

Learning and making novel predictions about others' preferences

Natalia Vélez¹, Yuan Chang Leong¹, Chelsey Pan, Jamil Zaki, & Hyowon Gweon

{nvelez, ycleong, chelspan, jzaki, gweon}@stanford.edu

Department of Psychology, 450 Serra Mall
Stanford, CA 94305 USA

Abstract

We often make decisions on behalf of others, such as picking out gifts or making restaurant recommendations. Yet, without direct access to others' preferences, our choices on behalf of others depend on what we think they like. Across two experiments, we examined *whether* and *how accurately* people are able to infer others' preferences by observing their choices. Our results suggest that people are capable of making reasonably accurate predictions about what others will choose next, given what they have chosen before. These results lay the groundwork to systematically study how people make novel predictions about others' preferences, and when different strategies might be appropriate.

Keywords: preference learning; social cognition; Theory of Mind; decision-making

Introduction

People often make choices to please others, such as buying gifts or making restaurant recommendations. The effortlessness of these mundane, everyday decisions belies their underlying complexity. Others' preferences may differ from our own, and we do not have complete knowledge of what they like. Therefore we must base our decisions on what we *think* the other person likes or wants. These choices are easy when we can simply give others what they have chosen before; even infants can cast aside their own preferences to give others foods that they clearly like (Repacholi & Gopnik, 1997; Doan et al., 2015).

However, choosing for others is rarely so simple. For example, suppose you are recommending a movie to a friend. You might remember some movies your friend has watched and liked before, but recommending exactly those movies would hardly be useful to her. Thus, it is often insufficient, even inappropriate, to simply choose what others have liked before. Instead, you would most likely consider movies that your friend has not seen before and choose the one you think she would like best. How do people make these novel choices on other people's behalf?

In many cases, people's own preferences provide a useful template for reasoning about other people's preferences. For example, you might recommend whatever movie *you* like best, under the assumption that your friend has similar tastes. Indeed, people tend to project their own desires and beliefs to those who are perceived to be similar (Ames, 2004). This is a useful strategy for predicting others' choices, especially when we have sparse, noisy, or ambiguous information about their preferences.

However, it is not a perfect strategy; other people's preferences do not always align perfectly with our own. Observing what others have liked before is another valuable source of information. Humans are capable of drawing powerful generalizations from sparse, noisy data (see Tenenbaum et al., 2011, for a review). Even young children draw systematic inferences about others' goals, preferences, and beliefs in ways that go beyond the observable evidence (Hamlin et al., 2007; Gweon et al., 2010; Kushnir et al., 2010). Previous work has formalized this process as "inverse planning," working backwards from others' observable actions to infer the unobservable mental states that generated them (Baker et al., 2009). Thus, people might use others' past choices to generate an abstract representation of their preferences, abstracting from the specific items that others have chosen to spot qualities that they might also enjoy in novel items. This would be akin to reasoning from a few movies that your friend has liked before—such as *Love Actually*, *Pretty Woman*, and *Sleepless in Seattle*—that she likes romantic comedies, and that she might enjoy other romantic comedies.

The ultimate goal of the present work is to better understand *how* people generalize from observed previous choices to make predictions about what others would choose next—as a first step, we ask *whether* and *how accurately* people can do this. Experiment 1 validates key features of our approach, and tests whether people can generalize from others' choices in a simplified context. In Experiment 2, participants (henceforth "observers") were faced with a more naturalistic—and much more challenging—task: they observed choices made by a previous participant (the "target") among one set of movies, and then predicted what the other person chose among a completely different set of movies.

Experimental Task: Choosing Novel Movies

The experimental task used here was designed to mirror the everyday task of recommending a movie to a friend (Figure 1). We created posters and plot synopses of novel movies that varied along three features: valence (positive or negative), setting (historical or futuristic), and genre (romantic or action). Each of the three features was varied orthogonally to generate 8 categories of movies; 4 movies were made for each category, resulting in 32 novel movies total. All movies were normed on Amazon Mechanical Turk by an independent group of raters ($n = 90$; data not shown) to ensure that each movie was categorized according to its intended features. The benefit of using novel movies is

¹ These authors contributed equally to this work.

twofold: first, we mitigated the potentially complex effects of prior knowledge and familiarity, as all movies were novel to all participants; second, we imposed some structure on the features of our movies, thus simplifying the learning problem.

In both experiments, participants watched the choices made by a target among a set of 16 movies, and then predicted what the same target would choose among a novel set of 16 movies. We considered two metrics for human performance. First, we directly compared observers' predictions to targets' real choices; this served as a ground-truth indicator of accuracy. Second, we compared observers' predictions to those of a mixed multinomial logit (MML) model. Model accuracies served as a benchmark for how well observers could be expected to do, given the choices they have seen the target make.

MML models have been used extensively in economics to model consumer choices (e.g., Train, 1980); more recently, they have also been applied to describe the development of preference understanding in young children (Lucas et al., 2014). MML models assume an agent's preferences are: (1) stable over time; and (2) defined over features of objects, and thus generalizable to other objects with similar features. When choosing one option out of a set, an agent's choices are probabilistically related to the utility, or attractiveness, of each option (Luce, 1977). Each option i is represented as a binary vector of its features (x_i) and a vector of weights (β) corresponding to the agent's preferences for individual features. The utility of option i (u_i) is the weighted sum of its features ($\beta \cdot x_i$), scaled by a free parameter T that describes the stochasticity of the agent's choices. (In all cases, T was a free parameter fit with a regularizing prior and the models were fit to maximize the maximum a posteriori estimate.) Taken together, the probability of an individual choosing option i from a pair of options J is defined as:

$$P(c_i | X, \beta) = \frac{\exp(\beta \cdot x_i / T)}{\sum_j \exp(\beta \cdot x_j / T)}$$

In the experiments below, the MML model served three important functions. First, the model was used to describe targets' own preferences: in Experiment 1a, we trained the model on one half of participants' own choices to extract the preference weights (β), and then used these learned weights to predict the second half of participants' responses. The model should predict targets' choices to the extent that participants' responses are reliable and that their preferences align well with the features we imposed on the stimuli,

Second, the model was used to evaluate *participants'* inferences about the target's preferences. In Experiment 1b, the model was trained on the first half of the target agent's choices—the same choices that human participants observed—and then tested on the participants' prediction about the target agents' choices among novel options. The model's accuracy reflects whether participants learned the

feature weights that describe the other person's choices, and to what extent they used these weights to predict what the other person would choose among a set of novel options.

Finally, the model was used as a descriptive tool to capture participants' strategies when choosing for others. Most notably, in Experiment 2, we extended the model to not only capture participants' inferences about the preferences of a target agent, but also how participants' observations interact with their own preferences when making novel predictions about the target's choices. Critically, throughout the paper, we used the model as a benchmark for human performance, rather than as a formal, computational characterization of the inferential processes involved in preference learning and generalization.

Experiment 1

Experiment 1, we tested two key assumptions of our approach. In Exp. 1a, we asked whether people's preferences are reliable and well defined over the features built into our stimuli; here, participants chose whichever movies they liked best, and we used the MML model to predict their choices. In contrast, in Exp. 1b, we were interested in whether people are able to learn about the preferences of others after observing a series of choices. Here, we simplified the learning problem by having participants learn about a simulated agent whose preferences can be perfectly described in our feature space, and who deterministically chooses the option with highest utility. Taken together, Exp. 1 serves to validate our overall approach and pilot a behavioral paradigm that can be used to study preference learning under noisier, richer conditions.

Experiment 1a: Choosing for self

Participants: 40 adults participated in an online behavioral experiment for pay through Amazon Mechanical Turk. All participants in this and subsequent experiments had U.S. IP addresses and provided informed consent in accordance with the IRB at Stanford University.

Procedure: The study was split into two rounds; in each round, participants saw half (16) of the 32 novel movies. Each round was composed of two tasks: Meet the Movies and Choose for Self (Figure 1a). In Meet the Movies, participants were shown the title, poster, and synopsis of each movie (Figure 1b). To ensure that participants were attending to and forming preferences for the movies, they rated how much they would like to watch each movie using a Likert Scale (1 – Not at all; 7 – Very much). In Choose for Self, participants were shown pairs of movie posters that they had just “met” and were asked to choose which of the two they would rather watch. There were 56 trials in this task (i.e., 56 pairs chosen from 16 movies), spanning all possible permutations of non-identical conditions.

Experiments 1b: Choosing for other

Participants: 50 participants were recruited for an online behavioral experiment through Amazon Mechanical Turk.

Procedure: Participants were told that they would observe and predict the choices of a target agent who had previously participated in Experiment 1a. Unbeknownst to the participants, the responses were generated by simulating the responses of an agent with pre-defined weights who deterministically chose the option with the highest utility.

As before, the study was split into two rounds, and each round began with Meet the Movies (Figure 1a). In the first round, participants observed the target’s choices (Observe Other; Figure 1b). Participants were shown pairs of movie posters, and a border appeared around the movie that the target had chosen. To ensure that participants attended to the task, we asked them to imitate the target’s choice by selecting the highlighted movie. Instead of using all 56 possible pairs, we excluded 8 pairs in which the artificial agent would be indifferent between the two movies; thus, participants saw 48 choices total.

In the second round (Choose for Other; Figure 1b), participants were again shown pairs of movies, but were instead asked to select the movie that they believed the target had chosen. Participants were not given trial-by-trial feedback, but they were informed that they would receive a bonus based on the number of correct responses in this task.

Note that the MML model had perfect information about each movie’s features, while human observers had no prior knowledge of the movies. Even though participants had a chance to “meet” the movies, we reasoned that this brief pre-exposure would be insufficient to eliminate their uncertainty about the dimensions of the feature space as well as the uncertainty about each movie’s features. Thus, in Experiment 1b, we provided keywords for each movie (e.g., “positive, historical, romantic”) during all tasks, making explicit the features of each movie. This ensured that the task for our human participants was comparable to the task imposed on our model, making the performance comparison more meaningful.

Results and Discussion

Experiment 1a: The model was trained on participants’ responses in the first round and tested in the second, and vice versa. Our measure of model accuracy (henceforth cross-validation accuracy) is the model’s average accuracy in predicting participants’ choices in each iteration; overall, the model accurately predicted participants’ choices in the test set (Cross-validation accuracy: $M(SD) = 0.54(0.11)$, tested against 0.50: $t(39) = 2.34$, $p = 0.02$; Fig. 2a).

Experiment 1b: Participants’ accuracy was near ceiling in the Observe Other task ($M(SD) = 0.98(0.08)$), suggesting that participants were alert and attentive during the task. Impressively, people showed fairly high accuracy even in the Choose for Other task, where participants had to use their previous observations to predict the target’s responses among a set of new movies, ($M(SD) = 0.77(0.22)$, test against 0.50: $t(49) = 24.27$, $p < 0.001$; Fig. 2b).

Overall, our model performed reasonably well at capturing people’s own preferences, suggesting that participants themselves have stable preferences that can be

inferred and predicted by our model. Importantly, participants were never told about the three dimensions or the features of each movie in Experiment 1a—nevertheless, people’s preferences were well described by the model, suggesting that people’s preferences align reasonably well with the feature space we have imposed on the stimuli.

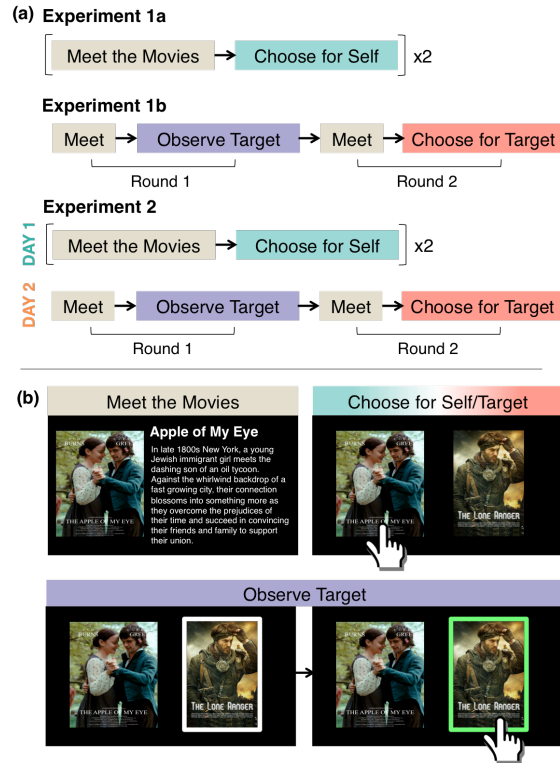


Figure 1: (a) Task order: In Experiment 1a, participants only chose for themselves; in Experiment 1b, participants only chose on behalf of another agent. Experiment 2 combines aspects of both of these into a multi-session experiment. (b) Schematic of tasks and example stimuli.

However, we note that the model is much worse at predicting real people’s preferences (Experiment 1a) than people are at predicting an artificial target agent’s preferences (Experiment 1b). These results serve complementary functions. On one hand, Exp. 1b provides an approximate upper bound for human performance, in the extreme case where the choice data provided are as clear and consistent as possible. Overall, we find that observers perform admirably when they are provided with good evidence. By contrast, model performance in Exp. 1a provides an estimate of the quality of the evidence available to observers when learning from real people’s choices. These results suggest that real people vary wildly in the extent to which their choices are consistent and aligned with the features built into the stimuli.

In Experiment 2, we aimed for a stronger test of people’s ability to learn and generalize about other’s preferences: having them observe a real person’s choices and asking them to predict their real choices. This is a much more

challenging task than learning from an idealized, deterministic target. However, given that people perform very well when given very clear information, we would still expect people to be able to learn and generalize from real person’s choices, at least to the extent that targets’ choices are consistent and informative.



Figure 2: Experiment 1 results. Left: Model cross-validation accuracy in Exp. 1a, indicating how well the model predicted each participant’s choices among the test set, based on weights inferred from participants’ choices among the training set. Right: Human performance in Choose for Target.

Experiment 2

As described above, Experiment 2 extended the previous experiments by having participants choose for themselves as well as for a target—here, the target was a real participant who participated earlier in the same experiment. This not only allowed us to test people’s ability to learn and generalize from noisy, messy choices of real participants, but also allowed us to examine the degree to which participants’ own preferences biased their choices on behalf of the target.

Procedure: 51 participants were recruited from the Stanford community for a two-day experiment. On day 1, participants completed the Meet the Movies and Choose for Self tasks, matching the procedure used in Experiment 1a. On day 2, participants first observed a target’s choices (Observe Target), then predicted the target’s choices among a different set of movies (Choose for Target). This session closely matched the procedure in Experiment 1b, with the critical difference that participants’ responses were yoked. That is, participant A’s responses during the first round of Choose for Self were presented to participant B during Observe Target, and participant B’s predictions in Choose for Target were tested against the actual choices that participant A made during the second round of Choose for Self. In sum, each participant observed and predicted the responses of another participant who had come before. Data from 2 participants were discarded because they did not return for the second day of the experiment.

Results and Discussion

First, the model was trained on participants’ binary choices in one round of Choose for Self and tested on the remaining round; as in Experiment 1, the model described

participants’ own preferences reasonably well (Cross-validation accuracy: $M(SD) = 0.61(0.12)$, test against 0.50: $t(48) = 6.4$, $p < 0.001$). If model accuracy is taken as a proxy for how consistent targets’ choices were and how closely they aligned to our stimulus features, then these results suggest that some targets’ choices were more informative than others, but that it is possible, on the whole, to generalize from the choices of targets in our sample.

Accordingly, participants’ predictions about the target’s choices on a novel set of movies also matched the target’s actual choices reasonably well ($M(SD) = 0.57(0.11)$, test against 0.50: $t(48) = 4.59$, $p < 0.001$, Fig. 3a) — despite the sparseness and noisiness of the target’s choices in both halves of the experiment. Impressively, participants’ performance was comparable to that of the model predicting the target’s choices based on the same observations (Model performance: $M(SD) = 0.59(0.13)$, paired t-test: $t(48) = -0.92$, $p = 0.361$). When the predictions of the MML model were compared to participants’ predictions, we found a reasonable correspondence ($M(SD) = 0.63$, test against 0.50: $t(48) = 7.1$, $p < 0.001$, Fig. 3a). Most importantly, we found a correlation between human and model accuracy in predicting the target’s choices ($r = 0.36$, $p = 0.01$; Fig 3b). Thus, even though 57% might not seem much higher than chance, our results suggest that observers seized the underlying structure in targets’ choices when possible.

So far we have considered a completely allocentric version of the MML model, which infers the targets’ preference weights from their previous choices and makes predictions based solely on these inferred weights. However, an alternative possibility is that one could make predictions by simply projecting one’s own preferences while completely ignoring others’ preferences, or that one could use a mix of these two strategies. In an exploratory analysis, we asked whether participants’ predictions about the choices of others were biased by their own preferences. We extended our MML model to make predictions based on a weighted average of self and learned target preferences, with the weight determined by free parameter η . Here, the model with $\eta=0$ is completely allocentric (i.e., made solely from others’ preferences), whereas the model with $\eta=1$ makes completely egocentric predictions (i.e., made solely from the participant’s own preferences).

We examined the distribution of best-fit values of η in our sample (Fig. 4). The distribution of best-fit values of η across participants peaks at $\eta = 0$, suggesting that many of the participants relied on their past observations of the target’s choices to make their predictions. There is also a second, smaller peak at $\eta = 1$, indicating that there is a second group of participants who relied solely on their own preferences when making decisions for others. However, there is a wide range of η values in our sample, suggesting that some participants combined self and other preferences to some degree when making their choices. Using a leave-one-trial-out cross-validation procedure, we compared the performance of the combined model to a model that made choices based only on target weights (other model) or only

on observer weights (self model). Notably, for 7 out of the 49 participants, the combined model predicts participants’ choices better than both the self and other model alone, suggesting that these participants were using both self and other preferences when making their choices.

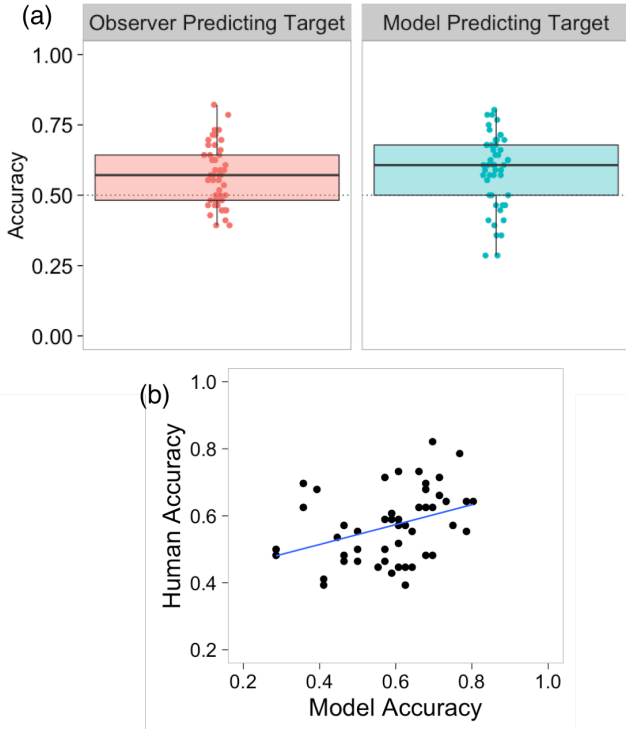


Figure 3: Experiment 2 results. (a) Human (left) and model (right) accuracy in Choose for Target. The model shown here was trained solely on the target’s choices, and thus corresponds to a totally allocentric observer. (b) Correlation between human and model accuracies.

Overall, we find initial evidence that people can learn impressively well from extremely noisy, sparse data to derive reasonable predictions about others’ choices, and that their accuracy is commensurate with the quality of the information gleaned from others’ past choices. We also find hints that people are not completely free of biases coming from their own preferences.

We note that there are two limitations of the paradigm in its current form, which will be addressed in future work. First, the interpretation of η is confounded when the observer and target have similar preferences. In such cases, the model would predict choice equally well when using self-preferences and when using other-preferences. The best-fit value of η would then depend on small differences in choice histories and are unlikely to reflect differences in participants’ strategies. Second, though many targets’ choices align closely with our stimulus dimensions and are consistent from one set of movies to the other, other targets’ choices were not at all informative, and in these cases our observers were set up to fail. While we randomly paired each observer to a target in the current work, future iterations of this paradigm will screen participants with

consistent, learnable preferences, and yoke participants to targets whose preferences are orthogonal to their own.




Figure 4: Distribution of η . Smaller η indicate that participants chose allocentrically (i.e., based on the target’s choices), and higher η suggest that participants chose egocentrically (i.e., based on their own preferences).

General Discussion

Across two experiments, we found that participants can generalize from the choices a target has made before to accurately predict what the target will choose next. In Experiment 1a, we first verified that the model captures people’s own preferences, suggesting that people have stable preferences that can be described by the dimensions we have imposed on the stimuli. Experiment 1b suggested that participants can accurately predict others’ choices in an idealized scenario, where they are observing the choices of an artificial target whose actions are deterministic and perfectly aligned with predetermined features of the stimuli.

Experiment 2 provided the main testing ground for our hypothesis: participants not only indicated their own preferences, but also grappled with the much harder task of learning from and predicting the noisy and often inconsistent choices of a real human participant. Despite the task being more challenging, we again found that participants made reasonably accurate predictions. Further, observers’ accuracy was commensurate with the consistency of the target’s choices, as measured by the accuracy of a simple MML model. In an exploratory analysis, we used this model to probe how strongly participants relied on their own preferences or on previous observations of the target’s choices in order to choose on their behalf.

Interestingly, we found that participants’ accuracy in making predictions about targets correlated with the accuracy of the MML model. Since the performance of the MML model is an index of consistency of a target’s preferences from one set of movies to the next, this result suggests that participants are better able to predict the target if the target provides consistent data for participants to learn about. What happens when the data provided are noisy? One possibility is that participants switch to an egocentric strategy of projecting their own preferences onto the target. Here, we demonstrate that participants use a mixture of egocentric and allocentric strategies; future work needs to be done to examine the factors determining how participants choose their strategy.

When we observe others’ choices, we make inferences about the hidden, internal states—such as preferences—that

motivated the choice. Using MML models that have been used in economics and cognitive science (Train, 1980; Lucas et al., 2009), we operationalize this process as learning the weights the other person attaches to underlying features that define the choice. While the model and our feature space are indeed too limited to capture the full richness and complexity of how people represent and learn about preferences, we believe that this simplified space is a good starting point for this investigation.

Given that we have found that people *can* learn and generalize from others' preferences, the natural next step is to explore *how* this is done. In the current work, we made the simplifying assumption that our stimuli varied along three binary dimensions; however, the feature space of people's actual movie preferences is much larger. In fact, for the same type of items, different people could use different sets of features to guide their choices—for example, one person might pay attention to a movie's reviews before choosing to watch it, while another person might decide which movies to watch based on the cast. As such, before learning which features another person values, observers have to first infer what features to learn about. This is a non-trivial problem, and it is similar to the structure learning problem explored in other domains of cognitive science (Gershman & Niv, 2010).

Another exciting extension of this work is to examine the degree to which social closeness affects people's predictions of others' choices. If participants systematically overestimate that the people closest to them are also more similar to them (Savitsky et al., 2011), then they might choose more egocentrically for a friend than for a stranger. However, closeness could very well play the opposite role: because people have had more opportunities to observe the choices of their close friends, they may choose more accurately for their friends than for a stranger. This direction also converges with prior neuroimaging work, which suggests that watching other people receive rewards engages neural systems involved in reward, and that this vicarious reward signal is influenced by perceived similarity between the observer and the target (Mobbs et al., 2009). However, little empirical work has directly tested the neural computations involved in making novel choices for others based on abstract, generalizable preferences, or how social closeness might modulate neural responses associated with preference learning.

The current work explored the deceptively mundane problem of predicting what others will like next based on what they have liked before. This is no small feat — other people's preferences may (or may not) differ substantially from our own, and their choices provide only a sparse and noisy reflection of their preferences. Motivated by both classical theories on egocentric biases (Ross et al., 1976) and more recent computational approaches to understand human learning as rational inductive inferences from sparse, noisy data (Tenenbaum et al., 2011), our experiments provide important empirical groundwork to better understand how this feat might be accomplished.

Acknowledgements

This work was supported by the Varieties of Understanding grant from the John Templeton foundation to HG and an NSF Graduate Research Fellowship to NV. We also thank our reviewers for their invaluable comments and feedback.

References

- Ames, D. R. (2004). Inside the mind reader's tool kit: projection and stereotyping in mental state inference. *Journal of personality and social psychology*, 87(3), 340.
- Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of personality and social psychology*, 87(3), 327.
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current opinion in neurobiology*, 20(2), 251-256.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, 107(20), 9066-9071.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557-559.
- Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, 21(8), 1134-1140.
- Mobbs, D., Yu, R., Meyer, M., Passamonti, L., Seymour, B., Calder, A. J., ... & Dalgleish, T. (2009). A key role for similarity in vicarious reward. *Science*, 324(5929).
- Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: evidence from 14- and 18-month-olds. *Developmental psychology*, 33(1), 12.
- Ross, L., Greene, D., & House, P. (1977). The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of experimental social psychology*, 13(3), 279-301.
- Savitsky, K., Keysar, B., Epley, N., Carter, T., & Swanson, A. (2011). The closeness-communication bias: Increased egocentrism among friends versus strangers. *Journal of Experimental Social Psychology*, 47(1), 269-273.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279-1285.
- Train, K. (1980). A structured logit model of auto ownership and model choice. *The Review of Economic Studies*, 47(2), 357-370.