INFORMATION CASCADES WITH INFORMATIVE RATINGS: AN EXPERIMENTAL TEST[†]

JOHN J. CONLON[*], PAUL J. HEALY[**] AND YEOCHANG YOON[‡]

ABSTRACT. We study behavior in an information cascades setting where previous buyers of a product leave noisy but informative ratings about its true quality. Although this increases the amount of public information available, Yoon (2015) shows that ratings can hurt welfare in some settings. In particular, the added public information causes cascades to form earlier, but if agents stop purchasing then there are no new ratings to break the inefficient cascade behavior. Although we find some evidence roughly in line with the theory, those results are swamped by a strong tendency for subjects to purchase even when public information suggests they should not.

Keywords: Information Cascades; Herding; Activity Bias

JEL Classification: D83; C92; D03.

I. INTRODUCTION

In an information cascade consumers ignore their private information about a product's quality and rely instead on the past purchase decisions of others (see Banerjee, 1992; Bikhchandani et al., 1992, e.g.). Consequently, these later purchase decisions provide no new information for subsequent potential buyers. If the first few consumers are acting on incorrect information, the entire market follows and welfare suffers.

Intuitively, this negative result might be avoided if early consumers could somehow reveal their private information. For example, if consumers were to leave informative feedback about their experience with the product, then erroneous cascades would be avoided. Many consumers use Internet-based ratings systems to voluntarily provide such feedback, so one might argue that the relevance of information cascades has been greatly reduced by the proliferation of these ratings systems.[1]

[†]The authors thank Brian Rogers and Yaron Azrieli for their helpful comments and suggestions.
[*] Federal Reserve Bank of New York: john.conlon@ny.frb.org. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System as a whole.
[**] Department of Economics, The Ohio State University; healy.52@osu.edu.
[‡]Institute of Economics and Finance, Nanjing Audit University; yoonyeochang@nau.edu.cn.
[1]Ratings websites are not only prevalent for Internet retail, but also allow consumers to review restaurants, doctors, stores, services, and even items purchased from brick-and-mortar stores (O'Guinn, 2015).

This argument assumes that such ratings systems are indeed welfare enhancing. Yoon (2015) evaluates this claim in a theoretical model and shows that in fact ratings systems can *decrease* welfare in certain scenarios. Thus, the welfare impact of Internet rating systems is a bit ambiguous.

In his setting consumers sequentially face a good of unknown quality (either high or low) and must choose either to 'purchase' or 'pass'. Each has a conditionally independent private signal of the true quality. Those that purchase automatically leave a costless and truthful rating about their experience, while those who pass do not. All future consumers see who purchased as well as the ratings of those who purchased. Ratings are simply binary messages: either 'good' or 'bad'. Individuals' experiences with the good are idiosyncratic—for example, some consumers will have a bad experience even though the product is high quality—so ratings are noisy, but informative.

For example, imagine that Ann dines at a low-quality restaurant but happens to receive excellent service from the waiter. Based on that, she leaves a truthful positive review of her experience. Bob sees that Ann purchased, reads Ann's positive review, and has his own private signal about the restaurant's quality. He updates his beliefs based on those three sources of information and chooses optimally whether to purchase or pass. If he chooses to purchase and has a bad experience, he will leave a negative rating. Eventually, the accumulated ratings data will provide strong evidence that the restaurant's true quality is low, inducing consumers to pass.

Now consider the opposite scenario: the restaurant is truly of high quality, but the first consumer (Ann) leaves a negative review. This might induce future consumers to pass, leading to a cascade in which nobody dines at the restaurant. In this case ratings cannot correct the cascade, because consumers must purchase to be able to provide a review. Thus, ratings can correct cascades in which consumers are incorrectly purchasing, but cannot correct cascades in which consumers are incorrectly *not* purchasing.

Finally, suppose that all three sources of information (purchase decisions, private signals, and ratings) are noisy, but the ratings data are the least noisy. Consequently, Bob weights Ann's review more than her purchase decision or his own private signal. In this case the presence of ratings can accelerate the formation of cascades: just one or two bad ratings can cause consumers to stop purchasing regardless of their private information. Now cascades are forming quickly, and ratings are only helpful in correcting errors when the cascade features purchasing behavior. Thus, ratings only help correct errors when the true quality is low. When the true quality is high, Yoon (2015) concludes that the presence of ratings can actually reduce consumer welfare.

Our primary goal is to test this counterintuitive prediction. We run a laboratory test of this theoretical setting. In the first treatment ("No Ratings") subjects choose to purchase or pass but leave no ratings. This is the standard information cascades setting, except that one decision (purchasing) gives a risky payoff, while the other (passing) does not. The second treatment ("Ratings") is identical to the first except subjects are forced to leave truthful but noisy ratings of the good when they purchase. To test the counterintuitive prediction that ratings can be harmful, we make both private information and ratings nearly uninformative, with ratings only slightly more informative. Under our parameters, Yoon (2015) predicts that, when the quality is high, we will see more incorrect cascades in the Ratings treatment than in the No Ratings treatment.

To our surprise, our primary result is unrelated to the theoretical predictions. We find that subjects in all settings exhibit a strong preference for purchasing. They often purchase even when the available public information strongly suggests the good is of low quality. We estimate an econometric model of individual behavior that allows for various biases and find that this is by far the most substantial deviation from risk-neutral profit maximization. And it significantly affects market outcomes. Without this bias we should see roughly equal frequencies of purchasing cascades and passing cascades in the No Ratings treatment. Instead, purchasing cascades are nearly four times as common. This bias has not been found in previous studies on information cascades perhaps because (to our knowledge) all previous studies presented subjects with two risky choices, rather than one risky choice and one safe choice. And it cannot be rationalized by risk aversion because the subjects are biased toward the riskier choice. Instead, it appears to be either a preference for small-stakes risk (much like casino gambling) perhaps combined with a preference to avoid the more 'boring' choice of passing.

We do find that the presence of ratings reduces this purchasing bias. Purchasing cascades are still far more common, but with ratings they are only 2.5 times as frequent as passing cascades. Finally, the presence of ratings does break cascades, as expected. This enables a second (or third or fourth) cascade to form within a single period. The overall result is more cascades (of either type) in each period, with cascades lasting for shorter periods of time.

Related Literature

Following the seminal theoretical works of Banerjee (1992) and Bikhchandani et al. (1992)—and the experimental results of Anderson and Holt (1997)—many extensions and applications of informational cascades have appeared in the literature. Herding behavior has been found in settings as diverse as interviewing job applicants (Kubler and Weizsacker, 2003), car purchasing (Grinblatt et al., 2008), and financial markets

(Shachat and Srinivasan, 2011, Drehmann et al. 2007, Cipriani and Guarino 2008). Celen and Kariv (2005) and Corazzini and Greiner (2007) look at the impact of reducing the information available to players, while Feltovich (2002) and Kubler and Weizsacker (2004) investigate settings in which players can pay to increase their information.

Quite a few studies investigate the extent to which players deviate from perfect Bayesian play. Hung and Plott (2001), Noth and Weber (2003), Kraemer et al. (2006), and Goeree et al. (2007) all find that players overweight their private signal relative to public information, though Dominitz and Hung (2009) do not replicate this result. Huck and Oechssler (2000) find that even students in advanced economics courses fail to successfully employ Bayes' rule in an exam setting, though Morone et al. (2007) find that agents improve their play once the theory behind the experiments has been explained to them. Celen et al. (2010) find that respondents react more strongly to public information when it is presented as advice rather than just as a player's previous decision.

Another branch of the literature looks at whether adding a stochastic component to players' decisions can explain their deviations from Bayesian rationality. Goeree and Holt (2004), Goeree et al. (2007), and Kubler and Weizsacker (2004) all develop similar models that include a random element in player's choices. When players are aware that others' decisions are stochastic, they should understand that those decisions are less informative. This would lead to an apparent underweighting of public information (and overweighting of private information), as has been observed. But Goeree et al. (2007) find that players overweight their private signal even more than the stochastic choice theory would predict. We replicate this result in our setting, modulo the activity bias.

Our paper contributes to the literature in several ways. This experiment is the first laboratory study to structure the decision problem similar to a true a purchasing decision: purchasing is risky because payoffs depend on the state of nature, while not purchasing guarantees a safe payoff. Though the theoretical predictions are identical to the classical setting of Banerjee (1992) and Bikhchandani et al. (1992), the different structure allows for (and, as we shall see, leads to) framing effects that do not appear in other settings. In particular, we find that subjects overwhelmingly prefer to take the risky choice. This 'purchasing bias' runs counter to the usual finding of risk aversion.

We also add a ratings component to the experiment in order to test the predictions of Yoon (2015). According to his model, noisy yet informative ratings have a differential impact on welfare depending on the (randomly realized) quality of the product. If the quality is low, ratings decrease the likelihood of an incorrect (purchasing) cascade, whereas if the quality is high ratings *increase* the likelihood of an incorrect (passing) cascade.

Finally, we estimate a stochastic choice model similar to that of Goeree et al. (2007), which incorporates three bias parameters to explain players' choices. In addition to a stochastic element, we allow for players to weight private information differently than public information (as in Mobius et al., 2014, e.g.) and for players to prefer one decision to another intrinsically. We find a strong bias towards purchasing, which if ignored would lead us to conclude that players overweight their private information. But when taking into account this purchasing bias the overweighting of private information becomes much smaller and does little to explain players' decisions.

II. Theory

The Setting

Our experimental design is based on the theoretical model presented by Yoon (2015). A set of individuals, indexed by $i \in N = \{1,2,\dots,n\}$, sequentially chooses whether or not to buy a product. The quality of the product is unknown, but is equally likely to be high or low. We denote the realized quality by $q \in \{-1,1\}$, with the interpretation that -1 represents bad quality and 1 represents good quality.

Each agent sequentially chooses whether or not to buy the good. The purchase decisions are denoted by $x_i \in \{0,1\}$, where $x_i = 1$ indicates that i purchased the good and $x_i = 0$ indicates that they did not purchase. All agents share a common value for the good equal to qv, which equals $v > 0$ when $q = 1$ and $-v < 0$ when $q = -1$. After purchase, consumers also experience a random, idiosyncratic shock θ_i to their utility. Payoffs are zero to agents who do not purchase. Thus, the *ex-post* utility of agent i is given by $u_i(x_i|q,\theta_i) = (qv+\theta_i)x_i$. We assume each θ_i is independently drawn from a symmetric, mean-zero distribution F that is identical across agents.[2]

Before the purchase/pass decision, agents receive a privately-observed random signal $s_i \in \{-1,1\}$ of the quality of the good. Each s_i is conditionally independent and is correct ($s_i = q$) with probability $p \in (0.5,1)$. Each agent also sees the purchase/pass decision of all prior agents, and this is commonly known. Thus, the information available to agent $i > 1$ at the time of purchase is given by $(x_1, x_2, \dots, x_{i-1}, s_i)$. Agent 1 observes only s_1.

The above description completely characterizes our No Ratings experimental treatment and is close to the model of Banerjee (1992) and Bikhchandani et al. (1992).[3] In our Ratings treatment, each consumer who purchased the good also posts a public statement of whether or not they experienced positive *ex-post* utility from the good. The rating is not a strategic decision: all consumers who purchased the good are forced to

[2]Formally, symmetry means $F(-\theta_i) = 1 - F(\theta_i)$ for all $\theta_i \geq 0$. This implies $F(v) \geq 1/2$ since $v > 0$.

[3]A slight difference is that consumers here experience risk and random utility shocks when purchasing, but not when passing.

post a rating, and all ratings must be truthful. Consumers who passed post an 'empty' rating. Formally, consumer i posts a rating $y_i \in \{-1,0,1\}$, where

$$y_i = \begin{cases} -1 & \text{if} \quad x_i = 1 \text{ and } u_i(1|q,\theta_i) < 0 \\ 0 & \text{if} \quad x_i = 0 \\ 1 & \text{if} \quad x_i = 1 \text{ and } u_i(1|q,\theta_i) \geq 0 \end{cases}$$

Later consumers see the (truthful) ratings of all earlier consumers, so the information available to consumer i with ratings becomes $((x_1,y_1),(x_2,,y_2),\ldots,(x_{i-1},y_{i-1}),s_i)$.[4]

Suppose that the support of θ_i is contained in $[-v,v)$, so that $qv+\theta_i$ is positive if and only if $q=1$. In that case the rating of any one consumer who purchased would be perfectly informative about the quality of the good. To eliminate this trivial case, we assume $F(v)<1$. Now when the quality is low ($q=-1$) a consumer who purchases will leave a negative rating if $\theta_i \leq v$, which occurs with probability $F(v)$. And when the quality is high ($q=1$) a consumer who purchases will leave a positive rating if $\theta_i \geq -v$, which occurs with probability $1-F(-v)=F(v)$. Thus, ratings of those who purchase act as conditionally independent public signals of the good's true quality, and each of those signals is accurate with probability $F(v)$. The 'empty' ratings of those who don't purchase are completely uninformative.

Herds, Cascades, and Equilibrium Predictions

We carefully distinguish between herds and information cascades, following both Smith and Sorensen (2000) and Celen and Kariv (2004). A *herd* is said to occur when consecutive agents make identical purchasing decisions. An *information cascade* (or, simply, a *cascade*) is a herd that happens because the purchase decisions of the preceding agents lead to posterior beliefs under which herding behavior is optimal.[5] In other words, a cascade is a type of rational herd.[6]

Suppose we observe in our data all players from i to i' making the same purchasing/passing decision. We will call it a *purchasing herd* if $x_j=1$ for all $j \in \{i,\ldots,i'\}$, and a *passing herd* if $x_j=0$ for all such j. A herd is *correct* if it is either a purchasing herd in the high state or a passing herd in the low state; otherwise it is *incorrect*. If $i'=n$ we say the herd is *terminal*; otherwise it is *transitory*. These descriptors also apply to cascades.

---

[4]Technically, the inclusion of x_j for $j<i$ is superfluous since y_j is sufficient for x_j.

[5]These concepts were made distinct by Smith and Sorensen (2000); see also Celen and Kariv (2004, 2005).

[6]Other types of herds could be rational, too. Smith and Sorensen (2000) study herds and cascades as equilibrium phenomena, so all herds and cascades are rational. In an experimental setting, however, some herds may not be rational.

Consider either treatment and suppose players are in equilibrium. When consumer i is about to move, let b_i represent the public belief that the quality is high (meaning, the Bayesian posterior that $q = 1$, given $((x_1, y_1), \dots, (x_{i-1}, y_{i-1}))$). This combines with i's private signal to generate a posterior belief r_i. Specifically, the Bayesian posterior is

$$r_i = \frac{b_i p}{b_i p + (1-b_i)(1-p)} I_{(s_i=1)} + \frac{b_i(1-p)}{b_i(1-p) + (1-b_i)p} I_{(s_i=-1)},$$

where $I_{(\cdot)}$ is an indicator function. Since θ_i is mean-zero and unknown at the time of purchase, consumers choose $x_i = 1$ if and only if $r_i \geq 1/2$. But if $b_i < 1-p$ then r_i will be less than 1/2 regardless of s_i. In that case, consumer i will not purchase even if her signal is positive. Similarly, if $b_i > p$ then $r_i > 1/2$ regardless of s_i, so i will purchase even if her signal is negative. Once this occurs, consumer $i+1$ will know that i's purchase was not informative, meaning $b_{i+1} = b_i$. Consequently, $i+1$ also ignores his private signal and chooses $x_{i+1} = x_i$. All future consumers $j > i$ will also choose $x_j = x_i$ regardless of their signals, giving an information cascade. In the data this will appear as a terminal herd.

For $i > 1$ let $D_i = \sum_{j<i}(2x_j - 1)$ be the number of consumers prior to i who purchased (encoded as +1 in the sum) minus the number who passed (encoded as −1).[7] In the No Ratings treatment, an information cascade occurs in equilibrium starting at consumer i if and only if $|D_i| \geq 2$. In particular, if $D_i \geq 2$ then a purchasing cascade occurs, and if $D_i \leq -2$ then a passing cascade occurs. No further updating occurs once a cascade begins, so both types of cascades are terminal. Cascades here may or may not be correct, and an incorrect purchasing cascade is just as likely to occur as an incorrect passing cascade.

Now consider the Ratings treatment. In a purchasing cascade consumers continue to leave ratings, causing the public belief to update. If the purchasing cascade is incorrect (the state is bad), then with positive probability enough negative feedback will accumulate so that $b_i < p$, at which point the cascade breaks. Thus, purchasing cascades may not be terminal. Even if the purchasing cascade is correct it is possible that enough negative ratings appear to break the cascade. And when ratings are relatively noisy (as in our experiment), this becomes more common. In passing cascades, however, the public belief does not update, so all passing cascades are terminal.

Although ratings reduce incorrect purchasing cascades in the low state, they can actually make incorrect passing cascades *more* frequent in the high state. To see this, suppose $p = 1/2 + \varepsilon$ for some small ε (private signals are almost completely uninformative, as are previous purchase decisions), and $F(v) = 1/2 + \varepsilon + \delta$ for some small δ (ratings are only slightly more informative than private signals). Suppose the quality is high and

[7] Technically D_i is a function of the observed history; we omit this as an argument to simplify exposition.

the public belief rises above p (which is $1/2+\varepsilon$). The group will enter a correct purchasing cascade. But ratings will continue to cause the public belief to update. And these ratings are often inaccurate, so it is reasonably likely that a string of negative ratings will drop the public belief back below $1/2+\varepsilon$, breaking the correct cascade.

If the good is truly low quality then ratings are helpful: correct cascades (passing cascades) are terminal, while incorrect cascades (purchasing cascades) are likely to be transitory. But if the good is high quality then ratings are harmful: correct cascades (purchasing cascades) may be transitory, while incorrect cascades (purchasing cascades) are always terminal. One might suspect that ratings help prevent incorrect cascades from forming in the first place, but recall that ratings disproportionately affect beliefs so it only takes a few incorrect ratings early on to trigger an incorrect cascade. Incorrect cascades are relatively less likely, but the overall incidence of cascades increases significantly with ratings.

In our experiment we test this case where private signals are fairly uninformative, and ratings information is only slightly more informative. Based on the above discussion and the calculations from Yoon (2015), we derive the following testable hypotheses. We state them as hypotheses about herds, not cascades, since we cannot distinguish the two in our data.

Hypothesis 1. When quality is low there will be fewer incorrect herds with ratings than without.

Hypothesis 2. When quality is high there will be *more* incorrect herds with ratings than without.

Hypothesis 3. Without ratings all herds will be terminal.

Hypothesis 4. With ratings all passing herds will be terminal, but purchasing herds will often be transitory.

III. Experimental Design

Our experimental design was based on that of Goeree et al (2007), and computerized using zTree (Fischbacher, 2007). We recruited undergraduates from Ohio State University using ORSEE (Greiner, 2015). In the experiment, the state of the world was represented as a jar that could either be red (the "Red Jar", representing the good state) or blue (the "Blue Jar", representing the bad state). In each round the color of the jar was equally likely to be red or blue, but subjects were not shown the chosen color until the conclusion of the round. Thus, we called this the "Unknown Jar".

If the Unknown Jar is the Red Jar, it contains five red and four blue balls. If it is the Blue Jar, it contains five blue and four red balls. Each subject is privately shown a single draw from the Unknown Jar, with replacement. This private signal s_i is therefore correct with probability $p = 5/9$.

On their turn, each subject is asked to choose between two options: "Jars" or "Pass". Choosing Jars gives the subject a lottery that pays 20 points if the Unknown Jar is Red and −20 points if the Unknown Jar is Blue. Thus, the Jars option represents purchasing the good ($x_i = 1$), which has potential payoffs of $v = 20$ and $-v = -20$. Choosing Pass ($x_i = 0$) gives the subject a guaranteed payoff of zero.

In the model, the true value is augmented by an idiosyncratic shock θ_i. This is implemented in the experiment by adding a second car called the "Points Jar". It contains three balls labeled "+30", one ball labeled "0", and three balls labeled "−30". Subjects privately draw one ball from the Points Jar and the amount on the drawn ball is added (or subtracted) from their earnings for the round. The probability mass function of the idiosyncratic shock is therefore given by $f(-30) = 3/7$, $f(0) = 1/7$, and $f(30) = 3/7$.

Since $v = 20$ we have that $F(v) = 4/7$, which is slightly greater than $p = 5/9$, though both are quite close to 1/2 and therefore fairly uninformative. Our experiment corresponds to the example described in the theory section wherein $p = 1/2 + \varepsilon$ and $F(v) = 1/2 + \varepsilon + \delta$ for small ε and δ.

Participants are randomly and anonymously assigned to groups of eight or nine each period and randomly assigned an ordering in that group (from 1 to n). They make their purchase/pass decisions sequentially and are able to see whether each participant before them chose Jars or Pass.

Those in the ratings treatment play the exact same game but are also able to see the previous decision-makers' ratings. Ratings were not a strategic choice variable; a previous participant's rating is "Positive" if her total points from both the Unknown Jar and Points Jar is positive. Otherwise her rating is "Negative". Recall that ratings match the sign of the state variable with probability $F(v) = 4/7$. If a previous participant chooses Pass then no rating is provided.

We ran four sessions of the No Ratings treatment and six sessions of the Ratings treatment. In each session there were either ten or twenty rounds.[8] Groups, orderings within groups, and the color of the Unknown Jar were randomly redrawn each round. There were either 16 or 18 participants per session—depending on how many subjects showed up—giving two groups of either 8 or 9 in every round. Participants were paid \$15 plus \$0.10 per point earned.

[8]Originally, we planned to have sessions consist of ten rounds each. However, during the first session we found that participants were able to complete rounds more quickly than we anticipated, so we increased the number of rounds to twenty for the subsequent sessions.

IV. RESULTS

Empirical Definition of Herds and Cascades

For our data analysis, we define a *herd* as any string of sequential identical purchase/pass decisions. Thus, a new herd begins whenever anyone chooses a different purchase/pass decision than their immediate predecessor. This definition allows for herds of length one, and guarantees that all decisions are a part of some herd.

The proper definition of a *cascade* would require information on subjects' beliefs, which are not available in this experiment. Instead, we calculate for every consumer i what the public belief b_i would be if they were Bayesian, and define player i to be in a cascade if (1) his purchase/pass decision matches that of his predecessor ($x_i = x_{i-1}$), and (2) either $b_i < 4/9$ and $x_i = 0$ (a passing cascade), or $b_i > 5/9$ and $x_i = 1$ (a purchasing cascade). When calculating the public belief b_i, we apply the assumption that previous players $j < i$ follow their private signal if $b_j \in [4/9, 5/9]$, that players who follow an existing cascade reveal no new information ($b_j = b_{j-1}$), and that players who deviate from an existing cascade are revealing their private signal. This is identical to how the public belief updates in equilibrium, with the intuitive refinement that any (probability-zero) deviation from a cascade is assumed to come from a player who followed their private signal.

A cascade is then identified as a sequence of consecutive players who choose the same action and have all been identified as in a cascade by the above criteria. Whenever a player does not follow his immediate predecessor the cascade ends. With ratings a purchasing cascade can also end because of negative ratings, even when no player deviates in their purchase/pass decision. Notice that every cascade is a strict subset of some herd in our data; indeed, the first player i in a given herd cannot be part of a cascade since $x_i \neq x_{i-1}$. We cannot have cascades of length n, though we can have cascades of length one. According to our definition, if player i deviates from an existing cascade but players $i+1$ and $i+2$ return play to the original cascade behavior, a new herd is said to begin at $i+1$ and a new cascade begins at $i+2$ (assuming b_{i+2} is still less than 4/9 or greater than 5/9).[9]

	Purchasing	Passing
No Ratings	0.82	0.22
Ratings	1.03	0.40

TABLE I. Cascades per period, by treatment and by type.

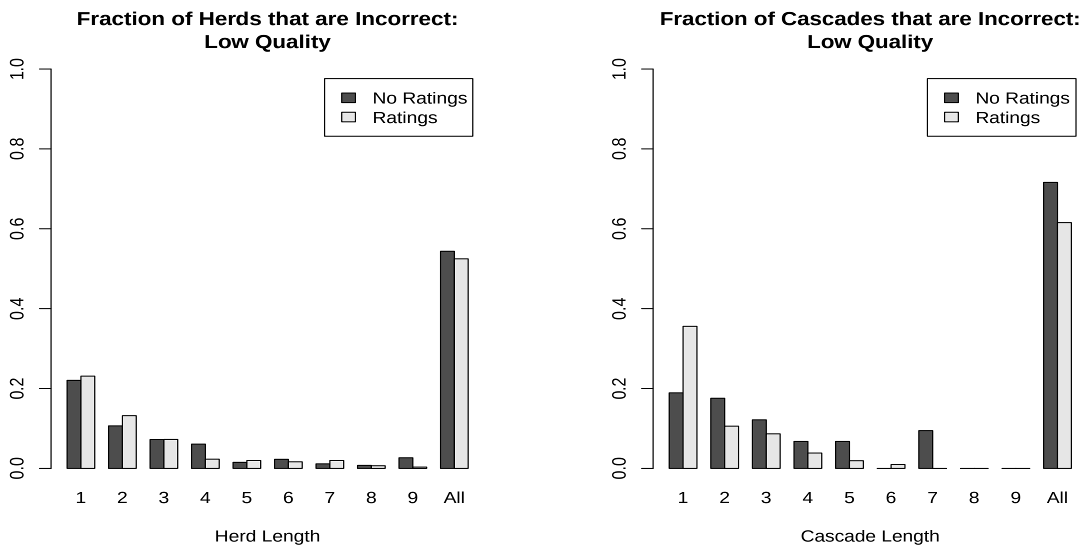

(A) (B)

FIGURE I. Fraction of (A) herds and (B) cascades that are incorrect when true quality is low.

Tests of the Main Hypotheses

To set the stage, we first measure the average number of cascades per period, broken down by treatment and by cascade type. The result is shown in Table I, which provides broad evidence for our three major findings: First, subjects purchase far more often than they pass. In the No Ratings treatment we should expect the frequencies of purchasing cascades and passing cascades to be roughly equal. Instead, purchasing cascades are nearly four times as common. Second, the presence of ratings reduces this bias. Purchasing cascades are still far more common, but now only 2.5 times as frequent as passing cascades. Finally, the presence of ratings breaks cascades, enabling more cascades to form per period.

Our first hypothesis from the theory is that, when quality is low, fewer (incorrect) purchasing herds will be observed in the ratings treatment because ratings will correct any such cascade. Figure I restricts attention to periods in which the true state was low

[9]Goeree et al. (2007) use a more stringent definition of cascades, requiring that they begin at the first i for which $b_i \in [0,4/9) \cup (5/9,1]$ and that they continue unbroken through player n. This is sensible since those authors do not study ratings; with ratings cascades can be broken (and reformed) in equilibrium. This difference inspires our definition.

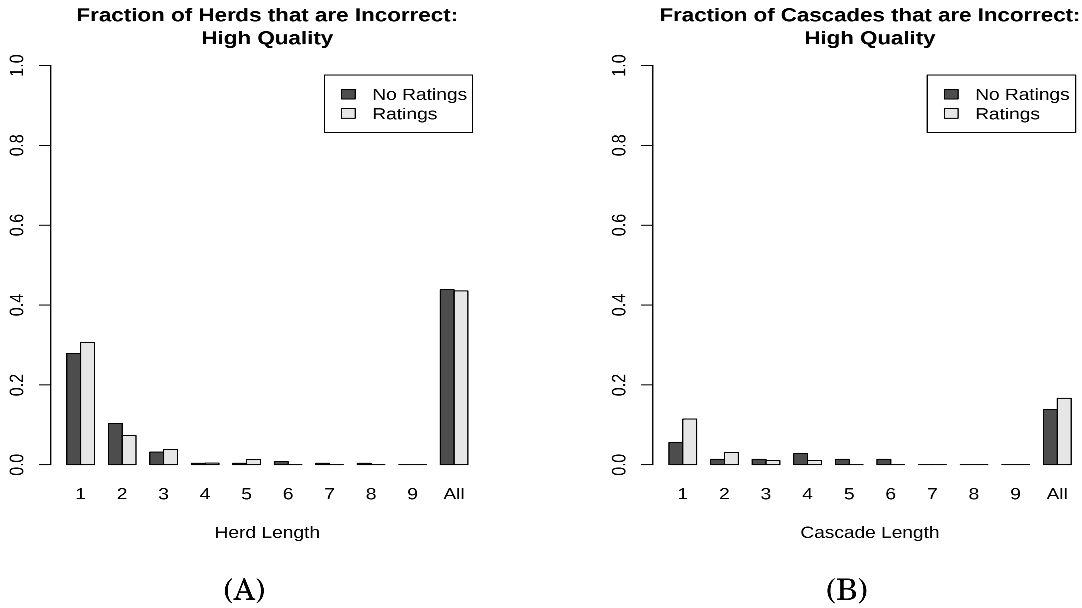

FIGURE II. Fraction of (A) herds and (B) cascades that are incorrect when true quality is high.

and shows the fraction of herds (panel A) or cascades (panel B) that were incorrect. The first nine bins show the frequency by herd or cascade length, while the tenth is the sum over all lengths. The overall difference does move in the direction of the hypothesis, but is not statistically significant either for herds (Fisher test p-value 0.673) or for cascades (p-value 0.200). For cascades this lack of significance comes because the frequency of length-one cascades is substantially higher with ratings. These length-one cascades occur because the public belief moves in larger increments with ratings—when moving either up or down—so there are more opportunities for the public belief b_i to rise above 5/9 with ratings. In that scenario, if subject i purchases but leaves a negative rating then $i+1$ is likely not to purchase, giving a cascade of length one. With ratings there is no negative public signal so subject $i+1$ is much more likely to purchase as well, extending the cascade.

A particularly surprising result is that the total fraction of incorrect herds and cascades (the last set of bars) is above one half for either treatment: If purchase decisions and ratings convey any information, incorrect herds should be less common than correct herds. The failure arises because subjects systematically purchase (or, choose "Jars") too frequently. When the true state is low, this purchasing bias leads to incorrect herds and cascades.

Our second hypothesis is that, counter-intuitively, ratings will *increase* the number of (incorrect) passing herds when the quality is high. In Figure II we see that these herds are not common—again because subjects purchase so often—and that differences

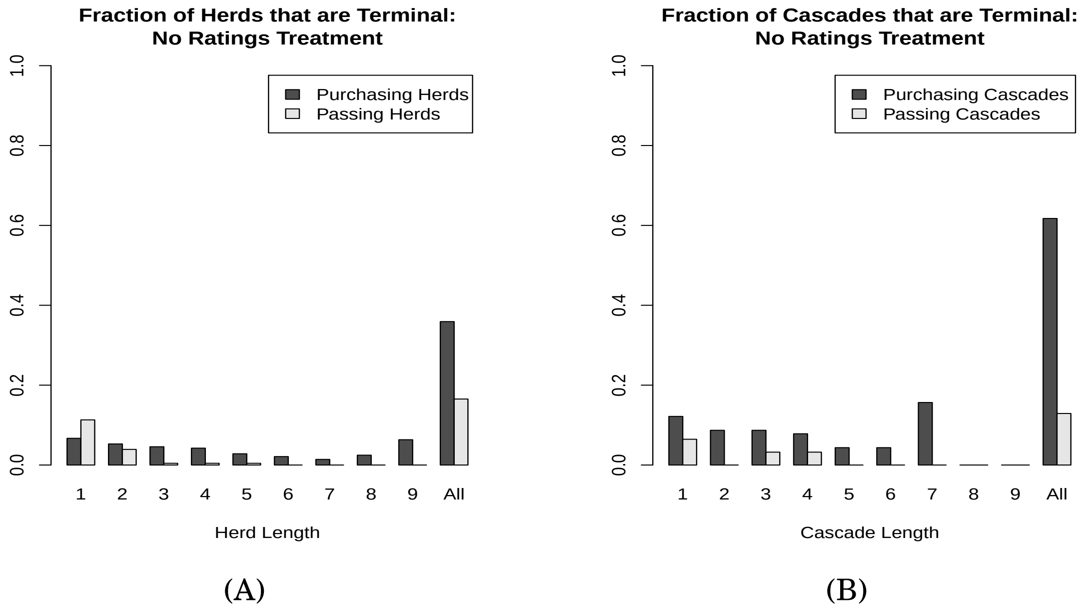

FIGURE III. Fraction of (A) herds and (B) cascades that are terminal in the treatment without ratings.

between treatments are small or non-existent. The Fisher test p-value for a treatment difference is 1.000 for herds and 0.671 for cascades. We also note that almost all of these incorrect passing herds (and cascades) are short-lived, rarely lasting more than three decisions. Again, this is evidence that subjects have a strong bias towards purchasing.

Our final two hypotheses focus on whether cascades are transitory or terminal. First, consider the No Ratings treatment. In theory, all cascades should be terminal without ratings because no new information is provided to break a cascade once it begins. To the contrary, panel (A) of Figure III reveals that only a small fraction of herds are terminal. This is in line with previous experimental results that suggest players choose stochastically, breaking cascades (Kubler and Weizsacker, 2004; Goeree et al., 2007, e.g.). Purchasing herds do tend to be terminal more often than passing herds (Fisher p-value < 0.001), again because subjects have a propensity to purchase. We obtain a similar result for cascades (p-value < 0.001), though purchasing cascades tend to be even longer-lived and are more often terminal than not. Overall, we clearly reject the hypothesis that cascades in the No Ratings treatment are all terminal.

In the Ratings treatment we hypothesize that passing cascades will be be terminal, but purchasing cascades will often be broken via ratings data. In Figure IV we see that the prediction is also not true. All types of herds and cascades are terminal less than 40% of the time. In fact, passing herds are terminal *less* often than purchasing herds (Fisher p-value 0.002). This is again explained by the purchasing bias, causing subjects to break out of passing herds but conform with purchasing herds.

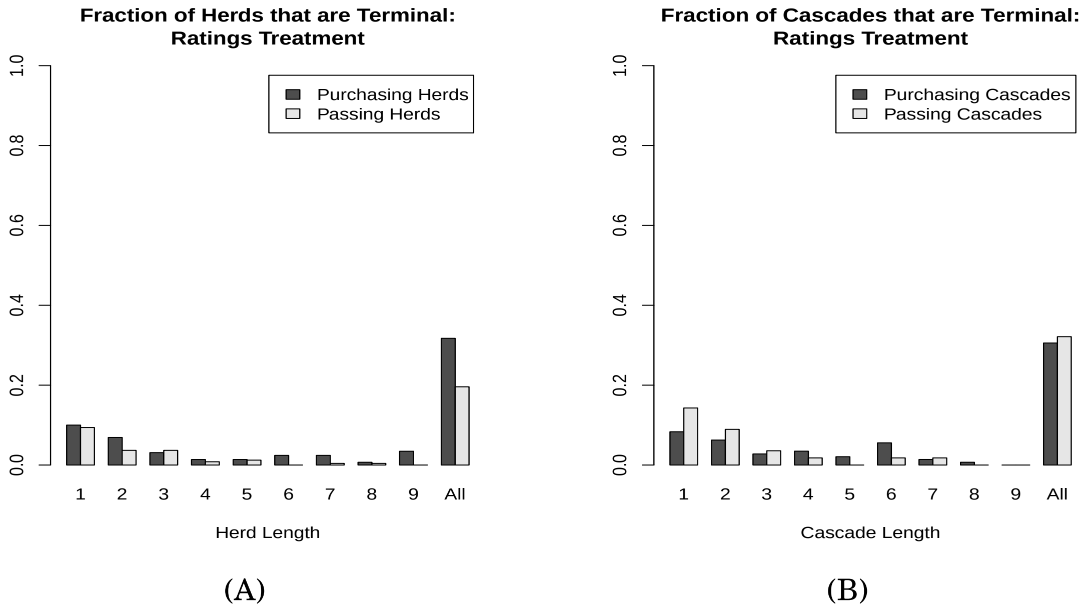

FIGURE IV. Fraction of (A) herds and (B) cascades that are terminal in the treatment with ratings.

Comparing treatments, we find no difference in the fraction of purchasing herds that are terminal (Fisher p-value 0.291) and no difference in the fraction of passing herds that are terminal (p-value 0.406). With cascades, however, we do see a difference. In the Ratings treatment there are significantly fewer purchasing cascades that are terminal than in the No Ratings treatment (p-value < 0.001). There are actually more terminal passing cascades in the Ratings treatment than in the No Ratings treatment, but the significance of this difference is marginal (p-value 0.071). The reduction in terminal purchasing cascades is exactly in line with theory, though the (insignificant) increase in terminal passing cascades is not.

We can expand on this last result by examining the distribution of purchasing and passing herd lengths for each treatment. This is shown (for herds and cascades) in Figure V. When looking at herds (panels A and C), the difference between treatments is not substantial. The average length of purchasing herds drops from 2.93 to 2.71, and the average length of passing herds drops from 1.69 to 1.64. Neither of these distributions is significantly different (Mann-Whitney p-values 0.332 and 0.861, respectively), so we conclude that the presence of ratings does not have a large impact on how frequently players mimic the actions of their immediate predecessor. But if we focus on cascades—which are a type of rational herd—we do see a treatment difference. Purchasing cascades drop in average length from 3.22 decisions without ratings to 2.24 decisions with ratings. In fact, the distribution of cascade lengths in the Ratings treatment is stochastically dominated by the distribution in the No Ratings treatment, if we exclude the

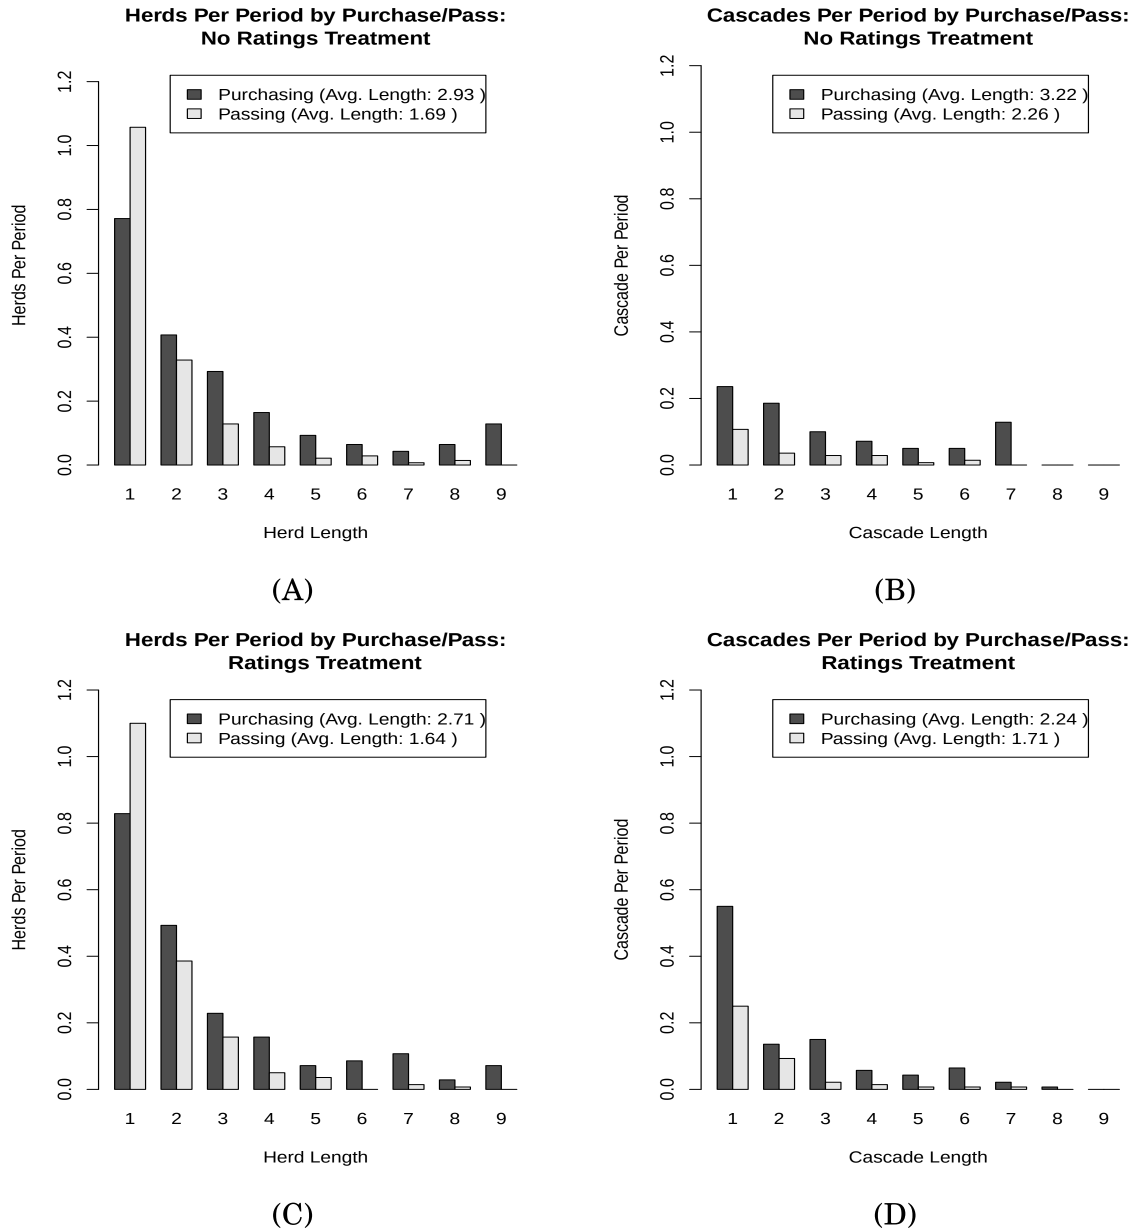

FIGURE V. Lengths of purchasing and passing herds and cascades, by treatment.

one cascade of length 8 in the Ratings treatment. A Mann-Whitney test confirms that these distributions are significantly different (p-value < 0.001). Passing cascade lengths drop from 2.26 on average to 1.71 on average—and the distributions have a stochastic dominance relationship if we exclude the one passing cascade of length seven—but the difference in the distributions is only marginally significant (Mann-Whitney p-value 0.092). Again, these results are qualitatively in line with the theoretical prediction: negative ratings eventually break long purchasing cascades, but do not have a large impact on passing cascades.

In the next subsection we turn to individual-level analysis. We will find that the reduction in purchasing cascades in the ratings treatment is largely driven by a reduction in the purchasing bias, apparently due to subjects focusing more on public information when ratings are introduced. This suggests that the theory's prediction was right, but for the wrong reason.

An Econometric Model of Individual Behavior

To study individual decisions we estimate an econometric model whose parameter estimates will allow us to identify the relative importance of various biases in subjects' decision-making. Its construction is similar to the models of Goeree and Holt (2004), Kubler and Weizsacker (2004), and Goeree et al. (2007), among others.

First, it has been observed that subjects often overweight their own private information. So we allow for non-Bayesian updating that incorporates this bias. Specifically, we assume that the posterior belief is given by

$$r_i = \frac{b_i p^\alpha}{b_i p^\alpha + (1-b_i)(1-p)^\alpha} I_{(s_i=1)} + \frac{b_i(1-p)^\alpha}{b_i(1-p)^\alpha + (1-b_i)p^\alpha} I_{(s_i=-1)},$$

where $\alpha > 0$ is a parameter that represents the weight on private information. If $\alpha = 1$ then we have the Bayesian formula. An agent with $\alpha > 1$ overweights private information, while an agent with $\alpha < 1$ underweights private information. Algebraic manipulation of this updating rule yields a convenient linear form in terms of log odds ratios, given by

$$\log\left(\frac{r_i}{1-r_i}\right) = \log\left(\frac{b_i}{1-b_i}\right) + \alpha\left[\log\left(\frac{p}{1-p}\right) I_{(s_i=1)} + \log\left(\frac{1-p}{p}\right) I_{(s_i=-1)}\right].$$

Second, we allow for a natural bias toward one decision or another. This is done simply by augmenting the utility function. Specifically, we alter the *ex-post* utility function of agent i to $u_i(x_i|q,\theta_i) = (qv + \theta_i + \beta)x_i$, where β represents a bias toward purchasing behavior.[10] If $\beta > 0$ then the agent is biased toward purchasing, and if $\beta < 0$ then the agent is biased toward passing. Our aggregate results suggest excessive purchasing, so we expect our estimate of β will be positive.

Finally, we allow for stochastic choice. Given beliefs r_i (which may be biased by α) and an *ex-post* utility function $u_i(x_i|q,\theta_i)$ (which may be biased by β), consumer i calculates the expected utility of purchasing and passing, which we denote here simply by EU_1 and EU_0, respectively. He then purchases the good with probability

$$Pr(x_i = 1|EU_1, EU_0, \lambda) = \frac{\exp(\lambda EU_1)}{\exp(\lambda EU_1) + \exp(\lambda EU_0)},$$

[10]Of course, the ratings players leave are still determined by their un-augmented *ex-post* utility.

where $\lambda \geq 0$ is a noise parameter that captures the subject's sensitivity to payoffs. The model predicts perfect best response as $\lambda \to \infty$, and uniform random choice when $\lambda = 0$. The maximum likelihood estimate of λ is sensitive to the scaling of payoffs; for the purposes of this estimation we normalize v to one.[11] With this normalization a risk-neutral subject with private belief r_i and purchasing bias β will prefer to purchase if $r_i - (1 - r_i) + \beta \geq 0$, or $r_i + \beta/2 \geq 1/2$. Thus, we can also interpret a bias of β as being equivalent to a $\beta/2$ increase in the subject's private belief about the true state being positive. It could also be interpreted as a proxy for risk preferences, with $\beta > 0$ indicating risk-loving behavior. We are forced to be agnostic about these varied interpretations as our experiment was not designed to tease apart their differences.

We assume all subjects employ the same parameter values and that this is common knowledge. The public belief is therefore adjusted accordingly. For example, a lower value of λ would imply that previous purchase decisions are less informative, causing the current player to weight them less in their own belief updating. If $\alpha = 0$ then his updating adjustment is fully Bayesian, given the noise in others' behavior. If $\alpha > 0$ then the subject puts even less weight on past play because of his own non-Bayesian bias.

Previous ratings data are truthful, so λ does not affect how ratings enter into the updated public belief. But α does. Summarizing, λ affects the randomness in choice and also shifts the weight on previous purchasing data, while α shifts the weights on both previous purchasing data and ratings data.

The purchasing bias β also alters the subjects' belief updating process. If previous subjects are prone to purchase ($\beta > 0$) then purchasing decisions are less informative than passing decisions. The reverse is true if they are prone to passing ($\beta < 0$). But β does not affect how previous ratings data are incorporated. Only α affects the weight placed on ratings data.

Our model is parameterized by the triple (α, β, λ). The unbiased equilibrium is a special case of our simple model where $\alpha = 1$, $\beta = 0$, and $\lambda \to \infty$. We estimate the parameters of our model via a grid search, finding the parameter vector that maximizes the log likelihood of our data. We do this for each treatment separately, and for both treatments combined. We also look at how the fit of the model is affected by 'shutting down' various parameters, to understand the relative importance of each in predicting actual behavior.

[11]Since the distribution of θ_i is symmetric its expectation is zero regardless of its scaling, so there is no need to normalize its value.

Econometric Model Parameter Estimates

The maximum likelihood estimates are shown in Table II. Recall that we ran more sessions of the Ratings treatment, so the log likelihood values are not directly comparable across treatments.

Specification 1 estimates the full model. First, estimates of β range from 0.10 to 0.16. This is a fairly strong purchasing bias, equivalent to a 5 to 8 percentage point increase in the subjects' private beliefs about the true state, regardless of their observed history.

Our estimated value of λ is roughly in line with Goeree et al. (2007), but our α is lower. This difference may stem from the purchasing bias described above: Rational subjects should place low weight on previous purchasing decisions in our experiment because subjects purchase far too often. If they fail to make this adjustment, they will appear to be putting too much weight on public information (and too little on private signals), causing α to decrease. When ratings are present α even drops below one. But notice that when we restrict $\alpha = 1$ (in specification 2) the estimates of β and the log likelihoods barely change. Thus, the full model's predictions are fairly insensitive to changes in α, and we cannot conclusively argue that subjects are non-Bayesian in our experiment.

The role of β appears more important: restricting β to zero (specification 3) noticeably alters α and decreases the predictive power of the model. This suggests that the purchasing bias is a much more robust and fundamental phenomenon in our data.

Comparing treatments (in the full model), we find that the presence of ratings reduces the purchasing bias and shifts subjects from overweighting private information ($\alpha > 1$) to overweighting public information ($\alpha < 1$). The reduction in the purchasing bias should reduce the overall incidence of purchasing cascades, though we do not expect the shift in α to have a large effect given its relatively unimportant role in prediction. This prediction is confirmed by the earlier results: We see somewhat fewer incorrect purchasing cascades with ratings (panel B of Figure I), and many fewer terminal purchasing cascades with ratings (comparing panel B of Figures III and IV).

V. DISCUSSION

Though our experiment was designed to study the counterintuitive predictions of Yoon (2015), our primary result is instead that subjects in this experiment had a strong propensity to purchase, even when public information suggested the good was of low quality. One might posit that subjects prefer to purchase out of altruism, because purchasing the good generates a rating which is useful for future consumers. But we see an even stronger propensity to purchase when no ratings are present, so this explanation can only explain part of the observed bias.

Specification		Both	No Ratings	Ratings
		Treatment		
1	α	0.87	1.25	0.68
	β	0.12	0.16	0.10
	λ	7.10	5.85	7.82
	Log L	-1901.2	-655.3	-1243.2
2	α	==Restricted to 1==		
	β	0.13	0.13	0.13
	λ	6.45	6.88	5.96
	Log L	-1901.3	-655.7	-1244.3
3	α	0.47	0.59	0.34
	β	==Restricted to 0==		
	λ	11.83	10.77	13.72
	Log L	-1973.4	-697.2	-1274.1
4	α	==Restricted to 1==		
	β	==Restricted to 0==		
	λ	7.16	7.63	6.63
	Log L	-1988.2	-701.1	-1285.7

TABLE II. Maximum likelihood estimates of the econometric model parameters, by treatment.

The apparent bias toward purchasing in the No Ratings treatment could be due to loss aversion. Subjects may enter the experiment with an expectation of positive payoffs, and therefore view a guaranteed payoff of zero (given by not purchasing) to be unappealing. Purchasing exposes the subject to the possibility of an even larger loss, but perhaps risk-seeking preferences in the domain of losses will lean subjects toward purchasing behavior. In other words, a prospect theory explanation may apply here, but would require fairly specific parameter values to capture this bias.

Recall, however, that all subjects received \$15 in addition to their earnings from the task. So any loss aversion explanation would require that the reference point for losses either be quite high (relative to standard experiment payments), or be independent of the fixed \$15 payment. We believe the latter is more plausible, as the experiment earnings are focal within the interface, and these can be negative.

An alternative explanation is that subjects simply find the purchase option to be more 'interesting' or 'active' than the pass option. This sort of activity bias has been documented in past studies as well, including Patt and Zeckhauser (2000); Lei et al. (2001); Bar-Eli et al. (2007); Carrillo and Palfrey (2011); and Atiker et al. (2011).

Our current design was not created to test these various explanations since we did not anticipate this result, so such an investigation remains an open question for future research.

REFERENCES

Anderson, L. R., Holt, C. A., 1997. Information cascades in the laboratory. American Economic Review 87, 847–862.

Atiker, E., Neilsen, W. S., Price, M. K., 2011. Activity bias and focal points in extensive games, University of Tennessee working paper.

Banerjee, A. V., 1992. A simple model of herd behavior. Quarterly Journal of Economics 107, 797–817.

Bar-Eli, M., Azar, O. H., Ritov, I., Keidar-Levin, Y., Schein, G., 2007. Action bias among elite soccer goalkeepers: The case of penalty kicks. Journal of Economic Psychology 28, 606–621.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100, 992–1026.

Carrillo, J. D., Palfrey, T. R., 2011. No trade. Games and Economic Behavior 71, 66–87.

Celen, B., Kariv, S., 2004. Distinguishing informational cascades from herd behavior in the laboratory. American Economic Review, 484–498.

Celen, B., Kariv, S., 2005. An experimental test of observational learning under imperfect information. Economic Theory 26 (3), 677–699.

Celen, B., Kariv, S., Schotter, A., 2010. An experimental test of advice and social learning. Management Science 56 (10), 1687–1701.

Cipriani, M., Guarino, A., 2008. Transaction costs and informational cascades in financial markets. Journal of Economic Behavior & Organization 68 (3), 581–592.

Corazzini, L., Greiner, B., 2007. Herding, social preferences and (non-) conformity. Economics Letters, 74–80.

Dominitz, J., Hung, A. A., 2009. Empirical models of discrete choice and belief updating in observational learning experiments. Journal of Economic Behavior & Organization 69 (2), 94–109.

Drehmann, M., Oechssler, J., Roider, A., 2007. Herding with and without payoff externalitiesâĂŤan internet experiment. International Journal of Industrial Organization 25 (2), 391–415.

Feltovich, N., 2002. Information cascades with endogenous signal precision. Working paper.

Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Experimental Economics 10 (2), 171–178.

Goeree, J. K., Holt, C. A., 2004. A model of noisy introspection. Games and Economic Behavior 46 (2), 356–382.

Goeree, J. K., Palfrey, T. R., Rogers, B. W., McKelvey, R. D., 2007. Self-correcting information cascades. The Review of Economic Studies 74 (3), 733–762.

Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with orsee. Journal of the Economic Science Association 1 (1), 114–125.

Grinblatt, M., Keloharju, M., Ikaheimo, S., November 2008. Social influence and consumption: Evidence from the automobile purchases of neighbors. The Review of Economics and Statistics 90 (4), 735–753.

Huck, S., Oechssler, J., 2000. Informational cascades in the laboratory: Do they occur for the right reasons? Journal of Economic Psycholog 21 (6), 661–671.

Hung, A. A., Plott, C. R., 2001. Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. The American economic review 91 (5), 1508–1520.

Kraemer, C., Noth, M., Weber, M., 2006. Information aggregation with costly information and random ordering: Experimental evidence. Journal of Economic Behavior & Organization 59 (3), 423–432.

Kubler, D., Weizsacker, G., 2003. Information cascades in the labor market. Journal of Economics 80 (3), 211–229.

Kubler, D., Weizsacker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. The Review of Economic Studies 71 (2), 425–441.

Lei, V., Noussair, C. N., Plott, C. R., 2001. Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. Econometrica 69, 831–859.

Mobius, M., Niederle, M., Niehaus, P., Rosenblat, T., 2014. Managing self confidence, University of Michigan Working Paper.

Morone, A., Fiore, A., Sandri, S., 2007. On the absorbability of herd behaviour and informational cascades: an experimental analysis. Dresden discussion paper in economics.

Noth, M., Weber, M., 2003. Information aggregation with random ordering: Cascades and overconfidence. The Economic Journal 113 (484), 166–189.

O'Guinn, T. C., 2015. Bricks and mortar: offline shopping in online america. Consumption Markets & Culture 18, 468–473.

Patt, A., Zeckhauser, R., 2000. Action bias and environmental decisions. Journal of Risk and Uncertainty 21, 45–72.

Shachat, J., Srinivasan, A., 2011. Informational price cascades and non-aggregation of asymmetric information in experimental asset markets. Working paper.

Smith, L., Sorensen, P., 2000. Pathological outcomes of observational learning. Econometrica 68 (2), 371–398.

Yoon, Y., 2015. Social learning with a ratings system, Ohio State University working paper.

Online Appendix

Not intended for publication

Appendix A. Transitions of Public Beliefs

In this section we calculate the transition of public beliefs, according to Bayes's rule. These calculations are adapted from Yoon (2015).

No Ratings Treatment: General Case

(1) When $b_i \in (1-p, p)$, the public belief is intermediate and the agent i's choice follows her own private signal s_i. In this case, her choice fully delivers her signal information to the later agents, i.e., $x_i = s_i$.

$$b_{i+1} = \begin{cases} \dfrac{b_i p}{b_i p + (1-b_i)(1-p)} & \text{if } x_i = 1 \\ \dfrac{b_i(1-p)}{b_i(1-p)+(1-b_i)p} & \text{if } x_i = 0 \end{cases}$$

(2) When $b_i = 1-p$, the agent i chooses $x_i = 0$ if her private signal is also $s_i = 0$. If $s_i = 1$, however, she is now indifferent between $x_i = 1$ and $x_i = 0$. In this case, the agent i chooses $x_i = 0$ with probability 1/2 even when $s_i = 1$. Hence, if $x_i = 1$, it simply implies that $s_i = 1$, but $x_i = 0$ does not always mean that $s_i = 0$.

$$b_{i+1} = \begin{cases} \dfrac{b_i p}{b_i p + (1-b_i)(1-p)} & \text{if } x_i = 1 \\ \dfrac{b_i(2-p)}{b_i(2-p)+(1-b_i)(1+p)} & \text{if } x_i = 0 \end{cases}$$

(3) When $b_i = p$, similarly, $x_i = 1$ does not fully deliver the agent i's private signal information while $x_i = 0$ implies $s_i = 0$.

$$b_{i+1} = \begin{cases} \dfrac{b_i(1+p)}{b_i(1+p)+(1-b_i)(2-p)} & \text{if } x_i = 1 \\ \dfrac{b_i(1-p)}{b_i(1-p)+(1-b_i)p} & \text{if } x_i = 0 \end{cases}$$

(4) When $b_i \in [0, 1-p) \cup (p, 1]$, the agent i simply follows the majority and ignore her private signal, s_i. Thus her action does not deliver any information.

$$b_{n+1} = b_n \text{ for } x_i \in \{0, 1\}.$$

Ratings Treatment: General Case

(1) When $b_i \in (1-p, p)$, the agent i's choice follows her private signal and she also gives ratings when she chooses $x_i = 1$.

$$b_{i+1} = \begin{cases} \dfrac{b_i p F(v)}{b_i p F(v) + (1-b_i)(1-p)F(-v)} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{b_i p F(-v)}{b_i p F(-v) + (1-b_i)(1-p)F(v)} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{b_i(1-p)}{b_i(1-p) + (1-b_i)p} & \text{if } x_i = 0 \end{cases}$$

(2) When $b_i = 1-p$, $x_i = 0$ does not fully deliver her private information while $x_i = 1$ implies $s_i = 1$ and the agent i provides rating informationo conditional on $s_i = 1$.

$$b_{i+1} = \begin{cases} \dfrac{b_i p F(v)}{b_i p F(v) + (1-b_i)(1-p)F(-v)} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{b_i p F(-v)}{b_i p F(-v) + (1-b_i)(1-p)F(v)} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{b_i(2-p)}{b_i(2-p) + (1-b_i)(1+p)} & \text{if } x_i = 0 \end{cases}$$

(3) When $b_i = p$, $x_i = 1$ only partially deliver the agent i's signal information.

$$b_{i+1} = \begin{cases} \dfrac{b_i(1+p)F(v)}{b_i(1+p)F(v) + (1-b_i)(2-p)F(-v)} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{b_i(1+p)F(-v)}{b_i(1+p)F(-v) + (1-b_i)(2-p)F(v)} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{b_i(1-p)}{b_i(1-p) + (1-b_i)p} & \text{if } x_i = 0 \end{cases}$$

(4) When $b_i \in [0, 1-p)$, agent i always choose $x_i = 0$ and no information is not provided to the later agents including ratings information.

$$b_{i+1} = b_i \text{ for } x_i \in \{0, 1\}$$

(5) When $b_i \in (p, 1]$, agent i always choose $x_i = 1$ and her choice does not deliver her private signal information. In this case, however, she gives a rating which deliver her ex-post payoff deliver information.

$$b_{i+1} = \begin{cases} \dfrac{b_i F(v)}{b_i F(v) + (1-b_i)F(-v)} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{b_i F(-v)}{b_i F(-v) + (1-b_i)F(v)} & \text{if } x_i = 1 \text{ \& } y_i = -1 \end{cases}$$

No Ratings Treatment: Experimental Parameters

In our experimental setting, we assume $b_0 = 1/2$, $p = 5/9$ and $F(v) = 4/7$. Therefore, the dynamics of the public beliefs are the following:

(1) When $b_i \in (\frac{4}{9}, \frac{5}{9})$,

$$b_{i+1} = \begin{cases} \dfrac{5b_i}{b_i + 4} & \text{if } x_i = 1 \\ \dfrac{4b_i}{5 - b_i} & \text{if } x_i = 0 \end{cases}$$

(2) When $b_i = 4/9$,

$$b_{i+1} = \begin{cases} \dfrac{5b_i}{b_i+4} = \dfrac{1}{2} & \text{if } x_i = 1 \\ \dfrac{13b_i}{14-b_i} = \dfrac{26}{61} & \text{if } x_i = 0 \end{cases}$$

(3) When $b_i = 5/9$,

$$b_{i+1} = \begin{cases} \dfrac{14b_i}{13+b_i} = \dfrac{35}{61} & \text{if } x_i = 1 \\ \dfrac{4b_i}{5-b_i} = \dfrac{1}{2} & \text{if } x_i = 0 \end{cases}$$

(4) When $b_i \in [0, \frac{4}{9}) \cup (\frac{5}{9}, 1]$,

$$b_{n+1} = b_n \text{ for } x_i \in \{0, 1\}.$$

Ratings Treatment: Experimental Parameters

(1) When $b_i \in (\frac{4}{9}, \frac{5}{9})$,

$$b_{i+1} = \begin{cases} \dfrac{5b_i}{3+2b_i} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{15b_i}{16-b_i} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{4b_i}{5-b_i} & \text{if } x_i = 0 \end{cases}$$

(2) When $b_i = 4/9$, .

$$b_{i+1} = \begin{cases} \dfrac{5b_i}{3+2b_i} = \dfrac{4}{7} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{15b_i}{16-b_i} = \dfrac{3}{7} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{13b_i}{14-b_i} = \dfrac{26}{61} & \text{if } x_i = 0 \end{cases}$$

(3) When $b_i = 5/9$,

$$b_{i+1} = \begin{cases} \dfrac{56b_i}{39+17b_i} = \dfrac{70}{109} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{42b_i}{52-10b_i} = \dfrac{105}{209} & \text{if } x_i = 1 \text{ \& } y_i = -1 \\ \dfrac{4b_i}{5-b_i} = \dfrac{1}{2} & \text{if } x_i = 0 \end{cases}$$

(4) When $b_i \in [0, \frac{4}{9})$,

$$b_{i+1} = b_i \text{ for } x_i \in \{0, 1\}$$

(5) When $b_i \in (\frac{5}{9}, 1]$,

$$b_{i+1} = \begin{cases} \dfrac{4b_i}{3+b_i} & \text{if } x_i = 1 \text{ \& } y_i = 1 \\ \dfrac{3b_i}{4-b_i} & \text{if } x_i = 1 \text{ \& } y_i = -1 \end{cases}$$

Appendix B. Experimental Instructions: No Ratings

This is an experiment in the economics of decision-making. If you listen carefully and make good decisions, you could earn a significant amount of money that will be paid to you in cash at the end of the experiment.

In this experiment, there is a jar called the "Unknown Jar" that is filled with 9 balls that are either red or blue. The jar is either the "red jar", meaning it contains 5 red balls and 4 blue balls, or the "blue jar", which contains 5 blue balls and 4 red balls. Thus, the "red jar" has slightly more red balls, and the "blue jar" has slightly more blue balls.

At the beginning of each round in the experiment, the computer will flip a coin to determine which jar to use – either the red jar or the blue jar. But you won't know which one is chosen. In the computer interface, you'll be shown a jar with 9 balls of unknown color (Figure 1). To give you a clue about which jar it is, you can then click on any one ball and its color will be revealed to you (Figure 2). The other 8 balls will not be revealed.

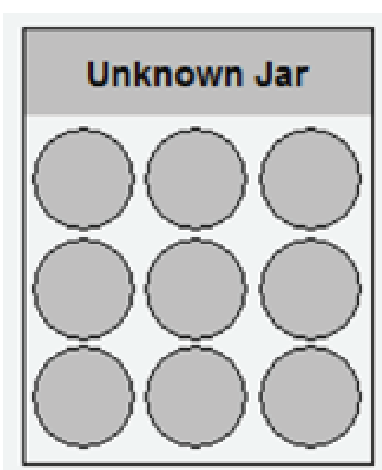

Figure 1: 9 balls of unknown color.

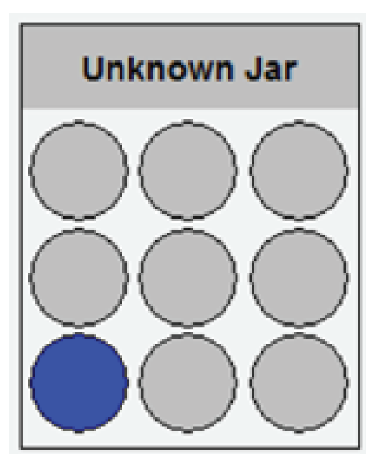

Figure 2: A blue ball was revealed.

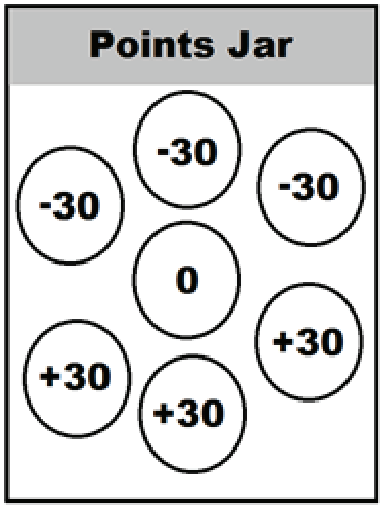

Figure 3: The "Points Jar".

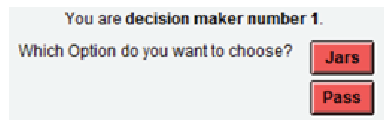

Figure 4: Choose "Jars" or "Pass"

You will then choose one ball from the "Points Jar". As shown in figure 3, the Points Jar contains 3 balls labeled "+30", 1 ball labeled "0", and 3 balls labeled "-30". The value of the "Points Ball" you drew is not yet revealed to you.

After drawing a ball from each jar, your job is to choose whether you want to get paid based on which jar the Unknown Jar is, or whether you prefer to "pass". You do this in the computer by clicking the button labeled "Jars" or the button labeled "Pass" (Figure 4).

If you choose "Jars", you receive 20 points if the unknown jar is the RED jar, and you lose 20 points if it is the BLUE jar. In addition, you will earn (or lose) the number of points shown on your Points Ball.

For example, suppose you choose "Jars". If the unknown jar is red and the Points Ball you drew is "+30", you earn 20 + 30 = 50 points. If the unknown jar is blue and your Points Ball is "0", then you earn -20 + 0 = -20 points. If we count all possible combinations, your total point earnings could be -50, -20, or 10 if the Unknown Jar is blue, and -10, 20, or 50 if the Unknown Jar is red.

If you choose to "Pass", then you earn zero points regardless of whether the Unknown Jar is red or blue, and the ball you drew from the Points Jar will not affect the number of points you earn. In summary, your choices are:

JARS: Receive 20 points if the unknown jar is RED. Lose 20 points if the unknown jar is BLUE. In addition, you earn (or lose) the number of points shown on the ball you drew from the Points Jar.

PASS: Receive 0 points.

You will be in a group of 9 subjects all making the same decision about the same jar. The other members of your group also get to choose one ball from both the Unknown Jar and the Points Jar and then choose Jars or Pass. The balls they draw may be the same as yours or different, depending on which they clicked. Thus, each of you may have different information about the Unknown Jar.

You will take turns making your decision. Each person in your group is assigned a "decision maker number", which tells the order in which you will make your decision. Everyone chooses a ball from each jar at the same time. Then, decision maker #1 chooses between Jars or Pass, followed by decision maker #2, and so on.

Decision Maker	Choice
1	Jars
2	Pass
3	Jars
4	Jars
5	Pass

Figure 5: The choices of decision makers 1 through 5.

Each person gets to see the history of what the earlier decision makers in their group chose. For example, decision maker #4 gets to see what decision makers #1, 2, and 3 chose. An example history that decision maker #6 might see is shown in Figure 5. This shows that decision maker #1 chose "Jars", decision maker #2 chose "Pass", and so on.

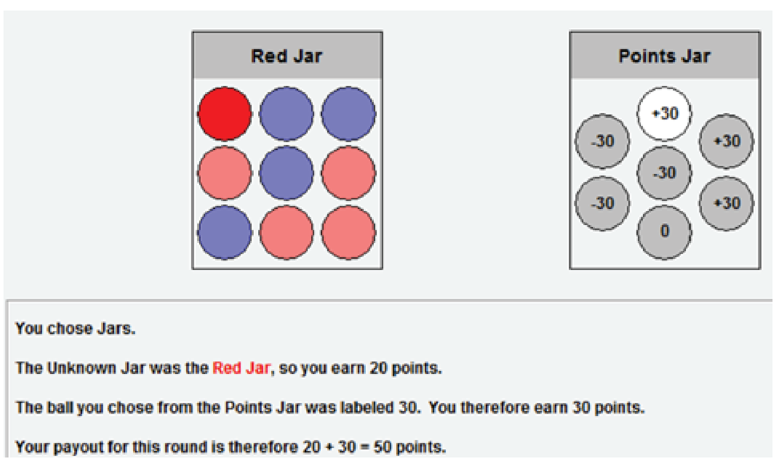

Figure 6: The jar is revealed and your earnings are calculated.

Once everyone in your group has made their choices, you will see whether the Unknown Jar is red or blue. If you chose "Jars", you will also see the contents of the Points Jar and the value of the ball you drew. Finally, you will see your earnings in points. For example, in Figure 6, the unknown jar was the Red jar. The player chose "Jars", so they earned +20 points. Their Points ball was "+30", so their total payoff is 20+30 = 50 points.

This completes one round of the experiment. You will play a total of 10 rounds. In each round everyone will be randomly assigned a new group, and will get a new decision maker number within that group. The computer will flip a coin again to decide whether the Unknown Jar will be the red jar or the blue jar. It will also randomly shuffle the contents of the Unknown Jar and the Points Jar.

At the end of the experiment, we will randomly pick one of the 10 rounds for actual payment. You will get paid $15 plus 10 cents for every point earned in that one chosen round (or, minus 10 cents for every point lost if your earnings for that round were negative). Since the maximum you can earn in any one round is 20+30=50 points, your maximum possible payment for the experiment is $15 + (50 × $0.10) = $20.00 The minimum you can earn in any round is -20-30 = -50 points, so your minimum possible payment for the experiment is $15 + (-50 × $0.10) = $10.00. (Your show-up fee of $5 is already included in these payments.)

Ground rules: Please do not communicate with any other participants, do not use your cell phone, and do not use any other software on the computer. If you have any questions, please raise your hand and the experimenter will assist you.

Are there any questions before we begin?

APPENDIX C. EXPERIMENTAL INSTRUCTIONS: NO RATINGS

This is an experiment in the economics of decision-making. If you listen carefully and make good decisions, you could earn a significant amount of money that will be paid to you in cash at the end of the experiment.

In this experiment, there is a jar called the "Unknown Jar" that is filled with 9 balls that are either red or blue. The jar is either the "red jar", meaning it contains 5 red balls and 4 blue balls, or the "blue jar", which contains 5 blue balls and 4 red balls. Thus, the "red jar" has slightly more red balls, and the "blue jar" has slightly more blue balls.

At the beginning of each round in the experiment, the computer will flip a coin to determine which jar to use – either the red jar or the blue jar. But you won't know which one is chosen. In the computer interface, you'll be shown a jar with 9 balls of unknown color (Figure 1). To give you a clue about which jar it is, you can then click on any one ball and its color will be revealed to you (Figure 2). The other 8 balls will not be revealed.

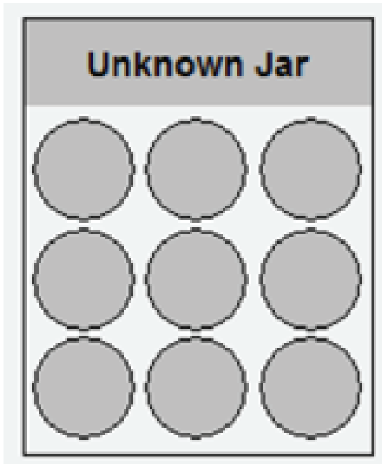

Figure 1: 9 balls of unknown color.

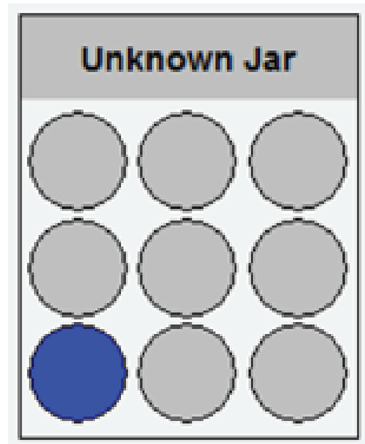

Figure 2: A blue ball was revealed.

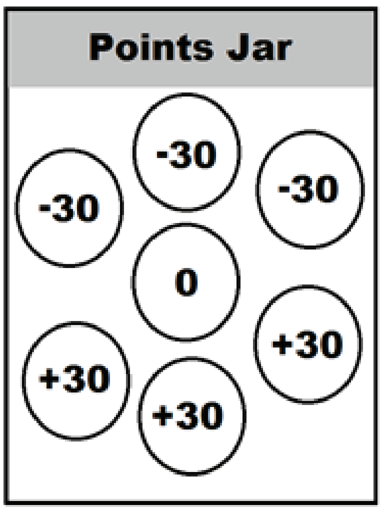

Figure 3: The “Points Jar”.

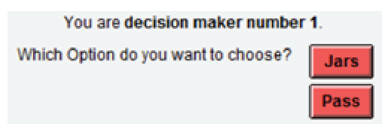

Figure 4: Choose “Jars” or “Pass”

You will then choose one ball from the “Points Jar”. As shown in figure 3, the Points Jar contains 3 balls labeled “+30”, 1 ball labeled “0”, and 3 balls labeled “-30”. The value of the “Points Ball” you draw is not yet revealed to you.

After drawing a ball from each jar, your job is to choose whether you want to get paid based on which jar the Unknown Jar is, or whether you prefer to “pass”. You do this in the computer by clicking the button labeled “Jars” or the button labeled “Pass” (Figure 4).

If you choose “Jars”, you receive 20 points if the unknown jar is the RED jar, and you lose 20 points if it is the BLUE jar. In addition, you will earn (or lose) the number of points shown on your Points Ball.

For example, suppose you choose “Jars”. If the unknown jar is red and the Points Ball you drew is “+30”, you earn 20 + 30 = 50 points. If the unknown jar is blue and your Points Ball is “0”, then you earn -20 + 0 = -20 points. If we count all possible combinations, your total point earnings could be -50, -20, or 10 if the Unknown Jar is blue, and -10, 20, or 50 if the Unknown Jar is red.

If you choose to "Pass", then you earn zero points regardless of whether the Unknown Jar is red or blue, and the ball you drew from the Points Jar will not affect the number of points you earn. In summary, your choices are:

JARS: Receive 20 points if the unknown jar is RED. Lose 20 points if the unknown jar is BLUE. In addition, you earn (or lose) the number of points shown on the ball you drew from the Points Jar (Figure 4).

PASS: Receive 0 points.

You will be in a group of 9 subjects all making the same decision about the same jar. The other members of your group also get to choose one ball from both the Unknown Jar and the Points Jar and then choose Jars or Pass. The balls they draw may be the same as yours or different, depending on which they clicked. Thus, each of you may have different information about the Unknown Jar.

You will take turns making your decision. Each person in your group is assigned a "decision maker number", which tells the order in which you will make your decision. First, everyone chooses a ball from each jar. Then, decision maker #1 chooses between Jars or Pass, followed by decision maker #2, and so on.

Decision Maker	Choice	Feedback
1	Jars	Positive
2	Pass	
3	Jars	Positive
4	Jars	Negative
5	Pass	

Figure 5: The choices & feedback of decision makers 1 through 5.

Each person gets to see the history of what the earlier decision makers in their group chose. For example, decision maker #4 gets to see what decision makers #1, 2, and 3 chose. An example history that decision maker #6 might see is shown in Figure 5. This shows that decision maker #1 chose "Jars", decision maker #2 chose "Pass", and so on. In the 3rd column of the history table, you can also see whether the earlier decision makers' point earnings were positive or negative. We call this their "feedback". Remember, if they chose Jars and the unknown jar is red, then their final payoff could be -10 (negative), 20 (positive), or 50 (positive). Since there's a 3/7 chance of drawing a "-30" Points Ball, there's a 3/7 chance that the person has negative feedback even though the jar was red. Similarly, if the unknown jar is blue, then their final payoff could be

-50 (negative), -20 (negative), or 10 (positive). Here, there's a 3/7 chance the person has positive feedback even though the jar was blue.

If a previous decision maker chose "Pass", then you know their payoff is zero. In that case, no feedback is provided.

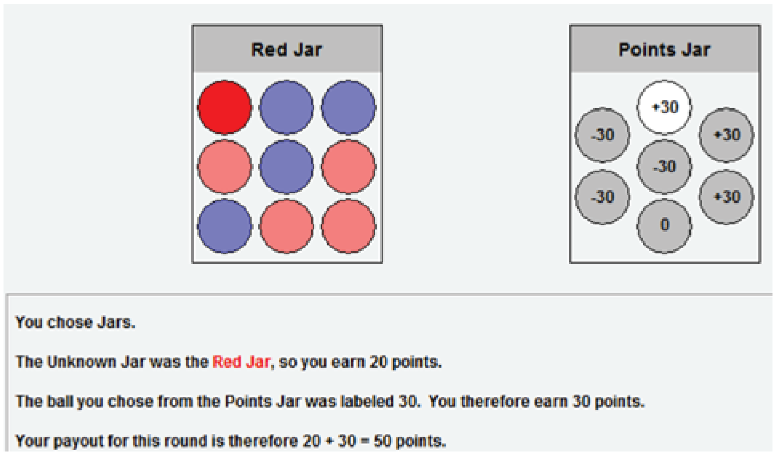

Figure 6: The jar is revealed and your earnings are calculated.

Once everyone in your group has made their choices, you will see whether the Unknown Jar is red or blue. If you chose "Jars", you will also see the contents of the Points Jar and the value of the ball you drew. Finally, you will see your earnings in points. For example, in Figure 6, the unknown jar was the Red jar. The player chose "Jars", so they earned +20 points. Their Points ball was "+30", so their total payoff is 20+30 = 50 points.

This completes one round of the experiment. You will play a total of 10 rounds. In each round everyone will be randomly assigned a new group, and will get a new decision maker number within that group. The computer will flip a coin again to decide whether the Unknown Jar will be the red jar or the blue jar. It will also randomly shuffle the contents of the Unknown Jar and the Points Jar.

At the end of the experiment, we will randomly pick one of the 10 rounds for actual payment. You will get paid \$15 plus 10 cents for every point earned in that one chosen round (or, minus 10 cents for every point lost if your earnings for that round were negative). Since the maximum you can earn in any one round is 20+30=50 points, your maximum possible payment for the experiment is \$15 + (50 × \$0.10) = \$20.00 The minimum you can earn in any round is -20-30 = -50 points, so your minimum possible payment for the experiment is \$15 + (-50 × \$0.10) = \$10.00. (Your show-up fee of \$5 is already included in these payments.)

Ground rules: Please do not communicate with any other participants, do not use your cell phone, and do not use any other software on the computer. If you have any questions, please raise your hand and the experimenter will assist you.

Are there any questions before we begin?