarily interpret decision time relative to the normative expectation associated with the choice outcome itself, rather than flexibly reweighting timing cues based on difficulty alone. As a targeted exploratory analysis, we examined whether decision difficulty influenced preference inferences when agents quickly chose the non–value-maximizing option. In this subset, easier decisions were indeed associated with especially strong inferred preferences ($\beta = 0.15$, $p = .044$), suggesting that decision structure provides a secondary cue when decision time is minimally informative.

Together, these results indicate that observers do not treat decision time as a generic cue to uncertainty. Instead, they integrate both decision speed and choice optimality to infer the strength of an agent's underlying preferences, attributing hesitation to weaker preferences particularly when an agent deviates from value-maximizing behavior (e.g., selecting a numerically smaller quantity).

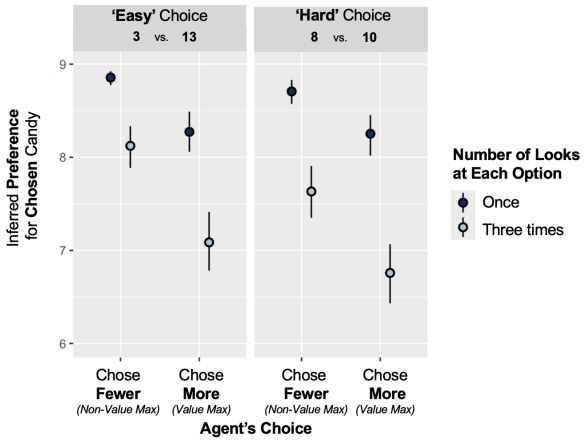

Figure 4. Plot showing people's preference inferences as a function of decision difficulty, the agent's number of looks, and the agent's choice. Pointranges show averages with 95% CIs.

Discussion

Experiment 2 demonstrates that observers use decision timing to infer agent-specific preferences. When both options were visible, participants systematically modulated their preference judgments as a function of how long the agent took to decide and also whether the choice reflected value-maximizing or non-value maximizing. Critically, identical choices supported different preference inferences depending on decision time and choice, suggesting that observers interpret timing relative to expectations about how easy a decision should be under neutral preferences (i.e. when to value maximize). Unexpectedly, these effects were not strongly moderated by objective decision difficulty, perhaps because the easy/hard manipulation did not sufficiently shift expectations about how long deliberation should take relative to the more salient cues of choice outcome and hesitation. Nonetheless, together, these findings indicate that observers do not treat hesitation as a generic signal of uncertainty, but reason about why uncertainty would arise, attributing extended decision time to features of the agent's subjective valuation, while integrating expectations from the structure of the choice itself.

General Discussion

Across two experiments, we investigated how observers use decision time, alongside observed choices, to infer unobservable aspects of others' minds. In Experiment 1, observers used decision speed to infer properties of the decision itself, estimating the value of an occluded option by interpreting how long an agent deliberated relative to the apparent difficulty of the choice. In Experiment 2, observers used the same timing cues to infer agent-specific preferences when the structure of the decision was fully known. Together, these findings suggest that people interpret decision time not in isolation, but relative to structured expectations about how long a decision should take, and invert these expectations to infer latent properties of both the world and the agent. More broadly, the present findings help connect work on intuitive utility-based social inference (e.g., Jara-Ettinger et al., 2016) with emerging work showing that observers use response dynamics to infer mental states (e.g., Gates et al., 2021; Bavard et al., 2024).

Crucially, observers' inferences were highly flexible. Participants did not treat decision speed as a simple monotonic cue (e.g., slower decisions implying greater uncertainty or weaker preferences). Instead, the meaning of extended deliberation depended on how long a decision was expected to take given the structure of the choice and the observed outcome. For example, in Experiment 1, longer deliberation led observers to infer closer value comparisons when an agent faced a strong visible alternative, but had a different implication when the visible option was weak. Similarly, in Experiment 2, hesitation was interpreted differently depending on whether a decision should have been easy or difficult under neutral preferences (i.e. value-maximizing), supporting more nuanced inferences about the agent's relative valuation of the options. This pattern suggests that observers appeal to decision speed relative to expectations, enabling fine-grained inferences about both conflict and value rather than relying on coarse heuristics.

Methodologically, we operationalized decision time using the agent's number of visual looks toward the options. While decision time and visual sampling are not identical constructs, they are closely coupled in naturalistic decision-making, and overt sampling behavior provides a transparent signal that observers readily interpret. Our claim is not that the number of looks is identical to decision time, but that observers treat overt sampling behavior as a proxy for deliberation. The inferences we measure therefore reveal observers' intuitive theory of decision dynamics, regardless of whether those dynamics are normatively optimal. Future work could dissociate these components more directly, for example by manipulating decision latency independently of visual attention, or by examining whether similar inferences arise from non-visual timing cues such as response delays or motor hesitations.

More broadly, these findings contribute to a growing body of work suggesting that social cognition involves reasoning not only about what others do, but about how their behavior is generated over time. Here we show novel evidence of a more nuanced intuitive theory of decision-making dynamics, where observers do not simply use decision time as a cue, but interpret timing relative to structured expectations about how long a decision should take. Such representations may support rapid, real-time social inference in everyday interactions, where timing information is often available even when outcomes or internal states are not.

References

Bavard, S., Stuchlý, E., Konovalov, A., & Gluth, S. (2024). Humans can infer social preferences from decision speed alone. *PLoS Biology*, 22(6), e3002686.

Berke, M., Tenenbaum, A., Sterling, B., & Jara-Ettinger, J. (2023). Thinking about thinking as rational computation. In *Proceedings of the Annual Conference of the Cognitive Science Society*.

Feigenson, L., Carey, S., & Hauser, M. (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological science*, *13*(2), 150-156.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in cognitive sciences*, 8(7), 307-314.

Frydman, C., & Krajbich, I. (2022). Using response times to infer others' private information: An application to information cascades. *Management Science*, 68(4), 2970-2986.

Gates, V., Callaway, F., Ho, M. K., & Griffiths, T. L. (2021). A rational model of people's inferences about others' preferences based on response times. *Cognition*, 217, 104885.

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

Halberda, J., Mazzocco, M. M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature*, 455(7213), 665-668.

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.

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.

Morris, B., Yurovsky, D., & Shaw, A. (2026). "Um…" Thinking out loud: Children infer the social meaning of speech disfluencies. *Child Development*, aacag062.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4), 873-922.

Richardson, E., & Keil, F. C. (2022). Thinking takes time: Children use agents' response times to infer the source, quality, and complexity of their knowledge. *Cognition*, 224, 105073.

Roberts, F., Margutti, P., & Takano, S. (2011). Judgments concerning the valence of inter-turn silence across speakers of American English, Italian, and Japanese. *Discourse Processes*, 48(5), 331-354.

Roberts, F., & Francis, A. L. (2013). Identifying a temporal threshold of tolerance for silent gaps after requests. *The Journal of the Acoustical Society of America*, 133(6), EL471-EL477.

Ziano, I., & Wang, D. (2021). Slow lies: Response delays promote perceptions of insincerity. *Journal of Personality and Social Psychology,* 120(6), 1457.