



*The World's Largest Open Access Agricultural & Applied Economics Digital Library*

**This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.**

**Help ensure our sustainability.**

Give to AgEcon Search

AgEcon Search  
<http://ageconsearch.umn.edu>  
[aesearch@umn.edu](mailto:aesearch@umn.edu)

*Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

*No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.*

**Information Cascades:  
Evidence from a Field Experiment with  
Financial Market Professionals**

*by*

Jonathan E. Alevy, Michael S. Haigh,  
and John A. List

WP 03-15

Department of Agricultural and Resource Economics  
The University of Maryland, College Park

# Information Cascades: Evidence from a Field Experiment with Financial Market Professionals

Jonathan E. Alevy, Michael S. Haigh, and John A. List  
University of Maryland, CFTC, and NBER

## **Abstract:**

In settings characterized by imperfect information about an underlying state of nature, but where inferences are made sequentially and are publicly observable, decisions may yield a “cascade” in which everyone herds on a single choice. While cascades potentially play a role in a variety of settings, from technology adoption to social processes such as mate selection, understanding cascade phenomena is imperative for financial markets. Previous empirical efforts studying cascade formation have used both naturally occurring data and laboratory experiments. In this paper, we combine one of the attractive elements of each line of research—observation of market professionals in a controlled environment—to push the investigation of cascade behavior into several new directions. Numerous empirical insights are obtained; perhaps most importantly, we find that market professionals behave quite differently than a control group of student subjects. In particular, market professionals, more so than students, base their decisions on the “quality” of the public signal, leading them to be more likely to disregard “bad” signals. And, unlike in the case with students, for market professionals, the propensity to be Bayesian does not differ significantly across the gain and loss domains. These results have important implications in both a positive and normative sense.

**Keywords:** Herd Behavior, Futures Traders, Experiments

**Manuscript Status:** First Draft: Dec 16<sup>th</sup> 2003

**Acknowledgements:** Thanks to John Di Clemente, former Managing Director of Research at the Chicago Board of Trade, for authorizing this study. Also special thanks to CBOT staff Dorothy Ackerman Anderson, Frederick Sturm, and Keith Schap for their incredible support on site. Valuable comments and suggestions were provided by seminar participants at Harvard University, University of Chicago, University of Maryland, the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting and Market Risk Management, in St. Louis, Missouri, April 2003, and the Economic Science Association International Conference, in Pittsburgh, Pennsylvania, July 2003. Liesl Koch and Vernon Smith provided comments on an earlier version of this manuscript.

## 1. Introduction

In economic environments where decision makers have imperfect information about the true state of the world, it can be rational to ignore one's own private information and make decisions based on what are believed to be more informative public signals. In particular, if decisions are made sequentially and the earlier decisions become public information, herding behavior or “information cascades” can result. Information cascades arise when individuals rationally choose identical actions despite having different private information. They may arise in a myriad of settings, including technology adoption, medical treatment choices, responses to environmental hazards, as well as decisions in financial markets, where bubbles and crashes may be examples of cascade behavior.<sup>1</sup>

Herding can be suboptimal, since the private information of herd followers is not revealed. As a result, a small amount of information revealed early in a sequence has a large impact. Cascades are thus idiosyncratic, with their welfare effects depending on chance revelation of information in early periods. One consequence is that cascades are often fragile as well, with abrupt shifts or reversals in direction when new information becomes available (Banerjee, 1992; Bikhchandani et al., 1992). Indeed, some argue that volatility induced by herding behavior can increase the fragility of financial markets and destabilize the broader market system (Eichengreen et al., 1998; Bikhchandani and Sharma, 2000).<sup>2</sup>

---

<sup>1</sup> It has been argued, also, that information cascades can explain a large variety of social behaviors such as fashion, customs, and rapid changes in political organization. Anderson (1994), Banerjee (1992), Bikhchandani et al. (1992; 1998), Kuran and Sunstein (1999), and Welch (1992) discuss a variety of examples.

<sup>2</sup> The distinction between technical and fundamental approaches to trading illustrates one avenue through which herding behavior may be introduced into financial markets. Technicians ignore underlying supply and demand conditions and assume that more informed traders' signal their understanding through price changes. Herding may also arise from agency problems, such as those associated with the management of pension or mutual funds when compensation derives from relative investment performance or from payoff externalities, such as those arising in bank runs (Bikhchandani and Sharma, 2000).

Empirical approaches to test for cascade behavior can be divided into two classes: regression-based tests using naturally occurring data and laboratory experiments using student subjects. Bikhchandani and Sharma (2000) review the extant regression-based results for herding in financial markets but note the difficulty of controlling for underlying fundamentals. A result of this difficulty, they argue, is that there is often “a lack of a direct link between the theoretical discussion of herding behavior and the empirical specifications used to test for herding.”<sup>3</sup> The laboratory environment allows for better control of public and private information and so explicit tests of theory are made more easily. Yet an important debate exists about the relevance of experimental findings from student subjects for understanding phenomena in the field. For example, Holt and Villamil (1986) provide numerous reasons to suspect that professional behavior may differ from non-professional behavior (training, regulation, etc.). Locke and Mann (2003) take the argument a step further by suggesting that any research that ignores the use of professional traders is likely to be received passively because “ordinary” individuals, as opposed to professional traders, are unlikely to have any substantial impact on market price because they are too far removed from the price discovery process. Bikhchandani and Sharma (2000, p. 13) argue similarly that “To examine herd behavior, one needs to find a group of participants that trade actively and act similarly.”

This study combines perhaps the most attractive aspects of the two classes of empirical research—observation of professionals in a controlled environment—to extend the literature in several new directions. First, the behavior of the market professionals from the floor of the Chicago Board of Trade (CBOT) is compared with that of college students in imperfect information environments. Second, considering the vast normative implications of recent work that has established the importance of domain (Kahneman and Tversky 1979; Shefrin and

---

<sup>3</sup> Fama (1998) discusses the interpretation of empirical results as evidence of irrational behavior.

Statman 1985; Tversky and Kahneman 1992; Odean 1998), we examine behavior of each group in both the gain and loss domains. Third, we examine whether, and to what extent, behavior is influenced by both private signal strength and the quality of previous public signals.

Empirical findings, gained from an examination of more than 1500 individual decisions, lend some interesting insights into cascade behavior. Most important, we find economically significant and robust differences in behavior across subject pools. Whereas both groups exhibit a considerable amount of Bayesian decision making across the various treatments, market professionals are much less apt to be influenced by the relevant domain. This finding is consonant with Locke and Mann (2003), Genesove and Mayer (2001), and List (2003; 2004), who find, in much different environments, that market experience is important in reducing deviations from utility theory that are associated with reference-dependent preferences. In addition, we find that market professionals base their decisions on the “quality” of the public signal to a greater extent than do students, leading them to enter fewer reverse cascades. We believe this result is novel to the literature and has important implications for financial markets.<sup>4</sup>

The remainder of the study is crafted as follows. Section 2 discusses our theory and experimental design. Section 3 describes our empirical results. Section 4 concludes.

## 2. Theory and Experimental Design

The cascade phenomenon has been discussed in various circles for several years (see Banerjee, 1992, for an economics-based model of herd behavior).<sup>5</sup> The underlying argument in the formal economic models is that cascade formation can be rational for Bayesian decision

---

<sup>4</sup> Evidence that professional traders exhibit non-rational behavior would increase support for behavioral approaches to asset pricing (Locke and Mann, 2003). See for example the model of Barberis et al. (1998) and Daniel et al. (2001).

<sup>5</sup> In 1841 Charles MacKay authored *Extraordinary Popular Delusions and the Madness of Crowds*, which provides interesting reading on financial market bubbles and crashes, as well as fads and fashions — what we would now call social learning.

makers. The cascade phenomenon, however, raises interesting questions beyond whether humans update information in a manner that is consistent with Bayes' rule.<sup>6</sup> The formation of informational cascades also forces one to address the question of how people think about the rationality of others, and specifically, how people respond to uncertainty about the quality of information that arises due to potential deviations from Bayesian rationality by others.

Anderson and Holt (1997) presented an interesting experimental study to investigate these issues, using a subject pool of undergraduate students. In their experiments each individual receives a private but noisy signal regarding the underlying state of nature in an exogenously determined order. After receiving their signal each subject publicly announces his or her belief about the state of nature in a sequential manner. Subsequent decision makers have the opportunity to observe the announcements of players who preceded them. Rational herding can develop when public announcements provide evidence for a specific state that overwhelms the informational content of the private draw.<sup>7</sup>

To ensure comparability of our results to the extant literature, we use experimental protocols that are closely related to the original work of Anderson and Holt (1997). Briefly, in our experimental investigation, we examine a situation in which two potential states of the world exist and are represented by two urns, Urn A and Urn B. The urns differ in their composition,

---

<sup>6</sup> The ability of humans to reason in a Bayesian manner seems to depend strongly on how new information is represented. Studies that present base rates as percentages often imply that we are poor "intuitive statisticians." Decisions tend to be more consistent with Bayesian rationality when individuals experience probability distributions through repeated exposure (see, e.g., Gigerenzer and Murray, 1987). Our experiment is consistent with protocols that have been shown to give Bayesian decision making its best chance.

<sup>7</sup> Anderson and Holt (1997) found that cascades formed in roughly 70 percent of the periods in which they were possible. Deviations from Bayesian cascade formation occurred most often when a simple counting rule gave a different indication of the underlying state than that of the Bayesian posterior. In these cases subjects tended to use the simpler counting rule. Extensions to the literature have introduced relevant complications to the cascade process, such as costly information, endogenous sequencing of choice order, and collective decision making (Hung and Plott, 2001; Kraemer et al., 2000; Kubler and Weizsäcker, 2002). In general, as complications are introduced, more significant deviations from normative rationality are observed.

containing different numbers of balls of type  $a$  and type  $b$ .<sup>8</sup> The unobserved “state of the world” in each period of play is the urn selected by a random process, the roll of a die. Each urn has a 50 percent chance of being selected. Subjects gain information about the state of the world by drawing a single ball out of an unmarked bag into which the contents of the selected urn have been transferred. This draw is made while the subject is isolated from the other experimental subjects so that the outcome of the draw remains private information. The subject then decides which urn he or she believes was selected by the random process and his or her decision is announced to the other subjects by the experimental monitor. Each subject’s inference, therefore, becomes public information for those who make decisions later in the period. The experimental sessions include 15 periods of this game in which each of the participants draws a ball and announces his or her urn choice. Sessions consist of either five or six players who participate in all 15 periods, but whose choice order in each period — either first, second, third, ...., sixth? is determined by a random draw.

To provide exogenous variation in the informational content of the private signal across treatments, we use two urn types. In the *symmetric* treatment, Urn A contained two type  $a$  balls and one type  $b$  ball, while Urn B contained two of type  $b$  and one of type  $a$ . In the *asymmetric* treatment four  $a$  balls are added to each urn so that Urn A contains six of type  $a$  and one of type  $b$ , while Urn B contains five of type  $a$  and two of type  $b$ . As a result, the  $a$  signal is weakened in the asymmetric treatment. The variation in signal strength is implemented in order to differentiate between Bayesian decision making and the use of a counting heuristic, in which the signal observed most often would indicate the most probable urn. In the symmetric treatment the

---

<sup>8</sup> The balls were golf balls of two types. Type  $a$  was a golf ball with visible stripes. Type  $b$  had a graphic of the University of Maryland terrapin mascot “Testudo” clearly visible.

counting heuristic and Bayesian decisions yield the same predictions; in the asymmetric condition they do not.

To provide exogenous variation in the relevant choice domain, we randomly placed subjects in either a *gain* or *loss* domain for all 15 periods. The treatment condition over gains and losses was implemented so that in gain (loss) space a correct (incorrect) inference about the underlying state resulted in positive (negative) earnings of \$1 for students and \$4 for the market professionals.<sup>9</sup> An incorrect (correct) choice in gain (loss) space resulted in no earnings. To provide similar monetary outcomes across treatments, in the loss treatments, students and market professionals were endowed with \$6.25 and \$25.00.<sup>10</sup> We believe this is the first study to vary the domain in cascade games.

Our final comparative static treatment variable concerns subject pools. Experimental subjects in a particular session consisted entirely of one of two subject types: students or market professionals. The experimental sessions with market professionals were conducted at the Chicago Board of Trade (CBOT) and the student data were gathered from undergraduates at the University of Maryland in College Park. The CBOT (student) subject pool included 55 (54) subjects recruited from the floor of CBOT (the university).<sup>11</sup>

Each experimental session consisted of a group of either five or six participants making decisions within the same treatment type over 15 periods. Table 1 summarizes the experimental

---

<sup>9</sup> CBOT officials suggested that designing a 30-minute game with an expected average payout of approximately \$30 was more than a reasonable approximation of an average trader's earnings for an equivalent amount of time on the floor. In our experiments the median earnings for the market professionals were slightly in excess of this amount and therefore likely to be salient.

<sup>10</sup> To ensure that subjects departed with positive money balances we had both subject pools participate in other unrelated games during the experimental session.

<sup>11</sup> In practice, the 55 subjects recruited from the floor of CBOT consisted of locals, brokers, clerks, and an exchange employee. We found no statistical difference among the types of floor participants with regard to Bayesian behavior and cascade formation and collectively call them the "market professionals." Homogeneity among the floor personnel is intuitively sensible since the average non-trader had accumulated approximately 9 years of floor experience and many reported experience as either brokers or local traders.

sessions that were conducted. The experimental design is a 2x2x2 factorial across urn type, either symmetric (S) or asymmetric (A); over gains (G) or losses (L); and with either college students (C) or market professionals (M).<sup>12</sup>

### ***Theoretical Predictions***

The formal examination of cascade behavior presented by Bikhchandani et al. (1992) assumes that individuals are Bayesian in their update of beliefs about underlying states of nature. Consider an environment in which there are two possible underlying states, *A* and *B*. Each individual receives a private signal, either *a* or *b*, indicating the probability of a state where  $p(A|a) > p(B|a)$  and  $p(A|b) < p(B|b)$ . Signal strength in the symmetric treatments implies that  $p(A|a) = p(B|b)$ . To understand the mechanics of cascade formation, we consider an example parameterized in accordance with our symmetric experimental treatment discussed above and presented in Figure 1. For the initial decision maker,  $p(a|A) = p(b|B) = 2/3$ , and therefore  $p(b|A) = p(a|B) = 1/3$ . Suppose, in fact, that the first draw is *a*. According to

Bayes' rule, the probability that the underlying state is *A* is given by

$$p(A|a) = \frac{p(a|A)p(A)}{p(a|A)p(A) + p(a|B)p(B)} = \frac{(2/3)(1/2)}{(2/3)(1/2) + (1/3)(1/2)} = \frac{2}{3}.$$

An expected utility

maximizer would, therefore, choose *A* as the state of nature since expected profits for announcing *A*,  $\mathbf{p}_A$ , exceed those for announcing *B*,  $\mathbf{p}_B$ .<sup>13</sup> Suppose that the second subject also draws signal *a* from the urn. Updating according to Bayes' rule yields

---

<sup>12</sup> Experimental instructions for the symmetric gain treatment for students are included in the appendix.

<sup>13</sup>  $\mathbf{p}_A - \mathbf{p}_B = \frac{\$W}{3}$  in the gain treatments, after an initial *a* signal, where  $\$W$  is the win amount, either \$1 for students or \$4 for professionals. Note also that treatments over gains and losses yield identical predictions (i.e., expected losses are minimized by picking the most probable urn).

$$p(A | A, a) = \frac{p(a | A)^2}{p(a | A)^2 + p(a | B)^2} = \frac{(2/3)^2}{(2/3)^2 + (1/3)^2} = \frac{4}{5}. \quad \text{Thus, two consecutive identical}$$

announcements yield a posterior probability of 0.80 in favor of the urn announced. As a result, the third decision maker should “follow the herd” and announce “A” regardless of his or her own private draw. This can be seen by examining the case where an opposing  $b$  signal is the private draw of the third player after two consecutive  $A$  announcements, yielding the posterior probability

$$p(A | A, A, b) = \frac{p(a | A)^2 p(b | A)}{p(a | A)^2 p(b | A) + p(a | B)^2 p(b | B)} = \frac{(2/3)^2 (1/3)}{(2/3)^2 (1/3) + (1/3)^2 (2/3)} = \frac{2}{3}. \quad \text{We}$$

denote a decision of this type—consistent with Bayesian rationality but in which one’s own private signal is ignored—a *cascade decision*.

Naturally the cascade decision may result in either a correct or incorrect inference about the underlying state. It is entirely possible, in the example above, that the true underlying state is B. Following the literature, we denote a cascade decision in which the wrong underlying state is announced a *reverse cascade*. Regardless of the underlying state, the theoretical analysis of cascade formation in the symmetric treatment implies that private information is uninformative whenever the number of public signals of one type exceeds the other by two or more.<sup>14</sup>

In our asymmetric treatment, presented in Figure 2, there are 6 (5)  $a$  signals and 1 (2)  $b$  signal in  $A$  ( $B$ ). The intuition behind formulation of Bayesian updating is similar to the symmetric treatment. Table 2 provides posterior probabilities for all possible sequences of draws

---

<sup>14</sup> The analysis of equilibrium behavior depends on the tie-breaking rule invoked when the posterior probability of an urn is 1/2. We follow Anderson and Holt (1997) in assuming that individuals who are indifferent announce their own signal. This is sensible if there is a nonzero probability of error in announcements. In our treatments this rule is followed 81 percent of the time. The type of tie-breaking rule chosen has no effect in the asymmetric treatment since a posterior probability of one-half is not possible. We point the reader to Koessler and Ziegelmeyer (2000) for further discussion of the impact of tie-breaking rules on equilibrium predictions.

for both the symmetric and asymmetric treatments. As an example, the two-thirds probability of urn  $A$  that arises after a single  $a$  draw in the symmetric treatment is matched in the asymmetric case only by 4 consecutive  $a$  draws.<sup>15</sup> One consequence of the change in signal strength is that herding on the  $B$  state should take place after a single  $b$  signal even with either one or two  $a$  signals in the game's history.

This difference presents an opportunity to distinguish, across treatments, Bayesian updating from a simpler choice heuristic. The main difference between the symmetric and asymmetric treatments is that the optimal Bayesian decision in the symmetric treatment corresponds to a simple counting rule: choosing the event with the most signals maximizes expected earnings. In the asymmetric case a number of sequences violate this counting rule in that it is optimal to choose  $B$  even when there are fewer  $b$  signals than  $a$  signals. The asymmetric treatment, therefore, allows us to distinguish Bayesian behavior from heuristics that may mimic Bayesian behavior in simple settings. The signal histories that are of interest in making this determination are  $(a, b) \in \{(2,1), (3,1), (3,2), (4,2)\}$  and are underlined in Table 2.

The treatments over gains and losses also raise questions that can be addressed by examining decision error. While expected utility implies that it is rational to always choose the most probable urn, previous experimental results suggest that there are deviations from this strategy. One explanation for this behavior derives from theories that incorporate decision costs into the choice of urns (Anderson and Holt, 1997). Consider the case where the posterior probability is only slightly in favor of one urn. Here the expected return from choosing the most probable urn is positive but small. Deviations from the Bayesian predictions may be rational if the cost of optimizing is greater than the expected benefits. Thus, we expect more Bayesian

---

<sup>15</sup> Another measure of the change in signal strength is that an initial  $a$  signal yields  $\text{Pr}(\text{urn}=A)=54.5$  percent in the asymmetric treatment and  $\text{Pr}(\text{urn}=A)=67$  percent in the symmetric treatment.

decisions when the posteriors move further away from one-half, and the expected utility of a correct choice increases. In this setting, we can investigate indirectly a prediction of prospect theory, by examining whether decisions are more or less Bayesian over gains than over losses of equal magnitude. The notion that losses loom larger than gains suggests more Bayesian decision making in the loss domain, if decision costs across the domains are assumed to be constant within a subject pool.<sup>16</sup>

### 3. Experimental Results

Table 3a presents descriptive statistics on Bayesian decision making and on the aggregate rate of cascade formation in our data. The experimental sessions yielded a total of 1,647 decisions, 1,293 (79 percent) of which were consistent with Bayesian rationality. Cascade decisions, defined as those that were Bayesian but in which the private signal was ignored, took place in 16 percent of the cases. In addition, one quarter of the cascades formed were reverse cascades, resulting in the wrong inferences about the underlying state. More revealing than the aggregate number of cascades is the proportion of cascade decisions made when it was possible to make one. Recall that a cascade decision is possible only when the private draw is inconsistent with the probability weight of the previous signals. In our data, cascade formation was possible in 466 decisions, representing 28 percent of the total. Cascades were realized in 262, or 56 percent, of the cases. These last two results are presented in the *potential* and *realized* cascades columns of Table 3a.

To shed further light on the experimental behavior we adopt a variety of parametric and nonparametric techniques and group our core results into five categories. Four of the categories

---

<sup>16</sup> To investigate properly this hypothesis we must control for the posterior probability, or rather the difference between the posterior and one-half, where expected utility maximizers are indifferent between urn types. We control for the posterior probability in the probit model below.

are associated with subject pool effects, and the fifth relates to a quality of decision making in which the student and market professionals behave similarly. With regard to subject pool effects we consider (1) differences in Bayesian decision making, (2) differences in cascade formation, (3) differences across the domains of gains and losses, and (4) differences within the market professional subject pool that are associated with individual trading experience. The fifth category of results relates to the exogenous alteration of signal strength through the use of the symmetric and asymmetric urns. The treatment effects provide insight into the descriptive validity of Bayesian rationality.

*Result 1. Market professionals are less Bayesian than students. Despite this behavioral discrepancy, earnings are not significantly different across subject pools.*

Table 3b reports similar statistics to those presented in Table 3A, but Table 3B disaggregates the data by treatment and subject type. In aggregate, 82 percent (76 percent) of the students' (market professionals') decisions are consistent with Bayesian decision making. Individual subjects ranged from 38 percent to 100 percent Bayesian, and of the seventeen subjects perfectly consistent with Bayesian rationality, twelve were students.<sup>17</sup> To investigate whether these differences in Bayesian behavior are statistically significant, we employ both unconditional and conditional tests. When using unconditional tests, we recognize the dependence among individual observations by using session level aggregates whenever possible to yield the most conservative estimates of treatment effects. Our unconditional test used to support Result 1 is a non-parametric Mann-Whitney U test, which suggests that the rate of Bayesian decision making differs across subject pools at a level of significance of  $p=.052$ .

To supplement the non-parametric test results, we estimate the following random effects probit model:

---

<sup>17</sup> Sixteen of the seventeen were in the symmetric urn treatment.

$$Baye_{it} = \beta' X_{it} + e_{it}, \quad e_{it} \sim N[0,1], \quad (1)$$

where  $Baye_{it}$  equals unity if agent  $i$  was a Bayesian in period  $t$ , and equals zero otherwise; and  $X_{it}$  includes variables predicted to influence play:  $order\_x$  (where  $x=2,..,6$ ), which is a categorical variable indicating where in the period of play the decision was made. The posterior probability is incorporated in the variable  $diff$  which is calculated as  $|prob(urn = A) - .5|$ , where the  $prob(urn = A)$  is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The absolute value of the difference of the posterior probability from one-half is an indication of the magnitude of the accrued public and private information. The  $diff$  variable thus varies from zero to one-half, increasing with the evidence for a specific urn type. Theories of noisy decision predict that the parameter on this variable is positive, with decisions more Bayesian as the posterior probability of a specific urn type increases. The variables  $gain$ ,  $sym$ , and  $trader$  are dichotomous and distinguish the experimental treatments, with  $gain$  equal to one for treatments over gains and zero over losses, and  $sym$  equal to one for the symmetric treatment and zero for the asymmetric treatment. In the specification pooling subject types,  $trader$  is equal to one for the market professionals and zero for the students. The variable  $heurist$  is a dummy variable equal to one for counting rule sequences and zero for all others. We specify  $e_{it} = u_{it} + \alpha_i$ , where the two components are independent and normally distributed with mean zero; thus,  $Var(e_{it}) = \sigma_u^2 + \sigma_\alpha^2$ . We estimate equation (1) using the maximum likelihood approach derived in Butler and Moffitt (1982).<sup>18</sup> Estimation of this particular model is quite complex, but is amenable to Hermite integration. To

---

<sup>18</sup> The likelihood function can be succinctly written as

$$L = \prod_i L_i = \int_{-\infty}^{\infty} (2\pi)^{-1/2} \prod_i \exp(-e_{it})^2 \phi(g_{it}q_{it}),$$

where  $g_{it} = 2nc_{it} - 1$ ; and  $q_{it} = \beta' X_{it} + [\text{corr}(e_{it}, e_{is}) / (1 - \text{corr}(e_{it}, e_{is}))]^{1/2} e_i$ .

estimate the model, we use an eight-point quadrature and follow Berndt et al. (1974) to compute the covariance matrix.<sup>19</sup>

Empirical results are presented in Table 4. Results from both a likelihood-ratio test and the *trader* dummy variable in the pooled regression model (panel a) indicate that there are differences in Bayesian behavior across subject pools: market professionals are less Bayesian than students.<sup>20</sup> An estimated coefficient of -0.201 suggests that traders are 4.8 percent less likely to be Bayesian compared to students. Considering the results from all three panels, it is clear that posterior probability (*diff* variable) also has important influences on Bayesian choice, as the larger is the divergence from a posterior of one-half, the more likely the choice will be consistent with Bayesian rationality.

Regressions based on subject type, presented in Table 4, panels b and c, reveal that students are more sensitive to some treatment effects than are the market professionals. Student subjects respond more dramatically to the urn symmetry. In addition, empirical results by subject pool indicate a decline in Bayesian behavior among market professionals who choose in the third through sixth position (students show no such effect). The magnitude of the effect is rather large, having from one-third to more than one-half of the effect of the posterior probability as represented in the *diff* variable (Table 4, panel c).

Overall, the maximum likelihood results confirm the results from the non-parametric tests. Despite their less Bayesian environment, market professionals do not perform significantly worse than the students with respect to the rate at which they choose the correct underlying state

---

<sup>19</sup> Results are robust to inclusion of a time trend for period or time dummies (categorical time dummy variables for each period of play). Below we discuss further our evidence of learning.

<sup>20</sup> A Chow test rejects the null hypothesis of no differences across the subject pools at the  $p = 0.01$  level.

(Mann-Whitney U-test on win percentage  $p= 0.29$ ), a point we take up again in our discussion of Result 2 on cascade formation.

In our third, and final, empirical technique to examine the extent of Bayesian decision making, we applied a version of the quantal response equilibrium (QRE), which incorporates boundedly rational decision making into game theoretic equilibrium considerations, to the data from the symmetric gain treatment. This approach, developed by Palfrey and McKelvey (1995), was first applied to the cascade setting by Anderson and Holt (1997) in a framework that assumes individuals have correct beliefs about the errors of others. The QRE results provide perhaps the clearest picture of the noisier environment that the market professionals face in the experimental setting due to the greater number of deviations from Bayesian rationality.

The QRE derives from the following logic: the probability of announcing an urn depends on the expected payment associated with choosing it. This logic is consistent with the results of our probit model, which reveals that as the posterior probability gets closer to one-half, decisions are less Bayesian. Alternatively, error-free Bayesian decisions require that the urn with a posterior probability greater than one-half would be chosen with probability 1. We calculate two probabilities associated with the QRE. The first reflects the posterior probability of an urn type. In what follows we discuss and report the probability that the type is Urn A. The QRE estimation yields a parameter that indicates the degree of departure from pure Bayesian decision making. The posterior and the estimated parameter allow us to calculate the second probability: the probability for a given history of play that the announcement by a subject is Urn A.

We implement the QRE by assuming a logistic choice rule:

$$pr(D_i = A) = \frac{\exp \mathbf{p}_i^A \mathbf{b}_i}{\exp \mathbf{p}_i^A \mathbf{b}_i + \exp \mathbf{p}_i^A \mathbf{b}_i} = \frac{1}{1 + \exp(\mathbf{p}_B - \mathbf{p}_A) \mathbf{b}_i}, \quad (2)$$

where  $\mathbf{p}_i^A = pr(A) * \$W + (1 - pr(A)) * 0$ ,

$$\mathbf{p}_i^B = (1 - pr(A)) * \$W + pr(A) * 0,$$

and so  $\mathbf{p}_i^B - \mathbf{p}_i^A = (1 - 2pr(A)) * \$W$ .

Notice in equation (2) that the estimated beta captures the extent of the noise in the decision-making process. As beta approaches zero, the probability of choosing Urn A approaches one-half, reflecting random choice. As beta grows large, the probability of choosing Urn A approaches 1 when  $\mathbf{p}_i^A > \mathbf{p}_i^B$ ,  $i \in 1, \dots, 6$ . Thus,  $\mathbf{b} = \infty$  implies perfect Bayesian rationality.

Contributions to the likelihood function from the first two decisions of the cascade game are

Decision 1:

$$LogL = \sum_{i=1}^n D_i \ln F(X_1 \mathbf{b}_1) + (1 - D_i) \ln (1 - F(X_1 \mathbf{b}_1))$$

Decision 2:

$$LogL = \sum_{i=n+1}^{2n} D_i \ln F(X_1(\bar{\mathbf{b}}_1) \mathbf{b}_2) + (1 - D_i) \ln (1 - F(X_1(\bar{\mathbf{b}}_1) \mathbf{b}_2))$$

where  $D_i$  is an urn  $A$  announcement and  $X$  is  $\mathbf{p}_i^B - \mathbf{p}_i^A$ .  $F$  assumes that errors follow a logistic distribution, and a scale parameter for the distribution,  $\mathbf{b}_i$ , is estimated from the data. The parameter derived from the first choice is used to calculate the probability  $D_i = A$  for all signal and announcement possibilities.<sup>21</sup> Note that for  $i = 2$  the independent variable, the difference in profits, is a function of the first period estimation through the parameter  $\bar{\mathbf{b}}_1$ . This is because the QRE assumes that decision makers incorporate the possibility of error in decision 1 into their estimate of the posterior probability of the urn type.

---

<sup>21</sup> In the first period there are four possibilities — an  $a$  or  $b$  signal can be followed by either an  $A$  announcement or a  $B$  announcement. In general there are  $2^{j+1}$  possibilities where  $j$  represents the choice order.

Parameter estimates and the implications for decision probabilities and posterior probabilities are presented in Tables 5 and 6. Estimation is conducted with the payments normalized to \$1 for both students and traders. Table 5 presents the parameters estimated for each choice order, and we again find significant differences between subject pools. In choice positions 1, 2, and 5, the students' parameters are larger, reflecting more Bayesian choices. In fact, in position 2 the students' decisions are all consistent with Bayesian updating, and we see a relatively large increase in the estimated beta as a result. Further, in all choice positions the parameters are significantly different from zero, indicating that urn choices differ significantly from what pure random choice would produce.

The posterior probabilities that result from incorporating decision error yield an indication of how cascade formation can fail. Table 6 presents the posterior probabilities and decision probabilities that reflect the adjustment for decision error. For comparison, the Bayesian posteriors and choice probabilities are also presented. The QRE results emphasize the fact that not only the numbers of each signal, but the order in which they are both revealed have an important impact on behavior. Note the posterior probability for order choice 3, the first choice where a cascade may form in the symmetric treatment, when the signal history is AAb (or BBa).<sup>22</sup> In this case the posterior has dropped dramatically from 0.67 for the most likely urn to 0.51 (0.59) for the market professionals (students). Thus, while cascade formation is still optimal for both groups, the noise in prior decisions dilutes the strength of the signals, with the market professionals facing essentially a random choice. Compare this to the ABa sequence, which has an identical Bayesian posterior probability. In this case, when one's own draw is

---

<sup>22</sup> Since the AAb and BBa are symmetric, Table 6 reports the results from these sequences as one choice history. The history reported is that of the AAb sequence. All other symmetric choice sequences are treated similarly.

consistent with the most likely urn, the posterior is 0.64, and the optimal decision is made almost uniformly by both subject pools.

While the evidence suggests that students are more Bayesian, we find that in fact the proportion of correct decisions, and thus earnings, are not significantly different across the subject pools.<sup>23</sup> This result is surprising since Bayesian rationality is predicated on maximizing earnings. To clarify this issue we delve deeper into the data and find the following:

*Result 2. In aggregate, the rate of cascade formation is not significantly different for market professionals and students, but market professionals enter into significantly fewer reverse cascades in the asymmetric treatments. This result follows from the ability of market professionals to gauge the quality of the signal.*

Aggregate differences in individual rates of cascade formation across subject types in the symmetric treatment are only marginally significant, with the Mann-Whitney statistic yielding significance at the  $p = 0.10$  level. We find an interesting difference between the subject pools, however, in the asymmetric treatments, after accounting for predictions in the underlying state. In particular, there are significant differences with regards to reverse cascade formation. Market professionals enter reverse cascades at a rate 29 percent below that of students, a difference that is statistically significant at the  $p = 0.01$  level.<sup>24</sup> We find no significant differences, however, in the rate at which the two groups enter cascades in which the correct underlying state is chosen ( $p = 0.31$ ). This result cannot be explained by our simple model of Bayesian decision making. Instead, we use a probit model to investigate the hypothesis that market professionals use information that the students ignore in order to avoid reverse cascades.

Summary empirical results from the random effects probit model that investigates cascade formation are reported in Table 7. The dependent variable is dichotomous and equal to

---

<sup>23</sup> A Mann-Whitney U test across subject pools of the rate at which individuals announce the correct state yields  $p = 0.30$ .

<sup>24</sup> This result remains significant when observations are aggregated more conservatively at the session level.

one when a cascade forms and zero when it does not. The dataset is restricted to the observations in which cascade formation is possible: those in which the posterior implies it is optimal to ignore one's own private signal. Failure of cascade formation thus implies that one relies on the private signal despite public evidence. Cascade formation varies significantly over the urn symmetry, with 81 percent of possible cascades realized in the symmetric treatment and only 48 percent realized in the asymmetric treatment (Table 3b). In the probit model, we focus on the asymmetric treatment, because the lower rate of cascade formation and Bayesian decision making in general allows us to test the hypothesis that traders are more successful at updating their play by observing the quality of decisions of others in previous periods. Empirical results of the symmetric treatment provide insufficient variability in decision quality.

As shown in Table 7, we include a variable, *othbys*, that indicates the aggregate performance of other subjects in terms of the proportion of Bayesian decisions they have made. This variable includes only observable information since the proportion of correct inferences includes only those that occurred in preceding periods. The measures for all relevant individuals, that is, those choosing prior to the particular observation in a period, are aggregated, giving a measure of the Bayesian rationality of individuals who precede each subject's choice. The variables *diff*, *heurist*, and *gain* are also included and are defined identically as in the previous model.

Empirical results are sharp. Although *othbys* is insignificant in the pooled specification, this model masks a large difference in how the two subject pools respond to others.<sup>25</sup> Specifically, in the asymmetric treatments, cascade formation for the market professionals is significantly associated with the quality of preceding decision makers, as hypothesized. What is

---

<sup>25</sup> The likelihood-ratio test suggests structural differences between the pooled and disaggregated models at the  $p = 0.04$  level.

unexpected is that student subjects respond in the opposite manner (hence the negative sign on the *othbys* parameter), becoming less willing to join a cascade when previous decision makers have been reliable.<sup>26</sup> Alternatively, the market professionals make much better use of available information on others' rationality.<sup>27</sup> One may wonder whether this result is due to market professionals having a greater level of previous interaction with one another than students have had, or if there is evidence of learning in the experiment. To explore this issue, we examined temporal behavioral patterns across the 15 periods of the experiment. The evidence is consistent with the view that market professionals learn over the 15 periods: comparing behavior from the first and last three rounds of a session, we find that market professionals significantly reduce the rate at which they join reverse cascades and increase the rate at which they join cascades with good outcomes ( $p < .01$  and  $p < 0.01$ ), respectively. In contrast, there are no significant changes in the rate of cascade for the students for either type of cascade ( $p = 0.80$  and  $p = 0.50$ ).

*Result 3. Bayesian behavior of the student population is affected by gain/loss domain, while market professionals are unaffected by the domain. We find evidence, however, that students are unaffected by the domain of earnings during latter experimental periods.*

Aggregate figures on gains and losses reported in Table 3a are roughly equivalent, with 80 percent of Bayesian decisions over gains and 77 percent over losses. Restricting attention to the asymmetric (A) treatments (see Table 8), however, shows that college students are less Bayesian in the gain treatment than in the loss treatment, while market professionals are unaffected by the domain of earnings (AGC v. ALC  $p=0.08$ ; AGM v. ALM  $p=0.61$ ).<sup>28</sup>

---

<sup>26</sup> A specification of the probit model on Bayesian behavior was estimated that included the *othbys* variable (detailed results omitted); *othbys* was not significant in this model, indicating that the quality of others' decisions affects the decisions of market professionals only in the relevant case where cascade formation is possible.

<sup>27</sup> Support for the significant differences between subject pools found in the parametric results is also found in non-parametric (Mann-Whitney) tests.

<sup>28</sup> p-values reported for the treatment level test found in Table 8 use observations aggregated at the individual participant level. As reported in Table 8, rates of cascade formation and reverse cascade formation do not differ over the domain of earnings.

Referring to Table 5, we find that the probit model yields results that are consistent with findings from the non-parametric tests. In the pooled data, the dummy variable *gain* is not significant at conventional levels, and remains so for the market professionals. For students, however, the parameter estimate is both significant ( $p = 0.04$ ) and negative ( $\beta = -0.289$ ), indicating that Bayesian behavior is increased in the loss domain. This result is consistent with the notion that, for the student population, losses loom larger than gains.

Students, however, are less affected by the earnings domain in later experimental periods. We investigate this insight by estimating separate probit models for the early and late periods in each session. In periods one through ten, students are roughly 8 percent less Bayesian in the gain treatments and the difference is statistically significant ( $p = 0.02$ ). In periods eleven through sixteen there is no significant difference in the rate of Bayesian behavior across the domains ( $p = 0.94$ ).

*Result 4. For market professionals, the reported intensity of their professional trading activity correlates positively with Bayesian decision making.*

Recall that we found no statistical differences among CBOT personnel with regard to the rate of Bayesian decision making or cascade formation, regardless of their exchange-trading role (broker, local, clerk, etc.). In this section we use additional survey data collected at the time of the experiments to determine if the underlying cause of Bayesian behavior differs across the trading subgroups. We find that trading intensity, defined as the number of contracts traded per day, is positively correlated with Bayesian behavior, although the effect is small.

The result is again based on a random effects probit specification and is reported in Table 9. We find that trading intensity (*intensity*) along with the counting rule sequences (*heurist*) have explanatory power. Other individual characteristics that included income and education are not significant for this group (results excluded to conserve space), nor are the treatment variables or

the posterior probability as represented by *diff*. By contrast, the trading support personnel (e.g., clerks) that supplied information on the intensity of their work yield results more similar to the student population. As with the students, *diff*, *heurist*, and *sym* are all significantly different than zero, although as with the traders, *intensity* is also significant, but small in magnitude. Our link between trading intensity and Bayesian rationality is consonant with the results of Locke and Mann (2003), Genesove and Mayer (2001), and List (2003, 2004), who find, in much different environments, that market experience is important in reducing deviations from utility theory that are associated with reference-dependent preferences.

*Result 5. Deviations from Bayesian norms are strongest when a counting rule makes a different prediction than Bayesian rationality.*

Table 3b shows that Bayesian decision making was dramatically reduced in the asymmetric treatment, with only 70 percent of decisions consistent with Bayes' rule. In the symmetric treatment, 91 percent of the decisions were Bayesian. In addition, cascade formation, when measured as the proportion of possible cascades, was significantly different across the symmetric and asymmetric urn types. These figures are presented in Table 3b in the column *Realized Cascades*, which shows that 81 percent of potential cascades were realized in the symmetric treatment but only 48 percent in the asymmetric case.<sup>29</sup>

In the discussion of some of the preceding results we have restricted attention to the asymmetric treatments. This was due to the fact that the high level of behavior consistent with Bayesian rationality in the symmetric treatments, 91 percent in aggregate, yielded little variability in the data. In the symmetric case, however, Bayesian rationality was consistent with following a simple counting rule: by choosing the urn with the most signals one followed Bayes'

---

<sup>29</sup> Because of the dependence among individual observations we use session level aggregates whenever possible to yield the most conservative estimates of treatment effects. The statistical test used is the Mann-Whitney U test. Significant differences in Bayesian behavior exist at  $p < .001$  across urn types.

rule. Alternatively, in the asymmetric treatment there are four sequences of draws in which the number of  $a$  signals is greater than the number of  $b$  signals but the Bayesian posterior implies urn B is the most likely state.

The counting heuristic suggests that individuals will choose the urn with the most signals, not necessarily that in which the posterior probability is greatest. This conjecture is investigated in Table 10, which presents all the available choice patterns in the asymmetric treatments. Those in which the counting rule and posteriors yield different predictions are highlighted in bold, and we call these *counting rule* sequences. The top number in the table represents the posterior probability and the bottom figure is the proportion of Bayesian decisions for each choice history. Statistical tests confirm what a visual scan of the data suggests: Bayesian behavior is significantly reduced in the *counting rule* sequences.<sup>30</sup> In fact, the four *counting rule* sequences have lower rates of Bayesian decision making than any of the other sequences, despite the fact that others have similar or smaller *diff* values. In aggregate, Bayesian behavior occurs at a rate of 41 percent in the “counting rule” choice histories and 81 percent in the remaining asymmetric treatments. This finding suggests that Bayesian updating is a poor approximation of the decision-making process used in the asymmetric treatment. Figure 3 illustrates this finding by presenting the proportion of Bayesian decisions for all observed histories of play as a function of the posterior probability. The counting rule sequences (lighter entries) are uniformly lower than the other choice histories, represented as black diamonds.

---

<sup>30</sup> We use a Wilcoxon matched pairs test with the proportion of Bayesian decisions the variable of interest, aggregated at the session level. The *diff* variable for the *counting rule* sequences are in the range from 0.0 to 0.2, and all other sequences with *diff* variables in this range are included for the paired comparison. Using data from the twelve asymmetric sessions we find that the *counting rule* sequences reflect less Bayesian decision making despite roughly equivalent *diff* scores at  $p < .01$ .

#### 4. Concluding Comments

The potential for herding behavior arising from informational cascades has been discussed in a wide variety of economic and social settings ranging from medical treatment choices to mate selection (Bikhchandani et al. 1998). Of particular interest is the potential that herding behavior has to destabilize financial markets. In this study we introduce market professionals from the CBOT floor to a controlled experimental environment in which cascade formation can be carefully studied. All previous empirical investigations of cascade behavior either examine data from naturally occurring markets where the researcher has no control over the data-generation process or use data from controlled laboratory studies with student subjects.

Our combination provides several interesting insights. Perhaps most importantly, we find that market professionals and student subjects behave differently in important and systematic ways: behavioral differences are found over the gain/loss domain and the rates of cascade formation. In general, our data suggest that traders do a better job of taking variability in the quality of others' decisions into account when choosing to rely on the information disclosed by others' actions. As a result, the market professionals end up in significantly fewer reverse cascades than the student population. In the CBOT group, we also find evidence that Bayesian behavior increases with an individual's trading activity. Overall, our findings highlight that there is much potentially useful behavioral economic and financial research to be done by implementing experimental protocols with non-standard subject pools.

## References

Anderson, L. R. Information Cascades: University of Virginia; 1994.

Anderson, L. R. and C. A. Holt. Information Cascades in the Laboratory. *The American Economic Review*. 1997; 87(5):847-62.

Banerjee, A. V. A Simple Model of Herd Behavior. *Quarterly Journal of Economics*. 1992; 107: 797-818.

Barberis, N, A. Scleifer and R. Vishny. A Model of Investor Sentiment. *Journal of Financial Economics*. 1998; 49: 307.343.

Berndt, Ernst, Bronwyn Hall, Robert Hall and Jerry Hausman, "Estimation and inference in nonlinear structural models," *Annals of Economic and Social Measurement*, (1974) 3: 653-65.

Bikhchandani, S.; D. Hirshleifer, and I. Welch. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*. 1998; 12(3): 151-170.

? A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*. 1992; 100(5): 992-1026.

Bikhchandani, S. and S. Sharma. Herd Behavior in Financial Markets: A Review. IMF Working Paper 48. 2000.

Butler, J.S. and R. Moffitt. A Computationally Efficient Quadrature Procedure for the One-Factor Multinomial Probit Model. *Econometrica*. 1982; 50(3): 761-64.

Daniel, K., D. Hirshleifer and A. Subrahmanyam. Overconfidence, Arbitrage, and Equilibrium Asset Pricing. *Journal of Finance*. 2001, 56: 921 – 65.

Eichengreen, B.; D. Mathieson; B. Chandha; A. Jansen; L. Kodres, and S. Sharma. Hedge Funds and Financial Market Dynamics. IMF Occasional Paper . 1998; 172.

Fama, E., Market Efficiency, Long-term Returns, and Behavioral Finance. *Journal of Financial Economics*. 1998; 49: 283-306.

Fox, C. R.; B.A. Rogers, and A. Tversky. Options Traders Exhibit Subadditive Decision Weights. *Journal of Risk and Uncertainty*. 1996; 13(1):5-17.

Genesove, D. and C. Mayer. Loss Aversion and Seller Behavior: Evidence from the Housing Market. *Quarterly Journal of Economics*. 2001; 116(4):1233-1260.

Gigerenzer, G. and D.J. Murray. Cognition as Intuitive Statistics. Hillsdale, NJ: Erlbaum (1997).

Holt, C.A. and A. Villamil. A Laboratory Experiment with a Single-Person Cobweb. *Atlantic Economic Journal*. 1986; 14(2):51-54.

Hung, A.A. and C. Plott. Information Cascades: Replication and an Extension to Majority Rule and **Conformity**-Rewarding Institutions. *American Economic Review*. 2001; 91(5):1508-1520.

Kahneman, D. and A. Tversky. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*. 1979; 47(2):263-291.

**Koessler, F., and A. Ziegelmeyer. "Tie-Breaking Rules and Informational Cascades: A Note." University Louis Pasteur Working Paper (2000).**

Kraemer, C. M. Noth and M. Weber. Information Aggregation with Costly Information and Random Ordering: Experimental Evidence. Manuscript. 2000.

Kubler, D. and G. Weizsacker. Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. Manuscript. 2002.

Kuran, T. and Sunstein, C. Availability Cascades and Risk Regulation. *Stanford Law Review*. 1999; 51(4):683-768.

List, J.A. Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace. *Econometrica* (2004), forthcoming.

List, J. A. Does Market Experience Eliminate Market Anomalies? *Quarterly Journal of Economics*. 2003; 118(1):41-71.

Locke, P. R. and S.C. Mann. Professional Trader Discipline and Trade Disposition. Working Paper, Texas Christian University. 2003.

MacKay, C. Extraordinary Popular Delusions and the Madness of Crowds. J. Wiley & Sons, New York, 1996.

Odean, T. Are Investors Reluctant to Realize their Losses? *Journal of Finance*. 1998; 53(5):1775-1798.

McKelvey, R. and T.R. Palfrey, Quantal Response Equilibria for Normal Form Games, *Games and Economic Behavior*, 1995; 10(1): 6-38.

Sarin, R. K. and Weber, M. Effects of Ambiguity in Market Experiments. *Management Science*. 1993; 39(5):602-615.

Shefrin, H. M. and M. Statman. The Disposition to Sell Winners too Early and Ride Losers too Long. *Journal of Finance*. 1985; 40:777-790.

Tversky, A. and D. Kahneman. Advances in Prospect Theory: Cumulative Representations of Uncertainty. *Journal of Risk and Uncertainty*. 1992; 5(4):297-323.

Welch, I. Sequential Sales, Learning, and Cascades. *Journal of Finance*. 1992; 47(2):695-732.

## Appendix: Experimental Instructions for Symmetric, Gain, Student (SGS) Treatment

Instructions:

In this experiment, you will be asked to decide from which of two urns balls are being drawn. We will begin by rolling a six-sided die. If the die roll yields a 1,2, or 3, we will draw from Urn A. If the roll of the die yields a 4,5, or 6, we will draw from Urn B. However, the roll of the die will be done behind a screen so that you will not know which urn has been chosen.

The urns differ in the following way:

Urn A (used if die is 1,2,or 3)	Urn B (used if die is 4,5,or 6)
2 Striped Balls	1 Striped Ball
1 Terp Ball	2 Terp Balls

Once an urn is determined by the roll of the die we will empty the contents of that urn into a container. (The container is always the same, regardless of which urn is being used.)

After the urn has been chosen each of you will come behind the screen one at a time and draw a ball from the container. The order in which you will draw has been determined randomly.

The result of your draw is your private information and **MUST NOT** be shared with other participants.

*After each draw, we will return the ball to the container before making the next private draw.* Each person will have one private draw, with the ball being replaced after each draw.

After each person has seen the results of their own draw, we will ask them to record the letter of the urn (A or B) that they think is more likely to have been used. When the first person to draw has indicated a letter, we will display that letter. After displaying the first person's decision, we will call out the next registration number, and the person with that number will draw a ball and record a letter (A or B). Again, their decision will be displayed on the overhead projector. This process will be repeated until everyone has made a draw and made a decision about which urn they believe is being used. After everyone has made a decision, the monitor will announce which of the urns was actually used. Everyone who chose the correct urn earns \$1. All others earn nothing.

This session will consist of 15 periods of the procedure just described.

Now I will describe the use of the record sheet, which is at the back of these instructions.

The results for each period are recorded on a separate row on the record sheet. Period numbers are listed on the left side of each row. Next to the period number record your draw (S or T) in column "**Own Draw**". In columns "**choice1**" through "**choice10**" record each participants decisions (A or B) as they are displayed. (If there are less than 10 players the last choice columns remain blank.) This means that when you are asked to make a decision about which Urn is being used the decisions of participants who have

drawn before you will be available. Write your decision in the appropriate column depending on the order in which you draw, and *circle your decision to distinguish it from other's decisions*.

When all participants have made their choices, the monitor will announce the letter of the Urn that was actually used. Record this letter in the column headed “**Urn**” for that period. If your circled decision matches the letter of the urn used, record your earnings of \$1 in the “**Payoff**” column. If your choice does not match the urn used record your earnings of \$0. You should keep track of your cumulative earnings in column “**Total Payoff**”.

Before we begin we will conduct a demonstration. During the demonstration, the roll of the die and the draw of the ball from the container will be publicly visible. When we move to Period 1 the roll of the die will be visible only to the monitor, and the draw will be visible only to the monitor and the person called behind the screen. Remember that urn A contains 2 striped balls and 1 terp ball. It is used if the throw of the die is 1,2, or 3. Urn B contains 1 striped ball and 2 terp balls, and is used if the throw of the die is 4,5, or 6.

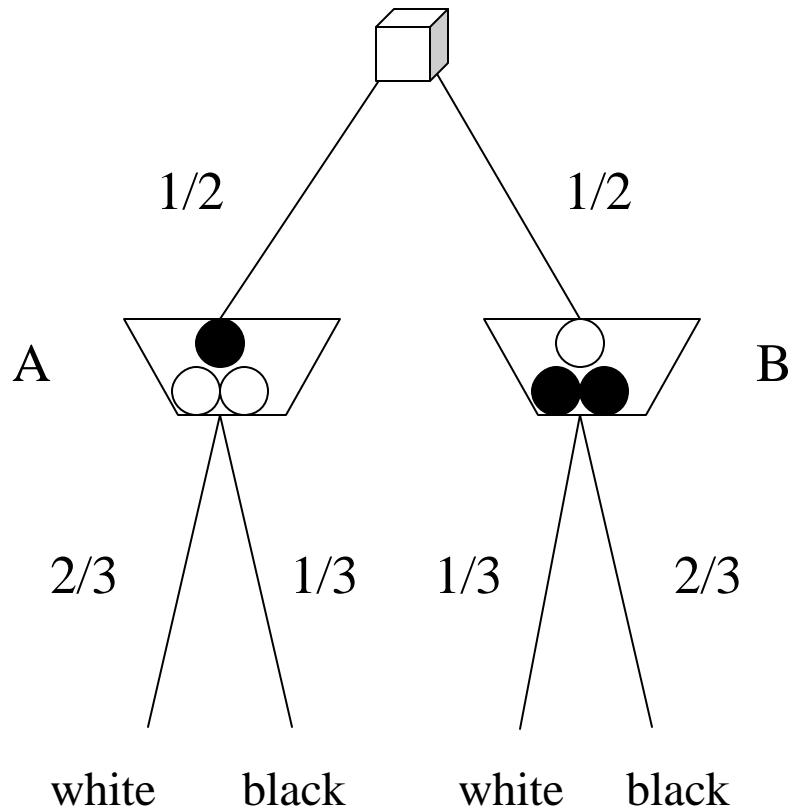
***Before we begin, be sure that your registration number is on the Record Sheet.***

Here is an overview of the procedure that will be followed in each period:

1. The monitor rolls the die to determine which urn is used and transfers balls from that urn to the container.
2. The monitor calls on a participant.
3. The participant goes behind the screen: (Be sure to bring your record sheet)
  - a) Makes a draw from the container
  - b) records the draw on record sheet in “**Own Draw**” column
  - c) records urn choice and *circles their choice*
4. The monitor displays the participant’s choice and the other participants record the urn choice on their record form.
5. Repeat steps 2 – 4 until all participants have made their choice.
6. The monitor reveals the urn used in that period by displaying the balls in the container.
7. Subjects record the urn used in that period and record their earnings, and their cumulative earnings.

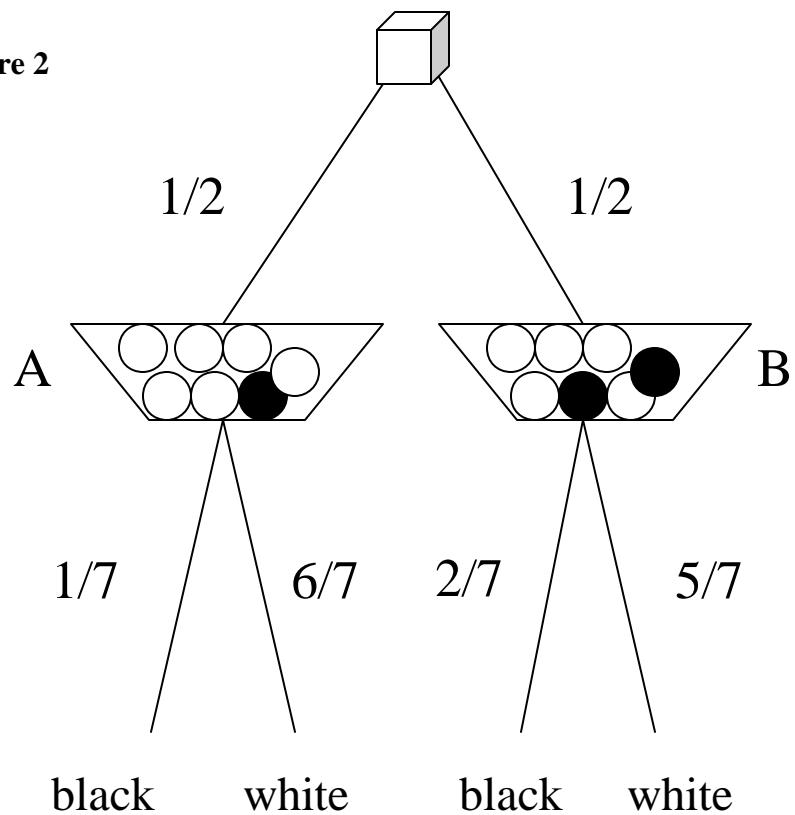
***Please refrain from conversation during all periods of play, and keep the information on your record sheet confidential.***

**Figure 1**



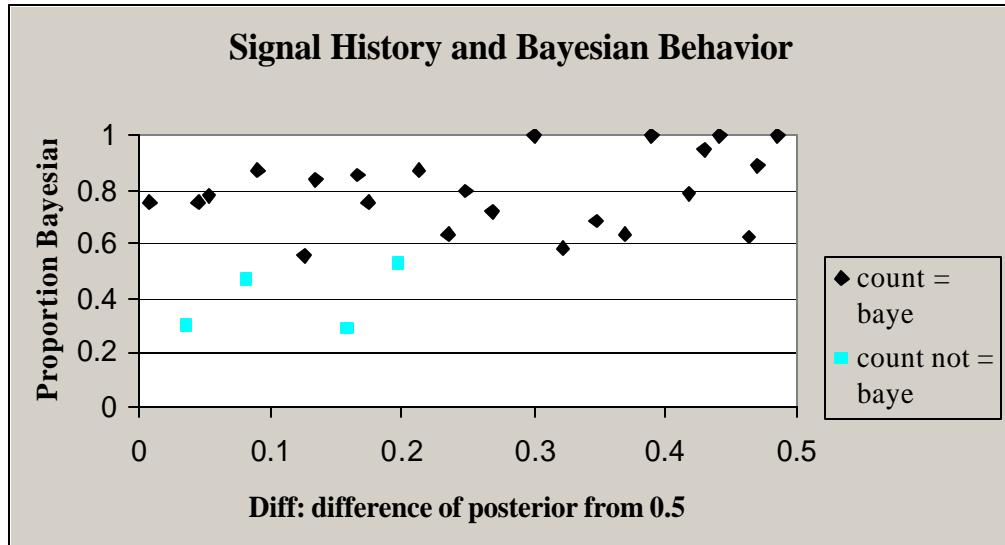
Symmetric Treatment

**Figure 2**



Asymmetric Treatment

**Figure 3: Counting Rule Heuristic**



The proportion of Bayesian decisions for every realized posterior probability is presented as a data point. The choice histories in which the counting rule and Bayesian posterior yield different predictions are presented as lighter squares. All other sequences are presented as black diamonds. Note that the sequences in which Bayesian behavior and the counting rule heuristic make different predictions have a uniformly lower proportion of Bayesian decisions than the others.

**Table 1. Experimental Design****Panel A: Ten Market Professional Sessions**

	Symmetric Urn		Asymmetric Urn	
	Gains	Losses	Gains	Losses
Number of Sessions	3	1	3	3
Participants in Session	5	5	One with 5, two with 6	6
Total number of Decisions	225	75	255	270
Average Earnings of Participants	\$43.20	-\$20.80	\$39.06	-\$22.89

**Panel B: Ten Student Sessions**

	Symmetric Urn		Asymmetric Urn	
	Gains	Losses	Gains	Losses
Number of Sessions	3	1	3	3
Participants in Session	One with 5, two with 6	5	One with 5, two with 6	5
Total number of Decisions	267	75	255	225
Average Earnings of Participants	\$11.61	-\$2.80	\$11.00	-\$6.40

Panel A (B) shows that Market Professionals (Students) were exposed to either the Symmetric or Asymmetric urn and played the game in either a gain or loss domain. The symmetric urn consisted of 3 balls — two *a* and one *b* in Urn A, and one *b* and two *a* in Urn B. The Asymmetric urn consisted of 7 balls — six *a* and one *b* in Urn A, and five *a* and two *b* in Urn B. The number of decisions is a function of the number of players, the number of games, and the number of periods in each game.

**Table 2. Posterior Probabilities: Symmetric (upper) and Asymmetric (lower, and *italic*) Urns**

a b	0	1	2	3	4	5	6
0	0.500 0.500	0.330 0.333	0.200 0.200	0.110 0.111	0.060 0.059	0.030 0.030	0.020 0.015
1	0.670 0.545	0.500 0.375	0.330 0.231	0.200 0.130	0.110 0.070	0.060 0.036	
2	0.800 0.590	0.670 0.419	0.500 0.265	0.330 0.153	0.200 0.083		
3	0.890 0.633	0.800 0.464	0.670 0.302	0.500 0.178			
4	0.940 0.633	0.890 0.509	0.800 0.341				
5	0.970 0.713	0.940 0.554					
6	0.980 0.749						

Entries represent the posterior probabilities for all possible sequences of draws for both symmetric (upper) and asymmetric (lower, and *italic*) treatments based on choice histories (a, b). The prior probability of an urn is 0.5 in (0,0). Underlined entries in the asymmetric urn are those consistent with counting heuristic sequences that will yield different predictions from the posterior prior probability.

**Table 3a. Aggregate Decision Making**

Panel A: Aggregate Decision Making (Combining Market Professional (M) and College Student (C) Treatments)						
C and M n = 1647		Bayesian	Cascades (total)	Reverse Cascades	Potential Cascades	Realized Cascades
	Proportion Number	0.79 1293	0.16 262	0.04 62	0.28 466	0.56 262/466

The *Bayesian* column represents the total number of decisions (and proportion) that were consistent with Bayesian updating. *Cascade* decisions (those which are Bayesian but private information ignored) and *reverse* cascades (same as cascades but the wrong inference of the underlying state takes place) occupy the next two columns. The *potential* cascades category represents the proportion (and number) of cascades that could have occurred when it was possible to make one, and the *realized* cascades category represents the proportion of those potential cascades that were actually realized. “n” = number of decisions.

**Table 3b. Disaggregated Decision Making across Treatments**

1. College Student Treatments (C)						
		Bayesian	Cascades (total)	Reverse Cascades	Potential Cascades	Realized Cascades
C n = 822	Proportion Number	.82 670	.18 150	.05 39	.30 247	.61 150/247
SGC n = 267	Proportion Number	.95 254	.17 46	.05 13	.19 51	.90 46/51
SLC n = 75	Proportion Number	.96 72	.07 6	.01 1	.08 6	.83 5/6
AGC n = 255	Proportion Number	.68 172	.25 64	.05 13	.46 117	.55 64/117
ALC n = 225	Proportion Number	.78 172	.16 35	.05 12	.32 73	.48 35/73
2. Market Professional Treatments (M)						
M n = 825	Proportion Number	.76 623	.14 112	.03 23	.27 219	.51 112/219
SGM n = 225	Proportion Number	.86 193	.14 32	.04 10	.19 43	.75 32/43
SLM n = 75	Proportion Number	.89 67	.17 13	.07 5	.24 18	.72 13/18
AGM n = 255	Proportion Number	.71 180	.13 34	.01 2	.28 72	.47 34/72
ALM n = 270	Proportion Number	.68 183	.12 33	.02 6	.32 86	.38 33/36

The *Bayesian* column represents the total number of decisions (and proportion) that were consistent with Bayesian updating. *Cascade* decisions (those which are Bayesian but private information ignored) and *reverse* cascades (same as cascades but the wrong inference of the underlying state takes place) occupy the next two columns. The *potential* cascades category represents the proportion (and number) of cascades that could have occurred when it was possible to make one, and the *realized* cascades category represents the proportion of those potential cascades that were actually realized. “n” = number of decisions. Treatment codes are S = symmetric, A = asymmetric, G = gain, L = loss, C = college student, M = market professional.

**Table 4. Bayesian Decisions: Probit Model**

Dependent variable: baye	4a: Pooled Model (combining market professionals and students) n = 1647			4b. Student Model n = 822			4c. Market Professionals Model n = 825		
Ind. Variables:	Coefficient	z stat	P> z	Coefficient	z stat	P> z	Coefficient	z stat	P> z
Diff	1.172	3.260	0.001	1.406	2.440	0.015	1.043	2.230	0.026
Heurist	-0.840	-6.580	0.000	-0.793	-3.970	0.000	-0.888	-5.330	0.000
Gain	-0.121	-1.180	0.237	-0.289	-2.050	0.040	0.046	0.330	0.740
Sym	0.515	4.600	0.000	0.902	5.820	0.000	0.190	1.240	0.215
Trader	-0.201	-2.050	0.040	-	-	-	-	-	-
order_2	-0.065	-0.510	0.608	0.125	0.690	0.491	-0.249	-1.410	0.157
order_3	-0.070	-0.510	0.608	0.204	1.000	0.316	-0.314	-1.690	0.091
order_4	-0.384	-2.880	0.004	-0.187	-0.940	0.347	-0.591	-3.220	0.001
order_5	-0.159	-1.130	0.258	0.032	0.150	0.882	-0.354	-1.870	0.062
order_6	-0.364	-2.080	0.037	-0.431	-1.560	0.118	-0.432	-1.870	0.062
	Log Likelihood: -774.054, Wald $\mathbf{c}_{(10)}^2$ = 137.79, Prob > $\mathbf{c}_{(10)}^2$ = 0.000			Log Likelihood: -339.236, Wald $\mathbf{c}_{(10)}^2$ = 84.33, Prob > $\mathbf{c}_{(10)}^2$ = 0.000			Log Likelihood: -427.283, Wald $\mathbf{c}_{(10)}^2$ = 68.91, Prob > $\mathbf{c}_{(10)}^2$ = 0.000		

The dichotomous dependent variable in all three probit models (pooled, student, and market professional) is coded one for a decision consistent with the Bayesian posterior and zero otherwise. Independent variables include *diff*, which is  $|prob(urn = A) - .5|$ , where the *prob(urn = A)* is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain*, *sym*, and *trader* (in the case of the pooled model) are dichotomous and distinguish the treatments. *Heurist* is a dummy variable equal to one for the counting rule sequences and zero for all others. *order\_x* (where  $x=2,..6$ ) is a categorical variable indicating where in the period of play the decision was made. The Wald statistic tests the null hypothesis that all coefficients are zero. A Chow test testing the pooled versus segregated models yields  $\mathbf{c}_9^2 = 20.98$ , which suggests significant differences between subject pools at the  $p = .0127$  level.

**Table 5. Quantal Response Equilibrium Parameter Estimates and Standard Errors: Symmetric Gain Treatment**

Subject Pool	$b_1$	$b_2$	$b_3$	$b_4$	$b_5$	$b_6$
Market Professionals	4.594* (1.170)	4.561** (1.271)	8.670 (2.818)	3.898 (1.099)	2.481*** (0.784)	
College Students	7.125 (1.568)	27.754 (10.018)	8.622 (2.425)	4.994 (1.305)	6.228 (1.773)	1.722 (0.675)

For each decision order the parameter estimates (standard errors are in parentheses) are presented for both subject pools. \*, \*\*, and \*\*\* represent differences between the subject pools at 10 percent, 5 percent, and 1 percent levels. All parameter estimates are significantly different from zero individually.

**Table 6. Posterior Probabilities and Choice Probabilities with QRE Decision Error for Both Market Professionals (M) and College Students (C)**

		Posterior			pr(D1=A)			Decision					
			M	C		M	C	M			C		
Choice Order	Signal History	Bayes	QRE		Bayes	QRE		A	B	Proportion A	A	B	Proportion A
1	A	0.67	0.667	0.667	1.00	0.822	0.915	37	8	0.822	43	4	0.822
2	Aa	0.80	0.757	0.779	1.00	0.912	0.999	23	3	0.885	27	0	1.000
2	Ab	0.50	0.436	0.469	0.50	0.358	0.150	4	15	0.211	3	17	0.150
3	AAa	0.89	0.806	0.854	1.00	0.995	0.998	14	0	1.000	17	1	0.944
3	AAb	0.67	0.509	0.593	1.00	0.539	0.833	7	6	0.538	12	0	1.000
3	ABA	0.67	0.649	0.638	1.00	0.929	0.916	10	0	1.000	8	0	1.000
3	ABB	0.33	0.316	0.306	0.00	0.039	0.034	1	7	0.125	0	9	0.000

Entries are the first three choices of the symmetric gain treatment, the posterior probability ( $\text{Pr}(\text{urn}=A)$ ), and decision probability ( $\text{Pr}(D_i=A)$ ) for each choice history that reflects the adjustment for decision error. For comparison, the Bayesian (Bayes) posteriors and choice probabilities are also presented for both students and market professionals. The actual decisions are also reported. Due to the symmetry of the treatments, urn A and B are treated equally.

**Table 7. Cascade Formation: Probit Model: Asymmetric Treatment**

Dependent variable: cascade	7a: Pooled Model (combining market professionals and students) n = 329			7b. Student Model n = 180			7c. Market Professionals Model n = 149		
Ind. Variables:	Coefficient	z stat	P> z	Coefficient	z stat	P> z	Coefficient	z stat	P> z
Diff	2.723	2.48	0.013	0.895	0.56	0.575	5.143	2.86	0.004
Heurist	-0.595	-2.13	0.034	-0.865	-2.13	0.033	-0.244	-0.56	0.576
Othbys	-0.503	-0.8	0.426	-2.835	-2.94	0.003	2.070	1.93	0.054
Gain	0.378	1.62	0.106	-0.030	-0.08	0.934	0.443	1.29	0.197
Trader	-0.044	-0.19	0.849	-	-	-	-	-	-
order_2	0.293	0.92	0.357	0.237	0.55	0.58	0.507	0.95	0.344
order_3	-0.098	-0.32	0.749	-0.081	-0.19	0.846	-0.003	-0.01	0.995
order_4	0.394	1.12	0.264	0.475	0.97	0.332	0.313	0.53	0.594
order_5	-0.471	-1.11	0.268	-0.372	-0.56	0.579	-0.450	-0.72	0.469
	Log Likelihood: -199.02, Wald = $\mathbf{c}_{(5)}^2$ 38.11, Prob > $\mathbf{c}_{(5)}^2$ = 0.000			Log Likelihood: -108.94, Wald = $\mathbf{c}_{(5)}^2$ 21.10, Prob > $\mathbf{c}_{(5)}^2$ = 0.007			Log Likelihood: -81.40, Wald = $\mathbf{c}_{(5)}^2$ 25.33, Prob > $\mathbf{c}_{(5)}^2$ = 0.001		

The dichotomous dependent variable in all three probit models (pooled, student, and market professional) is coded one for a cascade decision and zero otherwise. Independent variables include *diff*, which is  $|prob(urn = A) - .5|$ , where the *prob(urn = A)* is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain* and *trader* (in the case of the pooled model) are dichotomous and distinguish the treatment/subject type. *Othbys* indicates the aggregate performance of other subjects in terms of the proportion of Bayesian decisions they have made. *Heurist* is a dummy variable equal to one for the counting rule sequences and zero for all others. *Order\_x* is a categorical variable indicating where in the period of play the decision was made. Note: Because the *othbys* variable is not applicable for those first in order (they do not observe others' decisions), this is excluded. *Order\_6* is therefore excluded automatically to avoid problems associated with collinearity. The Wald statistic tests the null hypothesis that all coefficients are zero.

**Table 8. Non-Parametric Tests of Bayesian Behavior and Cascade Formation at the Individual Treatment Level**

Treatment	AGC	ALC	SGC	SLC	AGM	ALM	SGM	SLM
<b>AGC</b> <i>n = 17</i>	—	<b>0.085</b> 0.042 0.422	<b>0.000</b> 0.035 0.924	NA	<b>0.478</b> 0.008 0.028	NA	NA	NA
<b>ALC</b> <i>n = 15</i>		—	NA	<b>0.003</b> 0.120 0.114	NA	<b>0.106</b> 0.377 0.057	NA	NA
<b>SGC</b> <i>n = 17</i>			—	<b>0.898</b> 0.155 0.291	NA	NA	<b>0.005</b> 0.984 0.558	NA
<b>SLC</b> <i>n = 5</i>				—	NA	NA	NA	<b>0.154</b> 0.049 0.065
<b>AGM</b> <i>n = 17</i>					—	<b>0.607</b> 0.701 0.163	<b>0.004</b> 0.909 0.002	NA
<b>ALM</b> <i>n = 18</i>						—	NA	<b>0.008</b> 0.270 0.039
<b>SGM</b> <i>n = 15</i>							—	<b>0.476</b> 0.675 0.327
<b>SLM</b> <i>n = 5</i>								—

Treatment codes: First letter S or A indicates the *symmetric* or *asymmetric* treatment. Second letter is G or L for the domain of *gains* or *losses*. Third letter indicates subject type, where C = *college students* and M = *market professionals*. The top number (bold) represents the p-value associated with differences in Bayesian decisions across treatments. The middle (italic) number reports p-values for cascade formation. The lower number represents the p-values for reverse cascades. All are tested with the Mann-Whitney U test with an individual's aggregate choice being the observation level.

**Table 9. Trading Intensity**

Dependent variable: baye	9a: Trader subset of CBOT Market Professionals n = 255			9b. Non-trader subset of CBOT Market Professionals n = 450		
Ind. Variables:	Coefficient	z stat	P> z	Coefficient	z stat	P> z
Diff	0.118	-1.668	0.897	1.648	2.470	0.013
Heurist	-1.526	-2.194	0.000	-0.801	-3.900	0.000
Sym	0.099	-0.524	0.755	0.348	1.920	0.055
Gain	0.246	-0.427	0.474	0.005	0.030	0.976
Intensity	0.011	0.002	0.011	0.001	-1.920	0.055
order_2	0.273	-0.382	0.414	-0.284	-1.140	0.254
order_3	0.073	-0.587	0.829	-0.305	-1.150	0.252
order_4	-0.326	-1.015	0.354	-0.669	-2.680	0.007
order_5	0.144	-0.610	0.707	-0.390	-1.540	0.123
order_6	0.443	-0.514	0.364	-0.598	-1.990	0.047
	Log Likelihood: -103.50, Wald $\mathbf{c}_{(10)}^2 = 37.76$ , Prob $> \mathbf{c}_{(10)}^2 = 0.0002$			Log Likelihood: -215.88, Wald $\mathbf{c}_{(10)}^2 = 59.37$ , Prob $> \mathbf{c}_{(10)}^2 = 0.000$		

The dichotomous dependent variable in these two probit models is coded one for a decision consistent with the Bayesian posterior and zero otherwise. Independent variables include *diff*, which is  $|prob(urn = A) - .5|$ , where the *prob(urn = A)* is the posterior probability arising from the combination of public and private information at the disposal of each decision maker. The variables *gain* and *sym* (in the case of the pooled model) are dichotomous and distinguish the treatments. *Heurist* is a dummy variable equal to one for the counting rule sequences and zero for all others. *Order\_x* (where  $x=2,..6$ ) is a categorical variable indicating where in the period of play the decision was made. *Intensity* reflects the level of trading intensity among participants, measured as the number of contracts traded per day. The Wald statistic tests the null hypothesis that all coefficients are zero.

**Table 10. Posterior Probability Urn is A and Proportion of Bayesian Decisions (Counting Rule Sequences in Bold)**

a	b	0	1	2	3	4	5	6
0		0.33 <i>0.85</i>	0.20 <i>1.00</i>	0.11 <i>1.00</i>	0.06 <i>1.00</i>	0.03 <i>0.89</i>	0.02 <i>1.00</i>	
1		0.55 <i>0.76</i>	0.38 <i>0.56</i>	0.23 <i>0.72</i>	0.13 <i>0.64</i>	0.07 <i>0.95</i>	0.04 <i>0.63</i>	
2		0.59 <i>0.87</i>	<b>0.42</b> <i>0.46</i>	0.26 <i>0.63</i>	0.15 <i>0.69</i>	0.08 <i>0.79</i>		
3		0.63 <i>0.84</i>	<b>0.46</b> <i>0.30</i>	<b>0.30</b> <i>0.52</i>	0.18 <i>0.58</i>			
4		0.67 <i>0.76</i>	0.51 <i>0.76</i>	<b>0.34</b> <i>0.29</i>				
5		0.71 <i>0.87</i>	0.55 <i>0.78</i>					
6		0.75 <i>0.80</i>						

The number of A and B signals are given in the first row and column respectively. The pairs of numbers within an (a , b) pair represent the Bayesian posterior (upper number) and the proportion of Bayesian decisions (lower number and in *italics*). Those in bold type are the *counting heuristic* sequences. Thus (2,1) has a posterior probability of 42 percent that the urn is A (diff=0.08). Forty-six percent made the Bayesian decision in this case. By contrast the (2,0) sequence (in which diff=0.09) has a posterior probability of 0.59, and 87 percent of those decisions were Bayesian.