

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Modeling risky food sharing as rational communication about relationships

Permalink

<https://escholarship.org/uc/item/1066w77t>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Hung, Michelle Simona

Thomas, Ashley J

Radkani, Setayesh

et al.

Publication Date

2022

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

Modeling risky food sharing as rational communication about relationships

Michelle S. Hung (michl@mit.edu)

Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Ashley J. Thomas (ajthomas@mit.edu)

Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Setayesh Radkani (radkani@mit.edu)

Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Joshua B. Tenenbaum (jbt@mit.edu)

Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Rebecca Saxe (saxe@mit.edu)

Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

Abstract

The way two people choose to share food reveals how close their relationship is. Very close relationships alleviate the discomfort of exchanging saliva. We measure human inferences about relationships from observed food sharing actions with variable risks of saliva exchange; and then use a formal model of inverse planning to quantitatively capture these inferences. The model that best fits human judgments construes food sharing as a rational communicative social action, according to which actions are chosen both to maximize comfort given a relationship, and to communicate about the relationship itself.

Keywords: Bayesian modeling; Pragmatics; Reasoning; Social cognition; Theory of mind

Introduction

Imagine that you see two people you don't know (let's call them Sarah and Jane), sitting together at a table. They have a bowl of mac'n'cheese between them, and they are sharing it by alternately using a single spoon to take bites. What can you, an observer, learn from watching this interaction?

Cognitive scientists have long studied how humans can observe a person's actions, and infer their beliefs and desires (Premack & Woodruff, 1978; Wellman, 2017). Following this logic, people watching Sarah could infer that she wants some mac'n'cheese, and expects that sharing a spoon and bowl with Jane is an efficient way to get some mac'n'cheese. Yet this description leaves out a key inference that is apparent to most human observers: namely, Sarah and Jane must be in a very close relationship. Our goal in the current paper is to provide an initial characterization of how observing interactions allows observers to infer relationships, using as a case study, interactions involving sharing food.

Observed actions license inferences about beliefs and desires. Assuming the agents are rational actors, and thus tend to choose actions that maximize their utilities given their expectations, observers can use inverse planning to identify combinations of utilities (desires) and expectations (beliefs) that make the observed action most likely. Computational models that implement this intuition provide a close quantitative fit to human observers' inferences about an agent's goals, beliefs and perceptions (C. L. Baker, Tenenbaum, & Saxe,

2007; Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016). For example, if an agent takes a long deviation from the shortest path to their eventual goal, observers and models can infer that the agent was pursuing a subgoal, or initially expected to be able to achieve a preferred goal and then had to settle for a second-best goal (Jara-Ettinger et al., 2016; C. Baker, Saxe, & Tenenbaum, 2011).

In interpersonal interactions, each agent's beliefs and desires remain important, but many aspects of the interaction depend on the specific relationship between the two agents. Some aspects of the relationship can be described in terms of recursive utility functions (Powell, 2021; Kleiman-Weiner, Shaw, & Tenenbaum, 2017). For example, one aspect of a positive, or cooperative, interaction can be described as one agent having a positive utility on the second agent's utility. In this case, the first agent will choose actions that are likely to promote or achieve the second agent's goals. By contrast, in a conflictual or competitive interaction, the first agent may have a negative utility on the second agent's utility, and thus choose actions that obstruct or reduce the second agent's chances to achieve their goals. Computational models of recursive utility functions can give a close quantitative fit to human judgments of whether two observed agents are cooperating or competing (Ullman et al., 2009).

Yet there are key distinctions among human relationships that cannot be simply reduced to these recursive utility functions. One example is intimacy. Being in an intimate relationship is not simply a matter of highly valuing the other person's welfare, though high recursive utility is a common feature of intimacy. A more defining feature of intimate relationship is an experience of a shared essence, or "consubstantiation". Across cultures, intimate relationships are reflected in greater comfort with actions that substantively connect their bodies (Fiske & Haslam, 2005): exchanging bodily fluids (e.g., Carsten, 1995), touching body parts like face and genitals (e.g., Sorokowska et al., 2021), and synchronized body actions and appearances, like tattoos (e.g., Te Awakotuku, 2003). Sharing of saliva in particular is more common, expected, and acceptable among kin and intimate partners

than for less-close relationships (Margalit, 2017; Thomas, Woo, Nettle, Spelke, & Saxe, 2022; Fischler, 2011). Indeed, these actions not only reveal intimate relationships, but can be used to create and reinforce intimate relationships (Fiske & Haslam, 2005; Gallus, Reiff, Kamenica, & Fiske, 2021).

We hypothesize that, just as observers systematically infer individual mental states like beliefs and desires from observed actions in the environment, observers also systematically infer relationships from observed interactions between people. Furthermore, because people can anticipate these inferences, people may deliberately choose actions in order to cause observers to make desired inferences about relationships.

We focus on the example of consensual exchange of saliva, which occurs in the course of sharing food. For contemporary American adults, exchanging saliva during food sharing is seen as uncomfortable, or even disgusting (Tybur, Çınar, Karinen, & Perone, 2018), unless the people sharing the food are in an intimate relationship. In the current studies, we first (i) measure the quantitative relationship between saliva sharing risk, relationship closeness, and predicted comfort with 30 specific food sharing actions, and then (ii) test the hypothesis that after observing one person choose between two of these actions, observers make systematic inferences about the closeness of their relationship.

Next, we test whether a model of these inferences can quantitatively capture human inferences in this setting. The null hypothesis is that people do not infer relationship closeness from observed food sharing choices. Once we reject the null hypothesis, we consider two models of how observers could use the chosen action to infer the relationship. First, we made a Rational Choice model, adapted from the BToM framework (C. L. Baker et al., 2007). An ‘Actor’ chooses a way of sharing food that maximizes her comfort, given her belief about the relationship closeness. A ‘Receiver’ can infer her belief about the relationship, by inferring the relationship closeness which renders the chosen action most likely.

Building off of the Rational Choice model, we next modeled the choice of how to share food as a Rational Communicative Social Action (RCSA). Inspired by Rational Speech Act models of pragmatic speech (Frank & Goodman, 2012; Goodman & Frank, 2016), the RCSA model assumes that people choose social actions, in part, in order to cause desired inferences in an observing audience. Thus, a Communicative Actor chooses a way of sharing food to maximize her comfort and also to accurately communicate her belief about the relationship closeness to an observer (modeled as the Rational Choice Receiver). A Communicative Receiver infers that the Actor is choosing an action partly in order to communicate her belief about the closeness of the relationship.

Experiments

All studies were conducted online via Qualtrics and distributed through Prolific. In all experiments, recruited participants were at least 18 years of age, U.S. citizens, and fluent in

English. Studies were pre-registered at <https://osf.io/czdp/> (see studies 1 & 2). Stimuli and code are available at <https://github.com/michellehung7/bids/tree/cogsci2022>.

Experiment 1

To infer relationship closeness from food sharing choices, actors and receivers must share common knowledge that comfort with food sharing depends systematically on the interaction of the risk of saliva exchange and relationship closeness. Specifically, in Experiment 1, we set out to measure this common knowledge. We predicted that close relationships mitigate discomfort with saliva exchange.

Methods To test this hypothesis we collected three datasets. First, one group of human participants (Pilot Dataset 1, $N=160$) judged how likely saliva exchange would be, for each of thirty ways to share foods or drinks (10 foods, 3 actions each, sample vignette in Figure 3, left).

Second, an independent group of participants (Experiment 1a, $N=148$) read scenarios about two people who are sharing food or drink via one action (sample vignette in Figure 3, middle). The relationship was described as not close, somewhat close, close, or very close. Participants judged how comfortable the two people are (7-point Likert scale, from extremely uncomfortable to extremely comfortable). Each participant read 8 vignettes and 1 attention check.

Third, after confirming our hypothesis that comfort depends on an interaction of saliva sharing risk and relationship closeness in the planned sample, we collected additional responses to the same task, in order to have more stable estimates of the shared knowledge that grounds relationship inferences (Experiment 1b, $N=197$). Results reported here are from the combined dataset (Experiment 1, $N=345$).

We compared a linear model with relationship closeness and saliva sharing risk as fixed effects and subject as a random effect (Model 1), to a model that also included an interaction between relationship closeness and saliva sharing risk (Model 2). We pre-registered the confirmatory hypotheses that Model 2 would fit better than Model 1, indicated by a lower Akaike information criterion (AIC), and interaction term in Model 2 would be significant ($p < 0.05$).

Results and Discussion In the pilot data, the thirty actions varied substantially and systematically in the risk of saliva exchange (Figure 1). The lowest risk actions, like splitting a hot dog in half or dividing a pudding into separate bowls, have almost no risk of saliva exchange. By contrast, taking alternating bites from the same side of the hot dog, or alternating spoonfuls of pudding with the same spoon, is almost guaranteed to result in transfer of saliva.

The results of Experiment 1 (Figure 2) show that human participants know that food sharing involving more risk of saliva exchange is less comfortable, unless the people sharing the food are in a close relationship. The Akaike information criterion (AIC) of Model 2 was less than the AIC of Model 1 ($10388 < 10477$, $\chi^2 = 91.147$, $p < 2.2e - 16$), indicating




Figure 1: **Pilot dataset 1.** Risk of saliva sharing for each of three actions with each of ten foods.

that the addition of the interaction term in Model 2 improved model fit enough to compensate for the additional complexity of the model.

Likewise, the interaction term in Model 2 was significant ($\beta = 0.18, p < 2e - 16$). For two people in a very close relationship, sharing one spoon for the pudding is almost as comfortable as dividing the pudding into bowls. By contrast, for two people who are not close, sharing a spoon would be much less comfortable. Variation in the saliva sharing risk across foods and actions is systematically related to variation in comfort, given relationship closeness.




Figure 2: **Results of Experiment 1.** Comfort with food sharing actions, as a function of saliva sharing risk and relationship closeness. Each dot is one judgment (total N=2732).

Experiment 2

How comfortable people feel sharing food in particular ways depends on how close their relationship is. In Experiment 2, we measured how people use this common knowledge to infer relationship closeness from observed food sharing, and potentially to communicate about relationship closeness during food sharing.

Methods Human participants (Experiment 2, N=270) read vignettes about an Actor who chose a food sharing action, and then judged the relationship closeness that the Receiver would infer from that choice. Each participant saw one scenario per food (10 total), but the results from one food (sandwich) were lost due to experimenter error. This experiment was pre-registered, although our eventual analyses deviated from the pre-registered analysis plan. The analyses presented here are therefore exploratory, and will be confirmed in future studies.

Each vignette described two people who were eating together (sample vignette in Figure 3). The Receiver is described as uncertain about the Actor's belief about their relationship closeness. The Actor then chooses between two alternative ways to share the food. After each scenario, there was an attention check. Then participants judged, from the Receiver's point view, the Actor's belief about the relationship closeness.

Inferences about relationship closeness depend on the observed action, but may also depend on observers' priors about how close two people are, just given that they are eating the same food at the same time and place. To measure this prior, an independent group of participants (Pilot Dataset

Pilot Dataset 1	Experiment 1	Experiment 2
Imagine two people who are drinking a milkshake together in the following way: They drink the milkshake from one cup, alternately using the same straw. How likely is it that one person would end up with some of the other person's saliva in their mouth? extremely unlikely (1) – extremely likely (7)	Imagine two people who would describe their relationship as close. Imagine the following: One person pours half of their milkshake into a separate cup. Then, each person drinks from their own cup. How comfortable do you think the two people would be sharing a milkshake this way? extremely uncomfortable (1) — extremely comfortable (7)	Imagine that Sarah is out eating with Jane. Jane is not certain how close Sarah thinks their relationship is. On this occasion, Sarah and Jane are sharing a milkshake. Sarah is considering sharing the milkshake in one of the following two ways: (A) pouring half the milkshake into a separate cup, drinking from the original cup, and handing Jane the separate cup to drink from, or; (B) drinking the milkshake using a straw, then handing Jane the milkshake and a separate straw to drink from. Sarah chooses to share the milkshake with Jane by (A) pouring half the milkshake into a separate cup, drinking from the original cup, and handing Jane the separate cup to drink from. From Jane's point of view, how close does Sarah think their relationship is? not close (1) somewhat close (2) close (3) extremely close (4)

Figure 3: **Sample vignettes.** Left: Pilot dataset 1. Middle: Experiment 1. Right: Experiment 2.

2, N=197) estimated closeness from vignettes that simply describe two people eating at the same time and place, without specifying that they are sharing food or how they are eating.

Results and Discussion Without observing any particular way of sharing food, in Pilot Dataset 2, human participants considered eating together to be quite likely for all levels of relationship closeness, although somewhat less likely for people who are not at all close.

By contrast, when participants read about a specific way of sharing food, they made strong inferences about the relationship (Figure 4, left column). If the Actor chose an action with the highest saliva sharing risk, participants predicted that the Receiver would infer that the Actor believed they have a very close relationship. If the Actor chose an action with a moderate saliva sharing risk, participants predicted that the Receiver would infer that the Actor believed they have a close or somewhat close relationship. If the Actor chose an action with low saliva sharing risk, participants predicted that the Receiver would infer that the Actor believed they have a not close relationship. In addition to a strong influence of the chosen action, participants' judgments were also influenced by the rejected alternative action. In particular, the rejected alternative influenced human inferences about the intermediate, moderate saliva sharing action. When the Actor rejected a low saliva sharing risk, participants inferred that the Actor represented a closer relationship than when the Actor rejected a high saliva sharing risk.

Next, we sought to capture these patterns of human judgments in a framework based on inverse inference from an intuitive theory of rational communicative social action (Frank & Goodman, 2012).

Models

We compared three types of models of observers' inferences of relationship closeness in Experiment 2.

Null model

The null hypothesis is that human observers do not infer relationship closeness from observed food sharing interactions. In this case, the relationship closeness inferred after observing food sharing should not deviate from the inferences made with no observed food sharing (the relationship priors, $P(c)$).

Rational Choice Receiver in BToM Model

We adapted the Bayesian Theory of Mind model (C. L. Baker et al., 2007; Jara-Ettinger et al., 2016) to the current context.

A Rational Choice Actor A_0 selects a way of sharing food that maximizes her comfort, given her belief about the relationship. To estimate the subjective utility of each possible action, we fit the predicted comfort of the action from the results of Experiment 1, given the saliva sharing risk estimated in Pilot Dataset 1 and each level of relationship closeness. Comfort is translated into action probabilities via a soft max over the relative utilities of the chosen action and the alternative actions. Thus, the action probability distribution of A_0 is given by:

$$P_{A_0}(a|c) = \frac{e^{\beta \cdot U_{\text{comfort}}(a|c)}}{\sum_i e^{\beta \cdot U_{\text{comfort}}(a_i|c)}} \quad (1)$$

where a represents the chosen action, c , represents the relationship closeness, and β is a rationality parameter which represents the degree of determinism with which the actor does the utility maximizing action. This captures the intuition that food-sharing actors generally, but not always, do the more comfortable action.

In our pre-registered analyses, we planned to model a Rational Choice Actor that only considers the two alternative actions specified in the vignette. However, we subsequently added an additional model of a Rational Choice Actor that considers the utility of all three possible actions for sharing the food or drink. It is plausible that since each participant read 10 vignettes, each with two action alternatives, that participants implicitly considered the full range of actions in the experiment, even though in each vignette only two of those




Figure 4: **Results of Experiment 2.** Top row: Inferred emotional closeness as a function of the saliva sharing risk of the chosen action. Bottom row: Same results, as a function of the difference between the saliva sharing risks of the chosen and rejected alternative actions. Left column: Human judgments. Right columns: Predictions from four different models: the null model, a BToM model assuming rational choice between two available actions, a BToM model assuming rational choice between three possible actions, and an RCSA model.

alternatives were available. Thus, we modeled two variants of the Rational Choice Actor that differed in whether the Actor chose one action from a set of two available alternative actions, or from the three possible actions. Both versions of the Rational Actor use a soft max to noisily maximize utility with rationality parameter β .

Upon observing the Rational Choice Actor, the Rational Choice Receiver R_0 performs Bayesian inference by inverting the probability distribution of the Rational Choice Actor and then multiplying by the closeness prior $P(c)$. For each specific observed action, the Rational Choice Receiver estimates $P(c|a)$, the probability that the relationship closeness is c , given that the literal actor took action a . By Bayes' rule:

$$P_{R_0}(c|a) \propto P_{A_0}(a|c) \cdot P(c) \quad (2)$$

Two versions of the Rational Choice Receiver were matched to the corresponding Actor, interpreting the action as a choice among two, or three, alternative ways of sharing the food. For each version, we predicted the relationship closeness that the Receiver would infer, for each of the vignettes in Experiment 2.

Communicative Receiver in RCSA model

Next we considered a model of an Actor who gains additional utility from communicating her belief about the relationship to the Receiver.

The Communicative Actor A_1 chooses a food sharing action by balancing two utility functions. First, the Communicative Actor includes the same utility term as the three-choice Rational Choice actor: the utility of comfort, $U_{\text{comfort}}(a|c)$, for each possible action given relationship closeness.

Critically, the Communicative Actor has an additional utility, from having the Receiver correctly infer the Actor's belief about their relationship closeness. To estimate this utility, the Communicative Actor recursively models the Rational Choice Receiver described above (R_0). Then, as in RSA, the communicative utility term is $\log R_0(c|a)$, because A_1 seeks to minimize the surprisal of closeness c given action a for the Rational Choice Receiver. The relative weight of this communicative utility, versus the direct utility from the comfort of the action, is set by parameter α . The Communicative Actor's utility function is:

$$U_{A_1}(a|c) = \alpha \cdot \log R_0(c|a) + U_{\text{comfort}}(a|c). \quad (3)$$

Because the vignettes explicitly described the Actor as choosing between two explicit alternatives, we modeled an Actor that chooses to communicate a relationship closeness by rejecting an available action. Thus, the Communicative Actor compares the expected utility from the inference the Rational Choice Receiver would make from each of the two explicitly available action alternatives. The action probability distribution of A_1 is:

$$P_{A_1}(a|c) = \frac{e^{\beta \cdot U_{A_1}(a|c)}}{e^{\beta \cdot U_{A_1}(a|c)} + e^{\beta \cdot U_{A_1}(a_j|c)} + e^{\beta \cdot U_{\text{comfort}}(a_k|c)}} \quad (4)$$

where a_j is the rejected available option, a_k is the unavailable option, and $U_{A_1}(a|c)$ is given by Equation (3).

Upon observing the choice of the Communicative Actor, the Communicative Receiver R_1 computes the probability of closeness c given some action a by reasoning about the Communicative Actor A_1 . Just as the Rational Choice Receiver does inference by inverting the distribution of action given

relationship closeness as represented by the Rational Choice Actor, the Communicative Receiver does inference by inverting the distribution of actions given relationship closeness as represented by the Communicative Actor. By Bayes' rule, the probability of closeness c given action a is proportional to the probability that the Communicative Actor would choose action a , times the prior $P(c)$:

$$R_1(c|a) \propto P_{A_1}(a|c) \cdot P(c) \quad (5)$$

Model Parameter Fitting and Comparison

We pre-registered model comparisons using AIC, KL-divergence and r-squared to identify the single model that best fit the data (see <https://osf.io/czdpy/>; studies 1 & 2). However, our final model specifications deviated from the pre-registered versions, so the analyses presented below are exploratory and will be confirmed in future independent data.

The BToM model has one free parameter, β , which set the temperature of the softmax choice. The RCSA model has two free parameters: β and α , the weight that the Actor gives to the utility of communicating to the Receiver. We fit these parameters in the RCSA model by minimizing Mean Squared Error (MSE) between the model and the data, using a grid search with α, β between 0.1 and 5 in increments of 0.1. For both versions of Rational Choice Actors, MSE of the BToM model was minimized when $\beta = 1.8$. The Rational Choice Receiver embedded inside the RCSA model thus used $\beta = 1.8$. For the Communicative Receiver, MSE of the RCSA model was minimized when $\beta = 1.2$ and $\alpha = 0.9$.

Model Results

For each model, we considered the overall predicted closeness inferred, for each chosen action, as a function of that action's saliva sharing risk (Figure 4, top row); as well as the difference in the predicted inference, for each chosen action, as a function of the rejected alternative action (Figure 4, bottom row).

Null Model The null model fit the data very poorly (MSE=0.3866, KL=0.3539, $r^2=0.08$). Thus, we reject the null hypothesis that human observers do not infer relationship closeness from observed food sharing actions.

BToM Model The predictions of the BToM model, based on inverting an intuitive theory of a Rational Actor who aims to share food while maximizing comfort, were much closer to the human judgments. There was a qualitative difference between the models that considered only the two available alternatives, versus the full implied set of possible actions. The two-alternative Rational Choice Receiver inferred relationship closeness based on the contrast between the two available alternative, and thus so did not fully capture the human judgments (MSE=0.2012, KL=0.2866, $r^2=0.51$). On the other hand, the Rational Choice Receiver that considered all three possible action inferred increasing relationship closeness, as the risk of saliva sharing in the chosen action increased, and so fit the human judgments better (MSE=0.1544,

KL=0.2640, $r^2=0.80$). However, this model learned nothing from the rejected alternative.

In summary, the two versions of the Rational Choice receiver each captured aspects of the inference of the human observers, but could not fit the full pattern of judgments.

RCSA Model Adding an additional utility of communicating relationship closeness via the rejected action generated the best fit to human judgments (MSE=0.0979, KL=0.2130, $r^2=0.85$). This model, like humans, inferred different relationship closeness for the same chosen action, depending on which action was rejected. Also, although the least risky actions are comfortable in all relationships, they are only chosen by rejecting an action with greater saliva sharing. Thus, the Communicative Receiver infers a less close relationship when Actors choose the action with lowest saliva-sharing risk.

Discussion

From the way a person shares food, the recipient of the action can infer how close the sharer thinks their relationship is. Anticipating these inferences, people use actions while sharing food to communicate about their relationships. Here we provide a close quantitative fit to the pattern of human judgments by modeling food sharing as Rational Communicative Social Action (RCSA).

RCSA provides a principled framework for integrating the value of anticipated observer inferences (communication utility) into inverse planning in an intuitive theory of rational action. Similar models have been used in previous work to capture how the desire to appear impartial can influence a person's choice to distribute unequal financial payouts (Kleiman-Weiner et al., 2017); how the desire to appear considerate can influence the choice of negative expressions in polite speech (Yoon, Tessler, Goodman, & Frank, 2020); or how people trade off the values of learning new skills versus the desire to appear competent at familiar skills (Yoon, MacDonald, Asaba, Gweon, & Frank, 2018; Asaba & Gweon, 2019).

Here we use RCSA to model inferences about intimacy, based on observed food sharing activities. Even toddlers and infants recognize that consensually exchanging saliva is a reliable cue of an intimate relationship (Thomas et al., 2022). One open question concerns how similar these inferences would be, in different cultural contexts. Kinship and intimacy are conceived in terms of shared bodily substance across cultures (Fiske & Haslam, 2005). Yet there is also cross-national variation in specific behaviours that lead to saliva exchange, like mouth-to-mouth kissing (Suvilehto, Glerean, Dunbar, Hari, & Nummenmaa, 2015). In some cases, saliva sharing can occur in quotidian non-intimate interactions, such as via saliva-fermentation of beverages like *chicha* (Colehour et al., 2014). Thus, it is likely that there are culturally-specific constraints on how people use saliva-sharing interactions to infer, and communicate about, relationship closeness.

Using RCSA to capture inferences about food sharing thus illustrates the centrality of relationships to intuitive reasoning about social interactions (Kaufmann & Clément, 2014).

References

- Asaba, M., & Gweon, H. (2019, Dec). *Young children infer and manage what others think of the self*. PsyArXiv. Retrieved from psyarxiv.com/yxhv5 doi: 10.31234/osf.io/yxhv5
- Baker, C., Saxe, R., & Tenenbaum, J. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 33).
- Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (2007). Goal inference as inverse planning. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 29).
- Carsten, J. (1995). The substance of kinship and the heat of the hearth: feeding, personhood, and relatedness among malays in pulau langkawi. *American ethnologist*, 22(2), 223–241.
- Colehour, A. M., Meadow, J. F., Liebert, M. A., Cepon-Robins, T. J., Gildner, T. E., Urlacher, S. S., ... Sugiyama, L. S. (2014). Local domestication of lactic acid bacteria via cassava beer fermentation. *PeerJ*, 2, e479.
- Fischler, C. (2011). Commensality, society and culture. *Social science information*, 50(3-4), 528–548.
- Fiske, A. P., & Haslam, N. (2005). The four basic social bonds: Structures for coordinating interaction.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Gallus, J., Reiff, J., Kamenica, E., & Fiske, A. P. (2021). Relational incentives theory. *Psychological review*.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11), 818–829.
- 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.
- Kaufmann, L., & Clément, F. (2014). Wired for society: cognizing pathways to society and culture. *Topoi*, 33(2), 459–475.
- Kleiman-Weiner, M., Shaw, A., & Tenenbaum, J. (2017). Constructing social preferences from anticipated judgments: When impartial inequity is fair and why? In *Cogsci*.
- Margalit, A. (2017). On betrayal. In *On betrayal*. Harvard University Press.
- Powell, L. (2021). Adopted utility calculus: Origins of a concept of social affiliation.
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4), 515–526.
- Sorokowska, A., Saluja, S., Sorokowski, P., Frackowiak, T., Karwowski, M., Aavik, T., ... others (2021). Affective interpersonal touch in close relationships: A cross-cultural perspective. *Personality and Social Psychology Bulletin*, 0146167220988373.
- Suvilehto, J. T., Glerean, E., Dunbar, R. I., Hari, R., & Nummenmaa, L. (2015). Topography of social touching depends on emotional bonds between humans. *Proceedings of the National Academy of Sciences*, 112(45), 13811–13816.
- Te Awakotuku, N. (2003). Ta moko: Culture, body modification, and the psychology of identity. In *Tatau/tattoo: Embodied art and cultural exchange c. 1760-2000*.
- Thomas, A. J., Woo, B., Nettle, D., Spelke, E., & Saxe, R. (2022). Early concepts of intimacy: young humans use saliva sharing to infer close relationships. *Science*, 375(6578), 311–315.
- Tybur, J. M., Çınar, Ç., Karinen, A. K., & Perone, P. (2018). *Why do people vary in disgust?* (Vol. 373) (No. 1751). The Royal Society.
- Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum, J. (2009). Help or hinder: Bayesian models of social goal inference. *Advances in neural information processing systems*, 22.
- Wellman, H. M. (2017). The development of theory of mind: Historical reflections. *Child Development Perspectives*, 11(3), 207–214.
- Yoon, E. J., MacDonald, K., Asaba, M., Gweon, H., & Frank, M. C. (2018). Balancing informational and social goals in active learning. In *Cogsci*.
- Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (2020). Polite speech emerges from competing social goals. *Open Mind*, 4, 71–87.