

Heterogeneous Tastes and Social (Mis)Learning

Tristan Gagnon-Bartsch
Florida State

Benjamin Bushong*
Michigan State

Abstract

How do people learn from others' actions when those people may have differing tastes? We present data from two experiments in which properly extracting information from other people's actions requires an observer to account for how her predecessors' tastes may have influenced those actions. We find support for social learning that obeys some basic comparative statics predicted by the rational model. However, we also find significant and systematic departures. Participants seemingly over-infer from others' behavior when that behavior is weakly predictive of the underlying state and under-infer from others when their behavior is strongly predictive. This pattern of inferences is consistent with participants holding inaccurate beliefs about others where they over-weight the likelihood that others have tastes similar to their own. Information about others' tastes does not eliminate these biases in inferences.

JEL Classification: C91, D91, D84.

Keywords: Social learning, heterogeneous preferences, projection bias, experiment.

*E-mails: tgagnonbartsch@fsu.edu and bbushong@msu.edu. We thank Daniel Benjamin, David Gill, Victoria Prowse, Matthew Rabin, Antonio Rosato, Joshua Schwartzstein, Oren Ziv, and audiences at Chapman University, Florida State University, Michigan State University (Economics and Finance), Middlebury College, Purdue University, Texas A&M University, University of Nottingham, University of Oregon, and the 2023 ESA World and North America Meetings for comments. This research was reviewed by the Michigan State Human Research Protection Program.

1 Introduction

Others’ actions can provide a valuable source of information. But there are exceptionally few instances where people’s actions do not in some way depend on their idiosyncratic preferences. This poses a challenge for extracting information from their choices: proper inference requires us to account for how others’ actions were driven by private information versus their particular tastes. Did your neighbor buy a new model of car because she had information that it performs well, or did she simply have a taste for that style or brand? Is the new restaurant really that good, or are many people just partial to that type of cuisine? We designed two experiments to investigate how agents’ heterogeneous tastes—and their perceptions of others’ tastes—may impede information transmission.¹

The following example illustrates the basic logic of our experiment. Suppose a prospective home-buyer observes another interested party back out of a purchase agreement. This observer must disentangle whether the would-be buyer walked away because she received bad information about the quality of the house (e.g., a failed inspection) or because she simply soured on aesthetic features (e.g., its architectural style). If the observer over-estimates the likelihood that the other person still favors these features, he will conclude that the potential buyer likely received some negative private information about the house—otherwise the deal would have gone through. Such an observer’s beliefs about the quality of the house will then over-react to the other person’s action. In contrast, if the observer over-estimates the likelihood that the person soured on the aesthetics, his beliefs will under-react: he thinks the buyer would have walked away regardless of the inspection. This example highlights how the information conveyed by others’ actions depends on the observer’s perceptions of others’ tastes, and how inaccurate perceptions can distort social learning—one’s perception of the distribution of tastes dictates how they interpret others’ actions.

In this paper, we examine the degree to which people successfully engage in social learning when faced with naturally occurring heterogeneity in tastes. To do so, we employed a pair of three-stage field-in-the-lab experiments with sequential observational learning in which participants faced uncertainty about the nominal value of gift cards to various American businesses.

In the first stage, we introduced participants to the businesses used in our study. We then asked a few simple survey questions about participants’ tastes toward these businesses and their perceptions of others’ tastes. These surveyed preferences exhibited significant heterogeneity.

In the second stage, participants made a series of choices between a gift card to one of these businesses and a cash bonus $x \in \{\$30, \$40, \$50\}$. We call the pair of cash bonus and business

¹Some models of social learning and herding (e.g., [Banerjee, 1992](#); [Bikhchandani et al., 1992](#)) focus on idealized settings with common preferences to demonstrate how information may fail to aggregate even in the absence of realistic frictions such as heterogeneity in preferences. However, these frictions may be of first-order importance if they cause people to systematically misinterpret others’ private information.

faced, (x, k) , a “decision problem,” and participants faced a variety of decision problems. In each, participants were uncertain about the value of the card—it contained either \$20 or \$100—and were given a private, binary, noisy signal of this underlying state. For each business, participants made choices after both possible signal realizations.

In the third, “social-learning” stage, participants observed others’ privately-informed choices (either to accept the gift card or take the cash bonus) across a variety of decision problems. We then elicited their beliefs about which of the two signals their predecessor received. The fundamental challenge—which mirrors the home-buyer example above—is that these observers needed to consider, for instance, whether the prior actor rejected a card because they received a bad signal or because they simply disliked that particular business. We also elicited what the second-movers would choose if they were facing the same decision problem as the actor they observed. Second-movers were not given a private signal themselves and thus had to rely on information gleaned from the observed action. As noted above, our exploration of the role of heterogeneous tastes on information transmission is built on the straightforward observation that inferences (and choices) in the second stage should depend on participants’ beliefs about the tastes of the previous actor.

Our results offer four primary punchlines. First, despite the complexities of this learning environment, we find support for some social learning, and participants’ inferences obeyed a few simple comparative statics predicted by the rational model. After observing another person choose the gift card rather than the bonus cash, participants (i) correctly inferred that the other person likely received a positive signal about the value of the gift card, and (ii) made choices that largely align with what they would have done had they directly observed the signal themselves. Furthermore, participants’ inferences responded rationally to the magnitude of the cash bonus. For example, when this bonus was larger, observers who saw an actor take the card became more confident that the actor received a positive signal about the card. Finally, when armed with information about the actor’s taste, observers correctly exhibit a form of sophisticated inference wherein they put less weight on the actor having received a positive signal the more the actor enjoys the business in question.

But our second conclusion is that, despite these successes, we find significant and systematic departures from rational social learning. Our evidence can be broadly categorized as *under-inference* from behavior that ought to provide a strong signal about the underlying state and *over-inference* from behavior that provides a weak signal. For example, although almost all participants only chose the gift card after getting a positive signal about the card’s value, observers nearly universally failed to understand *how* informative the choice of taking the card was. In contrast, when a participant turned down the gift card in favor of the bonus cash, observers inferred that the likelihood of the positive signal was too low.² We provide a theoretical framework that highlights how

²This is consistent with the patterns documented by [Augenblick et al. \(2023\)](#) in both the lab and field. See also [Ba](#)

an observer’s (potentially inaccurate) perceptions of the distribution of preferences dictates the informativeness of others’ actions, and we show theoretically how belief updating depends on these perceptions. We provide multiple pieces of evidence—stemming from survey responses, elicited beliefs, and choices—suggesting that inaccurate perceptions of others’ tastes may drive inferences in our experiment.

Our third conclusion is that the degree to which a person infers from others’ behavior is closely related to their own taste. More specifically, participants in our experiment seemingly exhibited an egocentric bias in their perceptions of others’ tastes and subsequently inferred as if others shared their tastes. A key feature of our experimental design—which allows us to shed light on this egocentric bias—is that the relatively complicated space of preferences can be reduced based on how an actor responds to their private information. As mentioned above, participants in the second “actor” stage of our experiment saw a binary signal about the value of the gift card and then made a choice whether to take the card or bonus cash. Furthermore, participants made these privately-informed choices after both signal realizations. To analyze these decisions, we define three primary strategies (or mappings from signals to actions) for each decision problem: take the gift card regardless of the signal, take the gift card only after getting a positive signal, and take the bonus cash regardless of the signal.³

To highlight the degree to which an observer’s inferences depended on their tastes, we first show that participants’ choices—and thus their strategies—varied across the businesses in the study. Aggregating over these businesses (and the three bonus cash offerings), approximately half of participants took the gift card after a positive signal. In contrast, almost no participants took the card after a negative signal, meaning that choices were highly informative of the underlying signal. Thus, we primarily observe just two strategies among participants: following the signal or choosing the bonus cash regardless of signal. We next show that a person’s own strategy (which varied across decision problems) shaped their inference. Critically, participants played as both actor and, subsequently, observer. We can therefore analyze how an observer’s inferences from previous actors’ choices depended on his own strategy from the actor stage. We find that participants who employed a signal-dependent strategy—that is, they took the card only after the positive signal in a given decision problem—inferred more from others’ actions. We show that these taste-dependent inferences are consistent with an egocentric bias in which observers erroneously believed that others’ tastes were more similar to their own and thus overestimated the prevalence of their employed strategy.

Fourth and finally, we demonstrate that the errors in learning that we document persist even

et al. (2022) and Fan et al. (2023) for other recent papers that examine over- and under-inference.

³Note that the proportions actors employing each of these three strategies form a sufficient statistic for rational Bayesian updating in our setting. This allows us to simply calculate the fully-informed Bayesian benchmark in each decision problem in our experiment. We expand on this in Section 3.

when observers are given additional information about others. In a second experiment, we armed participants with highly informative signals about others’ tastes. More specifically, for each decision problem, the observer saw the subjective rating of the business (ranging from “negative” to “strongly positive”) from the person whom they were observing. Returning to the home-buyer example from above, this is akin to giving the observer some useful demographic information about the party who walked away. Despite this added information, inferences largely follow the patterns described above. We show that a few highly informative signals might improve inferences, but on the whole, inferences are perhaps worse than in our first experiment.

Our experiment is a significant departure from prior experiments on sequential social learning in which agents have common preferences. Following [Anderson and Holt \(1997\)](#), many subsequent experiments studied modifications of the canonical observational-learning setup yet maintained common or induced preferences; see, e.g., [Hung and Plott \(2001\)](#); [Nöth and Weber \(2003\)](#); [Kübler and Weizsäcker \(2004\)](#); [Çelen and Kariv \(2004\)](#); [Goeree et al. \(2007\)](#) and [Eyster et al. \(2018\)](#).⁴ As such, our primary contribution is to explore the role of heterogeneous tastes in such environments. Moreover, by examining choices over gift cards to popular businesses, we introduce naturalistic variation in tastes; that is, we (the experimenters) did not induce the tastes of participants. In this sense, our paper further contributes to this literature by taking a field-in-the-lab approach to studying belief updating when agents have heterogeneous preferences. In an additional departure from most of this literature, we elicit observers’ beliefs along with their actions in a social-learning setting.⁵ As our results in Section 4.3 demonstrate, this is a critical step that allows us to disentangle the degree to which accurate or biased social learning occurs.

Given that we find support for inaccurate beliefs that exaggerate similarity in tastes, we also contribute to a large body of research on “social projection bias” and the “false-consensus effect” which show that people tend to perceive their own tastes and attitudes as more common than they really are. The seminal study by [Ross et al. \(1977\)](#)—along with numerous studies that followed—find a positive correlation between subjects’ own stated preferences and their estimates of others’ preferences across many domains.⁶ While this correlation may be rational when there is uncertainty about others’ preferences ([Dawes, 1989, 1990](#); [Prelec, 2004](#)), later studies argue that these

⁴While the experimental design in [Goeree et al. \(2007\)](#) is based on common induced payouts, the analysis—with individual-specific preference parameters—could be interpreted as allowing for heterogeneous tastes in some sense.

⁵Other experiments in this area that also elicit beliefs—or involve finer action spaces that more precisely reveal beliefs—include [Çelen and Kariv \(2004\)](#), [Dominitz and Hung \(2009\)](#) and [De Filippis et al. \(2022\)](#).

⁶[Marks and Miller \(1987\)](#) review the false-consensus effect in 45 studies published in the decade following [Ross et al. \(1977\)](#), and [Mullen et al. \(1985\)](#) find robust evidence of the effect in a large meta-study. More recently, [Bursztyn and Yang \(2022\)](#) find that correlations consistent with the false-consensus effect are widespread in a meta-analysis of economics field studies. Evidence on the false-consensus effect also spans a broad range of domains, including political preferences (e.g., [Brown, 1982](#); [Delavande and Manski, 2012](#)), preferences over income redistribution (e.g., [Cruces et al., 2013](#)), risk preferences (e.g., [Faro and Rottenstreich, 2006](#)), and effort costs ([Bushong and Gagnon-Bartsch, 2023](#)).

perceptions reflect a systematic bias, whereby subjects weight their own preference too heavily relative to other information about population preferences (e.g., [Krueger and Clement, 1994](#)). In incentivized economics experiments, [Engelmann and Strobel \(2012\)](#) and [Ambuehl et al. \(2021\)](#) similarly find that a false-consensus bias remains if subjects must exert minimal effort to view information about others' choices. In a related vein, [Bushong and Gagnon-Bartsch \(2023\)](#) find that participants in a real-effort context project their current sense of fatigue onto others. While these papers show that people may mispredict others' tastes, our experiment examines how these mispredictions influence people's interpretation of others' actions and how those inferences shape subsequent choices.

Our experiment is also related to the theoretical literature on misspecified social learning. This literature considers several ways in which agents misunderstand how others' actions incorporate or reflect their private information. For instance, [Eyster and Rabin \(2010\)](#) and [Bohren \(2016\)](#), among others, examine how neglecting the redundancy of information in others' actions can lead society to grow convinced of a false state.⁷ Similar to our premise, [Bohren and Hauser \(2021\)](#) and [Frick et al. \(2020\)](#) analyze learning among agents who misperceive the distribution of strategies in the population they observe.⁸ [Gagnon-Bartsch et al. \(2021\)](#) relatedly study misinference in open-outcry auctions when bidders have misspecified models of others' tastes. And most closely related to our study, [Gagnon-Bartsch \(2016\)](#) and [Gagnon-Bartsch and Rosato \(2023\)](#) specifically examine social learning among agents who exaggerate commonalities in tastes. [Gagnon-Bartsch and Rosato \(2023\)](#) derives comparative statics that bear out in our data. Namely, when seeing an actor reject the card, participants who particularly liked the business became more convinced that the actor had bad information about the card's value.

More broadly, our results suggest the importance of better understanding how people engage in social learning in the presence of heterogeneity, especially given that this problem arises in many familiar contexts across fields. Consider a canonical setting for social learning: the choice of whether and where to attend university. It is well-documented that these choices are shaped by both information about future returns and by idiosyncratic tastes. Although the focus of much effort has been on disseminating information about the future returns to education, students' idiosyncratic tastes (e.g., a preference for strong athletic programs) may provide a friction to information aggregation. This may help explain why students' beliefs vary widely regarding whether and where to attend school and, conditional on attending a particular school, the returns to a given major or degree program ([Jensen, 2010](#); [Wiswall and Zafar, 2015](#); [Conlon, 2019](#); [Delavande and Zafar, 2019](#)). Beyond education, there are numerous domains where social learning plays a documented role

⁷[Chandrasekhar et al. \(2020\)](#) and [Eyster et al. \(2018\)](#) provide experimental evidence of such errors.

⁸See also [Smith and Sørensen \(2000\)](#) and [Goeree et al. \(2006\)](#) for rational models of social learning with heterogeneous preferences.

in consumer choice despite differing tastes or motives, including investment products (Bursztyn et al., 2014), insurance plans (Sorensen, 2006), agricultural technologies (Munshi, 2004), music (Hendricks et al., 2012), films (Moretti, 2011), and restaurants (Cai et al., 2009).

The paper proceeds as follows. In Section 2, we describe our experimental design. We then provide a theoretical framework in Section 3 and highlight some predictions of both rational and biased social learning. In Sections 4 and 5, we discuss the results from Experiments 1 and 2, respectively. We conclude and discuss some further applications of our results in Section 6.

2 Experimental Design

We conducted a pair of online experiments on Prolific with 455 total participants ($n = 229$ for Experiment 1; $n = 226$ for Experiment 2). Participants were restricted to those residing in the United States, and participants were required to have completed at least 100 previous tasks on Prolific to register for our study. In both experiments, participants made simple choices between cash and gift cards which were redeemable at various commonplace American businesses. Since the second experiment is a slight modification of the first, we begin by detailing our first experiment and then quickly describe the second by highlighting how it differs.

2.1 Experiment 1

All participants in our experiments participated as both an “actor” and an “observer.” Actors made a series of privately-informed decisions. Observers then saw choices made by actors, and we elicited the observers’ updated beliefs about the actor’s information conditional on their observed choices. Given that subjects participated in various roles, our experiment took place over several stages. There were three stages in total, and all participants completed these stages in the same order. See Figure 1 for a schematic of the experimental flow. The decisions in each stage of the experiment involved gift cards to seven American businesses: AMC Theatres, Amazon, Chick-fil-A, Home Depot, Old Navy, PetSmart, and Starbucks.

Stage 1: Survey Stage. In this stage, participants answered four questions about each of the seven businesses in our study. They first answered a simple yes/no to the following: “Have you or anyone in your household bought something from this business in the last five years?” Participants next provided a subjective rating of each shop in response to the cue: “Using the scale below [which took on four values ranging from ‘negative’ to ‘strongly positive’], describe your attitudes toward this business and its product(s).” Third, participants guessed the distribution of subjective ratings amongst all participants in the experiment. Fourth, we elicited participants’ valuation for a gift card for \$100. We did not incentivize responses in this stage.

Stage 2: Actor Stage. In the Actor Stage, participants made a series choices between gift cards and bonus cash. In these decision problems, they faced uncertainty over the nominal value of the gift card, which we denote by ω . The value ω varied independently across each decision trial, and it was either $\omega = \$100$ (which we call the “high-value state”) or $\omega = \$20$ (“low-value state”). In every trial, the two states were ex-ante equally likely and this was communicated to participants. Prior to each choice, actors were also provided with a partially informative binary signal $s \in \{l, h\}$ which accurately reflected the current state with probability $\frac{3}{4}$. That is, $\Pr(s = h|\omega = 100) = \Pr(s = l|\omega = 20) = \frac{3}{4}$ and $\Pr(s = l|\omega = 100) = \Pr(s = h|\omega = 20) = \frac{1}{4}$. After receiving a signal, participants chose between the gift card and a cash bonus of a known amount.

We varied the size of the cash bonus across the decision trials. In each trial, actors faced a choice between the gift card and a cash bonus of $x \in \{30, 40, 50\}$. To limit the total number of questions in the experiment, each participant faced just two values of x . All participants answered questions that involved $x = 40$. For a given business, participants faced an equal chance of additionally seeing $x = 30$ or $x = 50$. By varying the size of the bonus across questions, we implicitly altered the share of participants who might choose the gift card over the bonus. For instance, compare when the bonus was \$50 rather than \$40. Assuming a gift card is worth (weakly) less than its cash equivalent, it is apparent that risk-neutral or risk-averse participants facing either outside option would not take the gift card if they were sufficiently confident that the card was in the low-value state. But when facing a positive signal, some participants may switch from taking the card with $x = \$40$ to taking the cash when $x = \$50$.

Overall, participants faced many decision trials in the Actor Stage. For each of the seven businesses, they decided between the gift card and bonus cash upon receiving both the good signal ($s = 100$) and the bad signal ($s = 20$). And they repeated these decisions for two different amounts of bonus cash. Thus, each participant made 28 decisions in the Actor Stage (7 business \times 2 signal realizations \times 2 cash amounts). To incentivize these decisions, participants were told that one out of every twenty participants would have a choice implemented for real. If they were chosen, one of their decisions from the Actor Stage would be selected at random to determine their payoff from this stage (e.g., if they chose the gift card in the selected choice, then they would receive a gift card to that business for ω dollars).⁹

Stage 3: Observer Stage. In this stage, a participant saw choices that others made during the Actor Stage. For each observed choice, the participant guessed the probability that the actor received the high-value signal. More specifically, for each choice, the observer saw the business name, the amount of bonus money on offer and saw whether the actor took the gift card or the bonus money; the observer *did not* see the actor’s private signal. The observer then used a slider to indicate their

⁹Payoffs from all stages were determined and communicated to subjects after they completed the entire experiment. That is, participants were not given feedback during the experiment.

posterior belief that the actor received the high-value signal in that particular decision (see Figure 1). Each observer completed a series of such guesses, each time facing a randomly-drawn anonymous actor. Observers were paid according to the accuracy of each prediction using the binarized scoring rule.¹⁰

For each trial in this third stage, the observer also made a choice for themselves within the same decision problem faced by the actor. Since the observer did not receive a signal about the card, they had to use the information revealed from the actor’s decision to inform their choice. For instance, if the observer saw that an actor chose between \$40 and a card to Starbucks, then the observer also chose between \$40 and that same gift card the actor faced.¹¹ This aspect of our design mirrors the second round of canonical experimental designs for sequential-action social learning. These choices were incentivized identically to the choices in the Actor Stage.

To summarize, the timeline within each trial of the Observer Stage was as follows: (i) the observer saw the decision problem faced by an actor and saw the actor’s choice; (ii) the observer guessed the likelihood that this actor received the high-value signal; (iii) the observer stated what they would choose in the same decision problem.

While in most trials observers were told the identity of the shop prior to making a guess, we masked the shop identity in a small number of trials and replaced the salient logo in Figure 1 with a question mark. In this “unknown” case, observers were reminded that actors knew the shop identity when making decisions. Moreover, observers were (truthfully) told that the data revealed in unknown-shop trials was uninformative about the identity of the particular shop in any given trial.

In total, participants in the Observer Stage faced $28 + 6 = 34$ trials. Twenty eight of these trials result from each combination of business (7), the cash bonus the actor faced (2), and actor choice (2). The six remaining trials represent the “unknown shop” case for each combination of cash bonus $x \in \{30, 40, 50\}$ and actor choice. The order of the trials was randomized.

As we emphasize below, the inferences that observers drew from actors’ choices should have depended on the distribution of others’ preferences. Of course, observers may have been uncertain about this distribution. Thus, observing the choices of actors could, in theory, provide useful information about population preferences. For instance, if an observer was shown choice data in which many actors chose to accept the gift card over the cash bonus, the observer might have reasonably inferred that others quite enjoyed the shops in our study. To avoid influencing participants’ beliefs

¹⁰Following the guidance of [Danz et al. \(2022\)](#), our experimental instructions first described the observer’s objective in intuitive terms, and we relegated detailed information on the payment mechanism to a separate screen in the instructions. See the Online Appendix for the experimental instructions.

¹¹The value on the card faced by the observer was known to be perfectly correlated with that on the card faced by the actor. So, if the actor had received a high-value signal about their card, then the card in the observer’s decision problem also had a 75% chance of having the high value.

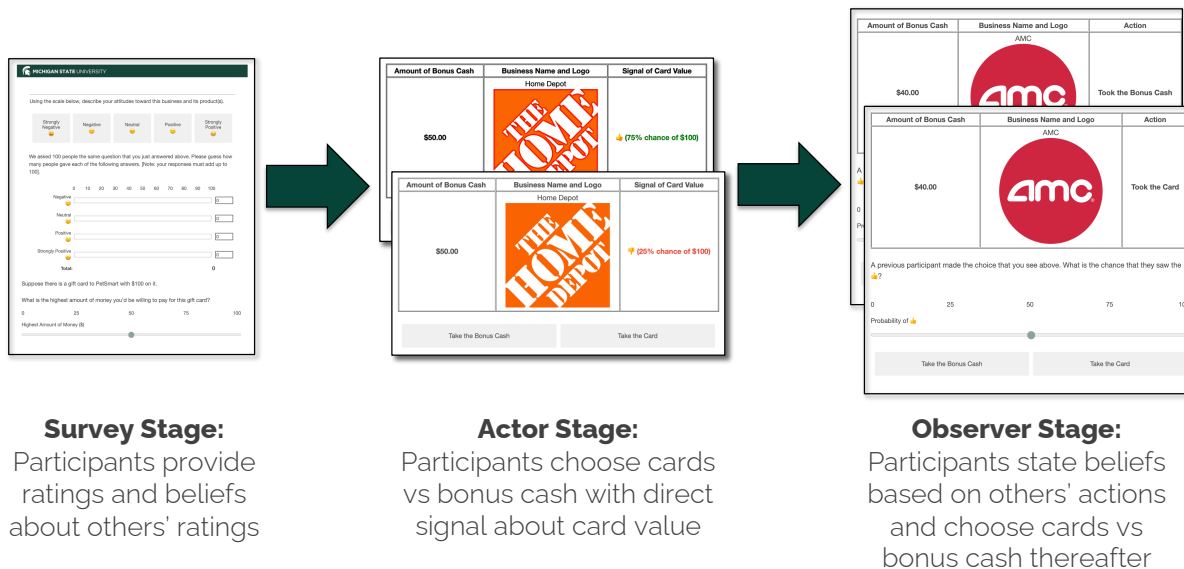


Figure 1: Overview of Experiments 1 and 2. First, we elicited basic preference information from participants. Next, we asked participants if they would choose a gift card or cash after receiving a noisy signal about the value of the card. Finally, we measured participants' inferences from others' choices.

in this manner, our experimental design took measures to limit this form of learning: the choice data we showed observers during the belief-elicitation trials were chosen in a predetermined way so that it revealed no information about actors' preferences. We informed observers about this during the instructions, emphasizing that the data they would see was *not* a representative sample of actors' choice behavior.¹²

Finally, after the three stages described above, our experiment concluded with a short series of demographic and attention-check questions. Figure 1 provides a flow chart of the three main stages of the experiment.

Discussion. The purpose of the Actor Stage was threefold. First, it yielded the privately-informed choices that uninformed observers learned from in the Observer Stage. Second, it provides us with a baseline measure of how participants responded to information. Choices in the Actor Stage reveal how people decided when they received direct signals about the gift-card's value, whereas the Observer Stage reveals how people decided when they received indirect infor-

¹²More specifically, the data that an observer encountered during the belief-elicitation trials was chosen such that for each shop k , the observer saw both possible actions for a given cash bonus x . That is, they saw exactly one actor accept the card and another reject it. This elicitation procedure is similar to the strategy method: we asked each observer for their updated beliefs conditional on each possible combination of the observables (i.e., shop identity, bonus cash, and actor choice). We described this procedure to observers prior to eliciting their beliefs, and thus they knew that they would see precisely 50% of actors accept the card and 50% reject it for each shop (including the "unknown" shop). Consequently, observers should not have used actors' choices to update their beliefs about population preferences.

mation filtered through others’ actions. Third, having all participants complete the Actor Stage means that participants in the Observer Stage had already been in the role of an actor. Thus, they had first-hand familiarity with actors’ choices and presumably understood how their decisions may have depended on their signals. An understanding of this mapping is crucial for drawing inference from others’ actions.

The Observer Stage allows us to assess how observers formed beliefs about the state based on others’ actions. Moreover, the choice data from observers allows us to analyze the degree to which observers acted on the information they gleaned from others’ decisions.

2.2 Experiment 2

Experiment 2 was identical to Experiment 1 except we varied the amount of information we provided participants in the Observer Stage prior to making their predictions about the actors’ signals. For each trial in the Observer Stage of Experiment 2, we told the observer the subjective rating of the actor who made the decision they were observing. That is, in addition to seeing an actor’s choice, observers also saw her subjective rating of the shop at hand. Furthermore, we restricted the bonus cash to $x = \$40$ in the Observer Stage of Experiment 2 to increase statistical power. Finally, to limit the number of questions, observers saw a random subset of all possible combinations of ratings, choices, and businesses.¹³ In total, participants made 48 inferences. Stages 1 and 2 of Experiment 2—along with all other aspects of Stage 3—were the same as in Experiment 1.

3 Theoretical Framework

We now provide a simple model of observational learning within our setting and formalize various predictions. The baseline model considers rational learning, and we discuss how it simply extends to capture forms of biased learning. We use this model to guide our empirical analyses in the following sections.

3.1 Observational Learning with Heterogeneous Tastes

Following our experimental design, we consider an observational-learning context with binary actions and a binary payoff-relevant state. We primarily focus on just two players: a privately-informed “actor” (Player 1) takes an action, and an “observer” (Player 2) sees this action. We analyze the observer’s updated belief about the state conditional on the actor’s action.

¹³As in Experiment 1, our experimental instructions described how the observed data was selected. The combinations were drawn uniformly so that the observed data provided no information about the aggregate choice behavior and preferences in the population (similar to the procedure discussed in Footnote 12).

Actions, States, and Decision Problems. In each decision problem, an actor must decide between two options: (i) a gift card specific to a particular shop or (ii) a monetary bonus of $\$x$. We denote the actor’s decision by $a \in \{A, R\}$, where A represents accepting the gift card and R represents rejecting it in favor of the cash. The nominal value of the gift card can be one of two amounts, $\omega \in \{h, l\}$, with $h > l$. Both states have an equal probability of occurring, and this is common knowledge.

We index the shops in our setting by k , and we label a *decision problem* by (x, k) . Hence, (x, k) describes the monetary bonus and the particular gift card that the actor chooses between.

Preferences. Players may have differing preferences over shops. Thus, even if the amount on a gift card were known, some may favor the gift card over the monetary bonus, while others may have the opposite preference. We denote the actor’s valuation for the gift card to shop k in state ω as $v_k(\omega)$, with her value of the bonus being simply x . We can interpret $v_k(\omega)$ as the cash amount at which the player would be indifferent between a guaranteed payment of that amount and the gift card in state ω . We assume that preferences are monotonic in ω — $v_k(h) > v_k(l)$ for all k —and that gift cards are weakly less desirable than their face value in cash— $v_k(\omega) \leq \omega$ for all k and ω .

Actor’s Private Information and Decision Rule. The actor receives a signal $s \in \{h, l\}$ correlated with the state. This signal is symmetric and has a precision of $\phi > 1/2$; that is, $\Pr(s = z | \omega = z) = \phi$ for each $z \in \{h, l\}$. Let $\mu(s)$ represent the actor’s belief in state $\omega = h$ given s . Under this structure, if the actor’s beliefs adhere to Bayes’ Rule, then $\mu(h) = \phi$ and $\mu(l) = 1 - \phi$. In our experimental design, we set $\phi = 3/4$.

We assume the actor’s decision rule maximizes her expected utility: she opts for the card iff $\mu(s)v_k(h) + [1 - \mu(s)]v_k(l) > x$. Given our binary-signal environment, this condition implies that the actor’s behavior in a given decision problem (x, k) can be summarized by one of three strategies:

- (i) *Signal-Independent Accept (SIA)*: choose the card for both signal realizations;
- (ii) *Signal-Independent Reject (SIR)*: choose the cash for both signal realizations;
- (iii) *Signal-Dependent (SD)*: choose the card if and only if $s = h$.¹⁴

Fixing the decision problem, we will categorize actors by the strategy they employ in that decision problem. We call this an actor’s “strategy type” (or simply “type” for brevity when there is no potential confusion). This type distills many aspects of a participant’s decision making—taste for the shop, risk attitudes, etc.—down to a simple mapping between signals and her choices.¹⁵

¹⁴An additional strategy is possible if we allow for irrational behavior: a *Mistaken* strategy is one in which the actor chooses the card if and only if $s = l$. Our predictions will focus on the case where such a strategy is not employed and players are aware of this.

¹⁵Although we use the term “strategy type” in reference to the mapping from signals to actions, it is not intended to

Observer’s Beliefs. For each decision problem (x, k) , the observer sees the action taken by a distinct anonymous actor. The observer then updates their belief about the likelihood that the actor received the signal $s = h$. We denote this updated belief by $\pi_{x,k}(a)$. Crucially, this belief relies on the observer’s beliefs about the actor’s preferences. A key feature of our design is that Bayesian updating in this environment depends solely on the likelihood the observer attaches to each of the strategy types mentioned above. Let $p_{x,k}$, $q_{x,k}$, and $1 - p_{x,k} - q_{x,k}$ represent the observer’s assessed probability that the actor employs a signal-independent-accept, signal-dependent, and a signal-independent-reject strategy, respectively, given shop k and cash bonus x . With these (potentially inaccurate) subjective probabilities, the observer’s beliefs about the likelihood that the actor received $s = h$ after applying Bayes’ Rule are:

$$\pi_{x,k}(A) = \frac{p_{x,k} + q_{x,k}}{2p_{x,k} + q_{x,k}} \quad \text{and} \quad \pi_{x,k}(R) = \frac{1 - p_{x,k} - q_{x,k}}{2(1 - p_{x,k}) - q_{x,k}}. \quad (1)$$

It’s important to note that $p_{k,x}$ and $q_{k,x}$ are sufficient to capture the observer’s beliefs: the observer need not separately consider risk attitudes or other aspects of taste (e.g., idiosyncratic preferences for the business or gift cards in general) beyond how they influence these proportions. Note further that by eliciting an observer’s beliefs when they see someone else accept the gift card in one trial and again in a separate trial when they see someone refuse it—i.e., by measuring both $\pi_{x,k}(A)$ and $\pi_{x,k}(R)$ while holding all other information constant—we can use the two equations in (1) to uncover the observer’s subjective assessment of $p_{x,k}$ and $q_{x,k}$.

Finally, our variation in the cash bonus across decision problems allows us to test a simple comparative static predicted by rational learning: an observer’s belief $\pi_{x,k}(a)$ should increase in x . Intuitively, when x is larger, an actor becomes less willing to select the gift card. Thus, seeing the actor choose the card more strongly indicates that the actor received the good signal, and $\pi_{x,k}(A)$ is therefore increasing in x . By contrast, seeing an actor reject the card when x is larger is less indicative of a bad signal, since the outside option has become more appealing. Thus, $\pi_{x,k}(R)$ is also increasing in x . We examine this prediction in Section 4.3.

3.2 How Beliefs About Others Shape Inferences

In this section we consider the possibility that people have inaccurate perceptions of others’ preferences and examine how these beliefs influence their stated inferences. In our setting with a binary state, these inaccurate perceptions of others’ preferences manifest as either apparent over- or under-reaction relative to the Bayesian benchmark. To measure these biases, we employ the framework from [Grether \(1980\)](#) as presented in [Benjamin \(2019\)](#). This involves examining whether the ob-

invoke any other aspects of strategic thinking such as a person’s level of strategic reasoning (as in level- k or cognitive-hierarchy models).

server's updated beliefs, $\pi_{x,k}(A)$ and $\pi_{x,k}(R)$, deviate systematically from the beliefs that would be expected under Bayesian updating with accurate information about the distribution of strategy types.

More formally, let $\bar{p}_{x,k}$ and $\bar{q}_{x,k}$ denote the true population frequencies of signal-independent-accept and signal-dependent types, respectively. Writing the updated beliefs in terms of log-odds, we can explore over- and under-reaction relative to the Bayesian benchmark with a set of two simple equations (which we estimate in the next section):

$$\ln \left(\frac{\pi_{x,k}(A)}{1 - \pi_{x,k}(A)} \right) = c_A \cdot \ln \left(\frac{\bar{p}_{x,k} + \bar{q}_{x,k}}{\bar{p}_{x,k}} \right), \quad (2)$$

$$\ln \left(\frac{\pi_{x,k}(R)}{1 - \pi_{x,k}(R)} \right) = c_R \cdot \ln \left(\frac{1 - \bar{p}_{x,k} - \bar{q}_{x,k}}{1 - \bar{p}_{x,k}} \right). \quad (3)$$

The left-hand side of Equations 2 and 3 is an observer's stated belief (in log-odds form) after observing an actor's choice. The right-hand side is the Bayesian belief scaled by c_A or c_R , which are constants measuring the degree of over- or under-reaction. After observing action $a \in \{A, R\}$, we say that an observer exhibits over-reaction ($c_a > 1$) when they place too much apparent weight on the new information, resulting in updated beliefs that are more extreme than warranted by the Bayesian benchmark. Conversely, under-reaction ($c_a < 1$) occurs when the observer does not update their beliefs enough based on the new information, leading to updated beliefs that are less extreme than the Bayesian benchmark would predict.

One reason we may see apparent deviations from the Bayesian benchmark is that people have inaccurate beliefs about the relative prevalence of types in the population. We can see this more clearly by writing $\pi_{x,k}(A)$ and $\pi_{x,k}(R)$ in Equations 2 and 3 in terms of an observer's perceived frequencies of types, as in Equation 1:

$$\ln \left(\frac{p_{x,k} + q_{x,k}}{p_{x,k}} \right) = c_A \cdot \ln \left(\frac{\bar{p}_{x,k} + \bar{q}_{x,k}}{\bar{p}_{x,k}} \right), \quad (4)$$

$$\ln \left(\frac{1 - p_{x,k} - q_{x,k}}{1 - p_{x,k}} \right) = c_R \cdot \ln \left(\frac{1 - \bar{p}_{x,k} - \bar{q}_{x,k}}{1 - \bar{p}_{x,k}} \right). \quad (5)$$

The following proposition highlights how inaccurate beliefs about the distribution of strategy types can lead to either over- or under-reaction in belief updating (relative to a fully-informed Bayesian):

Proposition 1. *Consider decision problem (x, k) . Let $\bar{p}_{x,k}$ and $\bar{q}_{x,k}$ denote the true population proportions of signal-independent-accept and signal-dependent strategy types, respectively. Consider a Bayesian observer who believes these proportions are $p_{x,k}$ and $q_{x,k}$, respectively. This observer infers as follows:*

1. When seeing action A , the observer's beliefs will appear to over-react if $\frac{q_{x,k}}{p_{x,k}} > \frac{\bar{q}_{x,k}}{\bar{p}_{x,k}}$ and under-react if the inequality is reversed.
2. When seeing action R , the observer's beliefs will appear to over-react if $\frac{q_{x,k}}{1-p_{x,k}} > \frac{\bar{q}_{x,k}}{1-\bar{p}_{x,k}}$ and under-react if the inequality is reversed.

The proposition—which follows directly from Equations 4 and 5—illustrates how inaccurate beliefs about the proportions of strategy types in the population can lead to over- and under-reaction when inferring from the actions of others. Depending on how the observer's perceived proportions compare to the true values, various combinations of over- and under-reaction can emerge among observers following actions A and R .

One plausible belief that participants may hold about the proportions of strategy types is an over-estimation of $p_{x,k}$. The parameters of our experiment were chosen so that the expected monetary value of the card after $s = l$ is $(1 - \phi)h + \phi l = 40$. This implies that, in theory, there should be few signal-independent-accept types, especially when $x \geq 40$ —that is, $\bar{p}_{x,k} \approx 0$. Given this feature of our design, the only way participants may misperceive $p_{x,k}$ is to overestimate it. Beyond this mechanical point, a number of concepts from both psychology and economics suggest such a misperception. For instance, people tend to overestimate or overweight small probabilities (see, e.g., [Kahneman and Tversky, 1979](#); [Viscusi, 1985](#); [Fischhoff and de Bruin, 1999](#)). Additionally, [Frederick \(2012\)](#) finds a widespread tendency for people to overestimate others' willingness to pay for goods. In our context this may manifest as overestimating the fraction of people willing to accept the card over the cash. Finally, participants may overestimate $p_{x,k}$ if they exaggerate the degree to which others are irrational (as in [Kneeland, 2015](#)). Proposition 1 reveals that even if the observer correctly perceived $\bar{q}_{x,k}$, then overestimating $p_{x,k}$ would cause the observer's beliefs to under-react when seeing $a = A$ and over-react when seeing $a = R$.

Another plausible belief is that participants wrongly view others' preferences as more similar to their own, in line with the large literature on social projection bias and the false-consensus effect discussed in the introduction. To model this particular pattern of biased beliefs about the population, we allow a participant's perception of the distribution of strategy types to depend on her own type. Let $\tau \in \{0, 1, 2\}$ denote a participant's employed strategy in decision problem (x, k) , where these values correspond to signal-independent-reject, signal-dependent and signal-independent-accept, respectively. We then let $p_{x,k}(\tau)$ and $q_{x,k}(\tau)$ denote a participant's perception of the type frequencies as a function of her own type.

We can then use Proposition 1 to study how social learning is affected by “projection bias”, where we use this term broadly to refer to inaccurate beliefs about the the distribution of types that overestimate the prevalence of types similar to one's own.¹⁶ Toward this end, it is useful to

¹⁶Although there are several plausible reasons why people may overestimate the prevalence of their own prefer-

consider what such a phenomenon implies about people’s perceptions of the distribution (i.e., what it implies about $p_{x,k}(\tau)$ and $q_{x,k}(\tau)$). First, it naturally suggests that participants overestimate the prevalence of their own type: $p_{x,k}(2) > \bar{p}_{x,k}$, $q_{x,k}(1) > \bar{q}_{x,k}$, and $p_{x,k}(0) + q_{x,k}(0) < \bar{p}_{x,k} + \bar{q}_{x,k}$. Second, it suggests that the perceived prevalence of a given type depends on how similar that type is to one’s own: people with types closer to τ perceive a higher prevalence of type τ than those with types farther from τ .¹⁷ This means, for instance, that signal-dependent types perceive a greater proportion of signal-dependent types than do signal-independent types—e.g., $q_{x,k}(1) > q_{x,k}(0)$ and $q_{x,k}(1) > q_{x,k}(2)$. And relative to those employing a signal-independent-reject strategy, signal-dependent types think there are more people employing a signal-independent-accept strategy—that is, $p_{x,k}(1) > p_{x,k}(0)$. Intuitively, since SIR types relatively dislike the business under consideration, they underestimate how much others enjoy it by a greater degree than signal-dependent types (who tend to like the business themselves relatively more).

The pattern of misperceptions described above allows us to (in some cases) predict which types of observers will under- or over-react more than others in a given decision problem. Consider, for instance, the beliefs of observers who are signal-dependent types. Applying the insights from Proposition 1, we can see that if $\frac{q_{x,k}(1)}{p_{x,k}(1)} > \frac{q_{x,k}(0)}{p_{x,k}(0)}$, signal-dependent types over-react by a greater degree than SIR types after observing action A . However, the aforementioned notion of projection bias does not make a crisp prediction about the ordering of these ratios of frequencies. By contrast, projection bias does yield an unambiguous prediction after observing action R . In that case, signal-dependent types overreact by a greater degree than SIR types after observing action R if $\frac{q_{x,k}(1)}{1-p_{x,k}(1)} > \frac{q_{x,k}(0)}{1-p_{x,k}(0)}$. Notice that this inequality always holds under projection bias because, relative to SIR types, signal-dependent types think both signal-dependent and SIA types are more prevalent.

Corollary 1. *Consider decision problem (x, k) . Suppose that $q_{x,k}(1) > q_{x,k}(0)$ and $p_{x,k}(1) > p_{x,k}(0)$. Then, when seeing action R , signal-dependent types unambiguously over-react more than signal-independent-reject types.*

In light of Corollary 1, our empirical analyses will frequently compare beliefs across strategic types as a way of exploring whether participants’ behavior is consistent with projection bias.

Observer’s Beliefs with Rating Information. The updated beliefs described above are for the case where the observer only sees the actor’s action, as in Experiment 1. In Experiment 2, we measure the observer’s updated beliefs about the actor’s information when they additionally see the actor’s rating for the shop in question. In Stage 1 of our experiments, we elicited the actor’s self-reported ratings for each business, denoted by $r \in \{1, 2, 3, 4\}$. A higher rating indicates a

ences (as we discuss further in Section 4.3), we use the term “projection bias” here to broadly capture this phenomenon in order to streamline our exposition.

¹⁷Our notion of “closeness” here is with respect to participants’ valuations for the gift card: SIR types have the lowest valuations, SIA types have the highest, and SD types are in between.

higher self-reported taste for shop, and thus it serves as a noisy signal of the actor’s underlying preferences. Accordingly, we let $\pi_{x,k}(a|r)$ denote the observer’s updated probability of $s = h$ conditional on the actor’s action and rating for business k .

Rational learning yields a stark prediction in this case: for either action $a = A, R$, the observer’s belief $\pi_{x,k}(a|r)$ is decreasing in r . For an intuition, suppose the actor chose action R . If the actor has a low rating for the business, this action is relatively uninformative since she likely would have rejected the card regardless of her signal. On the other hand, if the actor has a higher rating, action R more strongly indicates that the actor received the bad signal. Thus, the likelihood of $s = h$ is decreasing in r . If instead the actor chose A , this action is relatively uninformative when the actor has a high rating but is very informative when she has a low one—in the latter case she is likely to choose A only if she got the good signal. Again, the likelihood of $s = h$ is decreasing in r . Our analyses of Experiment 2 tests these comparative statics.

4 Results, Experiment 1

We organize our results by examining behavior over the three stages of the experiment. First, we provide some simple descriptive information about the tastes of participants gathered in the Survey Stage.¹⁸ We then explore the privately-informed choices of participants in the Actor Stage and show how they (i) made choices that accord with both their stated ratings towards the businesses and a weak aversion to risk; and (ii) utilized the information from their noisy signals. These choices form an important foundation for studying social learning, since they are correlated with participants’ stated preferences and they reflect their private information.

Finally, we turn to our main analyses, which examine participants’ behavior in the Observer Stage. We first show that observers’ inferences follow some of the comparative statics predicted by rational learning. However, we show that these inferences also differ in significant ways. In particular, observers systematically over- and under-infer from the actions they see. Moreover, we show that observers’ beliefs depended on their own tastes as predicted by a model of projection bias, and this dependency is not merely an artifact of cross-person heterogeneity or measurement error. We conclude by demonstrating that observers’ inferences have a significant effect on the choices they make after (mis)learning.

4.1 Survey Stage

We first provide evidence that participants indeed had heterogeneous tastes. Table 1 reports the average response to the survey questions from Stage 1 for each of the seven businesses. The

¹⁸We relegate descriptive statistics about our participants (i.e., demographics) to the Online Appendix.

variability in these responses reveal substantial heterogeneity in participants’ personal experience with the businesses, attitudes toward the businesses, and their willingness to pay (WTP) for a gift card. Hence, this setting appears well suited to study social learning under heterogeneous tastes. This heterogeneity is further explored in Panel (a) of Figure 2, which presents a histogram of participants’ ratings of the businesses on a 1-4 scale from “negative” to “strongly positive”.¹⁹

SUMMARY OF PARTICIPANTS’ PREFERENCES			
	ℙ(Personal Experience)	Rating	WTP
AMC	0.611 (0.488)	2.585 (0.723)	53.77 (28.13)
Amazon	0.983 (0.131)	2.830 (0.986)	72.77 (22.74)
Chick-Fil-A	0.760 (0.428)	2.603 (1.089)	54.66 (32.34)
Home Depot	0.860 (0.347)	2.729 (0.659)	63.84 (25.82)
Old Navy	0.559 (0.497)	2.476 (0.710)	54.13 (26.39)
PetSmart	0.607 (0.489)	2.603 (0.703)	53.55 (28.09)
Starbucks	0.865 (0.342)	2.533 (0.933)	57.14 (27.47)

Table 1: Mean Responses ($n = 229$) **from the Survey Stage.** “Personal Experience” corresponds to indicating that their household purchased something from the business in the last five years. “Rating” was reported on a four-point increasing scale. Willingness to Pay (“WTP”) for a \$100 card was reported using a slider from \$0 to \$100. Standard deviations in parentheses.

We also collected participants’ guesses about how others would rate the shops. Panel (b) of

¹⁹Figure A1 in the Appendix shows these ratings distributions for each individual business.

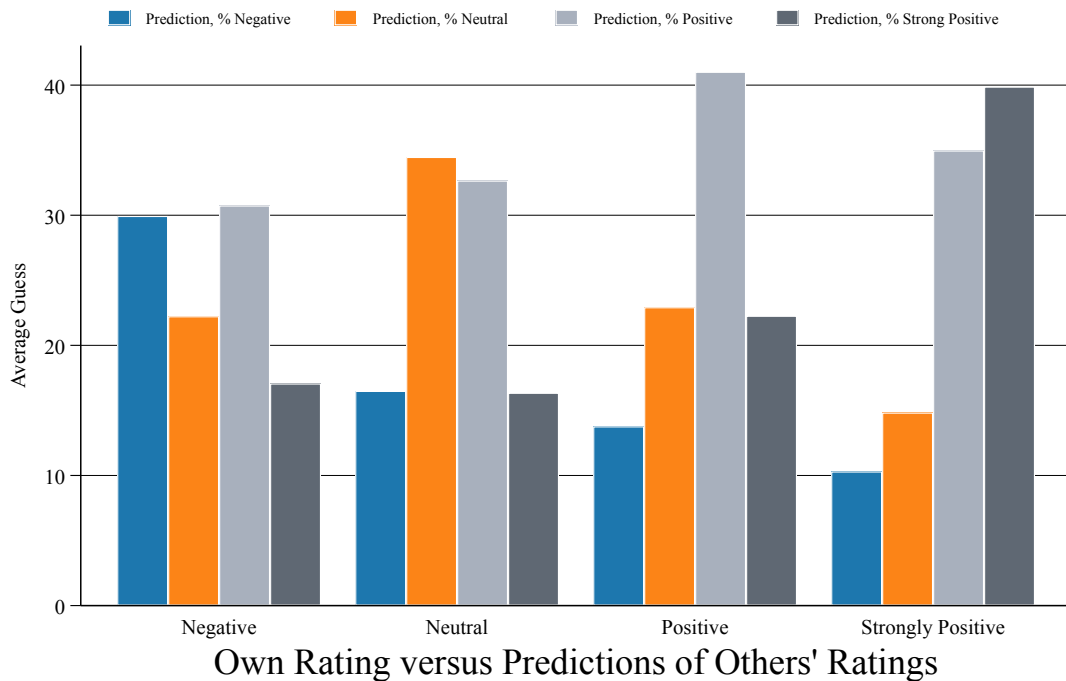
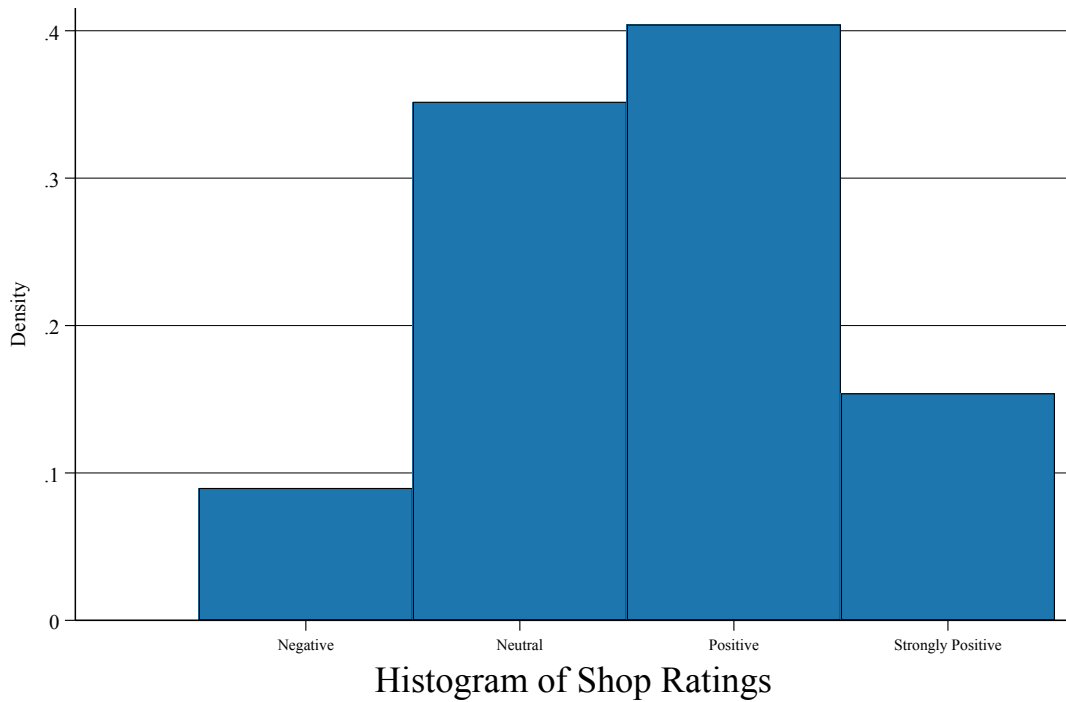


Figure 2: Ratings in the Survey Stage. Panel (a) shows the distribution of ratings aggregated across businesses. Panel (b) shows the perceived distribution of ratings conditional on participants' own rating. Perceptions of others' tastes were heavily skewed towards a person's own stated preference.

Figure 2 highlights that people’s perceptions of others’ ratings were biased towards their own. For example, while “negative” ratings represented only 9% of the overall data, participants who themselves rated a shop negatively believed that approximately 30% of others would agree with them. These strong distortions in perceptions of population ratings are observed across the spectrum of ratings. This is consistent with prior studies on the false-consensus effect noted in the introduction.

4.2 Actor Stage

We now turn to the privately-informed choice data from Stage 2. In this stage, participants faced a series of choices over gift cards and bonus cash. The left panel of Table 2 presents the aggregate choice data across all shops and participants, but disaggregated by the signal $s \in \{h, l\}$ the actor received and the cash bonus $x \in \{30, 40, 50\}$ they faced. After receiving the “good” signal (i.e., $s = h$), participants chose the gift card approximately 50% of the time across the various decision problems. In contrast, after receiving the “bad” signal ($s = l$), participants were very unlikely to choose the gift card. This data shows that participants used the signals to guide their choices and that they were typically not risk loving.

CHOICES AFTER SIGNAL $s = h$					DISTRIBUTION OF STRATEGY TYPES				
Cash Option					Cash Option				
	\$30	\$40	\$50	Total		\$30	\$40	\$50	Average
$a = R$	331	789	489	1609	Signal Ind. Reject	0.407	0.489	0.608	0.498
$a = A$	472	814	311	1597	Signal Dependent	0.509	0.492	0.384	0.469
CHOICES AFTER SIGNAL $s = l$					Signal Ind. Accept	0.078	0.016	0.005	0.029
$a = R$	736	1573	793	3102	Mistaken	0.005	0.003	0.004	0.004
$a = A$	67	30	7	104					

Table 2: Choices and Types from the Actor Stage of Experiment 1. *Participants made 14 choices for each signal realization across seven businesses. The left tables show the raw counts of these choices; the right table shows the associated distribution of strategy types.*

Furthermore, the percent of people who chose the gift card after the good signal was decreasing in the outside cash bonus, as demonstrated by moving across the columns in Table 2. For instance, rates of choosing the gift card after $s = h$ decrease from 59% to 51% to 39% when x increased

from \$30 to \$40 to \$50, respectively. This provides an important validation that our participants understood the choice environment and the relevant incentives.

Our choice data allows us to classify participants according to the three strategy types discussed in Section 3. We do this as follows. Recall that for each shop k and each bonus amount x , we observe each participant’s choice after both signal realizations. If the person chose the card for both signal realizations for a given (x, k) , we call them a “signal-independent-accept (SIA)” type in that decision problem; if they didn’t chose the card for either realization, we call them an “signal-independent-reject (SIR)” type; and if they chose the card only after receiving the good signal, we call them a “signal-dependent” type. A very small number of participants seemed to err by choosing the card only when they received the bad signal; we call them “mistaken” types. The right panel of Table 2 presents the frequencies of these types aggregated across businesses.

Following our theoretical setup from Section 3, we can further calculate $\bar{p}_{x,k}$ and $\bar{q}_{x,k}$, which denote the proportion of SIA and signal-dependent types in each decision problem (x, k) . We report these values in Table A1. As predicted, we find that (i) the proportion of SIA types is very low for all businesses, and (ii) the proportion of signal-dependent types is decreasing in the outside option, x . Averaging over all of the businesses in the study (as presented in Table 2), we find $\bar{p}_{30} = 0.078$, $\bar{q}_{30} = 0.509$, $\bar{p}_{40} = 0.016$, $\bar{q}_{40} = 0.492$, and $\bar{p}_{50} = 0.005$, $\bar{q}_{50} = 0.384$. Put succinctly, for all three outside options, there are very few participants who chose the gift cards when they received the bad signal.

We find that 22 (out of 229) participants never chose a gift card in any decision problem, while zero participants chose the gift card in every problem for both signals. Therefore, the vast majority of people had experience with both choosing the gift card and turning it down.²⁰ Importantly, our experimental design allows each participant to differ in their type across decision problems (i.e., across shops and bonus amounts). We leverage this within-participant variation in later analyses.

4.3 Observer Stage

Inferences Among Observers

We now analyze how observers formed inferences based on the choices they saw. Recall that observers were asked to infer the likelihood that an actor received the good signal after observing that actor’s choice. Table 3 shows the average of observers’ updated beliefs after seeing action $a \in \{A, R\}$. The left two columns show beliefs after the actor rejected the card ($a = R$), while the right two show beliefs after the actor accepted it ($a = A$). Participants’ belief updating obeys a

²⁰Nine participants were signal-dependent types in every decision problem. Two participants chose seemingly at random (and very quickly); we include these two participants from our main analyses, though our findings do not change depending on whether we exclude these participants or not.

basic comparative static: beliefs about the actor’s signal move significantly below the 50-50 prior when the actor rejects the card and move significantly above when they accept it. We now turn to some more nuanced features of our data.

AGGREGATE INFERENCES, ALL SHOPS, EXP 1

	After $a = R$		After $a = A$	
	Stated Belief	Benchmark	Stated Belief	Benchmark
Cash Option = 30	29.49 (1.36)	32.02 -	77.40 (0.96)	88.51 -
CashOption = 40	31.72 (1.28)	34.68 -	78.08 (0.87)	96.59 -
Cash Option = 50	33.76 (1.51)	38.59 -	81.27 (1.03)	98.42 -
Observations (column total)	3,206		3,206	

Standard errors (in parentheses) are clustered at the individual level.

Table 3: Average Inferences in Experiment 1. *Participants reported their belief about the likelihood that the actor received signal $s = h$ —reported as a number out of 100. The benchmark reflects what a fully informed Bayesian would report.*

The Effect of the Cash Option on Beliefs. We first examine how the cash on offer (x) shaped inferences. We find that the way observers account for x is consistent with the predictions of rational learning described in Section 3. For ease of comparison to the rational model, we define a “fully informed” benchmark as what a Bayesian would infer if they knew the population averages from our sample. We present this in columns 2 and 4 of Table 3. Note that a rational observer in our setting would understand that, as x increases, the card becomes relatively less attractive. Thus, when x is larger, there is more information revealed by action $a = A$, since it is more likely that the actor requires the good signal in order to choose the card over the cash. Beliefs following $a = A$ should therefore move closer to 1 as x increases. Participants’ beliefs follow this comparative static, as evident from comparing the rows of column 3 in Table 3 ($p < 0.001$, joint F -test for equality of rows in column 3). Similarly, less information is revealed by $a = R$ when x is larger, and beliefs should thus move closer to 0.5 following $a = R$ ($p = 0.002$, joint F -test for equality of rows in column 1). This is a key validation that observers in our experiment attended (at least

partially) to actors’ preferences and how their choices likely varied as their incentives changed.

The Responsiveness of Beliefs to Actions. While beliefs move in the rational direction, do they move by the rational amount? Recall that columns 2 and 4 of Table 3 show the Bayesian benchmark beliefs. The beliefs of participants in Table 3 departed from this benchmark in systematic ways. Most strikingly, observers (i) wildly under-inferred from observing action $a = A$ and (ii) over-inferred from action $a = R$.

However, the analysis in Table 3 aggregates over the (potentially very different) businesses in our study. To more formally examine the degree of over- or under-inference, we estimate Equations 2 and 3. This is a simple matter: given that we observe both people’s choices and inferences, we have the relevant data on hand. Specifically, we utilize actors’ choice behavior to calculate shop-by-shop proportions of SIR and signal-dependent types for each cash level— $\bar{p}_{x,k}$ and $\bar{q}_{x,k}$, presented in Table A1—to calculate the right-hand side of Equations 2 and 3. We then estimate the coefficients of over and under-reaction relative to the Bayesian benchmark using OLS.²¹

We present the results of this approach in Table 4. Column 2 shows that participants exhibited significant under-reaction across all businesses in the study after observing an actor accept the gift card ($p < 0.01$ for all businesses in rejecting test of $H_0 : c_A = 1$; Bonferroni adjusted for multiple hypothesis testing). We find mild heterogeneity in reactions across shops, with the notable outlier of Amazon, where participants under-reacted the least (that is, were closest to the Bayesian benchmark). This shop-specific result reflects an ex-ante design feature: a gift card to Amazon is fundamentally different than a gift card to the other businesses, since it is essentially cash. Given this, we expected that (i) participants’ valuations for an Amazon gift card would exhibit the least variation given its similarity to cash (see Table 1), and (ii) participants would be more likely to be signal-dependent types when facing Amazon gift cards relative to other businesses in the study (see Table A1). Put differently, our finding that participants reacted more appropriately when facing Amazon cards suggests that inferences may be more accurate when the confounding factor of heterogeneity is mitigated.

Furthermore, column 1 of Table 4 highlights that—after observing an actor reject the gift card—observers’ beliefs tended to *over-react*. Table 4 shows that observers’ beliefs over-reacted after seeing the actor reject the gift card for all shops except Amazon and Home Depot ($p < 0.01$ for other five businesses in rejecting test of $H_0 : c_A = 1$; Bonferroni adjusted). We also find strong evidence for under-reaction at Amazon ($p < 0.001$).²²

²¹We can also directly estimate these reaction coefficients for each cash level, business, and participant by simply dividing the left-hand side of these equations by the right.

²²The fact that observers’ beliefs when facing Home Depot were closer to the benchmark is consistent with our note above about the importance of heterogeneity on inferences: similar to Amazon, tastes at Home Depot were more homogeneous than those at the rest of the businesses in our study (see Figure A1).

REACTION COEFFICIENTS ACROSS BUSINESSES		
	After $a = R$	After $a = A$
AMC	1.67 (0.15)	0.38 (0.02)
Amazon	0.71 (0.05)	0.47 (0.02)
Chick-Fil-A	1.51 (0.12)	0.44 (0.02)
Home Depot	0.99 (0.09)	0.38 (0.02)
Old Navy	1.52 (0.15)	0.44 (0.02)
PetSmart	1.56 (0.13)	0.43 (0.02)
Starbucks	1.27 (0.12)	0.36 (0.02)
Observations (column total)	2984	2353

Standard errors (in parentheses) are clustered at the individual level.

Table 4: Reactions Relative to the Fully Informed Benchmark. *The table illustrates that participants’ inferences exhibited significant over-reaction after they saw a person reject the gift card (except when updating about an Amazon or Home Depot gift card) and significant under-reaction after seeing a person accept the gift card (for all businesses in the study). Sample-size variability stems from our inability to calculate benchmark log odds in decision problems where $\bar{p}_{x,k} = 0$.*

The Effect of a Person’s Own Strategy on Their Inferences. We now demonstrate that observers’ inferences also depended on their *own* choices from the Actor Stage. As described above, for each shop and bonus-cash amount, we can determine an observer’s own strategy type—signal-independent accept (SIA), signal-independent reject (SIR), or signal dependent—based on the strategy she deployed in Stage 2. We can then examine how an observer’s *own* type in a given de-

cision problem influenced the inference they drew from others’ choices in that same problem. For this exercise, we focus on the primary two types in our data—SIR and signal-dependent types—since (by design) they make up 96.5% of the data. Splitting inferences by these types reveals a striking truth: regardless of the observed action, signal dependent types inferred *more* from the actions of others. After seeing an actor choose $a = R$, signal-dependent observers stated that the chances the actor saw the good signal were 4.1% lower than SIR observers ($p = .037$; joint test across all bonus cash amounts). After seeing $a = A$, signal-dependent observers stated that these chances were 3.3% higher than SIR observers ($p = .001$; joint test across all bonus cash amounts).

The above result suggests that one’s own type directly influenced how they interpreted others’ actions. It is important to note, however, that this aggregate result could conceivably be driven by interpersonal differences in traits that are correlated with one’s type. For example, people who tended to be signal-dependent types may have been more engaged with the experiment and thus inferred differently from those who tended to be SIR types. To more carefully establish that a given observer’s inferences depended on her strategy type, we leverage the within-person variation in type. Specifically, since each participant’s type could change depending on the shop (k) and the cash amount (x), we can fix an observer i and ask how her inferences changed as a function of her type in a given decision problem. We employ a fixed-effects regression model to estimate this effect:

$$\pi_{i,x,k}(a) = \beta_0 + \beta_1\mathbb{I}(x = 40) + \beta_2\mathbb{I}(x = 50) + \beta_3\mathbb{I}(\text{Type} = \text{Sig Dep})_{i,x,k} + u_i + b_k + \epsilon_{i,x,k}$$

In this equation, $\pi_{i,x,k}(a)$ is the inference made by observer i after observing action a in decision problem (x, k) . We use dummies for the various cash options to allow for nonlinear effects. We use an indicator $\mathbb{I}(\text{Type} = \text{Signal Dep})_{i,x,k}$ to capture whether the observer i ’s own type was “signal dependent” when they themselves faced decision problem (x, k) in the Actor Stage. Participant-level and shop-level fixed effects are captured by u_i and b_k , respectively. To identify these coefficients, we focus on participants who exhibited variation in their type classifications—that is, they were classified as SIR in some decision problems and signal dependent in other problems. We leave out $x = 30$ and the SIR type as the omitted groups. Finally, we cluster standard errors at the individual level to account for correlation within each person.

Our findings, shown in Table 5, strongly suggest that a person’s type has both a statistically and calibrationally significant effect on her inferences. Consider, for example, when a participant observes $a = R$. We find that an observer’s own type matters about as much as a change in x from \$30 to \$50 (see column 1 of Table 5). The effect is slightly weaker but still significant when $a = A$ (see column 4 of Table 5). Since this panel approach controls for heterogeneity by identifying off of within-person variation, Table 5 is our most direct evidence that a person’s own type shapes her

inferences.

DETERMINANTS OF INFERENCES						
	Dep. Variable: Stated Belief					
	After $a = R$			After $a = A$		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	28.95 (0.69)	29.19 (1.06)	26.74 (1.43)	78.18 (0.58)	79.59 (0.82)	72.34 (2.36)
Cash Option = 40	1.31 (0.73)	1.47* (0.73)	1.37 (0.73)	0.95 (0.62)	0.94 (0.61)	0.71 (0.60)
Cash Option = 50	2.52* (1.03)	2.84** (1.02)	2.61* (1.02)	4.57*** (0.83)	4.54*** (0.82)	4.00*** (0.80)
Type = <i>Signal Dependent</i>	-2.56** (0.82)	-1.80* (0.81)	-1.81* (0.81)	1.93** (0.61)	1.87** (0.61)	1.84** (0.61)
Unknown-Shop Inference			0.08* (0.04)			0.10*** (0.03)
Shop-Level FEs	X	✓	✓	X	✓	✓
Observations	3101	3101	3101	3101	3101	3101

Estimated via panel regression with participant-level fixed effects. Standard errors (in parentheses) are clustered at the individual level. The coefficient on type is significant at $p = .0265$ (column 3) and $p = .0029$ (column 6). Sample restricted to decision problems where the observer was either *signal-independent-reject* or *signal-dependent* type. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Participant-Level Fixed-Effects Model of Inferences, Experiment 1. *Participants’ strategies (and thus their types) shaped their subsequent inferences (cols 1-2 and 4-5), even when controlling for additional idiosyncrasies using the unknown-shop inferences (cols 3 and 6).*

In columns 2-3 and 5-6 of Table 5, we extend this analysis in two ways. First, we demonstrate that our results are not merely an artifact of shop-specific heterogeneity by including fixed effects for each of the businesses in the study (columns 2 and 5). Second, we leverage the “unknown shop” trials where the specific business was hidden to the observer. Recall that participants made

inferences in unknown-shop decision problems for all combinations of a and x . We can therefore control for an observers' beliefs in the corresponding unknown-shop problem (columns 3 and 6), which allows us to further account for idiosyncratic differences beyond using simple fixed effects. Even after controlling for these factors, we continue to find that a person's own type shaped her inferences.

We find a similar role of type on inferences at each individual business when we revisit the more-structural estimation in Table 4 and separate results by type. We present this analysis in Table A3, where we present the reaction coefficients (from Equation 2 and 3) for each type. As suggested by Table 5, the shop-level coefficients show that signal-dependent observers under-inferred by less than their signal-independent-reject counterparts after seeing $a = A$. Moreover, signal-dependent observers over-inferred by more after seeing $a = R$. Intuitively, participants may have overestimated how much information others' choices revealed in decision problems where they themselves took actions that revealed their signal (i.e., in problems where they were a signal-dependent type). Taken together, the patterns described above match the prediction of Corollary 1.

Discussion of Results. Our data is consistent with inaccurate beliefs about others' types that consequently shape the interpretation of others' actions. Specifically, inaccurate beliefs about others' types explain *both* the over- and under-reactions we document in Table 4 and the type-dependent inferences of observers that we highlight in Table 5. Thus, our collective evidence accords with the predictions in Corollary 1. Along these lines, we push our data (perhaps beyond its limit) to back out observers' implicit beliefs about the distribution of types within the population as a function of the observer's *own* type. To do this, we utilize the expressions in Equation 1 to solve for the implied beliefs about each of the three main types. We then calculate these implied beliefs for each decision problem, and we average these results by the observer's type in that decision problem. We find that signal-independent-reject types (implicitly) believe that their own type makes up 31% of the population, that signal-dependent types make up 53%, and that signal-independent-accept types—a tiny component of the true population—comprise the remaining 16%. In contrast, the inferences of signal-dependent types suggest that they believe SIR types comprise 22% of the population, signal-dependent types (their own type) comprise 62%, and SIA types make up the remaining 16%.²³

More broadly, the pattern of under- and over-inference that we observe accords with recent work by [Augenblick et al. \(2023\)](#) who find that experimental participants under-infer from weak signals but over-infer from strong signals. Recall that our observers under-infer after seeing an actor accept the gift card. In our setting, action $a = A$ is a very strong indication that the actor received good

²³It is worth noting that our results are consistent with a variety of mechanisms that may cause a person to over-estimate the prevalence of their own tastes (e.g., projection bias, neglecting assortative matching based on tastes as in [Frick et al., 2022](#), or limited information as in [Dawes, 1989](#)).

information since extremely few people took the card otherwise. Thus, we too find under-reaction to “strong signals”—that is, under-inference after seeing a highly informative action. Moreover, our observers over-react after seeing an actor reject the gift card. The action $a = R$ is a much weaker indication of the actor’s private information, since many actors rejected the card even when they received a good signal. Thus, we also find over-reaction to “weak signals”—that is, over-inference after seeing a relatively uninformative action.

This pattern and interpretation also helps reconcile our findings with respect to Amazon. Given the features of tastes over Amazon gift cards discussed above, we would expect actions regarding Amazon to provide strong signals relative to other shops when $a = R$ and relatively weak signals when $a = A$ (since there are more SIA types for Amazon; see Table A1). This accords with our findings in Table 4 and provides further evidence of the pattern noted above.

Observers’ Choices after Social Learning

We now analyze the choices that observers made based on what they saw. This resembles the typical analysis of sequential actions in social-learning environments. As shown in Table 6, we find that these choices reflect significant social learning: participants themselves chose $a = A$ much more frequently after observing $a = A$ than after $a = R$. In fact, comparing Tables 2 and 6 highlights that the aggregated choices of participants in the second-mover position (i.e., observers) closely resemble the aggregated choices of participants in the first-mover position (i.e., actors).

CHOICE AFTER OBSERVING $a = A$					CHOICE AFTER OBSERVING $a = R$				
	Cash Option					Cash Option			
	\$30	\$40	\$50	Total		\$30	\$40	\$50	Total
$a = R$	352	773	447	1572	$a = R$	727	1492	756	2975
$a = A$	451	830	353	1634	$a = A$	76	111	44	231

Table 6: Counts of Choices for Observers, Experiment 1. *Observer’s choices in aggregate resemble the privately-informed decisions from Stage 2.*

The fact that the distribution of observer’s actions mirrors the distribution of privately-informed choices from Stage 2 may suggest a high degree of information transmission between actors and observers. However, our results above on observers’ beliefs suggest that information transmission

is flawed. We now utilize both choice data and belief data to better understand the drivers of participants' choices. To do so, we employ logit regressions with both subject-level and business-level fixed effects. Given our rich data, we include interaction terms between the observed action and all other variables to essentially fit separate choice models for each observed action (A vs. R). More formally, for any variable y and each $a \in \{A, R\}$, let $f_a(y) = \mathbb{I}\{\tilde{a} = a\} \times y$, where \tilde{a} is the observed action. Our econometric model of observer i 's choice, denoted by $a_{i,x,k}$, can then be summarized as:

$$\ln \left(\frac{\Pr(a_{i,x,k} = A|\tilde{a})}{1 - \Pr(a_{i,x,k} = A|\tilde{a})} \right) = \sum_{a \in \{A,R\}} [\beta_a f_a(\pi_{i,x,k}(a)) + \gamma_a f_a(Z_{i,x,k}) + f_a(b_k)] + u_i + \epsilon_{i,x,k}.$$

The vector $Z_{i,x,k}$ captures critical features of each decision problem including: the cash bonus x that the actor faced, the observer's rating of shop k , the observer's type in that same decision problem, and the observer's unknown-shop inference and choice after observing the same action \tilde{a} (for the same x). We present the results from this analysis in Table 7 and display marginal effects.

In column 1, we utilize a simple specification to demonstrate that stated beliefs are important in determining choice behavior of observers, as predicted by a rational model of social learning ($p < 0.001$). Those who infer a higher likelihood of $s = h$ should be more likely to accept the card, and our participants' behavior follows this basic feature of learning. We also show that the observer's preferences, as captured by her type from the Actor Stage, are (unsurprisingly) a significant driver of her behavior in the Observer Stage ($p < 0.001$). In column 2, we focus on how the observed action shaped choices. This mirrors our analysis from the Actor Stage (see Table A2), but since observers didn't receive the underlying signal, we instead control for the binary action they observed. We find that this estimation yields similar results to that of the Actor Stage.

While the previous results show some basic ways in which observer's choices may reflect rational social learning, a puzzle remains given that our belief data suggests systematic errors in information extraction. We thus simultaneously analyze the role of beliefs and observed actions to determine whether observers extracted information from the action beyond what is reflected in their stated beliefs. Column 3 of Table 7 shows that both the observer's stated beliefs and the observed action have a significant effect on choices. Two possibilities remain. First, it is possible that participants' stated beliefs didn't reflect their true underlying beliefs. Alternatively, perhaps observers' actions were swayed by some form of social influence or herding wherein the observer followed the action of the actor above and beyond the inference they drew from it.²⁴

²⁴The fact that the value of x does not seem to determine choices in column 2 (where we are not controlling for beliefs) is suggestive evidence in favor of this latter alternative.

DETERMINANTS OF OBSERVERS' CHOICES, EXP 1

	Dep. variable: $\mathbb{I}(\text{Chose Gift Card})$			
	(1)	(2)	(3)	(4)
Stated Belief	0.008*** (0.000)		0.004*** (0.000)	0.003*** (0.000)
Observed: $a = A$		0.435*** (0.020)	0.285*** (0.034)	0.203*** (0.027)
Cash Option = 40	-0.028* (0.013)	-0.004 (0.010)	-0.022 (0.012)	-0.008 (0.012)
Cash Option = 50	-0.089*** (0.018)	-0.016 (0.016)	-0.070*** (0.017)	-0.038* (0.017)
Rating = Neutral	0.036 (0.023)	0.050* (0.022)	0.035 (0.021)	0.030 (0.019)
Rating = Positive	0.030 (0.026)	0.061* (0.026)	0.036 (0.024)	0.033 (0.023)
Rating = Strongly Positive	0.068* (0.033)	0.121*** (0.033)	0.078* (0.032)	0.076* (0.030)
Type = <i>Signal Dependent</i>	0.180*** (0.017)	0.295*** (0.018)	0.174*** (0.017)	0.166*** (0.016)
WTP	0.002*** (0.000)	0.000 (0.000)	0.002*** (0.000)	0.002*** (0.000)
PersonalExperience	0.028 (0.016)	0.014 (0.018)	0.027 (0.014)	0.027* (0.014)
Unknown-Shop Inference				-0.001 (0.000)
Chose Card, Unknown Shop				0.209*** (0.036)
Shop-Level FEs	✓	✓	✓	✓
Observations	5332	6202	5332	5332

Estimated via fixed-effects logit with marginal effects displayed. Standard errors (in parentheses) are clustered at the individual level and calculated via delta method. Baseline categories are Observed: $a = R$, Cash Option = \$30, Rating = Negative, and Type = *Signal-Independent Reject*.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Observers' Choices, Experiment 1. Marginal effects from a fixed-effects logit model.

Finally, column 4 utilizes the behavior at the unknown-shop to control for both the average inference at the unknown shop and the choice when facing an unknown shop. By including the unknown-shop beliefs and choices—coarse proxies for individual differences in stated beliefs and risk attitudes, respectively—we control for two potentially important channels. We find that the direct effect of observing $a = A$ persists. Moreover, the magnitude is quite large.

Our collective body of evidence suggests participants may have made two different types of mistakes: observers may have chosen gift cards that they ought to reject, or may have rejected cards they ought to choose. We present suggestive evidence of both errors in in Table A5. There, we show that approximately 20% of participants who did not take the gift card for either signal in the Actor Stage subsequently took the card in the observer stage after seeing the action $a = A$. We also note that a high proportion of signal-dependent types followed the actor’s choice after $a = R$. Given that they over-infer from this action, this suggests some may have erroneously rejected the gift card due to this over-inference.

5 Results, Experiment 2

Our second experiment explores whether observers improved their inferences when given a noisy signal about the actor’s taste. This captures a key aspect of many real-world settings: people can sometimes observe characteristics or demographics of others when socially learning. To put this into practice, during the observer stage of the experiment, we told participants the subjective rating about the business— $r \in \{\text{Negative, Neutral, Positive, Strongly Positive}\}$ —of the person whose action they observed. The experiment was otherwise the same as Experiment 1.²⁵

We focus our analysis on the Observer Stage of Experiment 2. Table 8 shows the average of observers’ beliefs, $\pi(a|r)$, for each action $a \in \{A, R\}$ conditional on the observed rating r . Following the presentation of Table 3, we also provide as a benchmark what the fully-informed Bayesian would respond (calculated for each combination of business and observed rating, then averaged). The left side shows stated beliefs and these benchmarks when the actor rejected the card ($a = R$); the right side does so when the actor accepted it ($a = A$). As in Experiment 1, the general direction of updating is reasonable: beliefs about the actor’s signal move up after $a = A$ and down after $a = R$ (relative to the prior belief of 50%). Perhaps the most striking feature from Table 8 comes from comparing it to Table 3. Doing so reveals that, if anything, aggregate inferences were less accurate in Experiment 2 despite observers having more information about others.

²⁵The results from the Survey and Actor Stages of Experiment 2 closely match those of Experiment 1. For the sake of relative brevity, we present the results from these two stages in an Online Appendix.

AGGREGATE INFERENCES, ALL SHOPS, EXP 2

	After $a = R$		After $a = A$	
	Stated Belief	Benchmark	Stated Belief	Benchmark
Observed: Rating = Negative	41.98 (1.26)	39.66 -	76.44 (1.60)	100.00 -
Observed: Rating = Neutral	38.39 (1.19)	40.02 -	75.23 (1.06)	96.76 -
Observed: Rating = Positive	34.01 (1.60)	29.49 -	76.33 (0.85)	93.76 -
Observed: Rating = Strongly Positive	27.76 (1.80)	20.19 -	75.54 (1.08)	94.71 -
Observations (column total)	4003	-	3960	-

Standard errors (in parentheses) are clustered at the individual level. Benchmarks are calculated for each {business \times observed rating} combination.

Table 8: Inferences in Experiment 2. *After observing $a = R$, participants ($n = 226$) utilized ratings information to form their inferences, but under-inferred from highly informative positive ratings (columns 1 vs. 2). As in Experiment 1, participants under-inferred for all ratings after observing $a = A$ (columns 3 vs. 4).*

But how did observers’ beliefs respond to the actor’s rating? Note that this rating data is highly informative in the case of $a = R$: those who highly rate the business were much more likely to be signal-dependent types and hence their action is more revealing of the signal they received. We find that participants’ stated beliefs respond in accordance with this more subtle form of rational belief updating after observing $a = R$. Specifically, observers’ beliefs tend to decrease in the actor’s rating and thus exhibit the comparative static discussed in Section 3.2 ($p < 0.001$ for test of equality of descending rows in Table 8). This suggests that our participants (at least partially) considered an actor’s preferences when inferring from them. However, this pattern does not bear out when $a = A$.

Intriguingly, we find that—much like in Experiment 1—observers under-infer from $a = A$. This occurs for all observed ratings (compare columns 3 and 4 of Table 8). Moreover, although observers followed some qualitative patterns suggested by rational updating after observing $a =$

R , their stated beliefs are under-responsive to the observed ratings relative to the fully-informed benchmark (compare columns 1 and 2). The small difference between rows 1 and 4 of column 1 relative to the benchmarks highlights how observers failed to sufficiently adjust for the ratings information they saw. Beliefs ought to have moved significantly more than they did, given that the information was highly informative. This again bears out the pattern from Experiment 1: participants under-inferred from strong, highly-informative actions.²⁶ Our collection of evidence therefore provides a somewhat cautionary tale. If one thought that providing information would mitigate errors in social learning, our results suggests that it does not.

Observers' Choices after Social Learning. Finally, we reproduce our analysis from Experiment 1 wherein we examined the choices that observers made after seeing others' actions. Using a fixed-effects logit framework that closely mirrors that of Experiment 1, we find that observers in Experiment 2 exhibit a few behaviors that obey some rational comparative statics, but systematically depart from the rational benchmark in yet other ways.

First, we verify that choices exhibit the same basic features as documented in Experiment 1. Indeed, we find that observers' choices are naturally related to their stated beliefs (column 1 of Table 9), the action they observed (column 2), and their own type. In column 3, we show that seeing action A influences an observer's behavior above and beyond its effect on beliefs. As before, we interpret this as evidence for erroneous social learning akin to herding.

Second, we explore how choices were influenced by the observed rating. In column 2, we find that the choice behavior of observers is (perhaps counter intuitively) uncorrelated with the rating they saw. Given how informative this rating is, we find it surprising that observers seemingly under-utilized this information. Furthermore, when controlling for observers' stated beliefs, we find suggestive evidence that people took the card—regardless of the action observed—when observing those with particularly high ratings. This may be plausible if observers learned about their own tastes for a business from others' ratings. However, given that we control for personal experience with the business in these analyses, it seems reasonable to interpret this as a mistake stemming from either erroneous social learning or social influence rather than a result of optimal inference.

²⁶In Appendix Table A6, we show that one reason for this may be because participants' own rating still weighs heavily on their inferences. In that analysis we show that, fixing the participants' own rating, they infer from others' ratings in the rational direction. However, those with particularly strong ratings themselves wildly under-infer from the rating information.

DETERMINANTS OF OBSERVERS' CHOICES, EXP 2

	Dep. variable: $\mathbb{I}(\text{Chose Gift Card})$			
	(1)	(2)	(3)	(4)
Stated Belief	0.007*** (0.000)		0.004*** (0.000)	0.004*** (0.000)
Observed: $a = A$		0.418*** (0.019)	0.332*** (0.027)	0.207*** (0.028)
Observed: Rating = Neutral	0.023* (0.010)	0.007 (0.010)	0.009 (0.009)	0.013 (0.014)
Observed: Rating = Positive	0.034* (0.014)	0.023 (0.014)	0.017 (0.013)	0.002 (0.015)
Observed: Rating = Strongly Positive	0.050** (0.017)	0.025 (0.015)	0.029 (0.015)	0.054*** (0.016)
Type = <i>Signal Dependent</i>	0.180*** (0.020)	0.295*** (0.019)	0.165*** (0.020)	0.162*** (0.021)
WTA	0.001*** (0.000)	0.001 (0.000)	0.001** (0.000)	0.002*** (0.000)
PersonalExperience	0.029 (0.015)	0.020 (0.019)	0.035* (0.015)	0.031 (0.018)
Unknown-Shop Inference				-0.001 (0.000)
Chose Card, Unknown Shop				0.282*** (0.044)
Shop-Level FEs	✓	✓	✓	✓
Observations	6419	7760	6419	3850

Estimated via fixed-effects logit with marginal effects displayed. Standard errors (in parentheses) are clustered at the individual level and calculated via delta method. Baseline categories are Observed: $a = R$, Observed Rating = Negative, and Type = *Signal-Independent Reject*.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9: Observers' Choices in Experiment 2. Marginal effects from a fixed-effects logit model.

Finally, we conclude our analysis by utilizing the observer's behavior in the unknown-shop decision problems. As in Experiment 1, we find that the choice behavior in the unknown-shop

problems is highly correlated with choice behavior in the known-shop problems. While this analysis is limited by a severely reduced sample, our results in column 4 suggest that the same factors as in Experiment 1 primarily drive observers’ decisions: the stated belief, choice behavior in the Actor Stage, choice behavior in the unknown-shop trials, and the observed action.

Discussion of Results. Experiment 2 demonstrates that a (strong) informational intervention does not significantly alter the basic patterns in observers’ inferences that we saw in Experiment 1. However, we caution against directly comparing stated beliefs across Experiments 1 and 2.²⁷ Since our second experiment made the other person’s rating highly salient, we believe there may be fundamentally different psychologies in play across Experiments 1 and 2. Concretely, while we suspect our intervention was straightforward and easy for participants to interpret, Experiment 2 may have induced participants to compare themselves against the salient benchmark of the other person’s rating, as suggested by Table A6. Of course, such self-other comparisons were not possible in Experiment 1. Future work could further explore how making the differences between actors and observers salient interacts with social learning.

6 Conclusion

In this paper, we developed a field-in-the-lab experiment to explore social learning when agents have heterogeneous tastes. We find that, while social learning does occur, there are systematic deviations from fully-rational learning. By eliciting beliefs in addition to actions in a social-learning paradigm, we are able to uncover previously undiscussed patterns in social learning. Participants under-infer from highly-informative actions and over-infer from weakly-informative ones. We show that these misinferences are consistent with inaccurate beliefs about others’ preferences, perhaps stemming from projection bias. Understanding the deeper underlying mechanisms for misunderstanding others remains an open challenge.

We believe that our framework and results speak to some of the challenges faced by agents in a variety of social-learning contexts. For example, our results shed light on settings where actors’ objectives exhibit heterogeneity for reasons beyond simple differences in tastes, such as differing budget constraints or costs. Consider, for instance, a novel medical treatment that requires traveling to a special facility. When learning about this treatment based on others’ choices, a patient should account for the fact that some may opt for the status quo not because they have negative information about the novel treatment but simply because they can’t afford it or refuse to travel. Yet the frictions we document in this paper—wherein observers seemingly project their

²⁷Relatedly, Experiment 2 does not lend itself to the analytic approaches from Experiment 1. For example, when estimating Equation 2, since p and q both vary with the observed rating and the business, we are limited in our number of observations. As shown in Table A4, we face many cases with $p = 0$, which makes our approach to estimating c_A intractable.

own circumstances on the situation—may loom large. Likewise in agricultural contexts, farmers often have opportunities to learn from their neighbors about the benefits of new technologies, such as hybrid seeds and soil additives (Duflo et al., 2008; Conley and Udry, 2010). These products often work well for some soil types but not others, and farmers should therefore account for this variation when interpreting their neighbors’ choices (Munshi, 2004). As our evidence suggests, documented failures to learn in these environments might stem from agents failing to appreciate others’ differences and consequently misinterpreting their actions.

Moreover, our findings could help inform the design of social networks that best facilitate the spread of (factual) information. Harnessing such networks is of growing interest beyond the tech sphere; for example, organizations have tried to promote the adoption of agricultural technologies by connecting farmers in social networks (Fabregas et al., 2019). However, if people do not properly account for heterogeneity in such networks, then small homogeneous networks may outperform large diverse ones. Additionally, the results from our informational treatment in Experiment 2 suggest that interventions intended to improve social learning may need to go beyond simply providing information about others. We believe that better understanding such barriers to information transmission is central to better designing policies or networks aiming to promote learning.

References

- AMBUEHL, S., B. D. BERNHEIM, AND A. OCKENFELS (2021): “What Motivates Paternalism? An Experimental Study,” *American Economic Review*, 111, 787–830.
- ANDERSON, L. AND C. HOLT (1997): “Information Cascades in the Laboratory,” *American Economic Review*, 87, 847–62.
- AUGENBLICK, N., E. LAZARUS, AND M. THALER (2023): “Overinference from Weak Signals and Underinference from Strong Signals,” Mimeo.
- BA, C., J. A. BOHREN, AND A. IMAS (2022): “Over- and Underreaction to Information,” Mimeo.
- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- BENJAMIN, D. J. (2019): “Errors in probabilistic reasoning and judgment biases,” in *Handbook of Behavioral Economics: Applications and Foundations*, Elsevier, vol. 2, 69–186.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 992–1026.
- BOHREN, A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.

- BOHREN, J. A. AND D. HAUSER (2021): “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, 89, 3025–3077.
- BROWN, C. E. (1982): “A False Consensus Bias in 1980 Presidential Preferences,” *Journal of Social Psychology*, 118, 137–138.
- BURSZTYN, L., F. EDERER, B. FERMAN, AND N. YUCHTMAN (2014): “Understanding Mechanisms Underlying Peer Effects: Evidence From a Field Experiment on Financial Decisions,” *Econometrica*, 82, 1273–1301.
- BURSZTYN, L. AND D. Y. YANG (2022): “Misperceptions about Others,” *Annual Review of Economics*, 14, 425–452.
- BUSHONG, B. AND T. GAGNON-BARTSCH (2023): “Failures in Forecasting: An Experiment on Interpersonal Projection Bias,” *Management Science*, Forthcoming.
- CAI, H., Y. CHEN, AND H. FANG (2009): “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, 99, 864–82.
- ÇELEN, B. AND S. KARIV (2004): “Distinguishing Informational Cascades from Herd Behavior in the Laboratory,” *American Economic Review*, 94, 484–498.
- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): “Testing Models of Social Learning on Networks: Evidence From Two Experiments,” *Econometrica*, 88, 1–32.
- CONLEY, T. G. AND C. R. UDRY (2010): “Learning About a New Technology: Pineapple in Ghana,” *American Economic Review*, 100, 35–69.
- CONLON, J. J. (2019): “Major Malfunction,” *The Journal of Human Resources*, 56, 922–939.
- CRUCES, G., R. PEREZ-TRUGLIA, AND M. TETAZ (2013): “Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment,” *Journal of Public Economics*, 98, 100–112.
- DANZ, D., L. VESTERLUND, AND A. J. WILSON (2022): “Belief Elicitation and Behavioral Incentive Compatibility,” *American Economic Review*, 112, 2851–2883.
- DAWES, R. M. (1989): “Statistical Criteria for Establishing a Truly False Consensus Effect,” *Journal of Experimental Social Psychology*, 25, 1–17.
- (1990): “The Potential Nonfalsity of the False Consensus Effect,” in *Insights in Decision Making: A Tribute to Hillel J. Einhorn*, ed. by R. M. Hogarth, University of Chicago Press, 179–199.
- DE FILIPPIS, R., A. GUARINO, P. JEHIEL, AND T. KITAGAWA (2022): “Non-Bayesian Updating in a Social Learning Experiment,” *Journal of Economic Theory*, 199, 105188.
- DELANDE, A. AND C. F. MANSKI (2012): “Candidate Preferences and Expectations of Election Outcomes,” *Proceedings of the National Academy of Sciences*, 109, 3711–3715.

- DELAVANDE, A. AND B. ZAFAR (2019): “University Choice: The Role of Expected Earnings, Nonpecuniary Outcomes, and Financial Constraints,” *Journal of Political Economy*, 127, 2343–2393.
- DOMINITZ, J. AND A. A. HUNG (2009): “Empirical models of discrete choice and belief updating in observational learning experiments,” *Journal of Economic Behavior & Organization*, 69, 94–109, “Individual Decision-Making, Bayesian Estimation and Market Design: A Festschrift in honor of David Grether”.
- DUFLO, E., M. KREMER, AND J. ROBINSON (2008): “How High are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya,” *American Economic Review*, 98, 482–488.
- ENGELMANN, D. AND M. STROBEL (2012): “Deconstruction and Reconstruction of an Anomaly,” *Games and Economic Behavior*, 76, 678–689.
- EYSTER, E. AND M. RABIN (2010): “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–243.
- EYSTER, E., M. RABIN, AND G. WEIZSÄCKER (2018): “An Experiment on Social Mislearning,” Mimeo.
- FABREGAS, R., M. KREMER, AND F. SCHILBACH (2019): “Realizing the potential of digital development: The case of agricultural advice,” *Science*, 366, eaay3038.
- FAN, T., Y. LIANG, AND C. PENG (2023): “The Inference-Forecast Gap in Belief Updating,” Mimeo.
- FARO, D. AND Y. ROTTENSTREICH (2006): “Affect, Empathy, and Regressive Mispredictions of Others’ Preferences Under Risk,” *Management Science*, 52, 529–541.
- FISCHHOFF, B. AND W. B. DE BRUIN (1999): “Fifty–fifty = 50%?” *Journal of Behavioral Decision Making*, 12, 149–163.
- FREDERICK, S. (2012): “Overestimating Others’ Willingness to Pay,” *Journal of Consumer Research*, 40, 1–21.
- FRICK, M., R. IJIMA, AND Y. ISHII (2020): “Misinterpreting Others and the Fragility of Social Learning,” *Econometrica*, 88, 2281–2328.
- (2022): “Dispersed Behavior and Perceptions in Assortative Societies,” *American Economic Review*, 112, 3063–3105.
- GAGNON-BARTSCH, T. (2016): “Taste Projection in Models of Social Learning,” Mimeo.
- GAGNON-BARTSCH, T., M. PAGNOZZI, AND A. ROSATO (2021): “Projection of Private Values in Auctions,” *American Economic Review*, 111, 3256–3298.
- GAGNON-BARTSCH, T. AND A. ROSATO (2023): “Quality is in the Eye of the Beholder: Taste Projection in Markets with Observational Learning,” Mimeo.

- GOEREE, J., T. PALFREY, B. ROGERS, AND R. D. MCKELVEY (2007): “Self-Correcting Information Cascades,” *Review of Economic Studies*, 74, 733–762.
- GOEREE, J. K., T. R. PALFREY, AND B. W. ROGERS (2006): “Social Learning with Private and Common Values,” *Economic Theory*, 28, 245–264.
- GREYER, D. M. (1980): “Bayes Rule as a Descriptive Model: The Representativeness Heuristic,” *Quarterly Journal of Economics*, 95, 537–557.
- HENDRICKS, K., A. SORENSEN, AND T. WISEMAN (2012): “Observational Learning and Demand for Search Goods,” *American Economic Journal: Microeconomics*, 4, 1–31.
- HUNG, A. A. AND C. R. PLOTT (2001): “Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions,” *American Economic Review*, 91, 1508–1520.
- JENSEN, R. (2010): “The (Perceived) Returns to Education and the Demand for Schooling,” *The Quarterly Journal of Economics*, 125, 515–548.
- KAHNEMAN, D. AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.
- KNEELAND, T. (2015): “Identifying Higher-Order Rationality,” *Econometrica*, 83, 2065–2076.
- KRUEGER, J. AND R. W. CLEMENT (1994): “The Truly False Consensus Effect: An Ineradicable and Egocentric Bias in Social Perception,” *Journal of Personality and Social Psychology*, 67, 596–610.
- KÜBLER, D. AND G. WEIZSÄCKER (2004): “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *The Review of Economic Studies*, 71, 425–441.
- MARKS, G. AND N. MILLER (1987): “Ten Years of Research on the False-Consensus Effect: an Empirical and Theoretical Review,” *Psychological Bulletin*, 102, 72–90.
- MORETTI, E. (2011): “Social Learning and Peer Effects in Consumption: Evidence from Movie Sales,” *The Review of Economic Studies*, 78, 356–393.
- MULLEN, B., J. L. ATKINS, D. S. CHAMPION, C. EDWARDS, D. HARDY, J. E. STORY, AND M. VANDERKLOK (1985): “The False Consensus Effect: A Meta-Analysis of 115 Hypothesis Tests,” *Journal of Experimental Social Psychology*, 21, 262–283.
- MUNSHI, K. (2004): “Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution,” *Journal of Development Economics*, 73, 185–213.
- NÖTH, M. AND M. WEBER (2003): “Information Aggregation with Random Ordering: Cascades and Overconfidence,” *The Economic Journal*, 113, 166 – 189.
- PRELEC, D. (2004): “A Bayesian Truth Serum for Subjective Data,” *Science*, 306, 462–466.

- ROSS, L., D. GREENE, AND P. HOUSE (1977): “The “False Consensus Effect”: An Egocentric Bias in Social Perception and Attribution Processes,” *Journal of Experimental Social Psychology*, 13, 279–301.
- SMITH, L. AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–398.
- SORENSEN, A. (2006): “Social learning and health plan choice,” *The Rand journal of economics*, 37, 1–29.
- VISCUSI, W. K. (1985): “A Bayesian perspective on biases in risk perception,” *Economics Letters*, 17, 59–62.
- WISWALL, M. AND B. ZAFAR (2015): “How Do College Students Respond to Public Information about Earnings?” *Journal of Human Capital*, 9, 117–169.

Appendix

A Supplemental Analyses and Figures

In this Appendix, we present additional results that are supplemental to the main conclusions of the paper.




Figure A1: Histogram of Ratings by Shop, Experiment 1. While ratings varied significantly across shops, all shops were generally viewed positively. For example, the lowest-rated shop (Old Navy; center row, center of figure) had an average rating of 2.48 out of 4.

FREQUENCIES OF STRATEGY TYPES

	CashOption=30		CashOption=40		CashOption=50	
	\bar{p}_{30}	\bar{q}_{30}	\bar{p}_{40}	\bar{q}_{40}	\bar{p}_{50}	\bar{q}_{50}
AMC	0.05	0.42	0.00	0.39	0.00	0.31
Amazon	0.19	0.66	0.03	0.77	0.02	0.60
ChickFilA	0.06	0.50	0.01	0.44	0.00	0.37
HomeDepot	0.07	0.59	0.01	0.59	0.01	0.39
OldNavy	0.07	0.44	0.02	0.40	0.00	0.32
PetSmart	0.07	0.47	0.01	0.39	0.00	0.31
Starbucks	0.04	0.54	0.01	0.48	0.00	0.37

Table A1: Strategy Types in Experiment 1. Each column represents the sample averages for strategy types for signal-independent-accept (p) and signal-dependent (q) participants, given the specified bonus-cash level and business.

DETERMINANTS OF ACTORS' CHOICES

	Dep. variable: $\mathbb{I}(\text{Chose Gift Card})$		
	(1)	(2)	(3)
Good Signal ($s = H$)	0.512*** (0.009)	0.513*** (0.009)	0.513*** (0.009)
Cash Option = 40	-0.086*** (0.010)	-0.085*** (0.010)	-0.088*** (0.010)
Cash Option = 50	-0.161*** (0.014)	-0.161*** (0.014)	-0.165*** (0.014)
Rating = Neutral	0.040 (0.021)	-0.004 (0.023)	0.028 (0.022)
Rating = Positive	0.194*** (0.023)	0.044 (0.028)	0.085** (0.027)
Rating = Strongly Positive	0.285*** (0.026)	0.088** (0.032)	0.123*** (0.030)
PersonalExperience		0.062*** (0.017)	0.043* (0.017)
WTP		0.004*** (0.000)	0.003*** (0.000)
Shop-Level FEs	\times	\times	\checkmark
Observations	5823	5823	5823

Estimated via fixed-effects logit with marginal effects displayed. Standard errors (in parentheses) are clustered at the individual level. Omitted categories are Bad Signal, CashOption=30, and Rating = Negative.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A2: Modeling Actors' Choices in Experiment 1. *Using a panel-logit model with subject-level fixed effects, we explore the determinants of observers choosing the gift card. All models include interactions between signal and other independent variables.*

REACTION COEFFICIENTS ACROSS BUSINESSES AND TYPES

	After $a = R$		After $a = A$	
	<i>Signal Ind. Reject</i>	<i>Signal Dependent</i>	<i>Signal Ind. Reject</i>	<i>Signal Dependent</i>
AMC	1.58 (0.18)	1.84 (0.24)	0.37 (0.02)	0.41 (0.03)
Amazon	0.42 (0.10)	0.81 (0.07)	0.39 (0.04)	0.50 (0.02)
Chick-Fil-A	1.45 (0.15)	1.60 (0.18)	0.38 (0.03)	0.50 (0.03)
Home Depot	0.74 (0.13)	1.17 (0.11)	0.34 (0.02)	0.43 (0.02)
Old Navy	1.48 (0.17)	1.66 (0.26)	0.43 (0.03)	0.46 (0.04)
PetSmart	1.28 (0.18)	1.88 (0.19)	0.39 (0.03)	0.47 (0.03)
Starbucks	1.36 (0.14)	1.22 (0.17)	0.35 (0.03)	0.37 (0.02)
Observations (column total)	1507	1380	1101	1157

Standard errors (in parentheses) are clustered at the individual level.

Table A3: Shop-Level Reaction Coefficients for Two Primary Strategy Types, Experiment 1. *Participants' reactions varied depending on both their own type and the homogeneity of tastes for a given shop.*

FREQUENCIES OF STRATEGY TYPES								
	Negative		Neutral		Positive		Str Positive	
	\bar{p}_{40}	\bar{q}_{40}	\bar{p}_{40}	\bar{q}_{40}	\bar{p}_{40}	\bar{q}_{40}	\bar{p}_{40}	\bar{q}_{40}
AMC	0.00	0.20	0.01	0.25	0.04	0.46	0.00	0.88
Amazon	0.00	0.64	0.00	0.79	0.04	0.84	0.07	0.78
Chick-Fil-A	0.00	0.20	0.00	0.23	0.07	0.43	0.09	0.62
Home Depot	0.00	0.31	0.00	0.30	0.02	0.54	0.00	0.63
Old Navy	0.00	0.00	0.00	0.34	0.03	0.58	0.00	0.73
PetSmart	0.00	0.17	0.00	0.26	0.02	0.50	0.00	0.63
Starbucks	0.00	0.36	0.02	0.31	0.02	0.55	0.03	0.70

Table A4: Relationship Between Ratings and Strategy Types in Experiment 2 *Sample averages for signal-independent-accept (p) and signal-dependent (q) types for a given rating (columns) and business (rows).*

AFTER OBSERVING $a = R$		
STRATEGY IN ACTOR STAGE	CHOICE IN OBSERVER STAGE	
	$a = R$	$a = A$
Signal Independent Reject	32.51 (1.37) $n = 1548$	69.90 (3.05) $n = 39$
Signal Dependent	26.07 (1.45) $n = 1346$	57.92 (3.74) $n = 142$
AFTER OBSERVING $a = A$		
STRATEGY IN ACTOR STAGE	CHOICE IN OBSERVER STAGE	
	$a = R$	$a = A$
Signal Independent Reject	77.25 (1.25) $n = 1243$	77.72 (1.66) $n = 344$
Signal Dependent	74.39 (1.86) $n = 302$	82.19 (0.87) $n = 1196$

Standard errors (in parentheses) are clustered at the individual level.

Table A5: Average Inferences by Strategy Type in the Actor Stage and Choices in the Observer Stage. *Actions and inferences in the second stage are closely related, but a levels difference in inferences remains: signal-dependent types infer more for any observed action.*

INFERENCES AFTER $a = R$, EXPERIMENT 2

	OWN RATING			
	Negative	Neutral	Positive	Strong Positive
OBSERVED RATING				
Negative	39.00 (2.10)	43.31 (1.61)	42.37 (1.60)	39.91 (3.24)
Neutral	32.39 (2.28)	37.87 (1.49)	40.10 (1.59)	39.18 (2.74)
Positive	24.10 (2.95)	32.70 (1.87)	36.69 (1.92)	35.54 (3.70)
Strongly Positive	17.36 (2.62)	25.80 (2.11)	29.21 (2.30)	34.85 (3.83)
Observations (column total)	406	1340	1650	607

Standard errors (in parentheses) are clustered at the individual level.

Table A6: Inferences in Experiment 2 after Observing $a = R$, by Observed and Own Rating. *After seeing someone reject the gift card, participants utilized the ratings information to form their inferences. But their own ratings shaped the degree to which they utilized information about others.*