

Identifying social partners through indirect prosociality: a computational account

Isaac Davis¹, Ryan W. Carlson¹, Yarrow Dunham^{1,2}, and Julian
Jara-Ettinger^{1,2}

¹Department of Psychology, Yale University

²Wu Tsai Institute, Yale University

Abstract

The ability to identify people who are prosocial, supportive, and mindful of others is critical for choosing social partners. While past work has emphasized the information value of direct social interactions (such as watching someone help or hinder others), social tendencies can also be inferred from indirect evidence, such as how an agent considers others when making personal choices. Here we present a computational model of this capacity, grounded in a Bayesian framework for action understanding. Across three experiments we show that this model captures how people infer social preferences based on how agents' act when their choices indirectly affect others (Experiments 1a and 1b), and how people infer what an agent knows about others from knowledge of that agent's social preferences (Experiment 2). Critically, people's patterns of inferences could not be explained by simpler alternatives. These findings illuminate how people can discern potential social partners from indirect evidence of their prosociality, thus deepening our understanding of partner detection, and social cognition more broadly.

1 Introduction

A key challenge in selecting social partners is identifying people who are prosocial, cooperative, and mindful of others. Research suggests that people spontaneously detect others' potential value as social partners in terms of their capacity for empathy and prosociality (Fiske et al., 2007; Morelli et al., 2018), but how do people detect these qualities in others? Past work highlights how people draw on *direct* evidence of agents' prosocial inclinations, such as how they respond to requests for aid or social support. Research in this domain suggests that both adults and children can readily infer an agent's prosocial motives from their directly helpful actions (Hamlin and Wynn, 2011; Ullman et al., 2009; Kiley Hamlin et al., 2013), and use these inferences to form judgments about agents' moral character (Jara-Ettinger et al., 2020; Carlson et al., 2022; Hartman et al., 2022).

In actual social life, however, directly helpful actions are often a consequence of, rather than antecedent to, a social relationship. The person who helps their friend move apartments, for example, may do so because of their close personal relationship to that friend, rather than a general prosocial attitude toward others. In everyday life, it may therefore be difficult to observe direct evidence of prosociality unless one is already a member of (or at least privy to) the social relationship. How then do people detect potential social partners? We propose that part of the answer is a sensitivity to how agents make personal decisions in situations where their choices can indirectly affect others: that is, an agent's tendency to be mindful of how their decisions impact others' welfare (or ability to make choices) in everyday life.

The decision to act in a way that indirectly benefits others, sometimes called *social mindfulness* (Van Doesum et al., 2013; Van Lange and Van Doesum, 2015), can provide strong evidence of an agent's prosociality in two ways. First, it shows that the decider is using their theory of mind—the ability to reason about other people in terms of

unobservable mental states (Wellman, 2014; Frith and Frith, 2012)—to consider the preferences of others. Second, it signals a willingness to incur a personal cost (e.g.: forgoing a better outcome for oneself) to indirectly benefit someone else (e.g.: enabling them to better fulfill their own preferences) (Crockett et al., 2014; Fehr and Fischbacher, 2002; McClintock and Allison, 1989; Van Doesum et al., 2013; Van Lange, 1999). People who possess both of these qualities (those who are socially mindful) are more likely to care for and effectively support others, making them especially valuable social partners (Van Doesum et al., 2013, 2020). Thus, the capacity to detect prosociality from indirect evidence may be crucial for navigating human social life.

How do people detect socially mindful agents? Previous research suggests that, in social interactions, people interpret each others' behavior by inferring the mental states which best explain that behavior (Dennett, 1989; Gergely and Csibra, 2003a). In particular, inferences about agents' mental states are structured around an assumption that agents act to maximize their subjective utilities—the difference between the rewards they receive and the costs they incur (Jern et al., 2017; Lucas et al., 2014; Baker et al., 2017; Jara-Ettinger et al., 2016). Importantly, such models can also express social preferences in terms of recursive utilities, i.e.: the value an agent assigns to the outcomes of others (Jern and Kemp, 2014; Ullman et al., 2009). However, it remains unclear how people identify social partners from indirect (rather than direct) displays of prosociality.

Our work has two goals. First, we present a computational account of how people assess agents' social preferences from indirect evidence. We focus on settings where agents' choices can reveal whether they considered how their actions indirectly affect others, and how they value others' outcomes when making these choices. This model builds on prior work (Jern and Kemp, 2014), extending it to highlight the role of an agent's perceived knowledge or ignorance about the preferences of others, and how this shapes observers' assessments of the decider's prosociality and theory of mind.

Our second goal is to test our computational model in three experiments focusing on subtle displays of social mindfulness in sequential resource-allocation tasks. Scenarios like these frequently occur in everyday social life; for instance, when choosing a doughnut from a box in the office break room, one’s choice will affect which options remain for others. Thus, choices like these can reveal whether the decider deployed their theory of mind to try to indirectly benefit others, thereby providing indirect evidence of the decider’s prosociality. In Experiments 1a and 1b, we investigate how people interpret an agent’s personal choices in light of what the agent preferred, and crucially, what the agent believed others preferred. Our results demonstrate that people’s judgments about the decider’s social preferences (the degree to which the decider values others’ utility relative to their own) closely track with the predictions of a utility based model, and that simpler models which ignore the decider’s knowledge of others fail to explain these judgments. In Experiment 2, we demonstrate that the inverse is also true. That is, when a decider’s social preferences are known (for instance, that the agent is known to care about others), observers infer from an agent’s choices whether they knew the preferences of others, and what they believed those preferences were. This demonstrates not only that people can incorporate evidence that an agent is using theory of mind into their inferences of that agent’s social preferences, but also that people can infer what agents know or believe about each other based on their social behavior. These experiments provide converging support for a utility based model of how people identify prosocial partners from indirect evidence.

2 Computational framework

Our framework models scenarios in which an agent A (“actor”) makes a decision that indirectly affects another agent B (“bystander”). For our experiments, we used a

particular form of decision problem that allows agents to be indirectly prosocial (Van Doesum et al., 2013; Zhao et al., 2021). In these tasks, A is given first pick from a common supply of items of two types (e.g.: snacks of two different types), and permitted to take a fixed number of those items. A knows that B will pick second, and that B will therefore have access to whatever items A does not take. Based on A’s decision, we prompted participants to infer either the degree to which A cares about B (i.e. A’s “social preference” for B; Experiment 1), or A’s beliefs about B’s preferences for the available items (i.e. A’s “social knowledge” about B; Experiment 2).

We modeled participants’ inferences in both experiments using a single framework built on two core assumptions. The first is that social inference is structured around the expectation that agents act rationally to maximize their utilities—the differences between the costs they incur and the rewards they obtain (Gergely and Csibra, 2003b; Jara-Ettinger et al., 2016; Jern et al., 2017; Lucas et al., 2014). We encoded this assumption into a generative model of agent behavior (section 2.1). Our second assumption is that observers can invert this generative model to infer the underlying utilities and/or beliefs that best explain an observed action, and we implement this inversion as a form of Bayesian inference over the aforementioned generative model (section 2.2). Previous work has leveraged similar assumptions to explain a range of social inferences about goals (Baker et al., 2009), preferences (Lucas et al., 2014; Jern et al., 2017), beliefs (Baker et al., 2017; Tauber and Steyvers, 2011), competence (Jara-Ettinger et al., 2020), and social intentions (Ullman et al., 2009).

2.1 Generative model

A defining feature of our generative model is that A may care not only about how their decision affects them, but also how their decision affects others. We capture this notion using a recursive utility function. First, we assume each agent has an individual reward

function R , which captures the reward that the agent derives from an allotment of items. For example, if there are two types of treats (say, donuts and cupcakes), then $R_A([1, 1])$ is the total reward that A derives from having one donut and one cupcake, while $R_B([0, 2])$ is the total reward B derives from having zero donuts and two cupcakes. For clarity, we first discuss how our model uses reward functions to make decisions, and then turn to discuss the internal structure of these reward functions (section 2.1.1). To simplify notation, we will use d to denote A’s decision (i.e.: which items to keep), and $R_A(d)$ and $R_B(d)$ to denote the respective rewards that A and B receive from this decision (i.e.: $R_A(d)$ is A’s reward from the items they keep, while $R_B(d)$ is B’s reward from the items that remain after A’s decision).

Given these individual reward functions, we assume that the total utility $U(d)$ which A derives from the outcome of decision d is a weighted sum of A’s personal reward and B’s personal reward (following Jern and Kemp, 2014; Powell, 2021), i.e.:

$$U(d) = w_A R_A(d) + w_B R_B(d) \tag{2.1}$$

The pair of weights $W = (w_A, w_B)$ denotes A’s social preference, where $0 \leq w_A \leq 1$ captures A’s *egocentric* preference and $0 \leq w_B \leq 1$ captures A’s *allocentric* preference towards B (Murphy and Ackermann, 2011).¹ For example, $W = (1, 0)$ represents a case where A only values their own reward and is therefore completely self interested, while $W = (0, 1)$ reflects a completely altruistic A who only values B’s reward. Equation 2.1 thus reflects the total utility that A receives from decision d , including any utility A receives indirectly from B’s reward.

¹This formulation also allows using negative weights to capture A’s “dislike” for B, though our experiments focus only on self-interest versus other-interest, and do not invoke negative social preferences.

Note that equation 2.1 assumes A already knows B’s reward function. In reality, however, we often face varying degrees of uncertainty about the preferences of others: we may know what our close friends and family like or dislike, but have no idea about the preferences of strangers, or have only a vague idea about the preferences of a recent acquaintance (e.g.: A might think B prefers cupcakes over donuts, but not know the strength of the preference). To capture this uncertainty, we model A’s beliefs about B as a probability distribution $K(R_B)$ over possible reward functions. This distribution K reflects the degree of A’s knowledge about B’s preferences: if A knows B’s preferences exactly, then K is a point mass distribution centered on the correct reward function. If A has no knowledge whatsoever, then K is a uniform distribution over all possible reward functions. Importantly, K can also encode partial knowledge: if A believes that B prefers cupcakes over donuts, but does not know the strength of this preference, then K is a distribution over all reward functions with the property that $R_B(\text{cupcakes}) > R_B(\text{donuts})$.

Given this belief K , A’s total *expected* utility from decision d is given by

$$U(d) = w_A R_A(d) + w_B \mathbb{E}[R_B(d)] \tag{2.2}$$

where the expectation is taken with respect to A’s belief distribution $K(R_B)$. That is, $\mathbb{E}[R_B(d)] = \int xP(x)dx$, where $P(x)$ is the subjective probability that $R_B(d) = x$. Given this utility function, we assume that A will make decision d using a softmax decision function, so that A is more likely to make decisions that yield a higher total expected utility. The parameter value for this softmax function was tuned in an earlier pilot study (see Supplemental Materials).

2.1.1 Reward functions

In all of our experiments, each trial depicted a scenario in which A is given first pick from a common supply of food items of two types, and allowed to keep a fixed number of those items. A knows that B will pick second, and that B will therefore receive whatever items A does *not* take. Due to the quantitative structure of the task, and the fact that the resources being allocated may have a high rate of saturation (e.g.: one’s desire for additional cupcakes may decrease sharply after having a first cupcake), we incorporated discounting into the individual reward functions (Baucells and Sarin, 2010).

To this end, we introduced a discount parameter $0 < \rho < 1$, so that if A receives a reward r from eating one cupcake, A will receive $\rho * r$ from the second cupcake, $\rho^2 * r$ for the third, and so on. For each agent, we assumed a baseline reward value for each food item (r_1, r_2) . We then computed the reward that the agent receives from allotment $[X, Y]$ (i.e.: X of item 1 and Y of item 2) as

$$R([X, Y]) = r_1 * D(X, \rho) + r_2 * D(Y, \rho) \tag{2.3}$$

where $D(X, \rho)$ applies the discount rate function described above. Rather than attempting to estimate the discount rate from human data (as we expected it to vary across participants), we instead integrated this parameter out of the model using a $Beta(\alpha, \beta)$ prior distribution, and tuned the hyper-parameters α, β using an initial pilot study (See supplemental materials). This is equivalent to modeling social inference under the assumption that agents derive diminishing marginal utilities from additional food items, but without knowing each agent’s true discount rate. We demonstrate the importance of the discounting assumption by comparing our results against an alternative model with no discounting (i.e.: fixing $\rho = 1$). As we explain below, discounting is particularly important in cases where A is highly uncertain about B’s

preferences.

2.2 Inference

The generative model described above captures the proposed intuitive causal model that people use to reason about each other’s behavior. With this model, people can make predictions about how A will behave—expressed as a probability distribution

$P(d|R_A, K, W)$ —as a function of A’s personal preferences R_A , social preferences W , and beliefs K about the other agent’s preferences. Equipped with this model, an observer can leverage this distribution to infer an agent’s mental states using Bayesian inference.

Our experiments focused on two types of inference. In Experiment 1, participants observed A’s personal preferences (R_A), A’s beliefs about B’s preferences (K), and A’s decision d , and were then asked to infer A’s social preference $W = (w_A, w_B)$. In this situation, the observer can infer A’s social preference according to

$$P(W|d, R_A, K) \propto P(d|R_A, K, W)P(W|R_A, K) \quad (2.4)$$

In other words, the probability that A’s social preference is equal to W (given A’s personal preference, belief, and decision), is proportional to the probability that A would choose d (given A’s social preference) times the prior probability of A’s social preference.

In Experiment 2, participants observed A’s personal preferences, A’s social preference, and A’s decision, and were then asked to infer a) whether A knows which item B prefers and b) which item A thinks B prefers (assuming that A *does* know B’s preference). In this case, inferences about K can be captured by

$$P(K|d, R_A, W) \propto P(d|R_A, K, W)P(K|R_A, W) \quad (2.5)$$

When computing the prior distribution over A’s knowledge state $P(K|R_A, W)$, we encoded an assumption that A is more likely to attend to and remember B’s preference the more A cares about B. This is based on the principle that, in order to fulfill another person’s preferences, one must know what those preferences are in the first place. Thus, the more A cares about B (and therefore values fulfilling B’s preferences), the more motivation A has to attend to and remember B’s preferences. To this end, we define the prior $P(K|R_A, W)$ so that the probability that A is totally ignorant of B’s preferences is proportional to $w_A/(w_A + w_B)$ (See supplemental materials for the exact form of this prior). Note that this prior is only relevant for generating Experiment 2 predictions, as the term $P(K|R_A, W)$ does not appear in the equation for Experiment 1 (since A’s beliefs are fixed in the Experiment 1 tasks). We therefore used equations 2.4 and 2.5 to generate predictions for Experiments 1 and 2, respectively.

3 Experiments

In this section we present two model evaluations. First, we compared our model predictions against participant judgments of A’s social preference, based on how A’s decisions indirectly affect others (e.g., inferring that A is prosocial if they forego their favorite treat so that others may have it; Experiments 1a & 1b). Second, we compared our model predictions against participant judgments about A’s social knowledge, based on A’s social preferences and decisions (e.g., if we believe A to be mindful of others and observe A taking the last scone, this might suggest that A believes no one else likes scones; Experiment 2).

3.1 Experiment 1a

3.1.1 Participants

40 adult participants with US-based IP addresses were recruited via Amazon Mechanical Turk (mean age=40.1, S.D.=11.8). Four additional participants were recruited but excluded for failing at least one comprehension check question on each of two attempts. (See supplemental materials for survey questions and comprehension checks).

3.1.2 Stimuli

Stimuli consisted of 24 trials (see Fig. 1 for examples). Each trial contained a pictogram showing agents A and B, each with randomly chosen names, and a box containing brownies and cupcakes. Depending on the trial, the box contained either two of each treat (four total), or one of one type and two of the other (three total). The top of the screen showed a description explaining that A had to share two of the treats with B. A’s material preferences were depicted using a verbal description in a thought bubble (either “I prefer [brownies/cupcakes]” or “I like both treats equally”), as were A’s beliefs about B’s material preferences (either “B prefers [brownies/cupcakes],” “B likes both treats equally,” or “I don’t know what B prefers”). Finally, the pictogram showed a box containing whichever treats remained after A made their choice (either two brownies, two cupcakes, or one of each). Altogether, this created a combinatorial space varying the number of initial treats (3 or 4), A’s personal preference, A’s knowledge (or lack thereof) of B’s personal preference, and A’s choice. This yielded 6 unique “3-treat” trials, and 18 unique “4-treat” trials, for a total of 24 trials.

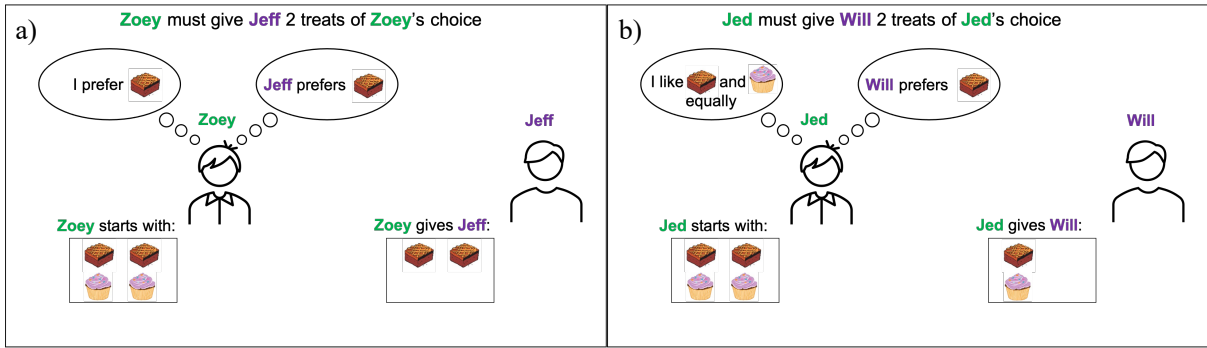


Figure 1: Two examples of stimuli from Experiment 1a. Agent A’s material preferences and beliefs about B’s preferences are shown in the two thought bubbles.

3.1.3 Procedure

Participants first read a cover story in which one agent (A) can take one (in the 3-treat case) or two (in the 4-treat case) treats from a box of cupcakes and brownies.

Participants were told that A knows B will get whichever treats A leaves behind, and were then taught how to interpret the information in each stimulus. Participants were then given a 6-question multiple-choice comprehension check (see Supplemental materials). Participants who failed at least one question on the first try were shown the instructions again, followed by a second attempt at the comprehension check questions. Those who also failed at least one question on the second attempt were excluded from the study. Participants that qualified for the task were shown all 24 trials in a random order upon completing the comprehension checks. For each trial, participants were asked “how much do you think A values [his/her] own preferences” and “how much do you think A values B’s preferences,” and gave each answer using a continuous sliding scale from 0 (“not at all”) to 5 (“a lot”).

3.1.4 Model predictions and alternate models

We generated model predictions using equation 2.4, which takes A’s personal reward function R_A , A’s belief K about B’s preferences, and A’s decision d , and outputs a

probability distribution $P(W|d, R_A, K)$ over A’s social preference W . To implement this, we needed to transform the verbal belief description in each stimulus (e.g.: “I think B prefers cupcakes over brownies”, or “I don’t know what B prefers”) into a probability distribution K over possible reward functions. If A is totally ignorant of B’s preferences, we treat K as a uniform distribution over all possible reward functions. If A believes that B prefers cupcakes, we condition this uniform distribution on the observation that $R_B(\text{cupcake}) > R_B(\text{brownie})$ (i.e.: that B gets more reward from cupcakes than brownies) using Bayesian posterior updating, and similarly if B prefers brownies (see Supplemental Materials for full details). Following this procedure, we used equation 2.4 to compute the distribution $P(W|d, R_A, K)$ for each trial, then used the expected value of this distribution as our prediction of A’s social preference.

To better understand our model predictions, we also implemented three alternative models, each designed to evaluate a different aspect of our main account. First, our computational model assumes that rewards are discounted (i.e., the more brownies you eat, the less enjoyable they become). We therefore implemented a “no-discounting” model, which is identical to the main model, with the difference that the rewards are not discounted.

Our last two models were designed to explore alternative accounts that do not require the full Theory of Mind reasoning present in our main model (see Supplemental Materials for full details). One possibility is that people infer social preferences in terms of how much an agent sacrifices themselves. For example, if A prefers brownies over donuts, giving both brownies to B would reveal a large “sacrifice” in A’s personal reward, regardless of B’s actual preference for brownies. Such an account would only require thinking about the difference between A’s highest possible reward and the true reward they obtained, without considering what A believes about B’s preferences (i.e.: the greater A’s sacrifice, the more prosocial they are). We call this model the

“self-sacrifice” model.

Conversely, it is possible that people assume that, the larger of a reward B gets, the more pro-social A is, regardless of A’s own reward. We therefore implemented a simple “other-outcome” model that equates B’s rewards with A’s social tendencies (i.e.: the greater B’s reward, the more prosocial A is). Classic literature on prosociality and altruism suggests that both of these factors are potential cues to an agent’s prosociality (Swap, 1991), and we included these models to evaluate whether either of these features alone is sufficient to capture participant inferences.

3.1.5 Results

For each of the 24 trials, participants provided separate judgments about self-weight (w_A in the model) and other-weight (w_B in the model), yielding 48 total inferences. Per our pre-registered analysis plan (see osf.io/23rq7) we Z-scored responses within participants and across trials (separately for each parameter), then averaged across participants for each trial. We then Z-scored the mean parameter values outputted by the model, and computed Pearson correlations with participant judgments. As shown in figure 2, main model predictions were highly and significantly correlated with participant responses ($r = .89$, 95% CI(.81, .94)), and were significantly more correlated with participant responses than all three alternate models (defined as a bootstrapped difference in correlations with 95% CI not crossing 0—see Supplemental Materials).

Figure 3 shows four trials that highlight our model’s qualitative behavior and how it matched human intuitions. These cases all show events where A prefers brownies over cupcakes, and encounters a box with two brownies and two cupcakes. When A did not know B’s preference, and A took a brownie and a cupcake (leaving B with a brownie and a cupcake), both participants and our model inferred that A was selfless and cared about B (first example). By contrast, when A took both brownies (leaving B with two

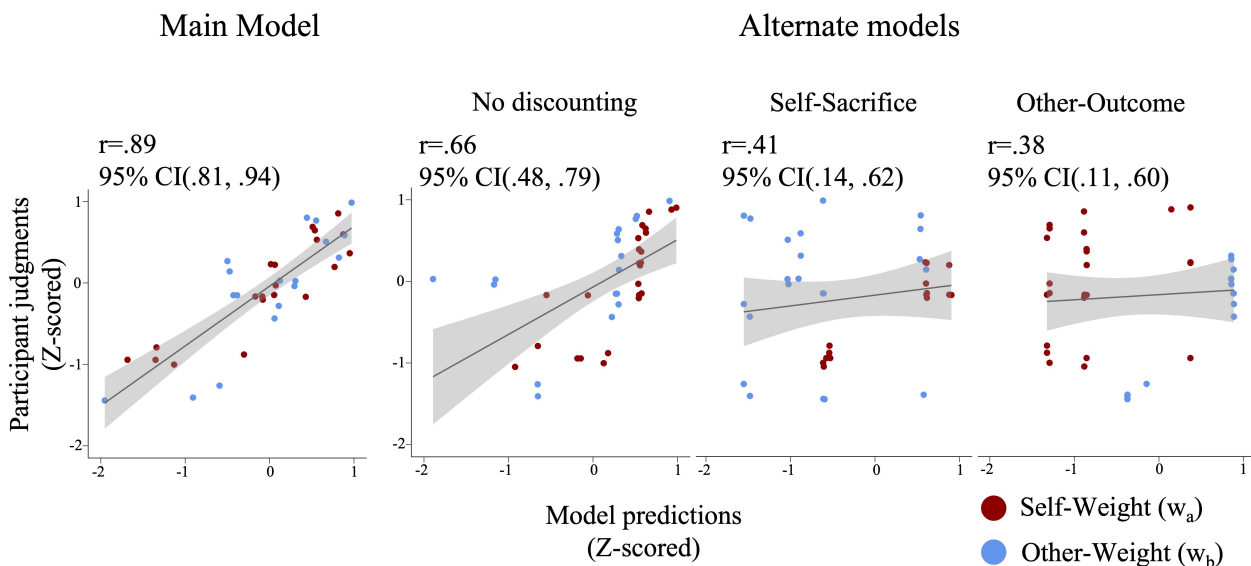


Figure 2: Results from Experiment 1a. Each subplot shows a comparison between predictions generated by one of the four models (including the three alternate models) against participant responses. Each point represents a judgment from a single trial: red dots correspond to self-weight (w_a) and blue dots correspond to other-weight (w_b). The x-axis depicts model predictions (z-scored) and the y-axis depicts participant responses (z-scored within participant and averaged across participants)

cupcakes), our model and participants both inferred that A cared highly about their own rewards and less so about others. This is consistent with previous findings on the “social mindfulness” phenomenon (Van Doesum et al., 2013; Zhao et al., 2021), whereby people associate prosocial attitudes with actions that preserve the greatest number of options for others (in this case, leaving one of each treat for the next person).

However, the third and fourth cases in Figure 3 show how preserving choice for others is not always considered the most socially mindful action. These two trials show identical scenarios to the first two, with the difference that A knew that B preferred cupcakes. When A made the “option-maximizing” decision (leaving one of each for B), participants and the model both inferred that A did not care about B. However, when A kept both brownies, participants and the model inferred that A cared about both B’s preferences and their own preferences, as this action maximized both agents’ rewards

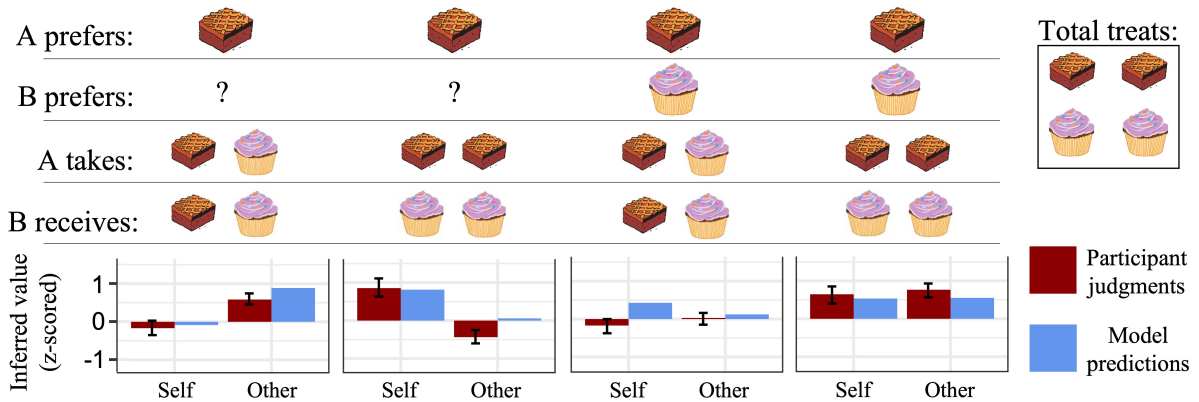


Figure 3: Demonstration of “social mindfulness” effect in both model inferences and participant responses for Experiment 1a. Each column corresponds to a single trial. Rows depict (from top to bottom): which type of treat A preferred, which type of treat A believed that B preferred (question marks indicate that A did not know), which two treats A took from the box, and which two treats were left for B after A’s choice. Bar charts show participant judgments (z-scored within participants and averaged across participants) compared against z-scored model predictions. Vertical bars show bootstrapped 95% confidence intervals

without conflict. Thus, our results suggest that the classic social mindfulness effect is a special case of this more general utility based reasoning, and that the effect only holds when A does not know B’s preference.

Figure 4 shows three trials that highlight the types of cases where the alternative models failed to explain participant judgments. In the first example, A liked both treats equally and did not know B’s preference. When A took both brownies, participants and the main model both inferred that A did not care about B, while the no-discounting model inferred that A was prosocial. Intuitively, this is because, when A is totally uncertain about B’s preferences, B is equally likely to prefer either treat. In the absence of discounting, it follows that B is equally likely to enjoy any combination of two treats: two brownies, two cupcakes, or one of each. Since all three choices provide the same expected reward to B, A’s choice does not reveal how much they care about B’s rewards. When the rewards are discounted, however, the second treat of the same type

contributes a diminished amount to B’s total rewards, so B is more likely to prefer one of each rather than two of the same type. Thus, when A is uncertain about B’s preference, A’s choice only reveals A’s social preference if we assume that rewards are discounted.

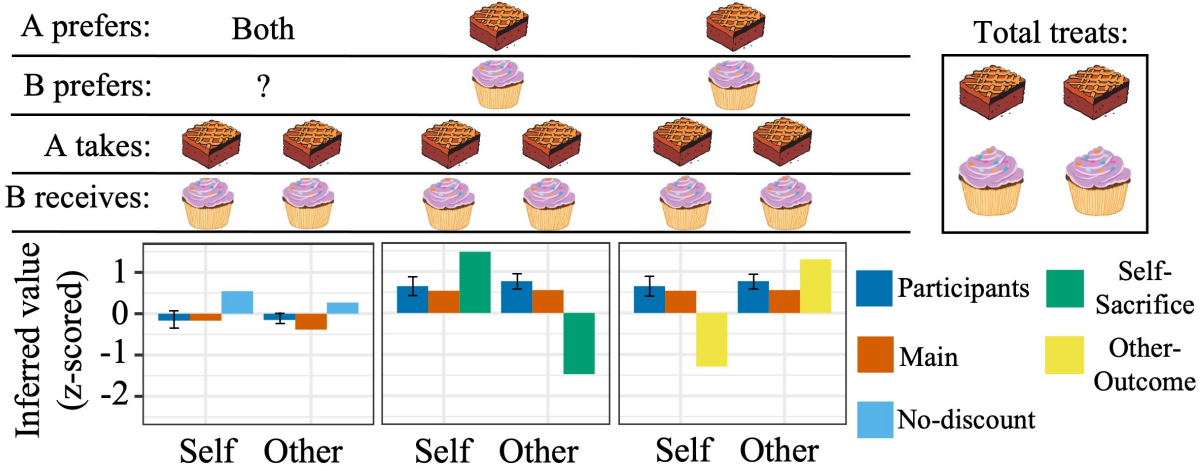


Figure 4: Three trials from Experiment 1a demonstrating the qualitative failures of each alternate model, respectively. Each column shows a single trial and a different alternate model. Rows depict (from top to bottom): which type of treat A preferred, which type of treat A believed that B preferred (question marks indicate that A did not know), which two treats A took from the box, and which two treats were left for B after A’s choice. Each bar chart shows z-scored predictions from one of the three alternate models compared against participant judgments (z-scored within participants and averaged across participants) and z-scored predictions from the main model. Vertical bars show bootstrapped 95% confidence intervals

The next trial (middle column; Figure 4) shows how the “self-sacrifice” model failed to explain participant intuitions. In this case, A preferred brownies and knew that B preferred cupcakes, so when A kept both brownies and left both cupcakes, this decision maximized both agents’ rewards without conflict. Participants and the main model therefore inferred that A likely cared about both their own reward and B’s reward. The self-sacrifice model, however, only reflects that A’s decision maximized their own personal reward, which means that A did not sacrifice anything with this decision. The self-sacrifice model therefore inferred that A is very selfish (i.e., $w_A \gg w_B$).

Finally, the last trial illustrates how the “other-outcome” model failed to explain participant intuitions. In this case, A preferred brownies and knew that B preferred cupcakes, so when A kept both brownies and left both cupcakes, this decision maximized both agents’ rewards without conflict. Participants and the main model therefore inferred that A likely cared about both their own reward and B’s reward about equally. However, because the outcome model focuses only on the total reward that B obtains, it inferred that A cared much more about B’s reward than their own. That is, despite correctly identifying that A’s decision maximized B’s reward, this model failed to account for the fact that A’s decision also maximized their own reward. Thus, these results demonstrate a) that participants attended to both the costs A incurred and the expected rewards that B received, and b) that participant inferences reflected an assumption of discounted utilities, particularly in cases where A does not know B’s preference.

3.2 Experiment 1b

Experiment 1a showed that people can infer social preferences by considering how an agent’s personal choices affect others. However, this experiment had two limitations. First, in Experiment 1a, the stimuli always read “A must give B 2 treats of A’s choice.” This wording therefore highlighted that A was directly making a choice that would affect both members. It is therefore possible that our model fit participant judgments only because participants were explicitly told that A was forced to choose which rewards B would get, which conflicts with our primary aim of evaluating social inference from indirect evidence of prosociality. To address this concern, we modified the cover story in Experiment 1b to a situation where A was clearly making a personal choice that indirectly affected B.

A second limitation of Experiment 1a is that it revealed people’s preferences using qualitative descriptions, which restricted us to only three preference categories

(preferring one treat or the other, or liking both equally). To test whether people can make these inferences from more granular information, the stimuli for Experiment 1b depicted preferences using quantitative choice histories, rather than verbal descriptions

3.2.1 Participants

40 adult participants with US-based IP addresses were recruited via Amazon Mechanical Turk (mean age=40.2, S.D.=10.3). 6 additional participants were recruited but excluded for failing at least one comprehension check question on each of two attempts.

3.2.2 Stimuli

The stimuli for Experiment 1b depicted the same variables as 1a. The key difference was that A’s personal preferences and social knowledge were depicted indirectly, using numerical choice histories, rather than direct verbal descriptions. To this end, participants were told that the agents often visit the same bakery to buy themselves treats, and will sometimes observe what other agents buy for themselves. In lieu of verbal preference and belief descriptions, each trial displayed A’s choice history over the previous 5 visits to the bakery, and what A had observed B buy for themselves over their previous 5 visits to the bakery (or a column of question marks, indicating that A did not visit the bakery at the same time as B and therefore did not see what B purchased).

Figure 5 shows examples of stimuli from this experiment.

By using quantitative choice histories over 5 different visits, Experiment 1b yielded significantly more possible trials than 1a. To select 25 trials, we chose 6 conditions for A’s choice history (one reflecting a strong preference for each treat, one reflecting a weak preference for each treat, and two reflecting ambiguous preferences between treats) and 7 conditions for A’s observation of B’s choice history (the same 6 as A, plus an “A did not observe B” ignorance condition). Combining this with the three choice conditions (two

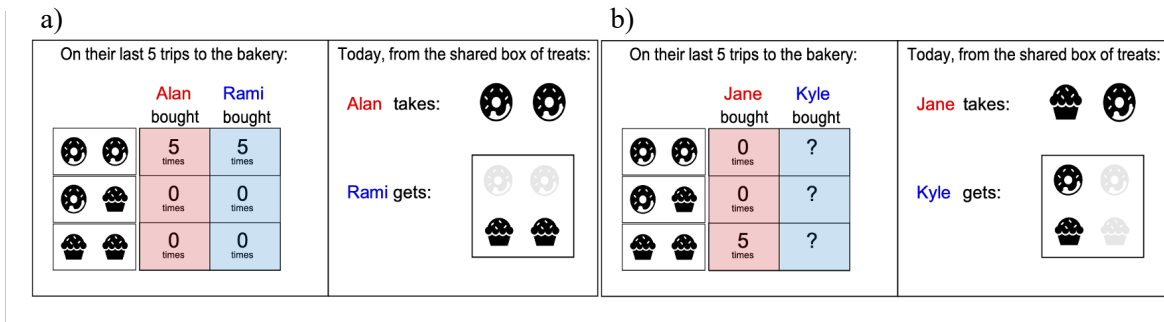


Figure 5: Two examples of stimuli from Experiment 1b. Agent A’s material preferences and beliefs about B’s preferences are depicted as numerical choice histories. Question marks indicate that A does not know the choice history of that agent.

of one treat, two of the other, or one of each) yielded 126 total conditions. We then generated model predictions for all 126 conditions, chose 5 social preference categories (very selfish, weakly selfish, roughly egalitarian, weakly altruistic, and very altruistic), and for each category, selected the 5 trials for which the model’s predictions most closely fit that category (e.g.: for “very selfish,” we chose the 5 trials for which the model inferred the 5 highest $w_A/(w_A + w_B)$ ratios). This yielded 25 total trials, with 5 for each social preference category.

3.2.3 Procedure

Experiment 1b was similar to Experiment 1a with a modified cover story: participants were introduced to an office setting where a manager had left a box of four snacks for two employees to enjoy. Participants were told that one agent (A) worked the morning shift, while the other (B) worked the afternoon shift, so that A would get to choose their two snacks first. Participants were additionally told that A knew that B works the afternoon shift.

Participants were then taught how to interpret the stimuli (Figure 5) and then given two chances to pass a 6 question comprehension check ensuring that they were paying attention and could correctly interpret the stimuli (see Supplemental materials for

details). Participants who failed at least one question on the first attempt were shown the instructions a second time, followed by a second attempt at the comprehension checks. Participants who failed one or more questions on the second attempt were excluded from the study. Participants who passed the questionnaire were then shown all 25 trials in a random order. For each trial, participants are asked to rate a) how much A cares about him/herself and b) how much A cares about B. Both responses were entered on a continuous sliding scale from 0 (doesn't care at all) to 5 (cares a lot).

3.2.4 Model predictions and alternative models

Model predictions for Experiment 1b were generated using the same procedure as Experiment 1a (see section 2.2 for full details). To implement the model, we needed to transform the numerical choice histories in the stimuli into probability distributions over possible reward functions. To this end, we transformed the “full ignorance” condition (in which A has no evidence of B’s preferences) into a uniform distribution over reward functions. The distribution for each non-ignorance condition was then computed by conditioning this uniform prior on the agent’s observed choices, using Bayesian posterior updating (see Supplemental Materials for full details). Finally, we generated three additional sets of predictions using the same three alternate models described in Experiment 1a (“no discounting,” “self-sacrifice,” and “other-outcome”).

3.2.5 Results

Data was analyzed in the same way as Experiment 1a, in accordance with our pre-registration (see osf.io/pn6zs). As shown in figure 6, our model closely tracked average participant responses, with correlations similar to the results of Experiment 1a ($r = .87$, 95%CI (.77, .92)). This reinforces our findings from Experiment 1a, and further suggests that people can integrate graded, quantitative information about preferences

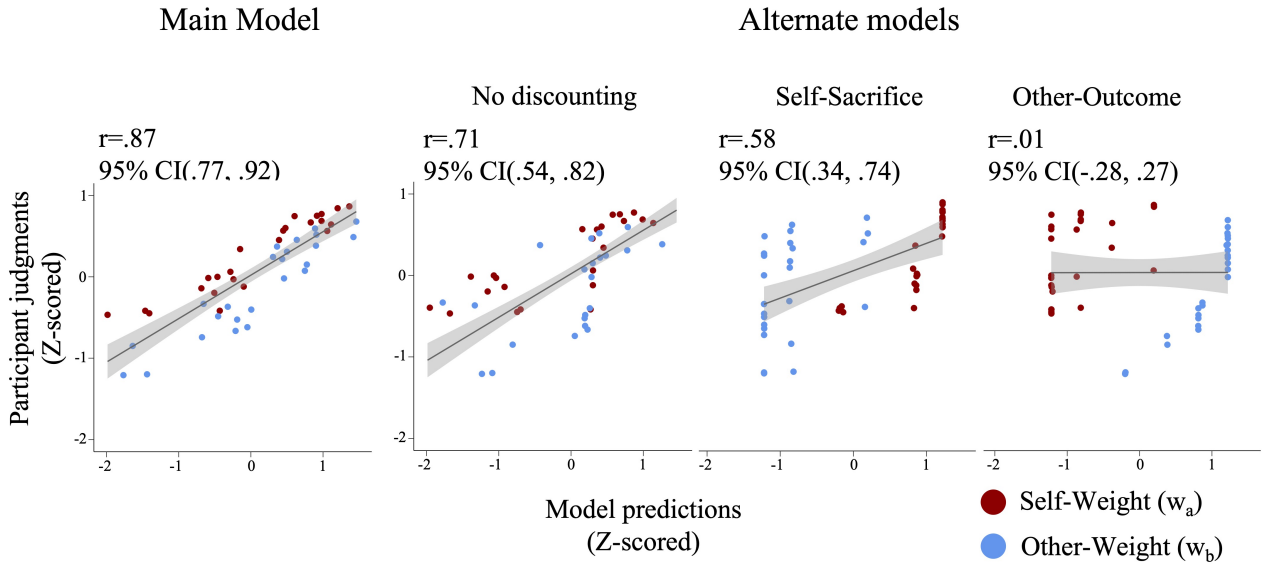


Figure 6: Results from Experiment 1b. Each subplot shows a comparison between predictions generated by one of the four models (including the three alternate models) against participant responses. Each point represents a judgment from a single trial: red dots correspond to self-weight (w_a) and blue dots correspond to other-weight (w_b). The x-axis depicts model predictions (z-scored) and the y-axis depicts participant responses (z-scored within participant and averaged across participants)

into their social inferences. Furthermore, whereas the stimuli from Experiment 1a depicted A giving two treats to B, the cover story and stimuli from Experiment 1b made it clear that A chose which treats to keep for themselves, rather choosing which treats to give to B. This demonstrates that people make similar inferences when A’s decision only affects B’s outcome indirectly.

Also similar to Experiment 1a, a bootstrapped difference in correlations revealed that all three alternate models were significantly less correlated with participant data than the main model, (“no discounting”: $r = .71$, 95%CI (.54, .82); “self-sacrifice”: $r = .58$, 95%CI (.34, .74); “other-outcome”: $r = .01$, 95%CI (-.28, .27)). These results further support that participant judgments reflected a discounted expected utility computation (at least in cases where A was ignorant of B’s preference), and were sensitive to both the cost that A incurred and the benefit that A indirectly yielded to B.

3.3 Experiment 2

Experiment 1 established that people can use information about an agent’s personal choices and social knowledge to infer their social preferences. However, our computational framework further predicts that people should be able to perform other kinds of inferences from the same generative model of agent behavior. In Experiment 2, we tested a reverse inference to the one from Experiment 1. Here we considered cases where an agent’s social preferences were known, and tested whether people could use this information to infer the agent’s social knowledge based on how they behave (e.g., if A is known to care about others, but takes a seemingly inconsiderate action, what does that reveal about their knowledge?).

3.3.1 Participants

40 adult participants with US-based IP addresses were recruited via Amazon Mechanical Turk (mean age=38.7, S.D.=10.6). Five additional participants were recruited but excluded for failing at least one comprehension check question on each of two attempts..

3.3.2 Stimuli

Experiment 2 comprised 25 trials, each of which depicted two agents A and B, and a box containing two burgers and two lettuce wraps (four items total) from which each agent could take two items. The stimuli depicted A’s choice of items to keep, A’s personal preference (“I prefer [burgers/veggie wraps] over [veggie wraps/burgers],” or “I like burgers and veggie wraps equally”), as well as A’s social preference towards B (“A cares about him/herself more than he/she cares about B,” “A cares about B more than he/she cares about him/herself,” and “A cares about him/herself and B about equally”). The stimuli also depicted A’s choice, but unlike the Experiment 1 stimuli, did not provide any information regarding A’s knowledge of B’s preference. See figure 7 for examples of

stimuli from Experiment 2.




Figure 7: Examples of stimuli from Experiment 2. Each stimulus depicts the chooser’s personal preference (either “I prefer [burgers/veggie wraps]” or “I like both equally”), social preference (which agent A cares about more), and choice (which two items to keep).

Each trial varied A’s personal preference (3 conditions), B’s social preference towards A (3 conditions), and A’s choice of food items (3 conditions: two of one item, two of the other, or one of each), yielding 27 total trials. Before collecting any data, we applied our inference model to each trial, computed the model’s likelihood of A’s choice in that trial, and eliminated any trial in which A’s choice was highly unlikely ($< 5\%$). We took this step to avoid showing participants any scenarios that lacked a coherent social interpretation (e.g.: A is stated to be selfish but behaves in an altruistic way). This process eliminated 2 trials, leaving the 25 included in Experiment 2.

3.3.3 Procedure

Participants were first shown a series of instructions providing a cover story and explaining how to interpret the stimuli. The cover story explained that A and B work in an office that provides lunch for its employees. Each morning the manager brings a box of food items (hamburgers and veggie wraps), and each employee is allowed to take two items. The story further explains that A’s lunch break is earlier than B’s, so that A gets first choice. After reading the instructions, participants were given an 8 question comprehension check. Participants who failed at least one question were shown the

instructions a second time, followed by a second attempt at the comprehension check. Participants who failed at least one question on the second attempt were excluded from the study.

All other participants were then shown all 25 trials in a random order. After each trial, participants were asked two questions. First, participants were asked whether they thought that A knew what B preferred, and responded using a continuous slider from 0 (definitely doesn't know) to 1 (definitely knows). Second, participants were told to suppose that A did, in fact, know what B preferred, and asked what A thought B preferred. Participants responded using a continuous slider that depicted one food item (hamburger) on the leftmost side of the scale, the other food item (veggie wrap) on the rightmost side of the scale, and "B likes both equally" in the center of the scale. Participant responses to the second question were converted to a real number in the interval $[-1, 1]$ for analysis.

3.3.4 Model predictions and alternate models

Predictions for Experiment 2 were generated using equation 2.5, which takes A's reward function R_A , A's social preference W , and A's decision d , and returns a probability distribution $P(K|d, W, R_A)$ over A's beliefs K . To implement this model, we converted the verbal descriptions of A's preferences into a probability distribution over reward functions following the same method from Experiment 1a. We used a similar method to convert the verbal description of A's social preferences (e.g.: "A cares more about B than about [him/her]self") into a probability distribution over social preferences: we started with a uniform prior distribution over all possible social preferences $W = (w_A, w_B)$, then used Bayesian posterior updating to condition on the observation that $w_A > w_B$ (for "A cares more about [him/her]self"), $w_A < w_B$ (for "A cares more about B"), or $w_A = w_B$ (for "A cares about both equally") (see supplemental materials for full details).

In addition to the main model, we generated predictions using three alternate models. First, we tested a “no-discounting” model, to evaluate whether people’s belief inferences also reflected an expectation of discounted rewards. Second, we wished to evaluate whether people were actually incorporating A’s reward function and decision into their inferences, rather than solely basing their inferences on A’s social preference (e.g.: assuming that A is more likely to know B’s preference when A cares about B, but ignoring all other information– see 2.2). To this end, our second alternate model (“prior-only”) generated predictions based solely on A’s social preference, without incorporating any other information in the stimulus. Finally, to evaluate whether people did, in fact, associate prosocial attitudes with an increased likelihood of knowing other agents preferences, we tested a “uniform-prior” model, which performs the same inference as the full model, but assumes that A is equally likely to know or not know B’s preference, regardless of A’s social preference towards B.

3.3.5 Results

Participant judgments and model predictions were processed in the same way as Experiments 1a and 1b, as preregistered (see osf.io/sf6bu). As shown in figure 8, main model predictions closely matched average participant judgments for both parameters ($P(\text{Knows?})$: $r = .92$, 95% CI(.83, .97), $P(\text{Belief})$: $r = .92$, 95% CI(.83, .97), combined: $r = .91$, 95% CI(.84, .95)), and were significantly more correlated with participant judgments than all three alternate models, defined as a bootstrapped difference in correlations with 95% CI not crossing 0 (see Supplemental Materials).

Participants were therefore able to leverage information about A’s personal preferences, social preferences, and decisions to infer both a) whether A has knowledge of B’s preference and b) what that knowledge is (assuming A has any) in a fashion that reflects expected utility computations. To better illustrate how our model captured

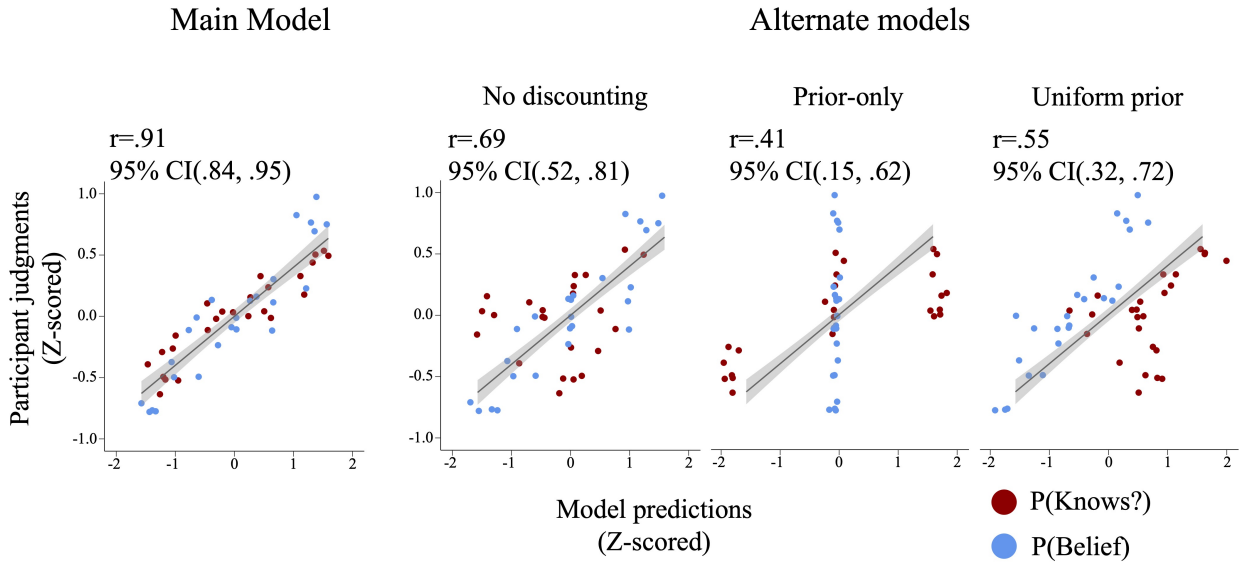


Figure 8: Results from Experiment 2. Each subplot shows a comparison between predictions generated by one of the four models (including the three alternate models) against participant responses. Each point represents a judgment from a single trial: red dots correspond to the probability that A knows B’s preference, and blue dots correspond A’s belief about B’s preference (assuming A *does* know). The x-axis depicts model predictions (z-scored) and the y-axis depicts participant responses (z-scored within participant and averaged across participants)

participant intuitions, figure 9 highlights four trials that demonstrate two key qualitative predictions of our model. The first two trials demonstrate how A’s social preference affected people’s inferences about A’s social knowledge. In these two events (first two cases in figure 9), A took the two burgers which she preferred, and left the two veggie wraps for B. When A was known to be selfish, A’s decision did not reflect their knowledge or ignorance about B’s preferences at all, since A would most likely make the selfish decision regardless of B’s preferences. In this case, the model relied solely on the knowledge prior (i.e.: the expectation that selfish agents are less likely to track the preferences of others), and thus inferred that a selfish A most likely did not know B’s preferences, consistent with participants’ inferences. By contrast, when A was known to be prosocial (second case in Figure 9), but took an apparently selfish action (keeping

both hamburgers), both participants and the model inferred a higher likelihood that A believed B preferred veggie wraps.

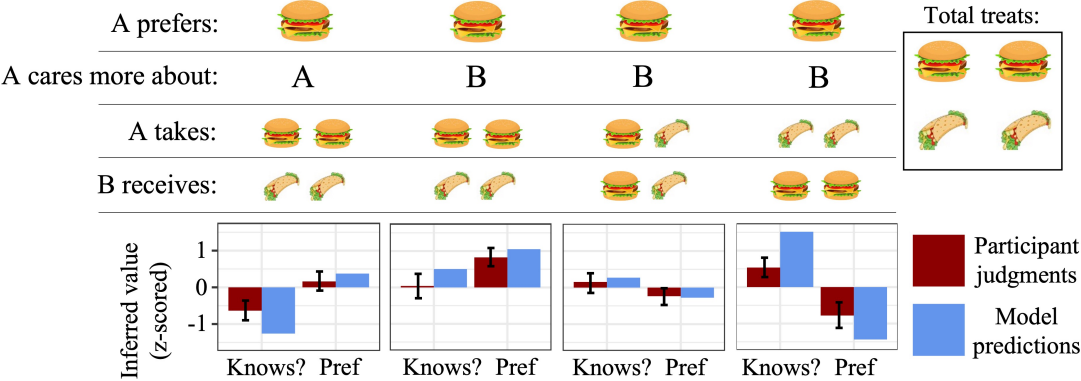


Figure 9: Four trials from Experiment 2 demonstrating two of the main qualitative phenomena captured by our main model. Each column corresponds to a single trial. Rows depict (from top to bottom): which type of treat A preferred, which type of treat A believed that B preferred (question marks indicate that A did not know), which two treats A took from the box, and which two treats were left for B after A’s choice. Each bar chart shows z-scored model predictions compared against participant judgments (z-scored within participants and averaged across participants). Vertical bars show bootstrapped 95% confidence intervals

This result might suggest that participants simply inferred that a prosocial A is generally more likely to know B’s preference, but the last two examples, wherein A was prosocial, showed this was not the case. When A took the “choice-maximizing” action, (i.e.: taking one of each– third case in figure 9), both the model and participants expressed uncertainty about A’s knowledge. Even though A was prosocial, which is associated with a higher overall likelihood of knowing B’s preference, A’s action suggested that A did not, in fact, know B’s preference, resulting in greater uncertainty about A’s social knowledge. However, when A took both of their less preferred snacks (veggie wraps) and left their more preferred snacks (hamburgers) for B (fourth case in figure 9), participants and the model both inferred that A knew what B preferred. Thus, participants leveraged an expectation of prosocial behavior to infer A’s knowledge or ignorance of B’s preference, as predicted by the model. This demonstrates that

participants could not only infer social preferences from actions and social knowledge, but also infer social knowledge from actions and social preferences.

The results of our three alternate models also emphasized how each of the core assumptions in our main model was necessary to account for participant intuitions. First, the “no-discounting” model ($r = .69$, 95% CI(.52, .81)) failed to predict participant inferences about A’s knowledge or ignorance of B’s preferences (i.e. $P(\text{Knows?})$). This is because, when A did not know B’s preferences, all three possible decisions yielded the same expected reward to B in the absence of discounting, so the “no-discounting” model could not infer whether or not A knows B’s preference (see results section of Experiment 1a for an extended explanation of this point). Second, the significantly lower correlation of the “prior-only” model ($r = .41$, 95%CI (.15, .62)) suggested that people were not just assuming that prosocial agents are more likely to know the preferences of others, but that they also incorporated information about agents’ rewards and decisions into their inferences. Finally, the “uniform prior” model, which assumed that A’s social preference has no effect on the likelihood that A knows B’s preference, had a significantly lower correlation with participant data than the main model ($r = .55$, 95%CI (.32, .72)). This suggests participants did expect that people with prosocial preferences were more likely to know the preferences of others. Thus, these results strongly suggest that a) participants could infer an agent’s social knowledge based on their social preferences and decisions, b) that these inferences reflected the same discounted utility computations underlying the social preference inferences from Experiment 1, and c) participants assumed that prosocial agents were more likely to know the preferences of others, but still attended to the costs they incurred and rewards they yielded when making these social judgments.

4 Discussion

How do we detect potential social partners who are kind, empathetic, and supportive in everyday social life? Here we examined the idea that people can detect these traits from indirect evidence of prosociality revealed by agents’ personal choices. We presented a computational model of this capacity, focusing on settings where agents’ personal choices can signal that they deployed theory of mind to indirectly benefit others. At the heart of this model is an assumption that people expect agents to maximize their own utilities, including the utility they assign to others’ welfare.

Across three experiments, we found converging support for our computational account. In Experiment 1a and 1b, we found that people leverage information about an agent’s personal preferences, choices, and knowledge of others’ preferences to infer that agent’s social preferences. We found that people can make these same evaluations using both direct, qualitative evidence of preferences (Experiment 1a) as well as indirect, quantitative evidence of preferences from choice histories (Experiment 1b). In Experiment 2, we found that the inverse is also true: people can infer an agent’s knowledge (or ignorance) of others’ preferences based on the agent’s social preferences and personal choices, and these inferences reflect the same underlying utility computations as our initial experiments. Together, this work supports the idea that people can draw on indirect evidence of prosociality (making personal choices that indirectly fulfill others’ preferences) to discern an agent’s social preferences. This provides strong support for a mentalistic, utility-based account of social inference, and emphasizes the key role of detecting when an agent is using theory of mind in the service of others.

A second, somewhat more surprising finding of our experiments was the role of discounting in producing the intuitive “social mindfulness” effect in these tasks. In particular, previous research suggests that, in the absence of any information about B’s

preferences, there is a strong intuition that the prosocial decision is the one that maximizes the number of available options for B (i.e.: for A to leave one of each type of item for B) (Van Doesum et al., 2013; Zhao et al., 2021). This effect was also found across both of our studies, despite the fact that B did not technically have a choice to make in our task (as B automatically got everything that A left behind). However, our analysis revealed that this social mindfulness effect was not predicted by an alternate model in which reward functions were not discounted. That is, when A was totally ignorant of B’s preference, then the expected reward to B was the same regardless of A’s choice, so the “no-discounting” model could not distinguish A’s social preference in these cases. Only when the individual reward functions were discounted did the social mindfulness effect emerge. This finding highlights the value of computational models in social psychology: while the social mindfulness intuition is fairly salient and immediate, it appears to be driven in these cases by a separate intuition regarding diminishing marginal utilities of consumable goods. This also points to potential future studies to further explore the role of discounting in social judgments. If, for example, we translate the task to another domain with no diminishing marginal utilities (e.g.: if I like one genre of book over another, I might enjoy a second book from the same genre just as much as the first one), we might expect qualitatively different intuitions about prosocial behavior in that context.

While these results are promising, they also have some limitations. First, our work focused on tasks in which choice were concrete and involved only a narrow set of choices. In real-world scenarios, however, people often encounter ‘tests’ of social mindfulness where the choice options themselves are not clearly visible, but must be represented abstractly at the time of choice. For instance, choosing to keep the seat next to one clear on a crowded bus, choosing to keep one’s desk clean in a shared office space, or choosing to shop at a local farmer’s market as opposed to a large chain supermarket. Generalizing

our framework to a wider range of more naturalistic task settings is therefore an important step for future research.

Additionally, the present work considered individual instances of choice outside of a broader social context. In analogous real life settings, numerous factors may color a person’s social inferences. One important factor not considered in the present work is how reputational concerns might such social judgments. In many situations, decisions that benefit others (even in subtle, indirect ways) are driven not just by altruistic motives, but also from a desire to improve (or maintain) a positive reputation with others. Indeed, existing work suggests that social behavior is highly influenced by reputational concerns (Ariely et al., 2009; Asaba and Gweon, 2022), and that people are sensitive to such influences when making inferences about why agents engage in prosocial behavior (Barasch et al., 2014; Carlson and Zaki, 2018). Thus, an important extension of our model would be to account for how observers consider reputational factors or other kinds of social pressures that might shape an agent’s behavior.

Another crucial factor is the perceiver’s own model of the social environment in which the interactions take place. For example, group identities and intergroup dynamics can powerfully shape people’s perceptions and expectations about how agents should (or do) treat each other (Lees and Cikara, 2020; Waytz et al., 2014; Rhodes, 2013). Moreover, social inferences surrounding a helpful act might be colored by very different expectations if an observer believes an actor and recipient belong to the same ingroup, versus if one is an ingroup member and one is an outgroup member. Accounting for the observer’s knowledge about the relevant social environment is another important next step for future research. Conversely, it may also be possible to perform the reverse inference- that is, to reason about the underlying structure of social groups by observing which agents do or do not act prosocially towards one another. Suppose, for example, an observer believes that ingroup members are more likely to behave prosocially towards

each other than towards outgroup members. The observer could apply a model like the one proposed here to infer dyadic social preferences between pairs of individuals, then extrapolate the broader social structures from these dyadic relations (Davis et al., 2022; Gershman and Cikara, 2020).

How do people detect potential social partners? Here we proposed that part of the solution is a sensitivity to how agents make personal decisions in situations where their choices might indirectly benefit others. Across a series of experiments, we showed that people not only leverage information about an agent’s personal choices to infer their social preferences, but also leverage information about an agent’s social preferences to infer their knowledge about the preferences of others. These findings showcase the promise of formal approaches for a deeper understanding of social psychology. Our future work will further explore the interactions between inferences about dyadic social relations and broader social structures, as well as the role of reputational concerns in such inferences. Indeed, by more closely examining how factors such as reputation, group dynamics, and choice representations are integrated into people’s social inferences, future work will yield richer accounts of how people discern others’ social preferences, and in turn, how people decide who to trust and befriend in social life.

References

- Ariely, D., A. Bracha, and S. Meier (2009). Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review* 99(1), 544–55.
- Asaba, M. and H. Gweon (2022). Young children infer and manage what others think about them. *Proceedings of the National Academy of Sciences* 119(32), e2105642119.
- Baker, C. L., J. Jara-Ettinger, R. Saxe, and J. B. Tenenbaum (2017). Rational

- quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour* 1(4), 0064.
- Baker, C. L., R. Saxe, and J. B. Tenenbaum (2009). Action understanding as inverse planning. *Cognition* 113(3), 329–349.
- Barasch, A., E. E. Levine, J. Z. Berman, and D. A. Small (2014). Selfish or selfless? on the signal value of emotion in altruistic behavior. *Journal of personality and social psychology* 107(3), 393.
- Baucells, M. and R. K. Sarin (2010). Predicting utility under satiation and habit formation. *Management Science* 56(2), 286–301.
- Carlson, R. W., Y. E. Bigman, K. Gray, M. J. Ferguson, and M. Crockett (2022). How inferred motives shape moral judgements. *Nature Reviews Psychology*, 1–11.
- Carlson, R. W. and J. Zaki (2018). Good deeds gone bad: Lay theories of altruism and selfishness. *Journal of Experimental Social Psychology* 75, 36–40.
- Crockett, M. J., Z. Kurth-Nelson, J. Z. Siegel, P. Dayan, and R. J. Dolan (2014). Harm to others outweighs harm to self in moral decision making. *Proceedings of the National Academy of Sciences* 111(48), 17320–17325.
- Davis, I., Y. Dunham, and J. Jara-Ettinger (2022). Inferring the internal structure of social collectives. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 44.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Fehr, E. and U. Fischbacher (2002). Why social preferences matter—the impact of non-selfish motives on competition, cooperation and incentives. *The economic journal* 112(478), C1–C33.

- Fiske, S. T., A. J. Cuddy, and P. Glick (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences* 11(2), 77–83.
- Frith, C. D. and U. Frith (2012). Mechanisms of social cognition. *Annual review of psychology* 63, 287–313.
- Gergely, G. and G. Csibra (2003a). Teleological reasoning in infancy: The naive theory of rational action. *Trends in cognitive sciences* 7(7), 287–292.
- Gergely, G. and G. Csibra (2003b). Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences* 7(7), 287–292.
- Gershman, S. J. and M. Cikara (2020). Social-structure learning. *Current Directions in Psychological Science* 29(5), 460–466.
- Hamlin, J. K. and K. Wynn (2011). Young infants prefer prosocial to antisocial others. *Cognitive development* 26(1), 30–39.
- Hartman, R., W. Blakey, and K. Gray (2022). Deconstructing moral character judgments. *Current Opinion in Psychology* 43, 205–212.
- Jara-Ettinger, J., H. Gweon, L. E. Schulz, and J. B. Tenenbaum (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in cognitive sciences* 20(8), 589–604.
- Jara-Ettinger, J., L. E. Schulz, and J. B. Tenenbaum (2020). The naive utility calculus as a unified, quantitative framework for action understanding. *Cognitive Psychology* 123, 101334.
- Jern, A. and C. Kemp (2014). Reasoning about social choices and social relationships. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 36.

- Jern, A., C. G. Lucas, and C. Kemp (2017). People learn other people’s preferences through inverse decision-making. *Cognition* 168, 46–64.
- Kiley Hamlin, J., T. Ullman, J. Tenenbaum, N. Goodman, and C. Baker (2013). The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Developmental science* 16(2), 209–226.
- Lees, J. and M. Cikara (2020). Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts. *Nature Human Behaviour* 4(3), 279–286.
- Lucas, C. G., T. L. Griffiths, F. Xu, C. Fawcett, A. Gopnik, T. Kushnir, L. Markson, and J. Hu (2014). The child as econometrician: A rational model of preference understanding in children. *PloS one* 9(3), e92160.
- McClintock, C. G. and S. T. Allison (1989). Social value orientation and helping behavior 1. *Journal of Applied Social Psychology* 19(4), 353–362.
- Morelli, S. A., Y. C. Leong, R. W. Carlson, M. Kullar, and J. Zaki (2018). Neural detection of socially valued community members. *Proceedings of the National Academy of Sciences* 115(32), 8149–8154.
- Murphy, R. O. and K. A. Ackermann (2011). A review of measurement methods for social preferences. *ETH Zurich Chair of Decision Theory and Behavioral Game Theory Working Paper*.
- Powell, L. J. (2021). Adopted utility calculus: Origins of a concept of social affiliation. *Perspectives on Psychological Science*, 17456916211048487.
- Rhodes, M. (2013). How two intuitive theories shape the development of social categorization. *Child Development Perspectives* 7(1), 12–16.

- Swap, W. C. (1991). When prosocial behavior becomes altruistic: An attributional analysis. *Current Psychology* 10(1), 49–64.
- Tauber, S. and M. Steyvers (2011). Using inverse planning and theory of mind for social goal inference. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 33.
- Ullman, T., C. Baker, O. Macindoe, O. Evans, N. Goodman, and J. B. Tenenbaum (2009). Help or hinder: Bayesian models of social goal inference. In *Advances in neural information processing systems*, pp. 1874–1882.
- Van Doesum, N. J., R. E. de Vries, A. A. Blokland, J. M. Hill, D. M. Kuhlman, A. W. Stivers, J. M. Tybur, and P. A. Van Lange (2020). Social mindfulness: Prosocial the active way. *The Journal of Positive Psychology* 15(2), 183–193.
- Van Doesum, N. J., D. A. Van Lange, and P. A. Van Lange (2013). Social mindfulness: skill and will to navigate the social world. *Journal of Personality and Social Psychology* 105(1), 86.
- Van Lange, P. A. (1999). The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation. *Journal of personality and social psychology* 77(2), 337.
- Van Lange, P. A. and N. J. Van Doesum (2015). Social mindfulness and social hostility. *Current opinion in behavioral sciences* 3, 18–24.
- Waytz, A., L. L. Young, and J. Ginges (2014). Motive attribution asymmetry for love vs. hate drives intractable conflict. *Proceedings of the National Academy of Sciences* 111(44), 15687–15692.

Wellman, H. M. (2014). *Making minds: How theory of mind develops*. Oxford University Press.

Zhao, X., X. Zhao, H. Gweon, and T. Kushnir (2021). Leaving a choice for others: Children's evaluations of considerate, socially-mindful actions. *Child development* 92(4), 1238–1253.