Herding and Cascades with Endogenous Price: An Experiment.

(Preliminary and Incomplete)

Christophe Bisière[1] Jean-Paul Décamps[2] Stefano Lovo[3]

Mai 2007

[1]Université de Toulouse 1
[2]Université de Toulouse 1 and Europlace Institute of Finance
[3]HEC School of Management

Financial support from Europlace Institute of Finance is gratefully acknowledged.

Abstract

This paper presents an experiment that simulates trading in financial market. The experiment design allows: i) to unambiguously identify herding, contrarian behavior and informational cascades phenomena; ii) to dissect the different components of subjects' decision process leading to trading decisions; iii) to explain how these phenomena and components are affected by the distribution of the asset fundamentals. Our data show that subjects are not risk neutral, they often do not update beliefs in a Bayesian way and display aversion to ambiguity. We show how this induces subjects to engage in herding and contrarian behavior therefore reducing market informational efficiency. An increase in the intrinsic uncertainty regarding the asset fundamentals magnifies these phenomena.

1 Introduction

One of the central roles of financial markets is to provide information about asset's fundamentals through the price system. This paper studies the capacity of the market to aggregate the pieces of information dispersed among market participants. We present the results of a series of experiments that simulate trading in financial markets where market participants are asymmetrically informed on asset fundamentals. Our objective is to understand how the informational content of the order flow changes with traders' risk attitude, with the strength of prior beliefs on fundamentals and with the intrinsic risk of the traded asset. In an influential paper Avery and Zemsky [1] show that when trading prices are endogenous and market participants are risk neutral, none of the above mentioned factors should have an impact on informed traders' behavior. As a consequence phenomena such as herding behavior, contrarian behavior and informational cascades are impossible. Our experimental data show that market participant are not risk neutral, subjects engage in herd and contrarian behavior, and informational cascades are possible.

The main objectives of experimental papers simulating trades in financial markets are to provide useful tests to market microstructure theories and help understanding investors' behavior in actual financial markets. Market microstructure models are typically built on the assumptions that informed traders are perfectly rational Bayesian agents and that this is common knowledge among market participants. Many experimental tests focus on microstructure theories that introduce the additional assumption of risk neutrality of market participants. This is the case for example of two important papers by Cipriani and Guarino [5] (CG henceforth) and Drehman, Oechssler and Roider [9] (DOR henceforth) that supported Avery and Zemsky view. They show in laboratory experiments that herding behavior and informational cascades are observed when prices are exogenously fixed but tend to disappear when trading prices are flexible. However, they also observe a significant amount of "irrational behaviors" some of which are explained by lack of common knowledge of subjects' rationality.

Our experiment is based on the theoretical papers by Décamps and Lovo [7], [8] that explicitly take into account market participants' risk attitude. We adopt an experimental design that allows i) to dissect the different components of subjects' decision process leading to trading decisions; ii) to unambiguously identifies herding, contrarian behavior and informational cascades; iii) to explain how these components and phenomena are affected by the distribution of the asset fundamentals. Differently from CG and DOR, our experiment is designed so that other subjects' rationality plays no role for determining a subject's optimal action, thus the lack of common knowledge of participants' rationality cannot explain deviation from what the theory predicts. Beside, we depart from CG and DOR's experimental studies in three additional features. First, the liquidation value of the asset includes two components: one that can be learned by aggregating all private information, and an additional shock on which agents have no information. Thus, the asset remains intrinsically risky even when all private information is disclosed. This set-up allows us to test the hypothesis of risk neutrality that is often assumed in previous experimental papers. Second, we observe agents' contingent choices for all possible realizations of the private signals, allowing to clearly identify situations where subjects do ignore their private information. Third, we run our experiment in two different formats so that subjects have to solve homeomorphic problems that nevertheless require different amounts of mental reasoning. More precisely in the first format, *market experiment* (ME), similarly to what happens for CG and DOR, subjects choose the quantity they want to trade in a market for a risky asset. When making their trading decision, subjects have to use public and private available information to figure out the probability distribution of the portfolio resulting from their trades. In the format we denote *lottery experiment* (LE), subjects are asked to choose among different lotteries where the distribution function of payoff is explicitly given. These lotteries are determined so that a fully rational Bayesian agent faces exactly the same problem in the two formats ME and LE. Therefore, such an agent should have exactly the same behavior in the two formats. The observation of subjects' choices in LE allows to understand their risk attitude (risk neutrality, risk aversion or risk loving). The comparison between a subject's answers in the two formats allows to detect non-Bayesian belief updating as well as other attitudes such as ambiguity aversion.

Our observations in LE show that many of what at a first glance may appear as "irrational behaviors" are not so if one takes into account risk attitude. Most subjects display risk aversion, many are risk lover but very few are risk neutral. As theoretically shown in Décamps and Lovo [7], [8] and observed in our experiment, when prior beliefs are strong, risk averse and risk lover subjects tend to ignore their private information. Risk lover subjects tend to engage in contrarian behavior while risk aversion induces agents to choose not to trade. The latter phenomenon is magnified when the intrinsic uncertainty regarding asset fundamentals increases.

Overall our data on in LE suggests that subjects tend to neglect their private information and that this tendency is exacerbated by the strength of the prior belief and by the intrinsic uncertainty regarding the asset fundamentals. From the comparison between ME and LE we can reject the hypothesis that agents are Bayesian. We rather find evidence of ambiguity aversion. Also in ME we observe a sensitive increase in herd behavior that can be explained with the fact that subjects' tendency to round probabilities is stronger when they have to mentally compute posterior probability. On the other hand, the percentage of non-informative, no-trade decisions, is lower in ME for extreme public beliefs, but higher when little is known about fundamentals. Therefore, order flow will be more informative in ME than in LE for extreme public beliefs, while it will be the opposite when public belief is less strong.

Thanks to the observation of contingent strategies at different levels of public belief, we can predict subjects' behavior when intervening at any different point in an hypothetical trading history. Thus, we can measure market informational efficiency by simulating an arbitrary large number of trading histories without having to physically run the experiment with subjects. Simulations show that the effect on pricing error is ambiguous. LE provides more efficient prices than ME in presence of intrinsic uncertainty. However, ME is more efficient when no unlearnable shock affects fundamentals.

The remainder of the paper is organized as follows. Section 2 presents the simplest version of Décamps and Lovo's theory, and its implication on agent's behavior and cascades. Section 3 presents the experimental design. Section 4 presents the results of the experiment. Section 5 concludes.

2 Theoretical framework

We consider a specification of the discrete time sequential trade model studied in Décamps and Lovo [7]: a single asset is exchanged for money among market makers and traders. We denote by $\tilde{v} = \tilde{V} + \tilde{\epsilon}$ the fundamental value of the asset. The random variable $\tilde{V}$ represents a realized shock on which market participants are asymmetrically informed. The random variable $\tilde{\epsilon}$ represents other shocks on fundamentals (for example future shocks) whose realization is unknown to everybody. We assume that $\tilde{V}$ takes value in $\{\underline{V}, \overline{V}\}$, where $\underline{V} < \overline{V}$ and $\tilde{\epsilon}$ takes value in $\{-\epsilon, +\epsilon\}$ with $\mathbb{P}(\tilde{\epsilon} = \epsilon) = \mathbb{P}(\tilde{V} = \overline{V}) = \frac{1}{2}$ and $\epsilon = (\overline{V} - \underline{V})/2$.[1] Each trader observes a conditionally independent and identically distributed private signal $\tilde{s}$ with possible values l, h. We assume $\mathbb{P}(l|\underline{V}) = \mathbb{P}(h|\overline{V}) = p$, with $p \in (1/2, 1)$ that implies that private signals are partially informative regarding $\tilde{V}$, but provides no information regarding $\tilde{\epsilon}$. We restrict tradable quantities to $\{-1, 0, 1\}$ that correspond to sell one unit, no trade and buy one unit respectively.

Let H_t be the history of trades (past quantities and prices) up to date $t-1$. All agents observe H_t and update their beliefs according to Bayes' rule. We denote $\pi_t = \mathbb{P}\left(\overline{V}|H_t\right)$ the public belief at time t and $\pi_t^s = \mathbb{P}\left(\overline{V}|H_t, s\right)$, a trader's belief at time t given a private signal $s \in \{l, h\}$. At any given period t a trader comes to the market, submits a market order and leaves the market. The trader expects a per

[1]Setting $\epsilon = (\overline{V} - \underline{V})/2$ implies that $\tilde{v}$ can take only three values: $\{\left(3\underline{V} - \overline{V}\right), (\overline{V} + \overline{V})/2, \left(3\overline{V} - \underline{V}\right)/2\}$. This has the advantage of reducing the complexity of subjects' maximization problems in the experiment.

share pricing schedule $P_t(.) : \{-1, 0, 1\} \longrightarrow \mathbb{R}$ and demands $Q^\star$ defined as:

$$Q^\star(P_t, H_t, s) := \arg \max_{Q \in \{-1,0,1\}} \mathbb{E}\left[u\left(m + x\tilde{v} + (\tilde{v} - P_t(Q))Q\right)|H_t, s\right], \tag{1}$$

where $u : \mathbb{R} \longrightarrow \mathbb{R}$ satisfies u' is positive and continuous. Note that risk neutrality, risk aversion and risk loving are contemplated in our analysis as we impose no restriction on the sign of u''. The variable m and x represent the trader's initial monetary wealth and initial inventory in the risky asset respectively. We denote by $Q^\star(P_t, H_t, \emptyset)$ the demand of an hypothetical trader that receives no private signal.

Risk neutral market makers compete to fill the trader's order without knowing the trader's signal and price the asset efficiently:

$$P_t(Q) := \mathbb{E}[\tilde{v}|H_t, Q^\star(P_t, H_t, \tilde{s}) = Q]. \tag{2}$$

That is, the market clearing price is equal to expectation of $\tilde{v}$ conditional on the information provided by past and current trades. An equilibrium is a situation where equations (1) and (2) are satisfied at any time t.

In the long run, the market is informational efficient if all the information dispersed in the economy is eventually incorporated into market prices. Considering that in our set-up private information only regards $\tilde{V}$, that $\mathbb{E}[\tilde{\epsilon}] = 0$ and that trading prices satisfy equation (2), we have informational efficiency if the trading price eventually converges to the realization of $\tilde{V}$. Hence the formal definitions:

Definition 1 *The market is informational efficient, if the random variables* $(\mathbb{E}[\tilde{v}|H_t])_{t\geq 0}$ *converges to* $\tilde{V}$ *almost surely when* t *tends to infinity.*

Definition 2 *A trader's order is said to be non-informative when it is not affected by the private signal the trader received, formally* $Q^\star(P_t, H_t, h) = Q^\star(P_t, H_t, l)$.

Definition 3 *An informational cascade occurs when the orders of all traders in the economy are non-informative.*

The model presented in this section is a special case of the economy analyzed in Décamps and Lovo [7] whose central result can be summarized as follows. If market makers and traders differ in their risk aversion, then as soon as the past history of trade provides sufficiently strong, but not complete, information regarding the realization of $\tilde{V}$, the equilibrium is unique and such that all traders submit non-informative orders. In this instance an information cascade occurs, price stops evolving and remains bounded away from the realization of $\tilde{V}$. As a consequence, market informational efficiency is impossible even in the presence of endogenous prices.[2] The following section provides an informal proof of this result and we refer to Décamps and Lovo [7] for the complete proof in a general framework.

2.1 An illustration of the behavior of a Bayesian expected utility maximizer

We discuss here a trader theoretical behavior in a numerical example that reflects the set-up of our experiment. Let $\underline{V} = 4$, $\overline{V} = 12$, $\epsilon = 4$ and $p = 0.65$. In this instance the fundamental value of the asset can take three values, i.e., $\widetilde{v} \in \{0, 8, 16\}$. For any finite history H_t let denote with g the *signals imbalance*, i.e., the difference between the number of private signals h and of private signals l received by past traders. Suppose for example that a history H_t drives agents to think that $g = 8$. This induces a public belief $\pi_t = 0.9930$.[3] Now, consider a trader coming to the market endowed with 0 amount of the risky asset and 12 units of money. This trader has the opportunity, but not the obligation, to buy or sell one unit of the asset at price $P_t = \mathbb{E}[\widetilde{v}|H_t] = 11.94$. The problem a Bayesian expected utility maximizer faces can be summarized as choices among lotteries in Tables 1, 2 and 3. The entries in the Tables represent the net gain of the trader consequent to the trading decision and contingent on the three possible realizations of the fundamental value $\widetilde{v}$, i.e., $tradedquantity * (\widetilde{v} - tradingprice) + 12$. In each table, the lines report the possible payoffs resulting to different trading decisions. The three tables only differ for the probabilities of getting the related payoffs. Table 1 represents the problem faced by a trader who received no private signal.

	0.35%	50.00%	49.65%	Expected value
Sell order	23.94	15.94	7.94	12
No trade	12.00	12.00	12.00	12
Buy order	0.06	8.06	16.06	12

Table 1: Lotteries when no signal is received

By definition of risk aversion, a risk averse trader will prefer the certain payment to the lottery, that is to say, "No trade" is the strictly preferred action. A risk-lover trader will typically strictly prefer

[2] A formal statement of this result is: there exist $\underline{\pi}$ and $\overline{\pi}$, with $0 < \underline{\pi} \leq \overline{\pi} < 1$ such that as soon as the public belief $\pi_t \notin (\underline{\pi}, \overline{\pi})$ then the equilibrium is unique and such that for all $s \geq t$: 1) An informational cascade occurs; 2) the equilibrium price schedule is $P_s(Q) = \mathbb{E}[\tilde{v}|H_t]$.

[3] From $\mathbb{P}(h|\overline{V}) = \mathbb{P}(l|\underline{V})$, it results that two histories lead to different public belief $\pi_\tau$and π'_τ only if $g \neq g'$.

either buying or selling to the other alternatives, whereas a risk neutral trader will be perfectly indifferent between the three actions.

Consider now the same trader but suppose he or she received a private signal $\widetilde{s}$ with precision $p = 0.65$. Will the private signal affect the trader's order? Table 2 represents the problem faced by a trader who received a private signal l whereas Table 3 represents the problem faced by a trader who received a private signal h. In other words, probabilities in tables 2 and 3 are obtained by the (Bayesian) updating of the public belief $\pi_t = 0.9930$ following private signal l and h respectively:

	0.65%	50.00%	49.35%	Expected value
Sell order	23.94	15.94	7.94	12.05
No trade	12.00	12.00	12.00	12.00
Buy order	0.06	8.06	16.06	11.95

Table 2: Lotteries when a signal l is received

	0.19%	50.00%	49.81%	Expected value
Sell order	23.94	15.94	7.94	11.97
No trade	12.00	12.00	12.00	12.00
Buy order	0.06	8.06	16.06	12.03

Table 3: Lotteries when a signal h is received

It is clear form the expected value column that a risk neutral trader will prefer to sell when $s = l$ and to buy when $s = h$. By contrast a sufficiently risk averse agent will still prefer "No trade" to the other options independently from the signal he received. Similarly a sufficiently risk lover trader will typically strictly prefer either buying or selling to the other alternatives but this choice will not be affected by the sign of the private signal. It is possible to rephrase the main result in Décamps and Lovo [7] as follows:

Proposition 1 *If before receiving a private signal a trader strictly prefers the action $Q^\star \in \{-1, 0, 1\}$ to the other two alternatives and π_t is sufficiently close to 0 or to 1, then after receiving any private signal $s \in \{l, h\}$ the trader will still prefer action $Q^\star$ to the alternative actions. Hence the trader's order is non-informative.*

Proof. First, note that as the public belief π_t approaches 0 or 1 the impact of a private signal $\widetilde{s}$ on an agent's belief becomes negligible. Formally, $\lim_{\pi_t \to 1} |\pi_t^s - \pi_t| = \lim_{\pi_t \to 0} |\pi_t^s - \pi_t| = 0$ for any

$s \in \{l, h\}$. Second, note that as $Q^\star$ is the unique element in $\{-1, 0, 1\}$ that solves equation (1) when $s = \emptyset$, it must be that for probabilities π_t^s close to π_t, action $Q^\star$ also solves equation (1) as the objective function is continuous in π_t^s. ■

Proposition 1 suggests that as soon as the public belief is sufficiently close to 0 or to 1, non-risk neutral agents will submit non-informative orders. More generally, the only traders that will always submit informative orders are those who ex-ante are indifferent among no-trading, buying or selling at the proposed price. In fact, only for these agents even an arbitrarily small change in belief induced by a private signal will affect the trading decision. If the probability of having such agents in the economy is null, then an informational cascade will eventually occur.

As an illustration of Proposition 1, consider Table 7. Columns labeled with LE are report an agent's contingent decision for different levels of g or, equivalently, for different level of public belief π, when the trading price is $P = \pi 12 + (1 - \pi)4$. The two letters in an entry represent the trading decision (Buy, Sell, or No-trade) after receiving signal l or signal h respectively. Different columns correspond to different level of risk aversion for a trader with CARA utility function. Risk averse agents will tend to hedge their positions. As soon as the public information is sufficiently strong, i.e., when $|g|$ is large, risk averse traders will prefer not to trade ignoring their private signal. Risk-lover agents will always prefer to trade but when the public information is sufficiently strong the sign of their private signal will have no effect on their trading decisions. Interestingly, it can be rational to adopt strategies different from the risk-neutral strategy S-B even in the framework studied in CG and DOR, i.e. when the additional risk linked to the $\widetilde{\epsilon}$ component is absent. Column LE in Table 8 presents the optimal strategies for different forms of utility function for $\epsilon = 0$.

Note that we assumed that agents can buy and sell at price $P(Q) = \mathbb{E}[\tilde{v}|H_t]$.[4] This simple pricing convention is of special interest for several reasons. First, this is the pricing rule adopted in the experiment, as it is clearly easier for subjects to determine their optimal demand when the unit price at which they can trade the asset does not depend on the sign of their own order. Second, this is the pricing rule adopted in CG and DOR. Thus by adopting the same rule we can better compare our results with theirs. Third, a quantity-sensitive price schedule, where the price at which a trader can buy is larger than the price at which the trader can sell, has the effect of reducing the incentive to trade and speculate on information. Thus, by fixing the price at $\mathbb{E}[\tilde{v}|H_t]$ for buy and sell orders, we increase the speculative component of traders' demand and this reduce the incentive to adopt non-informative orders and hence

[4]This is in contrast with the theory as the trading price should incorporate past and current information on fundamentals.

the possibility of cascade. In other words this pricing rule should bias our experiment in favor of market efficiency. Fourth, as shown in Décamps and Lovo [7], this will be the equilibrium price schedule as soon as a cascade occurs.

2.2 Testable implications

The previous discussion suggests several testable implications. First the use a traders makes of his private signal depends on the shape of his or her utility function. Risk neutral traders will buy or sell following a private signal h or l respectively. This happens for all levels of the public prior belief π. Tables 7 and 8 report the optimal choice for traders with different levels of risk aversion or risk loving. As a general rule non-risk neutral traders will submit non-informative order when π is extreme. Moreover by observing a trader contingent choice for different levels of π it is possible to recover his or her risk attitude. Second,the presences of intrinsic uncertainty implied by $\epsilon > 0$, does not affect risk neutral traders behavior but it does affect the one of non-risk neutral traders. Namely intrinsic uncertainty should increase non-informative trades for π extreme. Third, a trade in financial market corresponds to a specific lottery of payoffs. Thus, for a given level of the prior belief π and trading price $P = \pi 12 + (1 - \pi)4$, a trader's order should correspond to the lottery he or she prefers within a menu of three lotteries each one corresponding to trades $Q \in \{-1, 0, 1\}$. In other words a rational Bayesian trader should behave in the same way when facing maximization problem (1) and when choosing lotteries in a corresponding menu. Our experiment is meant to test these three hypothesis by detecting subjects risk attitude and testing for their Bayesian behavior. This allows to determining how these factors and the asset intrinsic uncertainty affect the order flow information content and market informational efficiency.

3 Experiment Design

Each subject participated to two formats of the experiment: a "Market Experiment" (ME) and a "Lottery Experiment" (LE). Each format consists of 3 warm-up rounds and 20 payoff relevant rounds. In both formats, each round reproduces the decision problem of a trader in the economy described in Section 2. Namely, for both formats, the problems that subjects had to solve in each round were equivalent to the decision of a trader having his unique opportunity to trade a risky asset after observing the trade of fully rational partially informed traders and after receiving an additional private signal on the value of the asset. Our experiment design differs from the ones in CG and DOR in two important perspectives.

First, in order to better understand how subjects use their private information, we adopted the "strategy approach". That is, instead of first disclosing the subject private signal and then asking the subject to choose an action, each subject was asked to declare a contingent strategy, that is, his or her choice contingent on receiving signal h and his or her choice contingent on receiving signal l. Then the subject's realized signal was disclosed and trade executed according to the subject's strategy. This approach has the advantage of allowing to detect non-informative orders that are the crucial ingredient for herding contrarian behavior and informational cascades in general. Moreover, by asking a subject's strategy for different levels of the public belief π_t, we are able to predict the behavior the subject would have when intervening at different stages of a trading history. This leads to the second important difference between our experiment and the previous works. In CG and DOR subjects intervene sequentially and hence have to interpret, and react to, the action of previous subjects. Instead, in each round of our experiment the relevant information provided by an *hypothetical* past trading history was made public through the announcement of a public belief π_t and also reflected into the trading price. Thus, other subjects strategies are completely irrelevant for a subject decision. Hence, what a subject believes regarding the rationality of other participant plays no role in determining his or her choice.Thus differently from what happens in many experimental works on herding, any deviation from what predicted by the theory cannot be ascribed to the lack of common knowledge regarding agents' rationality but rather to the actual lack of agents rationality.

Furthermore, by knowing the strategy of all subjects for different levels of the public belief, we are able to simulate an arbitrary number of trading histories and the corresponding evolution of price without having to run the actual experiments with the subjects. It is clear that this simulation exercise is worth because subject's behavior is not affected by factors contingent to the experiment such as their belief regarding other subjects' rationality.

3.1 Market Experiment

ME consists of 20 rounds where in each round each subject was asked the position he or she wanted to take in a given risky asset. Namely in a given round τ the subject is asked whether he or she wants to buy, to sell or not to trade a given risky asset, that we will denote asset τ. The fundamental value of asset τ is a random variable $\widetilde{v}_\tau = \widetilde{V}_\tau + \widetilde{\epsilon}$, as described in Section 2, where $\widetilde{V}_\tau \in \{4, 12\}$. The trading price for asset τ was fixed at $P_\tau = \pi_\tau \overline{V} + (1 - \pi_\tau)\underline{V}$. Both π_τ and P_τ are announced to the subject in round τ (see Figure 1). Moreover in each round each subject receives a private signal $\tilde{s} \in \{h, l\}$

with precision $p = 0.65$. Before being informed of his private signal and after observing π_τ and P_τ, the subject was asked to declare his or her desired trade conditional on receiving private signals h or l. The only difference among the 20 rounds was given by the probability π_τ and the corresponding trading price. For 17 of the 20 assets, the variables π_τ was determined in order to reflect the public belief obtained after observing 17 different histories of private signals. More precisely each of the first 17 assets corresponded to a different unbalance g varying from -8 to 8. Asset 18, 19, 20 were repetition of assets 1, 9 and 17 to control for the subject consistency.[5]

3.2 Lottery Experiment

LE is designed such that a rational Bayesian subject would find it perfectly equivalent to ME. In particular in LE subjects were asked exactly the same questions in exactly the same order as in the market experiment but with a different formulation. Instead of asking each subject the position to take into a given financial asset, the subject was asked to choose one item in a menu of lotteries. For each lottery in the menu it was specified the three possible outcomes and the corresponding probabilities. Similarly to the example of Table 2 and Table 3, each lottery in a menu corresponds to the random net gain obtained from selling, no-trade and buying one unit of asset τ given the private signal s. In order to parallel the "strategy approach" of the market experiment, at each stage, each subject was proposed two menus and asked to choose one lottery in each menu (see Figure 2). The only difference between the two menus proposed in a given round was in the probabilities attached to each payoff. This reflecting the different impact that a signal l and h would have on the probability a subject would attach to the different gains related to the trading decision. The lottery experiment consisted of 20 payoff relevant rounds each one including two menus. Overall each subject had to choose 40 lotteries among 40 menus. It was never mentioned to the subjects the relation between the two formats of the experiment neither the fact that from the perspective of a Bayesian rational subject the two experiments are perfectly equivalent.

3.3 Treatments

We implement one main treatment that involved 125 subjects (5 cohorts) and a control treatment that involved 42 subjects (2 cohorts). Treatments differ in the parameters defining the fundamental value of the asset. In the main treatment of the fundamentals where $\underline{V} = 4$, $\overline{V} = 12$, $\epsilon = 4$ that imply that $\widetilde{v}$ could take three values $\{0, 8, 16\}$. We will denote this treatment the *ϵ-based case*. In the control treatment we

[5] See for example Table 6 for a correspondence between the numbering of the asset, the corresponding π_τ and g.

eliminate the $\widetilde{\epsilon}$ component, i.e. $\underline{V} = 4$, $\overline{V} = 12$ and $\epsilon = 0$. The control treatment parallels the set-up in CG and DOR.

3.4 Subjects' Payoff

A subject's actual payoff was not based on his or her performance in all rounds of the experiment. Subjects' payoffs were determined on the basis of the gain on one round for the market experiment, and one round for the lottery experiment. The ex-post relevant rounds were randomly selected at the end of the experiment. We opted for this scheme as risk aversion plays a crucial role in this experiment. In particular, rewarding a subject on the basis of his or her average performance in the experiment would have substantially reduced the risk linked to a given decision and would not have parallel the theoretical framework.

More precisely, in order to determine a subject's payoff in the market experiment we proceeded as follows. First, one of the 20 rounds of the market experiment was randomly selected, say the round corresponding to asset τ. Second, the fundamental value of the asset was determined according to the announced distribution π_τ for this asset. Third, based on the realization of $\widetilde{V}_\tau$ the subject private signal was determined. Fourth, the subject desired transaction corresponding to this asset and this realization of the subject private signal was executed. Then the subject payoff was the sum of a fixed amount m plus an amount equal to the gain realized on this transaction. A subject payoff in the lottery experiment was determined by first, randomly selecting one of the 40 lotteries chosen by the subject. Then the subject's payoff was determined according to the probability distribution of the selected lottery.

3.5 Implementation

The experiment was run in HEC School of Management and Toulouse University. The 177 subjects were recruited from undergraduate finance classes in HEC and Toulouse University and had no previous experience in financial market and social learning experiments. In each session between 13 and 43 subjects participated as decision makers.

At the beginning of a session we gave written instruction that where also read aloud by an experiment administrator. Then two trial sub-sessions, involving the trade of three assets each, were run. Each of the trial sessions reproduced the trading mechanism in the two formats of the experiment. After the trial sub-sessions and before the first payoff relevant sub-session, subjects answered a questionnaire that tested their level of understanding of the rules of the experiment. Administrators answered all subjects' clarifying questions regarding the rules of the game until the questionnaire was distributed. Afterwards subjects were not allowed to ask additional questions and administrators ensured no form of communication among the subjects to take place. Throughout the experiment it was impossible for participant to observe other's screens. Each experiment lasted overall about 1 hours and a half. An average of €22.28 were paid to each subject. This amount includes a participation fee of €8 and earning for decision. Subjects were also rewarded with bonus points valid to improve their mark in the Financial Markets course. We discarded from our data-set the decisions of 6 subjects who gave more than 3 wrong answers out of the 14 questions in the questionnaire, as we considered these subjects had not understood the main rules of the experiment. The final number of observations was of $4,147$ in the ϵ-based case and $1,428$ in the control treatment.

4 Experimental Results

In the following we describe the main results of the experiment and we compare with the prediction of the theory as well as with the finding in CG and DOR. Our methodology allows us to identify how private and public information affect subjects' trading strategies. On the one hand, by observing each subject's trading decisions for both realizations of the private signal we can clearly identify when a subject private signal affects his or her action and when it does not. This is crucial when studying the informational efficiency property of the market. On the other hand, by observing agents' contingent strategy at different levels of the public belief we can analyze how the strength of public beliefs affects subjects sensitivity to private information. Furthermore, with the introduction of an intrinsically risky asset we study the impact of risk attitude on trading decision and on willingness to make use of private information.

We recall here that the prediction of the theory are that whereas risk neutral traders will always buy when receiving signal h and sell when receiving signal l, risk averse traders and risk lover traders will tend to ignore their signal when the public belief is extreme, i.e., when $|g|$ is large. For intermediate levels of the public belief the subject's behavior will depend on the curvature of his or her utility function as illustrated in Table 7 and Table 8. Note that it is possible to roughly recover the shape of a subject's utility function and risk attitude looking at the subject's answers in LE. In the following, we will identify a subject's strategy with two letters indicating the action chosen for signal l and h respectively. Namely S, N, and B stand for sell order, no trade and buy order respectively.[6] We adopt the following definition

[6]For example, strategy N-B corresponds to no-trade when receiving signal l and buy order when receiving signal h.

of herd and contrarian behavior.

Definition 4 *A subject engages in herd behavior if there exist $g^* > 0$ ($g^* < 0$) such that the subject adopts strategy is B-B (resp. S-S) for all $g \geq g^*$ (resp. $g \leq g^*$).*

Definition 5 *A subject engage in contrarian behavior if there exist $g^* > 0$ ($g^* < 0$) such that the subject adopts strategy is S-S (resp. B-B) for all $g \geq g^*$ (resp. $g \leq g^*$).*

In other words a subject engages in herd (contrarian) behavior if for example a sufficiently positive a priori induces him or her to buy (resp. sell) the asset independently of the realization of the private signal.

4.1 Lottery Experiment

In LE, subjects do not have to compute the possible payoffs related to a trading decision, and more importantly the probabilities attached to each possible outcome. This framework, closely matches the theoretical assumption that agents are Bayesian expected utility maximizers. Not surprisingly the outcome in LE is close to the theoretical predictions.

Non-informative trades. Non-informative trades are those strategies that are not affected by the sign of private signals, i.e., S-S, N-N, B-B. Our methodology allows us to identify and to study the nature of non-informative trades. This was not possible in the setting of the previous studies as only ex-post decisions were observed.

Table 9 and Table 10 report the frequency by order imbalance of informative and non-informative trades in the control LE and in the ϵ-based LE, respectively. Overall, we observe that the percentage of non-informative trades increases with $|g|$ and that trades are less informative in the ϵ-based case when compared to the control treatment. In other words subjects tend to ignore their private information more often when the public information is sufficiently strong. This phenomenon is amplified in the presence of additional uncertainty on the fundamentals of the asset. Namely, in the control treatment when $g = 0$, more than 45% of trades are non-informative and this percentage increases to 64% and 74% for $g = -8$ and $g = 8$ respectively. In the ϵ-based case, non-informative trades represent more than 47% of all trades when $g = 0$ and 84% and 83% for $g = -8$ and $g = 8$ respectively. This has clear implications for market informational efficiency. The flow of information coming to the market decreases with the strength of public belief and with the level of intrinsic uncertainty regarding the fundamental value of the asset.

It is interesting to disentangle the nature of non-informative trades. We distinguish situations in which the subject decides not to trade independently of his private signal (strategy N-N), from situations where the subject chooses a non-neutral trading position that is the same for both realizations of the private signal (strategies B-B and S-S).

The use of strategy N-N can be ascribed to risk aversion. When facing uncertainty regarding the fundamental value of the asset, and when the expected profit from trading is small, a sufficiently risk averse subject will prefer to take no position in the asset and will retain from trading. Uncertainty regarding the fundamentals of the asset is maximum when the public information is vague, i.e., when g is close to 0. Whereas for $|g|$ large, the uncertainty reduces as the public information clearly indicates whether the $\widetilde{V}$ component is 12 or 4. On the other hand, the expected profit from trading decreases with $|g|$ and this implies that an investor should be more prone to speculate on information when g is close to 0. Overall, using a simple mean-variance argument, one can predict that the smaller the difference between the expected gain from trading and the intrinsic risk of a trade, the larger should be the percentage of subjects choosing strategy N-N.

The predicted patterns of non-informative trades differs for the ϵ-based case and the control treatment. In the control treatment, uncertainty vanishes as $|g|$ increases. Consistently with the prediction of the theory, in the control treatment we observe a pick of strategy N-N (about 45%) for g close to 0, whereas the proportion of subjects choosing strategy N-N falls to about 30% when $|g|$ is close to 8. When considering the ϵ-based case, the presence of the component ϵ makes a trading decision intrinsically risky even when $|g|$ is large and little uncertainty is left on component $\widetilde{V}$. Note however that the expected profit from a trade tends to 0 as $|g|$ increases. Thus the theory predicts that when $|g|$ is large the proportion of subjects choosing strategy N-N should increase. Consistently with this prediction, the percentage of subjects choosing strategy N-N is about 37% for g around 0 and it increases to 66% and 55% for $g = -7$ and $g = 7$ respectively.

Non-informative strategies consisting in choosing the same trade independently of the private signal correspond to what the herding literature classifies as herding or contrarian behavior. Strategy S-S (resp. B-B) observed for g close to 8 (resp. g close to -8) corresponds to a sort of contrarian behavior as the agent trades against the strong information provided by the public belief. Risk-lover subjects will have such a behavior as illustrated in Table 7. In fact, by choosing a trading strategy that goes against the public information the subject is de facto choosing the lottery having the highest potential gain. Contrarian behavior is more likely when g is extreme. As for example when g is large it is sufficient a smaller degree of risk-loving to choose strategy S-S.[7]

Herding strategies, that consist in choosing S-S and B-B for g close to -8 and 8 respectively, can hardly been explained on the basis of subjects' risk attitude. In our framework such behavior can be ascribed to the subject's tendency of rounding probabilities and/or ignore events whose probability is sufficiently small. As it is apparent in Table 11 and Table 12 the percentage of herding strategies is larger for $|g|$ larger than 3 and it is higher in the control treatment. This phenomenon can be explained as follows. A subjects already having strong a priori belief on the event $\widetilde{V} = V$, will consider this event virtually sure. Hence the subject will buy whenever the asset price is below V and will sell if the price is above V. This irrespective of the realization of the private signal. This phenomenon is mitigated in the ϵ-based treatment. In this case, even if a subject neglects the residual uncertainty on $\widetilde{V}$, he or she will not ignore the uncertainty on $\widetilde{\epsilon}$ that makes any trade intrinsically risky.

The observed pattern for strategies S-S and B-B are symmetric. In the control treatment S-S and B-B are absent for g close to 0 and increase with $|g|$. The increase is stronger for contrarian trades rather than for herding. This suggests that the proportion of risk lover agents is superior to the proportion of subjects that tends to round given probabilities. Namely the percentage of S-S strategies (resp. B-B) rises to 29% for $g = 8$ (resp. $g = -8$) and to 19% (resp. 14%) for $g = -8$ (resp. $g = 8$). For the ϵ-based case, figures are similar: the percentage of S-S strategies (resp. B-B) is 8% (resp. 2%) for $g = 0$ and rises to 32% (resp. 25%) for $g = 8$ (resp. $g = -8$) and to 10% (resp. 12%) for $g = -8$ (resp. $g = 8$). It is important to stress that due to the specific design of our experiment, differently from what happens in CG and DOR, the presence of contrarian behavior cannot be ascribed to the lack of common knowledge of agents rationality but should directly be linked to subjects attitude toward risk.

Informative trades. Strategy S-B, consisting in buying when the private signal is low and selling when the private signal is high, is what CG and DOR refer to as the rational strategy. In fact, such a strategy is rational only if the subject is risk neutral. A risk neutral subject would choose strategy S-B independently of $|g|$ and independently of the treatment. Our experimental data suggests that very few subjects consistently choose strategy S-B. In the control treatment, the percentage of subjects using such strategy is maximum for $g = 0$ (33% of strategies) however no subject follows this strategy for $g \geq 6$ and only 2% of subjects still use this strategy for $g = -8$. The ϵ-based case presents a similar patterns. The percentage of subjects following strategy S-B is maximum for $g = 0$ (37%) but only ??% of subjects consistently use this strategy for all levels of the public belief. It is natural to observe a pick of S-B

[7]For example in Table 7 it is sufficient to compare the strategies for columns labeled LE for $\gamma = -0.26$ and $\gamma = -0.3$.

strategy when the expected profit from trade is maximum, i.e. $g = 0$. However, the fact that for $|g|$ large, subjects tend to prefer non-informative strategies is clearly in contradiction with the risk neutrality assumption.

Informative strategies S-N and N-B are used much more frequently than S-B. These strategies are optimal for risk averse agents, but whether these strategies are used with positive or negative g depends on the specific shape of the utility function. For example, in the control treatment, the strategy consisting in trading according to the signal whenever this confirms the public belief, but avoiding to trade otherwise, is consistent with a mean-variance utility function.[8] The strategy that consists in trading according to the signal when this contradicts the public belief and preventing from trading otherwise is optimal for a CARA utility agent.[9] Our observations for the control treatment is that strategy S-N (resp. N-B) represents about 25% of all strategies for $g \leq -2$ (resp. $g \geq 2$), but it represents only 7% or less for $g \geq 7$ (resp. $g \leq -5$). This suggests that mean-variance utility seems to explain more subjects' behavior than a CARA utility function.

For the ϵ-based treatment, both mean-variance utility and CARA utility are consistent with the use of S-N strategy (N-B strategy) for g (resp. $-g$) positive but small. Our observations for the ϵ-based treatment are consistent with these predictions. Strategy S-N (resp. N-B) picks at 27% (resp. 24%) of the strategies for $g = 1$ (resp. $g = -1$) but it represents less then 10% of the strategies for $g < 0$ and $g > 5$ (resp. $g < -5$ and $g > -1$).

Finally irrational strategies (B-N, N-S, B-S) represent less than 1.4% percent of trades in the control treatment and 2.3% percent of trades in the ϵ-based treatment.

4.2 Market Experiment

The prediction of the theory is that a subject should choose exactly the same strategies in LE and ME. This prediction is clearly rejected by our data. Overall in the the ϵ-based treatment (control treatment) only in 42.34% (33.33%) of the observations subjects' answers are the same for LE and ME. Only ??% (??%) of subjects answers in LE and ME where consistent for all level of g. In most of these cases subjects preferred strategy N-N for all levels of g in both formats. A closer look for comparing to subjects answers shows phenomena that are common to the two treatment. First, contrarian behavior trades (B-B for g negative and S-S for g positive) tend to disappear in format ME. Second, in comparison with format LE there is an increase in strategies consisting in following the signal whenever this confirms the public

[8]See for instance the first LE labeled column in Table 8.

[9]See for instance the second LE labeled column in Table 8.

history and not trading otherwise (i.e. S-N for g negative and N-B for g positive). Third strategies N-N are less frequent in ME compared to LE for for $|g|$ large but it is more frequent in ME for $g = 0$. For the ϵ-based experiment the frequency of strategies S-B increases in ME.

These differences result in a overall increase in the use of informative strategies for ME when $|g| > 0$, but a reduction for $g = 0$. In the ϵ- based treatment informative strategies (control treatment) pass from 28.26% (resp. 45.1%) of all choice in LE to 42.35% (resp. 70.17%) in ME. This suggests that when public information is small, the informational content of order flow is lower in ME compared to LE. However for strong prior belief, in ME information will be better signaled through subjects strategies.

In order to understand the source of these deviations from the theory we have to keep in mind the key difference between the two formats: in ME subjects have to mentally[10] compute the probability distribution of payoffs that results from the public belief and their private signals. In LE these distributions are explicitly provided. Hence in ME subjects typically do not have an exact knowledge of the distribution of the asset fundamentals. Uncertainty regarding the actual probability distribution of the asset fundamentals can induce subjects to display ambiguity aversion. This can explain why in ME subjects tend to prevent from trading when the sign of their private signal is opposite to the sign of the public history, while they tend to trade according to the signal whenever this confirms the public history.

Ambiguity aversion has been first pointed out with the famous Elsberg Paradox that show that individuals prefer to bet on lotteries for which the probability distribution is known. Gilboa and Schmeidler [10] produced a decision theory consistent with this behavior. They adopt an axiomatic approach that implies that decision makers evaluate each action according to the minimum expected utility it yields, where the minimum is take over the possible distribution functions of the fundamentals. In our framework the set of possible distribution functions that a subject considers as possible will change with the realization of the private signal. When the private signal is opposite to the public signal g, the set of possible distribution functions is relatively wide and the action maximizing the minimum expected utility is the no-trade action. When the private signal and the public signal go in the same direction the ambiguity is lower and the trading according to the signal can maximize the minimum expected utility.

4.3 Ex-post observations

In order to compare our results to those of CG and DOR (table 2 page 1434, section No history), we look at the ex-post decisions in our ME for the control treatment. This represents the set-up of the experiment

[10] Subject had no access to calculators during the experiment.

that is closest to the experiments in those papers.

Table 15 reports the frequency, overall and by order imbalance, of different strategies (namely, “Contrarian”, “Follow” and “No Trade”) in the control ME. In CG, 61% of subjects decide to follow their private information. In our experiment this percentage falls to 53%. The proportion of agents who decide not to trade is 41% in our experiment and 25% in CG. Interestingly contrarian behavior amounts to 14% percent in CG whereas it is less than 6% percent in our setting.

Several comments are in order. Contrarian behaviors are rationalized in CG by taking into account that subjects may have doubts about the rationality of their predecessors. As already explained, our design rules out this consideration and accordingly we find a lower percentage of contrarian behaviors.

In CG, the frequency of no trades increases with the absolute value of the trade imbalance. This observation does not correspond to the theoretical model that underlies CG’s paper. The authors suggest that a possible explanation is that subjects prefer not to trade when the trade imbalance is high because they face a high maximum loss. In our experiment we obtain the opposite pattern: we observe a pick of no trade decisions for $g = 0$ and a decrease of no trade decisions as $|g|$ increases. Precisely, we observe about 60% of no trade decisions for $|g| = 0$ while this percentage decreases to around 30% for $|g| = 8$. This result corroborates the analysis of the N-N strategy in the lottery experiment and, as explained in section 4.1, are in line with the theory.

4.4 Market Informational Efficiency

Measure of market informational efficiency is given by the evolution of the pricing error as trade unfolds. We define the pricing error as the percentage difference between the trading price and the price that would prevail if market makers could directly observe market participant’s private signals.[11] Informational efficiency is higher when the pricing error decreases at a faster rate.

In our experiment we do not observe trading histories as subjects do not trade sequentially the same asset. Nevertheless, thanks to the observation of subjects contingent strategies at different levels of the public belief, we know how each subject would trade after receiving any private signal when intervening at any different points in an hypothetical trading history. Thus, we can simulate virtual trading histories and compute the resulting pricing errors. This method has the advantage of allowing to simulate an arbitrary large number of trading histories without having to physically run the experiment with subjects. Note also that as subjects decision was completely disjoint from other subjects’ maximization problem,

[11]The percentage pricing error is defined as $|(Trading\, priceat\, t) - (Efficient\, priceat\, t)|\, /\, (Efficient\, price\, at\, t)$.

it is sensible to simulate trading histories where the virtual traders are subjects from different cohorts. Of course this method is immune to the effect of subjects' beliefs regarding other participant rationality.

When simulating trading histories we have to make assumptions on subjects arrivals and how the public belief and trading prices react to the flow of trade. More precisely in order to simulate a trading history we proceed as follows. Initially the fundamental value of the asset is randomly selected according to $\mathbb{P}(\tilde{\epsilon} = \epsilon) = \mathbb{P}(\tilde{V} = \overline{V}) = \frac{1}{2}$ and the initial public belief is fixed at $\pi_0 = 0.5$. In each trading round, first, one among the 9 possible trading strategies in Table 8 was randomly selected in a way that reflected the empirical frequencies of the different strategies in the experiments for $g = 0$. Note that these frequencies change with g, that reflects the level of public belief. Hence the probability of picking a given strategy evolved according to the change in the public belief. Second, the trader private signal was randomly selected with precision $p = 0.65$. Third, the trader order was selected according to trader's strategy and private signal. Finally, the public belief were updated and a new trading round started. In each round the trading price corresponded to the expected value of the asset given the public belief. We considered three updating rules for the public belief that we will denote R1, R2 and R3.

Rule R1 The assumption underlying this updating rule is that market makers know the strategy a trader uses but they do not observe the trader's private signal. As a consequence, public belief is unchanged whenever the trader strategy is non informative, i.e. when the trader strategy is either S-S, N-N, or B-B. For all other strategies the trader order is informative, and hence the public belief will be updated according to the actual realization of the trader private signal. This updating rule represents the best market makers could learn in the real market after observing an agent trading decision.

Rule R2 This rule is the closest to the assumptions of the microstructure theory and is based on the hypothesis that market makers do not know the identity of the trader but they know the average behavior of the population of traders. That is, for any given level of public beliefs, market makers know the frequency with which each strategy is adopted by traders. These frequency are given in tables 11 to 14. After observing a given action the posterior belief will change according to the probability that this order comes from someone who received a signal l or a signal h.[12] This pricing rule is also probably the more realistic in a market where trading is anonymous and where

[12] Note that the resulting public belief is not necessarily on the grid of beliefs used in the experiment. In these cases, we approximate a trader behavior with the one corresponding to the public belief on the grid that is closest to the actual public belief.

market makers infer from the flow of trade an unbiased statistic of the asset fundamentals.

Rule R3 This is the updating rule adopted by CG and DOR. A buy order is interpreted as signal h, a sell order is interpreted as signal s and no-trade leaves beliefs unchanged. This is the simplest pricing rule but also the one that is the less compatible with market makers rationality.

These three pricing rules allow us to compute different levels of prices sequence. R1 provides an estimation of the minimum possible pricing error, while R2 and R3 give an idea of maximum level for the pricing error. Overall we simulated about $12,000$ trading histories each one lasting a maximum of 40 trading rounds. Tables 4 and 5 report average pricing errors for histories of length 10, 20 and 40 periods, for each pricing rule and experiment format.

	LE			ME		
t	R1	R2	R3	R1	R2	R3
10	18.2%	24.2%	22.9%	19.6%	23.9%	24.8%
20	17.6%	26.8%	24.4%	18.3%	24.4%	28.0%
40	12.2%	23.6%	21.4%	11.4%	19.3%	30.0%

Table 4: Average pricing error in the control treatment

	LE			ME		
t	R1	R2	R3	R1	R2	R3
10	20.6%	30.0%	32.1%	25.3%	26.2%	24.8%
20	22.8%	36.2%	41.1%	33.0%	40.8%	37.3%
40	18.8%	36.3%	47.1%	20.3%	41.3%	34.8%

Table 5: Average pricing error in the ϵ-based treatment

Not surprisingly the most efficient updating rule is R1 as it assumes market makers can perfectly deduce a trader's signal from his order whenever the trader uses an informative strategy. For both formats and both treatments, the resulting pricing error is consistently lower for pricing rule R1 when compared to the other pricing rules. The comparison between R2 and R3 is ambiguous. While R2 provides lower pricing errors than R3 in the ϵ-based treatment of LE and in the control treatment of format ME, R2 performs better than R3 in the other situations. The impact of the pricing rule on the measure of pricing error can be important and increases with the length of histories. In some cases the average pricing error under R3 is almost three times the average error under R1.

Confirming the prediction of the theory price convergence is substantially faster in the control treatment. At the 40th round of trade the average price error in the ϵ-based is between one and a half and more than twice as large when compared to the pricing error in the control treatment. This happens because subjects tend to neglect their signal more often when the underlying uncertainty is higher.

The comparison between the two formats suggests that in the control treatment the difference in pricing error between ME and LE is not significant, while for the ϵ-based treatment prices are more efficient in LE compared to ME.

Interestingly, long run mispricing can be high, i.e. situations where the trading price permanently and significantly diverge from the efficient price are quite common. In the control treatment the percentage of histories that after 40 rounds present a pricing error of more than 66% is between 7% and 16% of the simulated histories for pricing rule R2 and R3 respectively. In the ϵ-based treatment these percentages increases to 15% and 27%.

5 Conclusion

We have developed an experiment that simulates trading in financial market. The specific format of our experiment allows to unambiguously detect herd and contrarian behavior as well as informational cascade and to analyze the main forces driving individuals to engage in these behaviors. We show that many of the so "called irrational" behavior are not so if one takes into account subjects' risk attitude. A subjects risk aversion or risk loving can induce to ignore private information and possibly engage in contrarian behavior. This reduces market informational efficiency and can generate informational cascade. Our data suggests that when subjects have to mentally compute the distribution of their payoff, they face some uncertainty about the actual chances of gain they face and react to this situation by displaying ambiguity aversion. These phenomena reduce market informational efficiency and are magnified when the intrinsic uncertainty regarding asset fundamentals is large.

Overall, our results suggest that private information regarding past shocks on fundamentals should be incorporated into prices more slowly when further shocks are expected or when uncertainty regarding other components of the asset is larger.

References

[1] **Avery, Christopher and Zemsky, Peter.** "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review,* 1998, 88(4), pp. 724–48.

[2] **Bikhchandani S., D. Hirshleifer, and I. Welch**, 1992, "A Theory of Fads, Fashion, Customs and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100, pp. 992-1026.

[3] **Celen, Bogachan and Kariv, Shachar.** "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review,* 2004, 94(3), pp. 484–98.

[4] **Cipriani, Marco and Guarino, Antonio.** "Herd Behavior and Contagion in Financial Markets." Unpublished Paper, 2001.

[5] **Cipriani, Marco and Guarino, Antonio.** "Herd Behavior in a Laboratory Financial Market." *American Economic Review,* 2005, 95(5), pp. 1427–1443.

[6] **Chari, V. V. and Kehoe, Patrick J.** "Financial Crisis as Herds: Overturning the critiques." *Journal of Economic Theory*, 2004, 119(1), pp. 128–150.

[7] **Decamps, Jean-Paul and Lovo, Stefano.** "Informational Cascades with Endogenous Prices: The Role of Risk Aversion", *Journal of Mathematical Economics*, 2006, (42)1, pp. 109–120.

[8] **Decamps, Jean-Paul and Lovo, Stefano.** "A Note on Risk Aversion and Herd Behavior in Financial Markets." *The Geneva Risk and Insurance Review*, 2006, (31)1, pp. 35–42.

[9] **Drehman, Mathias, Oechssler, Jorg and Roider Andreas.** "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment."

[10] **Gilboa, Itzhak and Shmeidler David.** "Maximum Expected Utility with non-Unique Prior" *Journal of Mathematical Economics*, 1989, 18, pp. 141–153.

[11] **Smith, Lones and Sorensen, Peter.** "Pathological Outcomes of Observational Learning." *Econometrica*, 2000, 68(2), pp. 371–98.

Asset	π_t	g_t
Asset 1	0.013	-7
Asset 2	0.023	-6
...	...	...
Asset 7	0.350	-1
Asset 8	0.500	0
Asset 9	0.650	$+1$
...	...	...
Asset 15	0.987	$+7$
Asset 16	0.013	-7
Asset 17	0.500	0
Asset 18	0.987	$+7$

Table 6: Asset number, public belief π_t, and balance of signals g_t

$U(x) = -\gamma e^{-\gamma x}$		$\gamma > 0.78$	$\gamma = 0.02$	$-0.25 < \gamma < 0.003$	$\gamma = -0.26$	$\gamma = -0.3$
g	π	LE	LE	LE	LE	LE
-8	0.002	N,N	N,N	S,B	B,B	B,B
-7	0.013	N,N	N,N	S,B	B,B	B,B
-6	0.023	N,N	N,N	S,B	B,B	B,B
-5	0.043	N,N	N,B	S,B	S,B	B,B
-4	0.078	N,N	S,B	S,B	S,B	B,B
-3	0.135	N,N	S,B	S,B	S,B	B,B
-2	0.225	N,N	S,B	S,B	S,B	S,B
-1	0.350	N,N	S,B	S,B	S,B	S,B
0	0.500	N,N	S,B	S,B	S,B	S,B
$+1$	0.650	N,N	S,B	S,B	S,B	S,B
$+2$	0.765	N,N	S,B	S,B	S,B	S,B
$+3$	0.865	N,N	S,B	S,B	S,B	S,S
$+4$	0.922	N,N	S,B	S,B	S,B	S,S
$+5$	0.957	N,N	S,N	S,B	S,B	S,S
$+6$	0.977	N,N	N,N	S,B	S,B	S,S
$+7$	0.987	N,N	N,N	S,B	S,S	S,S
$+8$	0.998	N,N	N,N	S,B	S,S	S,S

Table 7: Optimal strategies for an investor with CARA utility in the ϵ-based case

$U(x) =$		$E[x] - 0.2Var[x]$	$-0.15e^{-0.15x}$	$-0.25 < \gamma < 0.003$	$\gamma = -0.26$	$\gamma = -0.3$
g	π	LE	LE	LE	LE	LE
-8	0.002	S,N	N,B	S,B	B,B	B,B
-7	0.013	S,N	N,B	S,B	B,B	B,B
-6	0.023	S,N	N,B	S,B	S,B	B,B
-5	0.043	S,N	N,B	S,B	S,B	B,B
-4	0.078	S,N	N,B	S,B	S,B	B,B
-3	0.135	S,N	N,B	S,B	S,B	B,B
-2	0.225	N,N	N,B	S,B	S,B	S,B
-1	0.35	N,N	S,B	S,B	S,B	S,B
0	0.5	N,N	S,B	S,B	S,B	S,B
$+1$	0.65	N,N	S,B	S,B	S,B	S,B
$+2$	0.765	N,N	S,N	S,B	S,B	S,B
$+3$	0.865	N,B	S,N	S,B	S,B	S,S
$+4$	0.922	N,B	S,N	S,B	S,B	S,S
$+5$	0.957	N,B	S,N	S,B	S,B	S,S
$+6$	0.977	N,B	S,N	S,B	S,B	S,S
$+7$	0.987	N,B	S,N	S,B	S,S	S,S
$+8$	0.998	N,B	S,N	S,B	S,S	S,S

Table 8: Optimal strategies for different forms of utility function when $\epsilon = 0$

g	non-informative	informative
-8	64.29	35.71
-7	71.43	28.57
-6	64.29	35.71
-5	61.90	38.10
-4	40.48	59.52
-3	54.76	45.24
-2	45.24	54.76
-1	52.38	47.62
0	45.24	54.76
$+1$	40.48	59.52
$+2$	40.48	59.52
$+3$	45.24	54.76
$+4$	52.38	47.62
$+5$	57.14	42.86
$+6$	52.38	47.62
$+7$	71.43	28.57
$+8$	73.81	26.19

Table 9: Informative and non-informative trades (in %) in the control Lottery Experiment

g	non-informative	informative
-8	84.48	15.52
-7	85.50	14.50
-6	87.79	12.21
-5	77.86	22.14
-4	77.10	22.90
-3	64.89	35.11
-2	65.65	34.35
-1	60.31	39.69
0	47.33	52.67
$+1$	57.25	42.75
$+2$	69.47	30.53
$+3$	63.36	36.64
$+4$	72.52	27.48
$+5$	77.86	22.14
$+6$	79.39	20.61
$+7$	79.39	20.61
$+8$	82.76	17.24

Table 10: Informative and non-informative trades (in %) in the ϵ-based Lottery Experiment

g	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
-8	28.57	2.38	0.00	4.76	16.67	0.00	2.38	26.19	19.05
-7	30.95	0.00	0.00	4.76	28.57	0.00	0.00	23.81	11.90
-6	30.95	0.00	0.00	7.14	21.43	0.00	4.76	23.81	11.90
-5	23.81	2.38	0.00	7.14	26.19	0.00	9.52	19.05	11.90
-4	11.90	0.00	0.00	23.81	26.19	0.00	7.14	28.57	2.38
-3	7.14	0.00	0.00	11.90	33.33	0.00	0.00	33.33	14.29
-2	2.38	0.00	0.00	16.67	38.10	0.00	9.52	28.57	4.76
-1	0.00	0.00	0.00	14.29	50.00	0.00	21.43	11.90	2.38
0	0.00	0.00	0.00	9.52	45.24	0.00	33.33	11.90	0.00
$+1$	0.00	0.00	0.00	14.29	40.48	0.00	19.05	26.19	0.00
$+2$	7.14	0.00	0.00	28.57	33.33	0.00	11.90	19.05	0.00
$+3$	9.52	0.00	0.00	30.95	30.95	4.76	0.00	19.05	4.76
$+4$	7.14	0.00	0.00	23.81	28.57	2.38	7.14	14.29	16.67
$+5$	11.90	2.38	0.00	21.43	26.19	2.38	4.76	11.90	19.05
$+6$	11.90	0.00	0.00	30.95	19.05	2.38	0.00	14.29	21.43
$+7$	11.90	0.00	0.00	21.43	28.57	2.38	0.00	4.76	30.95
$+8$	14.29	0.00	0.00	21.43	30.95	2.38	0.00	2.38	28.57

Table 11: Conditional decisions (in %) in the control Lottery Experiment

g	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
-8	25.86	0.00	3.45	3.45	48.28	1.72	3.45	3.45	10.34
-7	11.45	0.00	0.00	3.82	66.41	2.29	5.34	3.05	7.63
-6	15.27	0.76	0.76	3.05	68.70	0.76	3.82	3.05	3.82
-5	12.98	0.76	1.53	9.16	60.31	0.76	4.58	5.34	4.58
-4	16.03	1.53	1.53	10.69	56.49	0.76	5.34	3.05	4.58
-3	13.74	0.76	0.76	20.61	45.80	1.53	3.82	7.63	5.34
-2	19.08	1.53	0.76	20.61	43.51	1.53	6.11	3.82	3.05
-1	16.03	0.76	0.00	25.19	41.22	0.00	9.16	4.58	3.05
0	2.29	0.00	0.00	5.34	37.40	0.00	37.40	9.92	7.63
$+1$	3.05	0.00	0.00	3.82	35.88	0.76	10.69	27.48	18.32
$+2$	3.82	1.53	0.76	3.05	47.33	0.00	9.16	16.03	18.32
$+3$	3.05	0.00	0.76	6.87	45.04	0.76	10.69	17.56	15.27
$+4$	1.53	0.00	0.76	3.05	47.33	0.76	7.63	15.27	23.66
$+5$	6.11	0.00	0.00	4.58	51.91	0.76	5.34	11.45	19.85
$+6$	6.87	0.76	0.00	5.34	51.91	3.05	3.05	8.40	20.61
$+7$	9.92	0.76	0.00	5.34	54.96	0.76	6.11	7.63	14.50
$+8$	12.07	0.00	5.17	3.45	37.93	1.72	5.17	1.72	32.76

Table 12: Conditional decisions (in %) in the $epsilon$-based Lottery Experiment

g	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
-8	2.38	0.00	0.00	11.90	11.90	0.00	9.52	50.00	14.29
-7	0.00	0.00	0.00	14.29	14.29	2.38	4.76	47.62	16.67
-6	0.00	0.00	0.00	11.90	14.29	2.38	4.76	50.00	16.67
-5	0.00	0.00	0.00	9.52	14.29	2.38	7.14	50.00	16.67
-4	0.00	0.00	0.00	9.52	11.90	4.76	7.14	64.29	2.38
-3	0.00	0.00	0.00	7.14	11.90	0.00	14.29	64.29	2.38
-2	0.00	0.00	0.00	7.14	23.81	2.38	19.05	47.62	0.00
-1	0.00	0.00	0.00	2.38	50.00	2.38	21.43	23.81	0.00
0	0.00	0.00	0.00	4.76	64.29	0.00	30.95	0.00	0.00
$+1$	2.38	0.00	0.00	35.71	40.48	0.00	16.67	4.76	0.00
$+2$	2.38	0.00	0.00	61.90	11.90	0.00	16.67	7.14	0.00
$+3$	4.76	0.00	0.00	50.00	14.29	0.00	23.81	7.14	0.00
$+4$	4.76	0.00	0.00	64.29	11.90	0.00	11.90	7.14	0.00
$+5$	7.14	2.38	0.00	57.14	14.29	0.00	7.14	11.90	0.00
$+6$	11.90	0.00	0.00	50.00	21.43	0.00	4.76	11.90	0.00
$+7$	11.90	0.00	0.00	47.62	23.81	0.00	4.76	11.90	0.00
$+8$	21.43	0.00	0.00	47.62	14.29	0.00	4.76	11.90	0.00

Table 13: Conditional decisions (in %) in the control Market Experiment

g	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
-8	1.72	3.45	6.90	12.07	31.03	0.00	20.69	13.79	10.34
-7	1.54	2.31	1.54	7.69	45.38	1.54	13.85	12.31	13.85
-6	1.54	2.31	1.54	8.46	46.15	2.31	12.31	11.54	13.85
-5	1.54	1.54	2.31	8.46	47.69	1.54	11.54	13.85	11.54
-4	1.54	2.31	2.31	8.46	50.77	1.54	11.54	12.31	9.23
-3	1.54	0.77	3.08	8.46	49.23	0.77	13.85	15.38	6.92
-2	1.54	2.31	4.62	11.54	49.23	2.31	11.54	14.62	2.31
-1	1.54	0.77	3.85	10.77	59.23	0.77	13.85	6.15	3.08
0	1.54	0.00	1.54	3.85	69.23	3.08	16.15	1.54	3.08
$+1$	1.54	0.77	0.77	16.92	51.54	2.31	17.69	4.62	3.85
$+2$	3.85	0.00	1.54	22.31	46.15	0.77	16.92	5.38	3.08
$+3$	4.62	1.54	0.00	26.15	43.08	4.62	11.54	6.15	2.31
$+4$	6.92	0.77	1.54	20.00	44.62	2.31	12.31	8.46	3.08
$+5$	11.54	0.00	0.77	17.69	43.08	2.31	11.54	9.23	3.85
$+6$	13.85	0.00	1.54	15.38	42.31	3.08	12.31	6.92	4.62
$+7$	13.85	0.77	0.77	18.46	38.46	4.62	10.00	9.23	3.85
$+8$	12.07	1.72	1.72	29.31	15.52	6.90	13.79	15.52	3.45

Table 14: Conditional decisions (in %) in the $epsilon$-based Market Experiment

g	Contrarian	Follow	No Trade
-8	11.54	61.54	26.92
-7	23.08	42.31	34.62
-6	11.54	61.54	26.92
-5	3.85	53.85	42.31
-4	3.85	46.15	50.00
-3	3.85	69.23	26.92
-2	3.85	50.00	46.15
-1	0.00	34.62	65.38
0	0.00	38.46	61.54
$+1$	0.00	38.46	61.54
$+2$	0.00	53.85	46.15
$+3$	3.85	65.38	30.77
$+4$	0.00	61.54	38.46
$+5$	0.00	65.38	34.62
$+6$	7.69	50.00	42.31
$+7$	11.54	53.85	34.62
$+8$	11.54	53.85	34.62
All	5.66	52.94	41.40

Table 15: Ex-post decision (in %) in the control Market Experiment

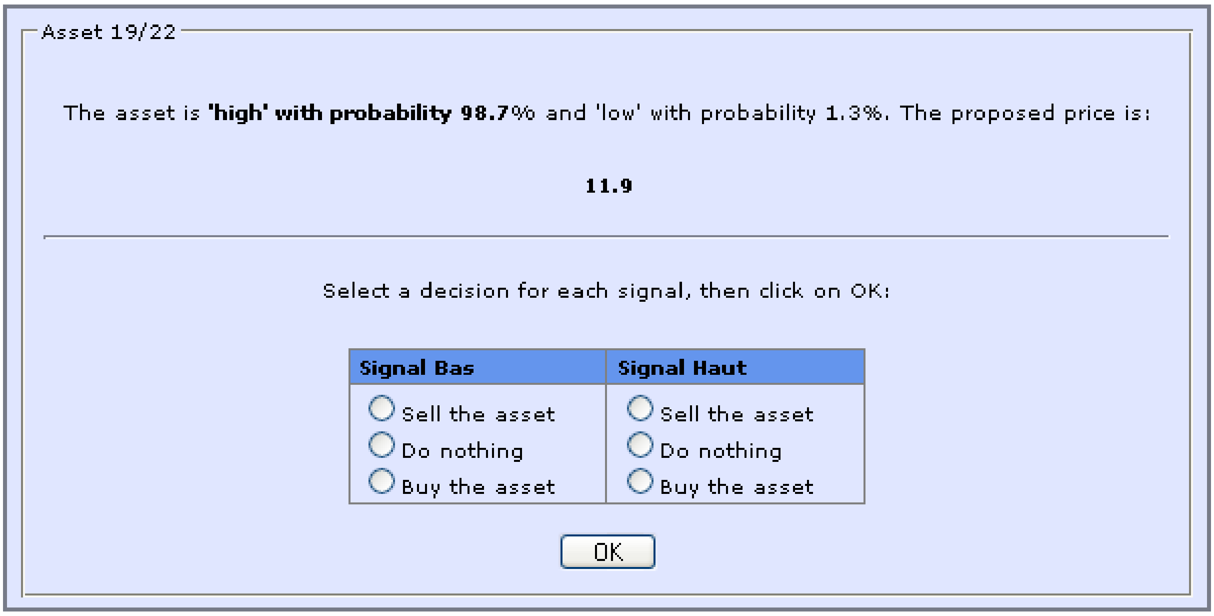

Figure 1: Screen in a Market Experiment

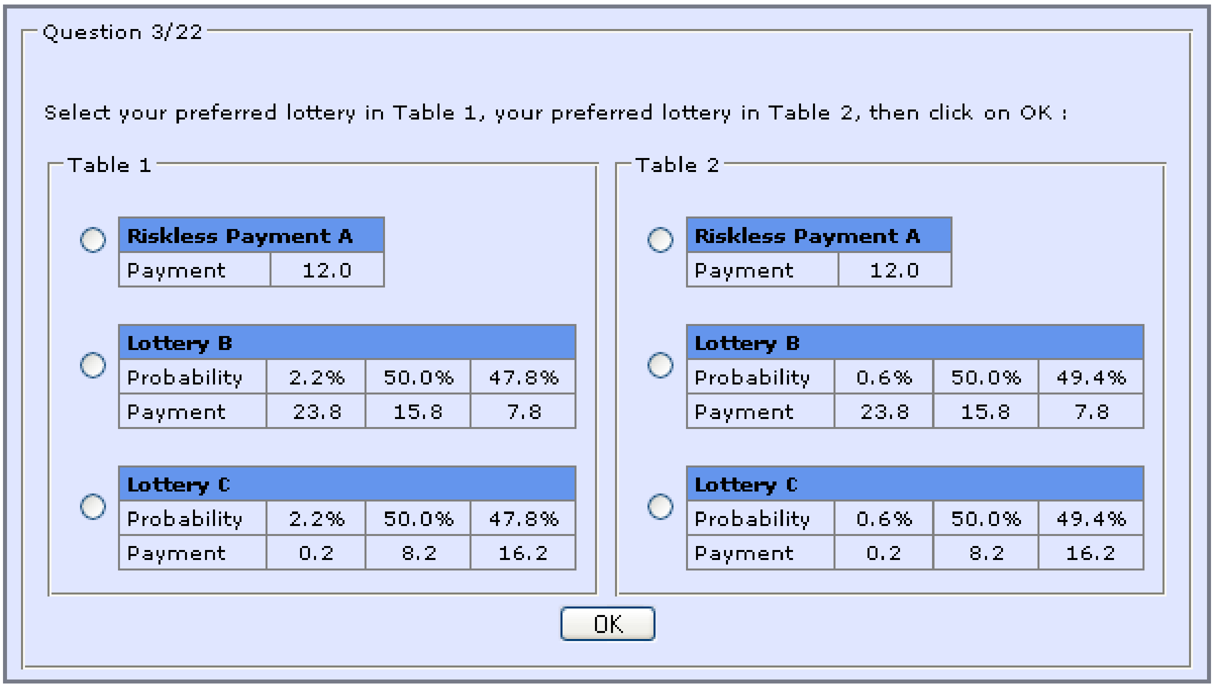

Figure 2: Screen in a Lottery Experiment