

# Risk Attitude, Beliefs Updating and the Information Content of Trades: An Experiment.

Christophe Bisière<sup>1</sup>

Jean-Paul Décamps<sup>2</sup>

Stefano Lovo<sup>3</sup>

This version: May 19, 2009

<sup>1</sup>Toulouse School of Economics (IDEI) and Institut d'Administration des Entreprises (CRG), Manufacture des Tabacs, 21, Allée de Brienne, 31000 Toulouse, France. Email: bisiere@univ-tlse1.fr. Phone: +33 (0)5 61 12 85 76.

<sup>2</sup>Toulouse School of Economics (GREMAQ-IDEI), Manufacture des Tabacs, 21, Allée de Brienne, 31000 Toulouse, France. Email: decamps@cict.fr. Phone: +33 (0)5 61 12 85 99.

<sup>3</sup>HEC School of Management, Paris, and GRECHEC, Finance and Economics Department, 1 rue de la Libération, 78351, Jouy en Josas, France. Email: lovo@hec.fr. Phone: +33 (0)1 39 67 74 62.

Décamps is an Academic Fellow at the Europlace Institute of Finance. This research benefited from the support of the Europlace Institute of Finance. Christophe Bisière and Jean-Paul Décamps gratefully thank the Chaire “Marché des risques et création de valeurs, fondation du risque/Scor”. Stefano Lovo gratefully thanks the HEC Foundation for financial support. The authors wish to thank Véronique Bessièrè, Bruno Biais, Sébastien Pouget and especially Rosemarie Nagel and Hersh Shefrin for their thoughtful feedback and suggestions. We also thank the participants at the AFSE meeting on Experimental Economics, Lyon, May 2007; the workshop on Individual and Collective Decision Making, Paris, December 2007; the Paris AFFI international meeting, Paris, December 2008; the International Financial Research Forum of the Europlace Institute of Finance (EIF), Paris, December 2008; the finance seminars at the Universities of Barcelona, Copenhagen, Louvain, Paris 9, Toulouse. We remain solely responsible for the content of this paper.

## **Abstract**

We conduct a series of experiments that simulate trading in financial markets and which allows us to identify the different effects that subjects' risk attitudes and belief updating rules have on the information content of the order flow. We find that there are very few risk-neutral subjects and that subjects displaying risk aversion or risk-loving tend to ignore private information when their prior beliefs on the asset fundamentals are strong. Consequently, private information struggles penetrating trading prices. We find evidence of non-Bayesian belief updating (confirmation bias and under-confidence). This reduces (improves) market efficiency when subjects' prior beliefs are weak (strong).

# 1 Introduction

This paper presents the results of a series of experiments that aim to measure how and to what extent investors' demand is affected by their private information regarding assets fundamentals. More specifically, we analyze how the information content of the order flow varies in relation to the intrinsic uncertainty regarding the fundamental value of the asset, and the strength of common prior beliefs on this fundamental. Since market ability to generate informational efficient prices essentially depends on the sensitivity of investors' demand to their private information, subjects' behavior in our experiment can be translated into implications regarding market informational efficiency. We find that the market is likely to fail aggregating private information when market prior beliefs are strong and/or when all private information cannot resolve uncertainty regarding the fundamentals. The experiment design we adopt allows to distinguish between the different impacts that risk attitudes and belief updating rules have on subjects' demand. Our data suggest that informational inefficiency is due to the virtual absence of risk neutral behavior but is mitigated by the presence of confirmation bias in the way subjects update their beliefs.

Consider the case of a trader who decides on the position to take on a financial asset and who observes a price reflecting public prior belief on the asset fundamentals. Any additional private information the trader may have will be construed to determine his or her posterior belief on the cash flows that the asset will generate. Since the seminal paper by Glosten and Milgrom (1985), it has often been taken for granted that if the trading price reflects public information fairly, then the position taken by a trader will reveal the nature of his or her private information, i.e., he or she will buy after receiving favorable private information on the asset perspective, or decide to sell if the private information is unfavorable. In this instance, orders will always reflect private information and as a result, prices will eventually assimilate all information dispersed in the economy. Indeed, this is what happens if one assumes that both privately informed traders and market makers, in charge of setting prices, are risk-neutral Bayesian investors. Outside the risk-neutral world, however, hedging risk matters and the direction of a trade does not necessarily corresponds to the sign of the trader's private information. As a matter of fact, Décamps and Lovo (2006a,b) show that when market makers and traders differ in risk aversion, trades do not necessarily disclose private information and informational efficiency fails. This suggests that when interpreting the information content of a trade one cannot abstract from the study of the trader's risk attitude and, more generally, from the decision process underlying the trader's choice.

Since the seminal paper by Anderson and Holt (1997), in most of the informational cascade and market efficiency experimental literature the analysis of subjects behavior is based on the assumption that subjects are risk-neutral.<sup>1</sup> Quite to the contrary, in this paper we conduct a series of experiments that allows to directly relate risk attitude, prior public beliefs and bias in interpretation of private information, to the information content of trades. This is possible thanks to four key features in our experimental design.

1. In order to detect non-Bayesian belief updating, we carried out our experiment in two different formats so that subjects had to solve equivalent problems that involved various degrees

---

<sup>1</sup>See also Huck and Oechssler (2000), Kübler and Weizsäcker (2004), Çelen and Kariv (2004), Cipriani and Guarino (2005) and Drehman et al. (2005).

of reasoning. In the format we denote *Market Experiment* (ME henceforth), subjects choose the quantity of a risky asset that they want to trade in a given market. When making their trading decisions, subjects have access to public and private information on the fundamental value of the asset. Public information consists of the *ex ante* distribution of the asset value, corresponds to the “market prior belief” and is reflected in the asset trading price. Private information consists of a private signal correlated with the true value of the asset. Subjects are asked to declare their preferred position (buy, sell, no-trade) for different levels of prior beliefs. In the format we denote *Lottery Experiment* (LE henceforth), subjects are asked to choose from among a series of different lotteries where the distribution function of payoffs is explicitly stated. More precisely, subjects choose the preferred lotteries in different menus corresponding to different levels of market prior belief in the ME. Each menu consists of three lotteries corresponding to a decision to buy, to sell or not to traded in the ME. Payoffs and probabilities in the lotteries reflect the value of a portfolio resulting from trading decisions in ME for different levels of prior beliefs. In other words, lotteries in LE are determined so that a Bayesian subject would face the exact same decision problem in both formats and therefore would make in ME the same choice as in LE. The observation of subjects’ choices in LE allows us to understand their risk attitude. The comparison of subject’s choices in the two formats enables us to detect deviations from the expected-utility or Bayesian-updating paradigms.

2. In both formats the information content of the order flow is directly observable. More precisely, by using a “strategy method”, we observe subjects choices in all potential realizations of the private signals and for different levels of prior beliefs. In doing this, we identify situations where subjects do ignore their private information and take the same action irrespectively of the sign of their information.
3. In our main treatment, the liquidation value of the asset comprises two components: one that can be learned by aggregating all private information, and an additional shock on which subjects have no private information. The latter component reflects the fact that real world investors are aware that future shocks might affect a stock’s value even if no private or public information about the sign of these shocks is available at the time of their investment. This set-up allows us to measure the effects of intrinsic risk on informational efficiency and in addition provides a further test for the hypothesis of risk-neutrality.
4. Subjects do not interact, thus lack of common knowledge of participants’ rationality cannot explain deviation from what the theory predicts.

Our observation in the LE shows that, contrary to what is implicitly assumed in most experimental papers on market efficiency, subjects are not risk neutral. We show that the behavior of about two-thirds of the subjects in LE is compatible with CARA and/or CRR utility functions, with subjects’ risk attitude ranging from a high degree of risk aversion to risk-loving. No subject can be classified as close to risk neutral. Taking into account risk attitude allows to explain subjects decision to abstain from trading, a behavior that is regularly observed in experimental literature, but which is hard to explain under the risk neutral assumption. The absence of risk neutrality has

clear implications on how the order flow information content varies with the strength of market prior beliefs. We show that the stronger the prior beliefs regarding the actual asset fundamental value, the smaller the proportion of subjects who make use of private information to determine the direction of their trade. Thus, the information impounded in the order flow shrinks when the market has a strong prior belief on what the asset value is. We find that the presence of a non-learnable shock on the asset fundamentals further deters subjects from using private information. This further reduces market ability to assimilate information on the learnable component of the asset value. These observations are broadly consistent with the theoretical prediction of Décamps and Lovo (2006a,b): when the market is sufficiently convinced about a stock prospect (positive or negative), only a few pieces of information dispersed in the economy will reach the market even when this information, if revealed, would lead to a sharp correction in the stock price.

Subjects' behavior in ME is different from the one we observe in LE. After running a control experiment we find little evidence that this difference is due to the different "framing" of the two formats. More precisely, in comparison with LE, for strong prior belief the ME presents an increase of strategies that consist of following the private signal when this confirms the prior belief, and of no-trading when the private signal and prior beliefs challenge each other. For weak prior belief, we observe an increase of no trading decisions, no matter what the private signal is. Starting from the utility function implied by a subject behavior in LE, we can measure the bias in belief updating that is implied by the subject's behavior in ME. We find evidence of confirmation bias for strong public belief and underconfidence bias for weak public belief. That is, in ME, for strong public belief, subjects tend to overweight (underweight) the information content of a private signal when this confirms (contradicts) prior beliefs while for weak public belief, subjects tend to underweight their private signal.

For each of the subjects, we observe the contingent trading strategy in LE and ME, for different levels of prior public belief and different realizations of the private signal. Thus, while in our experiment subjects did not trade sequentially, we can simulate an arbitrary large number of sequential trading histories where the behavior of virtual traders reflects the actual behavior of the pool of subjects in our experiment. These simulations generate sequences of trading prices that we use to measure market informational efficiency in a more direct way. We find that the virtual absence of risk neutral behavior significantly slows down the convergence of price to fundamentals. In ME, the information content of the order flow is lower than in LE when public prior belief is weak, while it is higher than in LE for stronger public beliefs. Thus, the effect of non-Bayesian updating is to improve information efficiency when the market is clearly bullish or clearly bearish, but to reduce efficiency when the market has no precise orientation.

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 agents' behavior. Section 3 presents the experimental design. Section 4 presents the results of the experiment.

## **2 Theoretical framework**

In this section we first describe the theoretical setting borrowed from Décamps and Lovo (2006a). Second, we illustrate the model with several numerical examples to help to explain the main pre-

dictions we tested in our experiment.

## 2.1 The model structure

We consider a discrete time sequential trade model à la Glosten and Milgrom (1985): 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, where  $\tilde{V}$  and  $\tilde{\epsilon}$  are independently distributed. 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 everyone. We assume that  $\tilde{V}$  takes value in  $\{\underline{V}, \bar{V}\}$ , where  $\underline{V} < \bar{V}$  and  $\tilde{\epsilon}$  takes value in  $\{-\epsilon, +\epsilon\}$  with  $\mathbb{P}(\tilde{\epsilon} = \epsilon) = \mathbb{P}(\tilde{V} = \bar{V}) = \frac{1}{2}$ . Each trader observes a conditionally independent and identically distributed private signal  $\tilde{s}$  with possible values  $l$  and  $h$ . We assume  $\mathbb{P}(l|\underline{V}) = \mathbb{P}(h|\bar{V}) = p$ , with  $p \in (1/2, 1)$  that implies that private signals are partially informative regarding  $\tilde{V}$ , but provide no information regarding  $\tilde{\epsilon}$ .

Time is discrete. At any period  $t$  a trader enters the market facing a unique opportunity to buy or sell one unit of the risky asset at the trading prices posted by market makers. We denote with  $H_t$  the history of trades (past quantities and prices) up to date  $t - 1$ . All agents observe  $H_t$  and update their beliefs on  $\tilde{V}$ . We denote  $\pi_t = \mathbb{P}(\bar{V}|H_t)$  the public belief at time  $t$  and  $\pi_t^s = \mathbb{P}(\bar{V}|H_t, s)$ , a trader's belief at time  $t$  given a private signal  $s \in \{l, h\}$ . Note that, because private signal precision is bounded, the closer the prior  $\pi_t$  is to 1 (or to 0) the smaller the change in belief  $|\pi_t^s - \pi_t|$  induced by the private signal will be. For this reason we adopt the following convention. We say that a prior belief  $\pi_t$  is *strong* when  $|\pi_t - 0.5|$  is large. Also, we call a public belief  $\pi_t$  that is larger than 0.5 a *positive prior*, and a public belief  $\pi_t$  that is lower than 0.5 a *negative prior*.

The demand of a trader with utility function  $u$  is:

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

where  $P_t(\cdot) : \{-1, 0, 1\} \rightarrow \mathbb{R}$  is the pricing schedule proposed by market makers. We assume  $u'$  positive and continuous but we impose no restriction on  $u''$ . Thus, our analysis takes risk-neutrality, risk-aversion and risk-loving into consideration. The variables  $m$  and  $x$  represent the trader's initial monetary wealth and his or her initial inventory in the risky asset, respectively.

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

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

All agents are Bayesian. An equilibrium is a situation where equations (1) and (2) are met at any time  $t$ . Private information enters prices when market makers can construe it from trading decisions. However, if a trader's demand is invariable with his or her private signal, nothing can be inferred from his or her order. Formally,

**Definition 1** *A trader's order is said to be non-informative when it is not affected by the private signal the trader received, i.e.,  $Q^*(P_t, H_t, h) = Q^*(P_t, H_t, l)$ .*

The larger the percentage of traders submitting non-informative orders, the smaller the flow of information that can be incorporated into the asset price.

Avery and Zemsky (1998) show that when informed traders are risk-neutral, the order flow provides a constant stream of information. In this instance, the order flow will never stop providing information and the trading price eventually converges to  $\tilde{V}$ .<sup>2</sup> Décamps and Lovo (2006a,b) show that if market makers and traders differ in their risk aversion, and if the agents' set of actions is discrete,<sup>3</sup> 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. This implies that price will stay bounded away from the realization of  $\tilde{V}$ . While we refer to Avery and Zemsky (1998) and Décamps and Lovo (2006a,b) for the formal proof of these statements, in the following section we will illustrate these findings with a numerical example that reflects the set up of our experiment.

## 2.2 An illustration of the behavior of a Bayesian expected utility maximizer

Consider the following parameters' values  $\underline{V} = 4$ ,  $\bar{V} = 12$ ,  $\epsilon = 4$  and  $p = 0.65$ . In this instance the fundamental value of the asset can take three values, i.e.,  $\tilde{v} \in \{0, 8, 16\}$ . In this illustration and throughout the paper, we will assume that agents can buy and sell at a price set at the expected asset value, conditional on the public information available at time  $t$ , i.e.,  $P(Q) = \mathbb{E}[\tilde{v}|H_t]$ .<sup>4</sup>

Let us focus on an extreme public belief  $\pi_t = \mathbb{P}(\bar{V}|H_t) = 0.9930$  corresponding to a trading price of  $P_t = \mathbb{E}[\tilde{v}|H_t] = 11.94$  and consider a Bayesian expected utility maximizer endowed with  $x = 0$  amount of the risky asset and  $m = 12$  units of money. The problem that a trader of this type faces can be represented as a choice in a menu of lotteries described in tables 1, 2 and 3. The entries in the tables report the possible payoffs resulting from the three possible trading decisions and the three possible realizations of the fundamental value  $\tilde{v}$ , i.e., *traded quantity*  $\times$   $(\tilde{v} - \text{trading price}) + 12$ . Tables 1, 2 and 3 differ only in the probabilities of obtaining the payoffs in each column.

Table 1 represents the problem faced by a trader who received no private signal. By definition, a risk-averse trader will prefer the certain payment 12 to the other two lotteries. That is to say, "No trade" is the strictly preferred action. A risk-loving trader will strictly prefer selling to the other alternatives, whereas a risk-neutral trader will be perfectly indifferent with regard to the three actions.

Now consider the same trader, but suppose that he or she received a private signal  $\tilde{s}$  with precision  $p = 0.65$ . Will the private signal affect the trader's order? Probabilities in tables 2 and 3 are obtained by the Bayesian updating of the public belief  $\pi_t = 0.9930$  following private signal  $l$

<sup>2</sup>The same feature applies in Glosten and Milgrom (1985) and Easley and O'Hara (1992)

<sup>3</sup>If agents were able to trade a continuum of quantities, risk aversion alone would not be enough to generate market inefficiency. See for instance Glosten (1989) and Vives (1995).

<sup>4</sup>This pricing convention is simpler than the one predicted by the theory where buy and sell orders are not necessarily executed at the same price. This is the pricing rule adopted in our experiment. By fixing the price at  $\mathbb{E}[\tilde{v}|H_t]$  for buy and sell orders, we increase the trader's expected profit from speculation and this reduces the incentive to adopt non-informative orders. In other words, this pricing rule should bias the results of our experiment in favor of market efficiency.

Table 1: Lotteries when no private signal is received

	$\mathbb{P}(\tilde{v} = 0) = 0.35\%$	$\mathbb{P}(\tilde{v} = 8) = 50.00\%$	$\mathbb{P}(\tilde{v} = 16) = 49.65\%$	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

and  $h$ , respectively. Thus, tables 2 and 3 represent the choices available to a trader who received a private signal  $l$  and  $h$ , respectively.

Table 2: Lotteries when a signal  $l$  is received

	$\mathbb{P}(\tilde{v} = 0) = 0.65\%$	$\mathbb{P}(\tilde{v} = 8) = 50.00\%$	$\mathbb{P}(\tilde{v} = 16) = 49.35\%$	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 3: Lotteries when a signal  $h$  is received

	$\mathbb{P}(\tilde{v} = 0) = 0.19\%$	$\mathbb{P}(\tilde{v} = 8) = 50.00\%$	$\mathbb{P}(\tilde{v} = 16) = 49.81\%$	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

It is clear from the expected value column that a risk-neutral trader will prefer to sell when  $s = l$  and to buy when  $s = h$ . In contrast, a sufficiently risk-averse agent will still prefer “No trade” to the other options available, independently from the signal he or she received. Similarly, a sufficiently risk-loving trader will strictly prefer selling to the other alternatives, and this choice will not be affected by the sign of the private signal. Hence, a trader who is either sufficiently risk averse or sufficiently risk-loving will submit a non-informative order.

This example suggests that the two ingredients that can generate non-informative orders are, on the one hand, the absence of risk neutrality and, on the other hand, the fact that a private signal affect posterior belief just slightly. Considering that when prior belief is extreme (that is when  $\pi_t$  is sufficiently close to 0 or to 1) a partially informative signal will only slightly affect posterior belief, we deduce that traders who are not risk-neutral will submit non-informative orders as soon

as prior public beliefs are sufficiently strong. On the other hand, a risk-neutral trader's order will incorporate the private signal, regardless of the strength of prior beliefs.<sup>5</sup>

The impact that risk attitude and prior belief has on trading strategies is further analyzed in Table 5 for CRR and CARA utility functions. This table presents the optimal contingent trading strategies for different levels of risk attitude and different level of public belief  $\pi$ . The first column of Table 5 reports the different level of *signal imbalance*  $g$  defined as the difference between the number of positive and negative private signals that the market has inferred from previous trades. The one-to-one mapping between signal imbalance  $g$  and public belief  $\pi$  will make the displaying of the results easier. Positive priors correspond to  $g > 0$ , negative priors correspond to  $g < 0$  and strong priors correspond to  $|g|$  close to 8. In Table 5, the random variables  $\tilde{V}$  and  $\tilde{\epsilon}$  take value in  $\{4, 12\}$  and  $\{-4, 4\}$ , respectively. In order to facilitate the analysis, we identify a contingent trading 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.<sup>6</sup>

Several comments are in order. Note first that traders strategies are symmetric with respect to  $|g|$ . For example, if a trader contingent strategy is S-N for a given level of  $g$ , then it will be N-B for signal imbalance  $-g$ .<sup>7</sup> Second, very risk-averse traders ( $\gamma > 0.078$  and  $\alpha < -0.85$  in Table 5) always choose strategy N-N, whatever the public belief. Third, strategy S-B is optimal for all the levels of public prior belief we considered only when the trader risk attitude is sufficiently close to risk-neutrality. This strategy is also optimal for traders with intermediate levels of risk aversion (or risk-loving) but only when public belief is weak (i.e.,  $|g|$  small). However, as soon as public prior belief is sufficiently strong (i.e.,  $|g|$  large), these traders will submit non-informative orders. More precisely, risk-averse traders will prefer not to trade and will ignore their private signal. Risk-loving traders (i.e.  $\gamma < -0.25$  and  $\alpha > 4.7$  in Table 5) will choose to buy when the prior is strong and negative ( $g$  close to  $-8$ ) and to sell when the prior is strong and positive ( $g$  close to 8), but in both cases they will ignore their private signal.

These remarks have a number of empirical implications at the individual level as well as at the aggregate level.

**Implication 1** *An expected utility maximizer contingent trading strategy is symmetric with respect to  $|g|$ .*

**Implication 2** *By observing an expected utility maximizer contingent trading strategy for different levels of prior belief  $\pi$ , one can estimate the trader's risk attitude. In particular, a risk-neutral trader will choose strategy S-B for all levels of  $\pi$ .*

**Implication 3** *In an economy composed of Bayesian traders that are expected-utility maximizers but differ in their risk attitude, the stronger the public belief, the higher the frequency of non-informative orders and the lower the order flow information content.*

<sup>5</sup>In fact, risk-neutrality implies *ex ante* indifference with respect to the three trading options. In this case, even if a private signal causes an arbitrarily small change in belief, this will be sufficient to swing the sign of the trader's order.

<sup>6</sup>For example, strategy N-B corresponds to no-trade when receiving signal  $l$  and buy order when receiving signal  $h$ .

<sup>7</sup>In fact, the maximization problem faced by an agent with prior  $\pi$  and private signal  $l$  ( $h$ ) is symmetric to the maximization problem corresponding to prior  $1 - \pi$  and private signal  $h$  (resp.  $l$ ). The symmetry simply requires switching sell orders into buy orders.

It is interesting to link non-informative strategies with what the herding literature classifies as herding or contrarian behavior.<sup>8</sup> A trader 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. Formally,

**Definition 2** *A subject engages in herd behavior if there exist  $\pi^* > 0.5$  ( $\pi^* < 0.5$ ) such that the subject adopts strategy B-B (resp. S-S) when  $\pi_t \geq \pi^*$  (resp.  $\pi_t \leq \pi^*$ ).*

**Definition 3** *A subject engages in contrarian behavior if there exist  $\pi^* > 0.5$  ( $\pi^* < 0.5$ ) such that the subject adopts strategy S-S (resp. B-B) for all  $\pi_t \geq \pi^*$  (resp.  $\pi_t \leq \pi^*$ ).*

Table 5 suggests that contrarian behavior should be related to risk-loving attitude, while herd behavior is not consistent with CARA or CRR utility functions and Bayesian updating.

**Implication 4** *In an economy composed of Bayesian traders that are expected-utility maximizers, contrarian behavior arises in the presence of risk-loving traders but herd behavior is not observed.*

## 2.3 Testing Bayesian updating rule

In the previous section, we assumed Bayesian updating to illustrate how subjects' contingent trading strategies are affected by their risk attitude. However subjects need not update beliefs using Bayes rule. In this case, a subject's actual behavior for a given risk attitude will differ from the one described above. In order to separate the effect of risk attitude from the effect of non-Bayesian updating, each subject participated to two formats of the experiment: the *Lottery Experiment* (LE) and the *Market Experiment* (ME). Both formats reproduce the decision problem of a trader in the economy described in Section 2.1. However, in LE questions are stated in the same form as in the menu of lotteries in tables 2 and 3. By explicitly providing subjects with the distribution function of payoffs, we ensure that the belief updating rule plays no role in their decision. In ME, subjects are first informed of the prior  $\pi$  and of the accuracy of their private signal and then asked to declare their preferred trading position contingent to the realization of the private signal. In other words, while in LE posterior probability is explicitly provided, in ME subjects have all elements necessary to derive posterior probability. The two formats are designed so that a rational Bayesian expected utility maximizer would find them perfectly equivalent.

**Implication 5** *The behavior of a Bayesian expected-utility maximizer in ME is the same as in LE.*

---

<sup>8</sup>The literature on rational herding starts with the seminal papers by Bikhchandani and Hirshleifer (1992), Welch (1992) and Banerjee (1992). Herding with endogenous prices has been recently studied in a series of papers, including Avery and Zemsky (1998), Lee (1998), Chari and Kehoe (2004), Décamps and Lovo (2006a,b). See, for instance, Hirshleifer and Teoh (2003) for a survey on herd behavior in capital markets and Chamley (2004) for an extensive study on rational herding.

### 3 Experiment Design

We performed our experiment under two different—but in some ways equivalent—formats: ME and LE. Each subject participated in both formats.<sup>9</sup> Our main treatment matches the theoretical setup described in Section 2.1 where the random variables  $\tilde{V}$  and  $\tilde{\epsilon}$  take value in  $\{4, 12\}$  and  $\{-4, 4\}$ , respectively. We also conducted two control treatments: the *Simplified Market Experiment* (SME henceforth) and the *No-Unlearnable Risk treatment* (NUR treatment henceforth) that will be detailed in sections 4.2.1 and 4.3, respectively. Below is a detailed description of the Market Experiment (ME), the Lottery Experiment (LE), the subjects’ payoff and the implementation.

**Market Experiment** The Market Experiment consisted of a series of 17 questions or “rounds”. In a given round  $\tau$  each subject was asked whether he or she wanted to buy, to sell or not to trade a given risky asset, that we will denote asset  $\tau$ . As described in Section 2, the fundamental value of asset  $\tau$  is a random variable  $\tilde{v}_\tau = \tilde{V}_\tau + \tilde{\epsilon}$  with  $\tilde{V}_\tau \in \{4, 12\}$  and  $\tilde{\epsilon} \in \{-4, 4\}$ . The trading price for asset  $\tau$  was fixed at  $P_\tau = \pi_\tau 12 + (1 - \pi_\tau)4$ . Both  $\pi_\tau$  and  $P_\tau$  were made known to the subjects in round  $\tau$  (see a screen shot in Figure 7). Moreover, in each round each subject received a private signal  $\tilde{s} \in \{h, l\}$  with precision  $p = 0.65$ . Before being informed of the private signal and after observing  $\pi_\tau$  and  $P_\tau$ , each subject was asked to declare his or her desired trade conditional on receiving private signals  $h$  or  $l$ . The only difference among the rounds was provided by the probability  $\pi_\tau$  and the corresponding trading price  $P_\tau$ . For the 17 assets, the variables  $\pi_\tau$  were determined so that they reflected the public belief obtained after observing 17 different private signal histories. More precisely, each of the 17 assets corresponded to a different unbalance  $g$  varying from  $-8$  to  $+8$ .<sup>10 11</sup>

**Lottery Experiment** The Lottery Experiment, LE, is designed so that a rational Bayesian subject would find it perfectly equivalent to ME: subjects were asked exactly the same questions in exactly the same order but with a different formulation. Instead of asking the subjects the position they would take on a given financial asset, they were asked to choose one item in a menu of lotteries. Three possible outcomes and the corresponding probabilities were specified for each lottery in the menu. Similarly to the example of Table 2 and Table 3, each lottery in a menu corresponded to the random net gain obtained from selling, no-trade and buying one unit of asset  $\tau$  given the private signal  $s$ . In order to match the strategy method implemented in ME, each subject was offered two menus at each stage and asked to choose one lottery in each menu (see Figure 8). The only difference between the two menus proposed in a given round was in the probabilities attached to each payoff. This reflected the different impact that a signal  $l$  and  $h$  would have on a subject (Bayesian) posterior probabilities. LE consisted of 17 payoff-relevant rounds, each one comprising two menus.<sup>12</sup> Overall, each subject had to choose 34 lotteries in 34 menus. The relation between

---

<sup>9</sup>LE preceded or followed ME, depending on the cohort. The order in which subjects participated in the two formats had no qualitative effect on the behavior observed. Hence, our results refer to the aggregate data across cohorts.

<sup>10</sup>Only from  $-7$  to  $+7$  for subjects in cohorts 1 and 2, that is, 15 assets.

<sup>11</sup>See for example Table 5 for a correspondence between the numbering of the asset and the corresponding  $\pi_\tau$  and  $g$ .

<sup>12</sup>Only 15 rounds for cohorts 1 and 2.

the two formats of the experiment was never mentioned to the subjects, nor was the fact that the two experiments were equivalent from the perspective of a Bayesian rational subject.

**Implementation** The experiment was carried out at HEC Paris and Toulouse University. We recruited 227 subjects from undergraduate finance classes. The subjects had no previous experience in financial market experiments. Between 10 and 43 subjects participated in each session as decision makers. The main treatment involved 134 subjects (5 cohorts) while the two control treatments concerned 93 subjects (2 cohorts for the NUR treatment and 3 cohorts for SME). At the beginning of a session, we gave written instruction which were also read aloud by an experiment administrator. Then two trial sub-sessions, each involving the trade of three assets, 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 the subjects' questions regarding the rules of the game up until the distribution of the questionnaire. After this, subjects were not allowed to ask additional questions and administrators ensured that no form of communication among the subjects took place. Throughout the experiment participants were unable to observe each other's screens. Each experiment lasted about an hour and a half. An average of €19.87 was paid to each subject. Subjects were also rewarded with bonus points enabling them to raise their marks in the Financial Markets course.<sup>13</sup> Subjects' payoffs were determined on the basis of the gain on one round only, both for the market experiment and the lottery experiment. These rounds were randomly selected at the end of the experiment.<sup>14</sup> We discarded from our dataset the decisions of 5 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 4,132 for 130 subjects in the main treatment, 1,428 for 42 subjects in the NUR treatment and 850 for 50 subjects in the SME treatment.

## 4 Experimental Results

This section describes the main results of the experiment. Sections 4.1 and 4.2 examine the results obtained in LE and ME, respectively. Section 4.3 discusses the effects of intrinsic uncertainty. Section 4.4 is devoted to the analysis of market informational efficiency.

In this section, a major concern is how decisions change with the strength of prior belief. For the sake of clarity, we report the results by class of signal imbalance, defined as follows:

---

<sup>13</sup>This device is standard in the literature (see for instance Camerer and Hogarth, 1999; Williams, 2008; Biais et al., 2005) and allow to incentivize participants in their experiments without distorting their risk attitude.

<sup>14</sup>See the experiment instructions for a precise description of the algorithm determining a subject's payoffs.

Class of signal imbalance	Strength of belief
$g \in [-8, -5]$	Strong negative belief
$g \in [-4, -1]$	Weak negative belief
$g = 0$	Neutral belief
$g \in [+1, +4]$	Weak positive belief
$g \in [+5, +8]$	Strong positive belief

## 4.1 Lottery Experiment

In LE, probabilities attached to each possible event are explicitly provided. Thus, for this format an expected utility maximizer’s decision depends only on the shape of his or her utility function and not on the way he or she interprets public and private information. As a consequence, LE provides a simple framework for judging whether the subjects’ behavior can be explained by the expected utility assumption. This format also allows us to measure subjects’ risk attitude.

Before proceeding with the detailed analysis of the data, we will focus on the main results of the Lottery Experiment for the main treatment. First, we find that none of the participants can be considered as risk neutral subjects. Second, the observed impact of risk attitude and prior belief on trading strategies fits the theoretical properties summarized in Table 5. This is highlighted in Figure 1 page 13, which provides a general overview of the distribution of subjects’ strategies in LE as a function of different levels  $g$  of signal imbalance. We observe a symmetry of the contingent orders with respect to  $|g|$ . There is a peak of strategies S-B for  $g = 0$ . The proportion of strategies N-N increase with  $|g|$ . Strategies B-B decrease with  $g$  while strategies S-S increase with  $g$ . We observe a peak of strategies N-B for weakly negative priors (that is  $g \in [-4, -1]$ ) and, symmetrically, a peak of strategy S-N for weak positive priors. The fraction of non-informative strategies (i.e., sum of N-N, B-B and S-S) as a function of  $g$  displays a U-shape. That is, the information content of the order flow decreases when prior beliefs are strong. We now turn to the detailed analysis.

### 4.1.1 Symmetry of choices

As we pointed out in Section 2, the decision problem a subject faces for a given level of public belief  $\pi$  is the same as the one he or she faces when public belief is  $1 - \pi$ . This feature is apparent in LE as the lottery in the menus corresponding to prior  $\pi$  and to prior  $1 - \pi$  are the same but are presented in a different order. Thus, rationality requires that a subject’s preferred lottery does not depend on how lotteries are ranked in the menu. For each subject we compute a “symmetry score” by comparing the preferred lottery for a given prior  $\pi$  with the one chosen for prior  $1 - \pi$ . The score gives us the proportion of a subject’s choices that respects the symmetry rule.<sup>15</sup> Thus, the closer the score is to 1, the more the subject’s behavior is compatible with rationality. The median symmetry score is 0.81 with 75% of subjects having scored higher than 0.57.<sup>16</sup> Overall, these data

<sup>15</sup>Precisely, for each subject we count 0.5 each time the conditional decision is symmetric for public beliefs  $\pi$  and  $1 - \pi$ . The symmetry score of a subject is then obtained by dividing this count by the count corresponding to a fully symmetric strategy profile.

<sup>16</sup>The cumulative distribution of the symmetry score is given in Figure 10.

suggest that the subjects' behavior is consistent with Implication 1 and does not contradict this basic test of rationality. Nevertheless, there is a small fraction of subjects (8.46%) that clearly behaves inconsistently and displays a symmetry score of less than 0.33.

#### 4.1.2 Risk attitude

In order to infer subjects' risk attitude, we compare each subject's actual behavior in LE with the behaviors predicted by CARA and CRR utility functions for different levels of risk aversion. Each subject is then assigned the utility function (CARA or CRR) and a risk aversion parameter that more accurately matches the subject's observed behavior.<sup>17</sup> After that we assign a matching score to each subject between the observed behavior and closest theoretical behavior. The closer the matching score is to 1 the better the subject's behavior can be explained with a CARA or a CRR utility function. For 47 of the 130 subjects of the main treatment, the matching score is less than 0.75. The median matching score of the remaining 83 subjects is 0.97. Thus, for 63.85% of the subjects we found a utility function that explains at least 75% of the subjects' choices.<sup>18</sup>

The distribution of risk aversion is reported in Table 6 for CARA and CRR utility functions. Subjects displaying a strong enough risk aversion to induce N-N in at least 75% of the situations represent 31.54% of the population. The percentage of risk-loving subjects is 11.54%, while 20.77% of the subjects display intermediate levels of risk aversion. The remaining 36.15% has a matching score below 0.75. Surprisingly none of the participants in the main treatment come close to behaving in a risk neutral way. This finding is in sharp contradiction with the assumption that the subjects are risk neutral.

#### 4.1.3 Informative orders and public information

In order to understand whether and how the strength of prior belief affects the information content of the order flow, we study how the proportion of subjects using non-informative orders (corresponding to strategies N-N, S-S, and B-B) changes with the signal imbalance.

Table 7 reports the frequency of informative trades by strength of belief. In LE, we observe that the percentage of informative orders decreases with the strength of belief. That is, the proportion of non-informative orders is a U-shaped function of the strength of belief. In other words, subjects tend to ignore their private information more often when the public information is sufficiently strong. Namely, informative orders represent 52.31% of all trades for neutral prior belief and this percentage falls to 16.29% and 20.54% for strong negative and strong positive prior belief,

---

<sup>17</sup>For each subject, the matching score is computed as follows. Let us start with the CARA function. A given value of the risk parameter of a CARA utility function implies a certain theoretical behavior, i.e. a set of conditional decisions, one for each level of signal imbalance. For each of these theoretical decisions, we compute a matching score with the observed behavior, by counting 0.5 for each side of the conditional decision that matches the observed behavior, and dividing this count by the number of signal imbalances. The score is thus a number between 0 (no matches) and 1 (perfect match). The theoretical behavior with the highest score gives us the aversion CARA utility function parameter assigned to the subject. We then repeat this procedure for a CRR utility function, and obtain the aversion parameter for this class of function. Finally, the subject's highest overall matching score determines an aversion parameter and a class of function (CARA or CRR) that best reflect his behavior.

<sup>18</sup>The cumulative distribution of the matching score is given in Figure 11.

respectively. This is consistent with Implication 3 on market informational efficiency: the flow of information decreases with the strength of public belief.

#### 4.1.4 Herding and contrarian behavior

Figure 1 displays the proportion of each strategy, by class of belief.<sup>19</sup> Corresponding percentages are reported in Table 8.

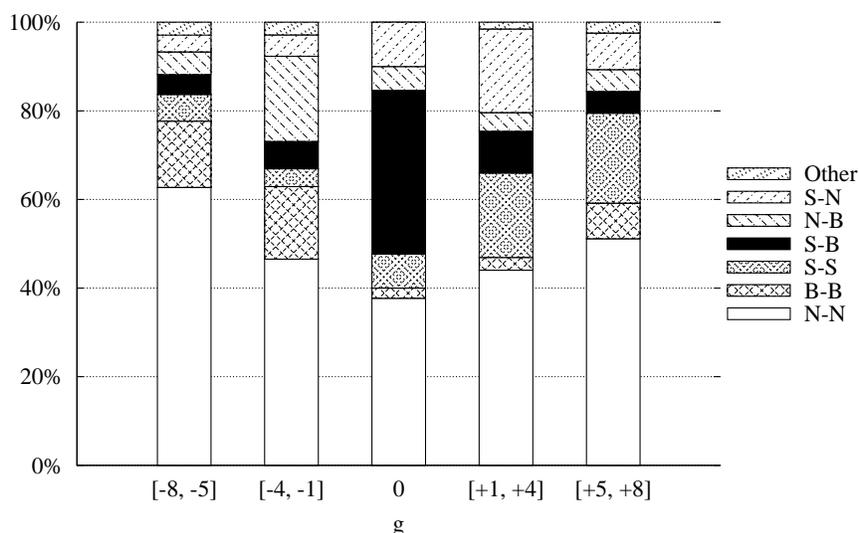


Figure 1: Distribution of subjects' strategies in Lottery Experiment (main treatment)

Strategies S-S (B-B) are related to contrarian behavior when they occur for positive (negative) priors. We observe both buy and sell contrarian behavior. The frequency of strategies S-S for neutral priors is 7.69%, and is much higher for strong positive priors (20.31%). Similarly, the frequency of strategies B-B rises from 2.31% for neutral priors, to 14.96% for strong negative priors. According to Implication 4, contrarian behavior can be attributed to the fraction of subjects that are risk-loving. It is important to stress that because of the specific design of our experiment the presence of contrarian behavior cannot be ascribed to the lack of common knowledge of agents' rationality but should be linked directly to the subjects' risk attitude.

Strategies S-S (B-B) are related to herd behavior when they are associated with negative (positive) priors. While our theory gives no justification for herding strategies, a small fraction of the subjects engaged nevertheless in herd behavior. In particular, sell herding (S-S for strong negative prior belief) amounts to 6.03% and buy herding (B-B for strong positive prior belief) to 8.04%.

<sup>19</sup>Strategy "Other" includes B-N, B-S, and N-S.

## 4.2 Market Experiment

Implication 5 of the theory suggests that subjects will choose exactly the same conditional strategies in LE and in ME. This prediction is clearly rejected by our data. Overall, subjects' answers are the same for the LE and the ME only in 42.64% of the observations. For only 27.69% of subjects, the answers in LE and ME were identical in at least 75% of the questions. In most of these cases subjects preferred strategy N-N for all levels of prior belief in both formats. More specifically, for every level of signal imbalance  $g$  we run a Bhapkar test of marginal homogeneity and rejects at the 1% significance level the hypothesis that the format (LE or ME) has no effect on the frequency of conditional decisions.<sup>20</sup> Figure 2 and the corresponding Table 10, summarize the distribution of the subjects' strategies.

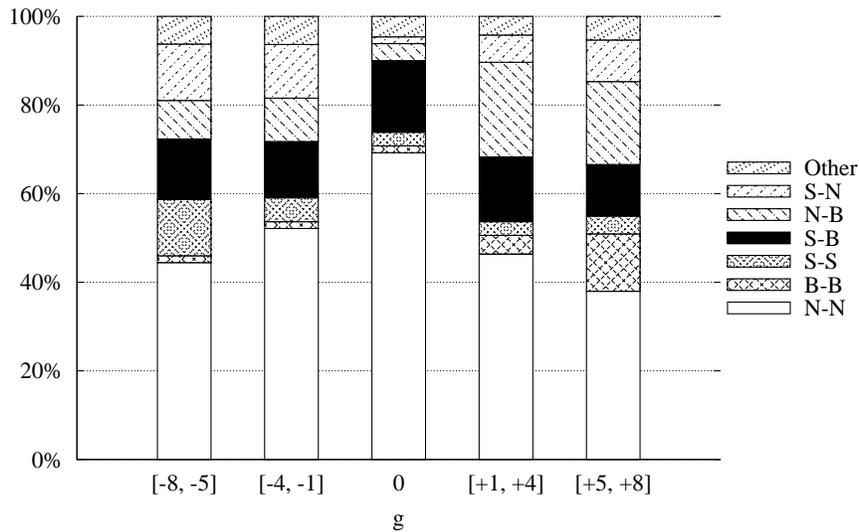


Figure 2: Distribution of subjects' strategies in Market Experiment main treatment

The only common pattern with LE is the symmetry of the subjects' choices. The rest of the observations differ. First, contrarian trades (B-B for negative priors and S-S for positive priors) tend to disappear in the ME. Second, strategies consisting in following the signal whenever this confirms the public history, and in not trading otherwise (i.e. S-N for negative priors and N-B for positive priors) are more frequent in ME than in LE. Third, for strong (resp. neutral) priors, strategies N-N are less frequent (resp. more frequent) in ME than in LE. Fourth, the frequency of strategies S-B increases in ME.<sup>21</sup> Fifth, herd behavior increases in ME.

The effect of the ME format on the order flow information content is ambiguous. As illustrated in Table 7, informative strategies rise from 28.17% of all choices in LE to 42.35% in ME. In comparison with LE, the frequency of informative strategies in ME increases for strong belief, but decreases for neutral belief. This suggests that when public information is weak, the order

<sup>20</sup>For a presentation of tests of marginal homogeneity, we refer to Davis and Holt (1993) and Agresti (2002).

<sup>21</sup>With the exception of  $g = 0$ .

flow information content is lower in ME compared to LE. However, for strong prior belief in ME, private information will be better signaled through subjects' strategies. As a result, the non-informative contingent strategies N-N, B-B and S-S as a function of beliefs display a humped shape as illustrated in Figure 2.

#### 4.2.1 Non-Bayesian updating or framing effect?

There are at least two possible explanations for the discordance in observed behaviors in the two formats. First, subjects may behave inconsistently because of the difference in the framing of LE and ME. For instance, while there is no direct reference to financial markets in the way LE is presented, in ME subjects are asked to take trading decision in financial assets. This framing could trigger heuristic behaviors in ME that are absent in LE.<sup>22</sup> Another possible explanation for the difference in the subjects' behavior in the two formats is that in LE probabilities are explicitly provided, while in ME subjects have to interpret public and private information when forming their decision. Subjects who do not conform to Bayes rule could behave differently in ME and LE.

To test for the framing hypothesis, we ran a control format that we called Simplified Market Experiment (SME). This format takes the frame of the ME with the difference that subjects do not have to interpret private and public information. Namely, subjects had the opportunity to trade a financial asset at a given price, similarly to ME. However, instead of providing subjects with prior belief and private signals, we directly supplied them with the posterior probability that  $\tilde{V} = 12$ . These probabilities were computed as follows. For a given asset  $\tau$ , we took into account the prior belief corresponding to the proposed trading price and updated this belief –using Bayes rule– following either a signal  $l$  or a signal  $h$ . Figure 9 shows the screen layout presented to subjects. Also, for this format, the predicted behavior of a Bayesian expected utility maximizer is identical to the one in LE or ME.

Results for this experiment are reported in Figure 3 and Table 12. Overall, the subjects' behavior in SME is closer to the one we observed in LE than to the one we observed in ME. In particular, the order flow information content decreases as the prior belief becomes stronger, similarly to what happens in LE. Also, strategies S-B (N-N) are more (resp. less) frequent for a neutral prior, as observed in LE. In one respect, however, behavior in SME is closer to that in ME than in LE: herd behavior is more frequent in SME than in LE. Taken as a whole, SME suggests that, with the exception of the insurgence of herd behavior, the framing has little impact on the subjects' trading decisions.

In order to analyze the hypothesis of non-Bayesian updating, we proceed as follows. We focus on the 83 subjects for which there exists a CARA or a CRR utility function that explains at least 75% of their choices in the LE of the main treatment. Given a subject  $i$  matching utility function, a level of  $g$  and a private signal  $s$ , we look for the set of posterior beliefs that generates the subject's behavior in ME for that  $g$  and that  $s$ . Within this set of posterior beliefs, we focus on the one closest to the Bayesian posterior belief and we denote it  $\hat{\pi}_i(g, s)$ . Repeating the process for both realizations of the private signals in all levels of  $g$ , we can correlate a subject  $i$  with a point-wise function  $\hat{\pi}_i(g, s)$  mapping prior belief  $g \in \{-8, -7, \dots, +7, +8\}$  and private signals  $s \in \{l, h\}$

---

<sup>22</sup>Some examples of heuristic behavior: buying (selling) when the prior belief is strong and positive (resp. negative); always trading according to the signal; buying when the price is low and selling when the price is high, etc.

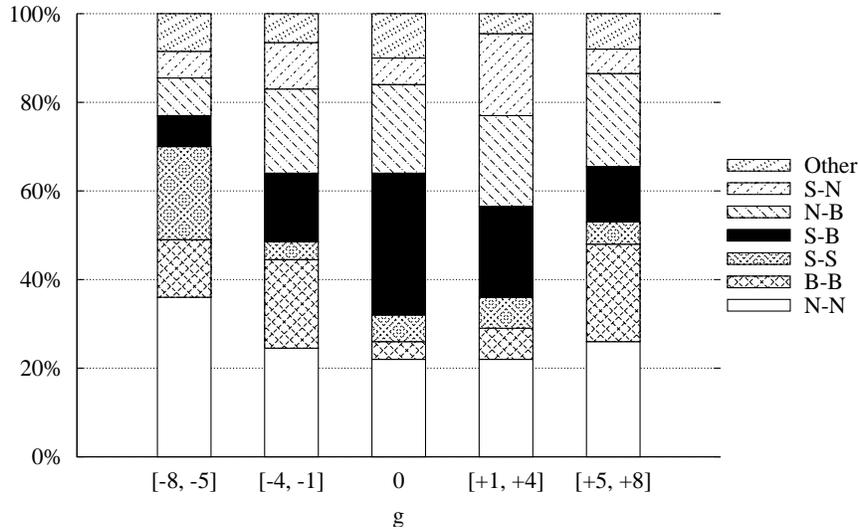


Figure 3: Distribution of subjects' strategies in Simplified Market Experiment

into posterior beliefs. We define subject  $i$ 's "Bayesian-score" as the fraction of the 34 posterior beliefs  $\hat{\pi}_i(g, s)$  that equal the corresponding Bayesian posterior beliefs. We will say that a subject is non-Bayesian if his or her Bayesian-score is lower than 0.75. According to this definition, 29 of the 83 subjects considered for this analysis are non-Bayesian. Figure 4 represents the average bias in posterior beliefs for subjects that are non-Bayesian.<sup>23</sup>

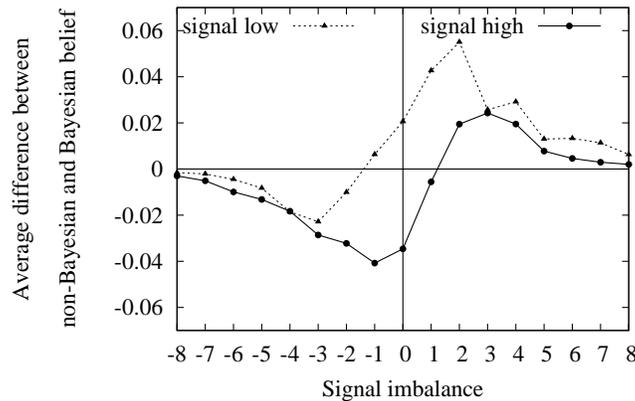


Figure 4: Bias in belief updating rule

For relatively strong public belief ( $g \geq 2$  and  $g \leq -2$ ), we find that on average non-Bayesian subjects display confirmation bias and tend to interpret private signals in a way that either confirms

<sup>23</sup>For some subjects, the CARA and CRR utility functions explain the behavior in LE equally well. In these cases, the utility function used to determine  $\hat{\pi}_i(g, s)$  is the one minimizing the average bias on posterior beliefs.

prior belief or does not challenge it. Namely Figure 4 shows that, when  $g \geq 2$  ( $g \leq -2$ ), subjects' behaviors in ME can be explained with posterior beliefs that are higher (lower) than posteriors derived using the Bayes rule. As an illustration, consider a sufficiently positive public belief, i.e.  $g \geq 2$ . The impact of a signal *high* is reinforced by the public belief and traders are more inclined to buy when they receive a signal of this type. In addition, the impact of a signal *low* is compensated by the positive public belief. It follows that subjects receiving a signal *low* are more inclined not to trade or even to buy. Consequently, contrarian behaviors disappear, herding behaviors appear and we observe a larger proportion of N-B strategies than in LE. The argument is symmetric for strong negative prior beliefs.

Interestingly, for public belief  $g$  between  $-1$  and  $1$ , the analysis differs. In that case, we find evidence of underconfidence, namely that subjects undervalue their private signal. As shown in Figure 4, for a prior belief of around  $0.5$ , the behavior of subjects with signal *low* (*high*) in ME can be explained by posterior beliefs that are higher (lower) than posteriors derived using the Bayes rule. Underweighting the private signal creates a sort of additional uncertainty that leads to a peak of N-N contingent orders.

### 4.3 Effect of intrinsic uncertainty

We now describe the results of the NUR treatment in which the intrinsic risk component  $\tilde{\epsilon}$  is absent, i.e.,  $\underline{V} = 4, \overline{V} = 12, \epsilon = 0$ . By comparing it with our main treatment, we seek to better understand the effect that a non-learnable risky component  $\tilde{\epsilon}$  has on the information content of the order flow.

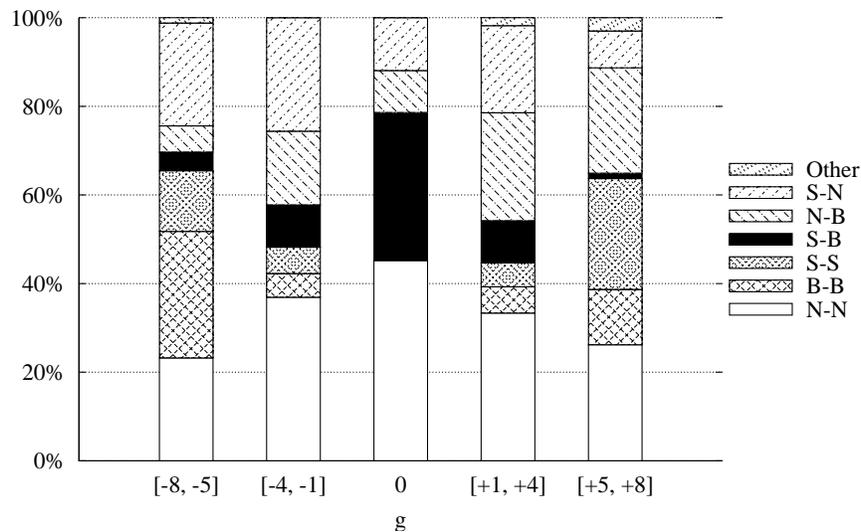


Figure 5: Distribution of subjects' strategies in Lottery Experiment (NUR treatment)

The main features concerning the distribution of subjects' strategies obtained in the experiments are the same as in the main treatment, on the whole. In particular, non-informative contingent strategies as a function of beliefs display a U-shape in LE and a humped shape in ME. Figure 5

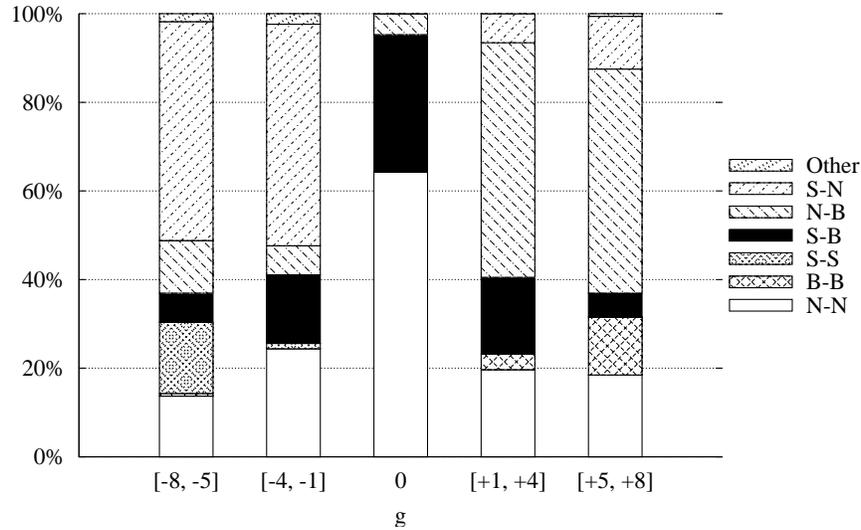


Figure 6: Distribution of subjects' strategies in Market Experiment (NUR treatment)

and 6 synthesize these observations. There are, however, two noticeable differences between the two treatments. First, the frequency of informative orders for a strong negative and strong positive prior belief is higher for the NUR treatment. For strong negative priors, informative orders in LE rise from the 16.29% of the main treatment to 34.52%, and for strong positive priors, from 20.54% to 36.31%.<sup>24</sup> Thus, the absence of the non-learnable component  $\tilde{\epsilon}$  increases subjects' sensitivity to private information. In other words, the flow of information decreases with the level of intrinsic uncertainty regarding the fundamental value of the asset. Second, in LE, the frequency of herd behavior is more marked in the NUR treatment than in the main treatment.<sup>25</sup> Specifically, sell herding (S-S for strong negative prior belief) in the NUR treatment amounts to 13.69% and buy herding (B-B for strong positive prior belief) to 12.50%. In the main treatment, these percentages fall to 6.04% and 8.04%. A possible explanation could be the subjects' tendency to round the probabilities of the lotteries and/or to ignore events whose probability is sufficiently small. In the control treatment, the event  $\tilde{V} = V$  is considered by traders with strong beliefs as virtually sure when  $|g|$  is large. Hence, independently of the realization of the private signal, subjects will buy (sell) when their belief is sufficiently positive (resp. negative). This phenomenon is mitigated in the main treatment where the additional uncertainty on  $\tilde{\epsilon}$  makes any trade intrinsically risky even when  $\tilde{V}$  is virtually known. Section 4.4, devoted to the market informational efficiency, will go into the analysis of the effect of intrinsic uncertainty in greater detail.

<sup>24</sup>See Table 7.

<sup>25</sup>See tables 8 and 9.

## 4.4 Market Informational Efficiency

In the previous section we have shown that, in LE, the proportion of non-informative orders is a U-shaped function of the signal imbalance, indicating that the information flow decreases as prior beliefs become stronger. Subjects' behavior in ME displays the opposite pattern, with the percentage of non-informative orders decreasing with the strength of public belief. In this section we try to measure the actual impact of these figures on the price dynamics and on market information efficiency. Within a sequential trade framework, market informational efficiency can be measured by the evolution of the pricing error defined as the difference between the actual price and the full information price, i.e., the price that would prevail had market makers directly observed past traders' private information. In our experiment, however, we do not observe trading histories as the subjects did not trade sequentially. Nevertheless, we can simulate an arbitrary large number of virtual trading histories and measure the average pricing error for histories of different lengths. To do this, we exploit our observations of contingent trading strategies at different levels of public belief. In order to generate virtual trading histories, we assume that virtual subjects randomly come to the market to trade once and behave as the real subjects did in the actual experiment.<sup>26</sup> As in our experiment, trading prices always reflect prior objective probabilities, our simulations reflect situations where for any given past history of traders, subjects believe that the asset is correctly priced by market makers. After each trading round, public belief and price are updated in a way that reflects the assumptions of the theory in Section 2. Namely, the price updating rule is based on the hypothesis that market makers and traders do not know the identity of past traders. However, they have a correct estimation of 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 trading strategy is adopted by traders. These frequencies are those observed in the experiment and are summarized in tables 8 to 11.<sup>27</sup> After observing a given action, the public belief will change according to the Bayesian probability that the order comes from someone who received a signal  $l$  or a signal  $h$ .<sup>28</sup> The trading price is updated accordingly. We will denote this pricing rule R1.

We simulated about 5,000 trading histories per treatment and format, with each trading history covering a maximum of 20 trading rounds. Figure 12 reports the evolution of the average pricing error in the main treatment. The pricing error in ME is consistently higher when compared to that of LE. After 20 trades, the average (median) pricing error is 30% (25.14%) in LE and 34% (34.6%) in ME. Figure 14 reports the distribution of pricing errors at the 20<sup>th</sup> round of trade. For LE (ME) in the main treatment, we find that in 14.52% (13.94%) of histories the pricing error is less than 10%,

---

<sup>26</sup>We used the following algorithm to generate a virtual trading history. At the beginning of the trading history the value of  $\tilde{V}$  is randomly determined according to  $\mathbb{P}(\tilde{V} = \bar{V}) = \frac{1}{2}$  and the initial public belief is fixed at  $\pi_0 = 0.5$ . First, in each trading round, one among the nine possible trading strategies is randomly selected in a way that reflects the empirical frequencies observed in the experiments. Note that these frequencies change with the level of public belief, as well as with the treatment and format used for the simulation. Second, a private signal is randomly determined so that it is consistent with probability 65%. The virtual trader's order is the one corresponding to the trading strategy and private signal determined in the previous two steps. Finally, the public belief is updated and a new virtual trading round starts.

<sup>27</sup>It is worth noting that these frequencies take into account all the traders including the 36.15% of subjects for whom we cannot assign a utility function.

<sup>28</sup>For the simulation, this new belief is approximated to the closest point on the grid of belief used for the experiment.

while in 8.34% (8.68%) of histories the pricing error is higher than 70%, suggesting that in both formats information cascades in the “wrong direction” are not unusual. Figures 13 and 15 report the evolution of the average pricing error in the NUR treatment. The absence of the additional risk  $\tilde{\epsilon}$  improves the information content of the order flow leading to an average (median) pricing error at the 20<sup>th</sup> round of 24.4% (14.57%) and 22% (16.92%) in LE and ME, respectively. In comparison with the main treatment, the average pricing error and the frequency of wrong cascades are lower. Interestingly, while simulations based on subjects’ behavior in LE provide more efficient prices in the short run, ME generates more efficient prices in the long run.

Overall, our simulations suggest that market information efficiency is reduced in the presence of additional risk regarding the asset fundamentals and that this phenomenon is amplified by non-Bayesian behavior, at least in the short run. In fact, the public and private information regarding the  $\tilde{V}$  component is the same in the main treatment and the control treatment. However the presence of an additional non-learnable component  $\tilde{\epsilon}$ , coupled with the virtual absence of risk-neutral traders, has the effect of reducing the market ability to learn  $\tilde{V}$ . Non-Bayesian behavior has an ambiguous effect. On the one hand, it reduces the information content of the order flow when public belief is weak, and as a result prices and public belief tend to stagnate at first. On the other hand, if public belief reaches some strength, non-Bayesian behavior tends to increase the information flow.

## 5 Conclusion

In this paper, we report results of an experiment that simulates trading in the financial market. We adopted two formats for our experiment: the Lottery Experiment and the Market Experiment. This allowed us to unambiguously measure the information content of the order flow and to disentangle the impact that risk attitude and non-Bayesian updating have on it. We show that many of the so called “irrational” behaviors are not so if one takes into account subjects’ risk attitude. The design of LE enabled us to measure subjects’ risk attitude. We found that CARA and CRR utility functions can explain the behavior of about two-third of the subjects in LE. While no subject has a risk neutral behavior, about half of the subject display risk aversion and subjects displaying risk-loving represent about one-ninth of the population. These risk attitudes have the effect of reducing the information content of the order flow when market participants have strong prior beliefs on the asset fundamentals. Contrary to what is predicted by the theory, subjects behaved differently in ME and LE. This discrepancy can be ascribed to non-Bayesian belief updating and to a lesser degree to framing effects. More specifically, by considering that subjects’ utility functions are consistent with their behaviors in LE, we find that confirmation bias and underconfidence seems to fit the subjects’ behavior in ME. The effect of non-Bayesian belief updating on market efficiency is ambiguous. Non-Bayesian updating reduces (improves) the information flow when the subject’s prior belief is weak (strong).

## References

Agresti, A., 2002. Categorical Data Analysis, 2nd Edition. John Wiley and Sons.

- Anderson, L., Holt, C., 1997. Information cascades in the laboratory. *American Economic Review* 87, 847–862.
- Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review* 88, 724–748.
- Banerjee, A., 1992. A simple model of herd behavior. *Quarterly Journal of Economics* 107, 797–817.
- Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgmental overconfidence, self monitoring and trading performance in an experimental financial market. *Review of Economic Studies* 72, 287–313.
- Bikhchandani, S., Hirshleifer, D. Welch, I., 1992. A theory of fads, fashion, customs and cultural change as informational cascades. *Journal of Political Economy* 100, 992–1026.
- Camerer, C., Hogarth, R., 1999. The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty* 19, 7–42.
- Çelen, B., Kariv, S., 2004. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review* 94, 484–498.
- Chamley, C., 2004. *Rational herds: Economic models of social learning*. Cambridge University Press, Cambridge.
- Chari, V., Kehoe, P., 2004. Financial crisis as herds: Overturning the critiques. *Journal of Economic Theory* 119, 128–150.
- Cipriani, M., Guarino, A., 2005. Herd behavior in a laboratory financial market. *American Economic Review* 95, 1427–1443.
- Davis, D., Holt, C., 1993. *Experimental Economics*. Princeton University Press, Princeton.
- Décamps, J.-P., Lovo, S., 2006a. Informational cascades with endogenous prices: The role of risk aversion. *Journal of Mathematical Economics* 42, 109–120.
- Décamps, J.-P., Lovo, S., 2006b. A note on risk aversion and herd behavior in financial markets. *The Geneva Risk and Insurance Review* 31, 35–42.
- Drehman, M., Oechssler, J., Roider, A., 2005. Herding and contrarian behavior in financial markets: An internet experiment. *American Economic Review* 95, 1403–1426.
- Easley, D., O’Hara, M., 1992. Time and the process of security price adjustment. *Journal of Finance* 47, 577–605.
- Glosten, L., 1989. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business* 62, 211–235.

- Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed agents. *Journal of Financial Economics* 14, 71–100.
- Hirshleifer, D., Teoh, S., 2003. Herd behaviour and cascading in capital markets: a review and synthesis. *European Financial Management* 9, 25–66.
- Huck, S., Oechssler, J., 2000. Information cascades in the laboratory: Do they occur for the right reasons? *Journal of Economic Psychology* 21, 661–671.
- Kübler, D., Weizsäcker, G., 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies* 71, 425–441.
- Lee, I., 1998. Market crashes and informational avalanches. *Review of Economic Studies* 65, 741–759.
- Vives, X., 1995. The speed of information revelation in a financial market mechanism. *Journal of Economic Theory* 67, 178–204.
- Welch, I., 1992. Sequential sales, learning, and cascades. *Journal of Finance* 47, 695–732.
- Williams, A., 2008. Price bubbles in large financial asset markets. In: Plott, C. R., Smith, V. L. (Eds.), *Handbook of Experimental Economics Results*. Vol. 1. North Holland, Amsterdam, pp. 242–246.

Table 4: Cohorts and participants in the experiment

Cohort #	Treatment	Order	# subjects	# discarded
1	Main	LE-ME	43	2
2	Main	LE-ME	32	1
3	NUR	LE-ME	16	0
4	NUR	ME-LE	26	0
5	Main	ME-LE	18	1
6	Main	ME-LE	21	0
7	Main	ME-LE	20	0
8	SME	-	10	0
9	SME	-	21	0
10	SME	-	20	1

Table 5: Optimal strategies for an investor with CRR and CARA utility

CRR:	$U(x) = \frac{x^{(1+\alpha)}}{(1+\alpha)}$	$\alpha < -0.85$	$\alpha = -0.25$	$-0.034 < \alpha < 4.7$	$\alpha = 4.83$	$\alpha = 5.91$
CARA:	$U(x) = -\gamma e^{-\gamma x}$	$\gamma > 0.078$	$\gamma = 0.02$	$-0.25 < \gamma < 0.003$	$\gamma = -0.26$	$\gamma = -0.32$
$g$	$\pi$					
-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,S	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 6: Subject's risk attitude

	CARA $U(x) = -\gamma e^{-\gamma x}$	CRR $U(x) = \frac{x^{(1+\alpha)}}{(1+\alpha)}$	Number of subjects	%	Average matching score
High risk averse	$\gamma > 0.078$	$\alpha < -0.85$	41	31.54	97.47
Medium risk averse	$0.005 < \gamma < 0.078$	$-0.85 < \alpha < -0.065$	27	20.77	87.62
Close to risk neutral	$-0.25 < \gamma < 0.005$	$-0.065 < \alpha < 4.7$	0	0.00	—
Risk-loving	$\gamma < -0.25$	$\alpha > 4.7$	15	11.54	89.82
Not classed			47	36.15	61.44
All			130	100.00	81.52

Table 7: Informative orders (in %) in the main treatment and in the NUR control treatments

Signal imbalance $g$	Main treatment:		NUR treatment:	
	ME	LE	ME	LE
$g \in [-8, -5]$	41.29	16.29	69.64	34.52
$g \in [-4, -1]$	40.96	33.08	74.40	51.79
$g = 0$	26.15	52.31	35.71	54.76
$g \in [+1, +4]$	46.35	34.04	76.79	55.36
$g \in [+5, +8]$	45.09	20.54	68.45	36.31
All	42.35	28.17	70.17	45.10

Table 8: Conditional decisions (in %) in the Lottery Experiment (main treatment)

Signal imbalance $g$	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
$g \in [-8, -5]$	14.96	0.45	1.12	5.13	62.72	1.34	4.46	3.79	6.03
$g \in [-4, -1]$	16.35	1.15	0.77	19.23	46.54	0.96	6.15	4.81	4.04
$g = 0$	2.31	0.00	0.00	5.38	37.69	0.00	36.92	10.00	7.69
$g \in [+1, +4]$	2.88	0.38	0.58	4.23	44.04	0.58	9.42	18.85	19.04
$g \in [+5, +8]$	8.04	0.22	0.67	4.91	51.12	1.56	4.91	8.26	20.31

Table 9: Conditional decisions (in %) in the Lottery Experiment (NUR treatment)

Signal imbalance $g$	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
$g \in [-8, -5]$	28.57	1.19	0.00	5.95	23.21	0.00	4.17	23.21	13.69
$g \in [-4, -1]$	5.36	0.00	0.00	16.67	36.90	0.00	9.52	25.60	5.95
$g = 0$	0.00	0.00	0.00	9.52	45.24	0.00	33.33	11.90	0.00
$g \in [+1, +4]$	5.95	0.00	0.00	24.40	33.33	1.79	9.52	19.64	5.36
$g \in [+5, +8]$	12.50	0.60	0.00	23.81	26.19	2.38	1.19	8.33	25.00

Table 10: Conditional decisions (in %) in the Market Experiment (main treatment)

Signal imbalance $g$	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
$g \in [-8, -5]$	1.56	2.23	2.46	8.71	44.42	1.56	13.62	12.72	12.72
$g \in [-4, -1]$	1.54	1.54	3.46	9.81	52.12	1.35	12.69	12.12	5.38
$g = 0$	1.54	0.00	1.54	3.85	69.23	3.08	16.15	1.54	3.08
$g \in [+1, +4]$	4.23	0.77	0.96	21.35	46.35	2.50	14.62	6.15	3.08
$g \in [+5, +8]$	12.95	0.45	1.12	18.75	37.95	3.79	11.61	9.38	4.02

Table 11: Conditional decisions (in %) in the Market Experiment (NUR treatment)

Signal imbalance $g$	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
$g \in [-8, -5]$	0.60	0.00	0.00	11.90	13.69	1.79	6.55	49.40	16.07
$g \in [-4, -1]$	0.00	0.00	0.00	6.55	24.40	2.38	15.48	50.00	1.19
$g = 0$	0.00	0.00	0.00	4.76	64.29	0.00	30.95	0.00	0.00
$g \in [+1, +4]$	3.57	0.00	0.00	52.98	19.64	0.00	17.26	6.55	0.00
$g \in [+5, +8]$	13.10	0.60	0.00	50.60	18.45	0.00	5.36	11.90	0.00

Table 12: Conditional decisions (in %) in the SME treatment

Signal imbalance $g$	B-B	B-N	B-S	N-B	N-N	N-S	S-B	S-N	S-S
$g \in [-8, -5]$	13.00	4.50	2.50	8.50	36.00	1.50	7.00	6.00	21.00
$g \in [-4, -1]$	20.00	2.50	1.00	19.00	24.50	3.00	15.50	10.50	4.00
$g = 0$	4.00	6.00	2.00	20.00	22.00	2.00	32.00	6.00	6.00
$g \in [+1, +4]$	7.00	0.00	1.50	20.50	22.00	3.00	20.50	18.50	7.00
$g \in [+5, +8]$	22.00	2.00	4.50	21.00	26.00	1.50	12.50	5.50	5.00

Asset 3/17

The asset is **'high' with probability 97.6%** and 'low' with probability 2.4%. The proposed price is:  
**11.8**

---

Select a decision for each signal, then click on OK:

Signal Bas	Signal Haut
<input type="radio"/> Sell the asset	<input type="radio"/> Sell the asset
<input type="radio"/> Do nothing	<input type="radio"/> Do nothing
<input type="radio"/> Buy the asset	<input type="radio"/> Buy the asset

OK

Figure 7: Screen layout in a Market Experiment (main treatment)

Question 3/17

Select your preferred lottery in Table 1, your preferred lottery in Table 2, then click on OK :

Table 1

**Riskless Payment A**

Payment	12.0
---------	------

**Lottery B**

Probability	2.2%	50.0%	47.8%
Payment	23.8	15.8	7.8

**Lottery C**

Probability	2.2%	50.0%	47.8%
Payment	0.2	8.2	16.2

Table 2

**Riskless Payment A**

Payment	12.0
---------	------

**Lottery B**

Probability	0.6%	50.0%	49.4%
Payment	23.8	15.8	7.8

**Lottery C**

Probability	0.6%	50.0%	49.4%
Payment	0.2	8.2	16.2

Figure 8: Screen layout in a Lottery Experiment (main treatment)

Asset 3/17

The proposed price is:  
**11.8**

---

**[1]** What would be your decision if the asset were 'high' with probability 95.7%?

Sell the asset  
 Do nothing  
 Buy the asset

---

**[2]** What would be your decision if the asset were 'high' with probability 98.7%?

Sell the asset  
 Do nothing  
 Buy the asset

Figure 9: Screen layout in a Simplified Market Experiment

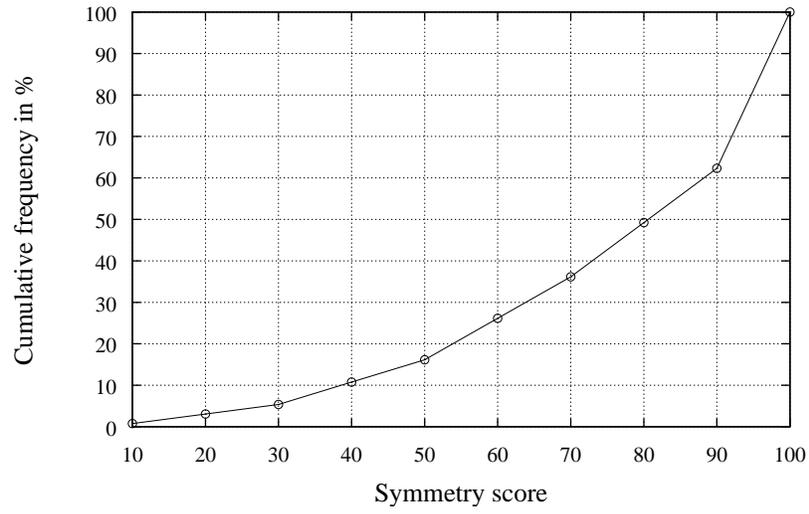


Figure 10: Cumulative distribution of the symmetry score (main treatment)

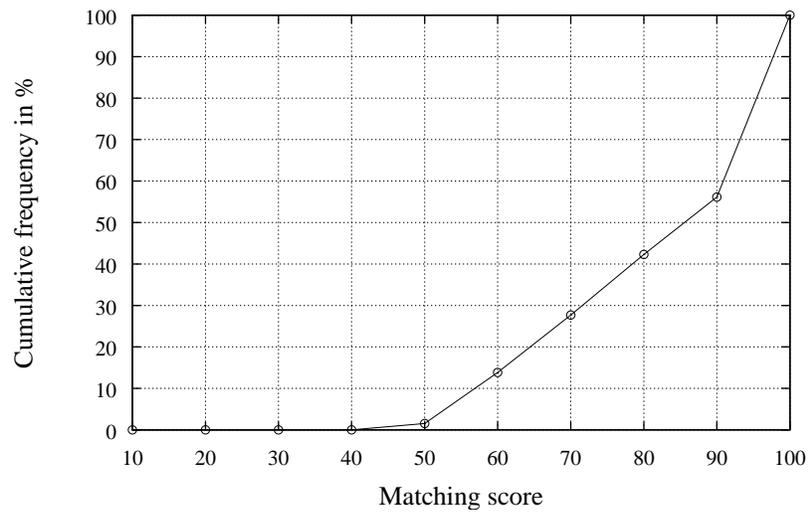


Figure 11: Cumulative distribution of the matching score (main treatment)

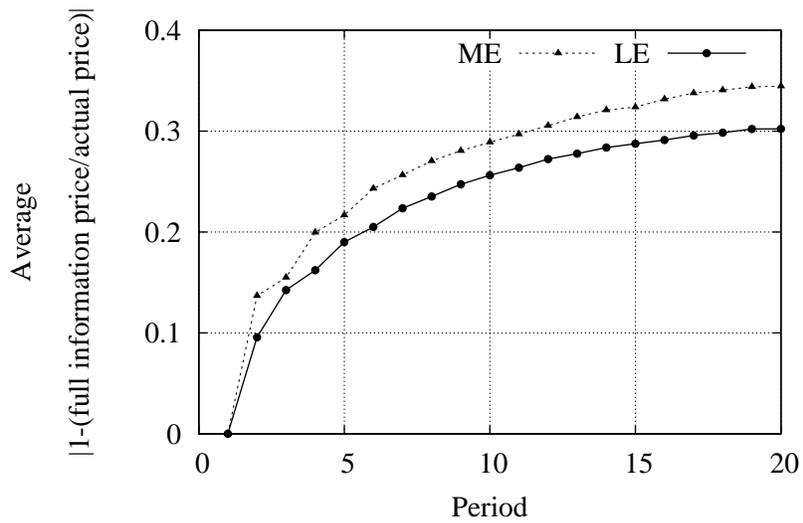


Figure 12: Evolution of the pricing error in the main treatment

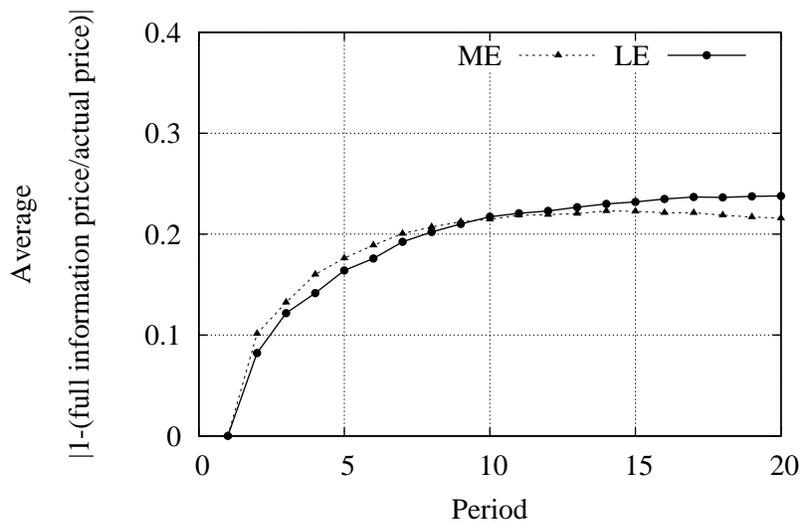


Figure 13: Evolution of the pricing error in the NUR treatment

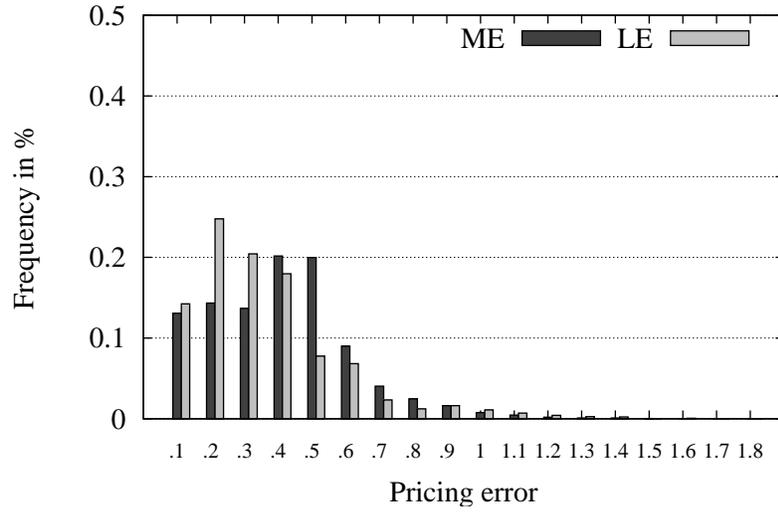


Figure 14: Distribution of pricing error at the 20<sup>th</sup> round in the main treatment

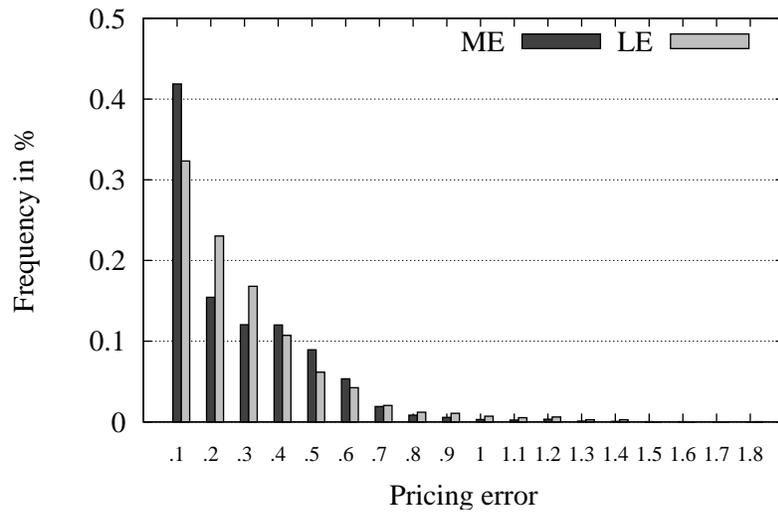


Figure 15: Distribution of pricing error at the 20<sup>th</sup> round in the NUR treatment