

On the Information Content of the Order Flow: An Experiment.

Christophe Bisière [**presenting author**]¹ Jean-Paul Décamps² Stefano Lovo³

This version: May 2008

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

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

³HEC 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.

Financial support from Europlace Institute of Finance is gratefully acknowledged. The authors thank Bruno Biais, Sébastien Pouget and especially Hersh Shefrin for thoughtful discussions and suggestions. We also thank participants at the AFSE meeting on Experimental Economics, Lyon, Mai 2007; the workshop on Individual and Collective decision Making, Paris December 2007; the finance seminars at universities of Barcelona, Copenhagen, Paris 9 and Toulouse. We remain solely responsible for the content of this paper.

Abstract

We report results of a series of experiments that simulates trading in financial market. The specific format of our experiment allows to unambiguously measure the information content of the order flow and to disentangle the impact that risk attitudes and belief updating rules have on market informational efficiency. On the one hand, we show that many of the so called “irrational” behaviors are not so if one takes into account subjects’ risk attitude. On the other hand we find evidence of non-Bayesian updating of beliefs. Risk neutral subjects are rare and subjects displaying risk aversion or risk loving tend to ignore private information when their prior beliefs on the asset fundamentals are strong. This behavior implies that when the market has a sharp opinion on an asset fundamental value, the private information dispersed in the economy struggles to enter trading prices. Non-Bayesian belief updating has an ambiguous effect on market efficiency as it reduces (improves) the information flow when subject prior belief is weak (strong).

1 Introduction

This paper presents the results of a series of experiments that simulate trading in financial markets where market participants are asymmetrically informed. Our objective is to understand how and to what extent investors' private information on asset fundamentals can be incorporated into the asset trading price. To this purpose, we analyze how the information content of the order flow varies with subject's risk attitude, with the intrinsic uncertainty regarding the asset fundamental value and with the strength of common prior beliefs on this fundamental. We are able to disentangle the impact that risk attitudes and belief updating rules have on market informational efficiency. Overall, we find that market informational efficiency decreases when subjects have strong prior beliefs on the asset fundamentals and when the intrinsic risk of the asset increases.

Consider a trader who chooses the position to take in a given financial asset. To begin with, the trader forms prior belief on the asset fundamentals by construing public information coming from newspapers, past trading prices, observable actions of other market participants, etc. Then, the interpretation of whatever additional private information in the trader's hands will determine his or her posterior beliefs on the cash-flows the asset will generate. Since the seminar paper by Glosten and Milgrom [14], it is often taken for granted that if the trading price fairly reflects public information, then the position taken by the trader will reveal the sign of his or her private information, i.e., he or she will buy after receiving favorable private information on the asset perspective, whereas a sell decision will follow unfavorable private information. In this instance, traders' orders will always reflect traders' private information and hence eventually, prices will assimilate all information dispersed in the economy. In an influential theoretical paper, Avery and Zemsky [2] (AZ henceforth) showed that this is, indeed, 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 market efficiency is problematic. As a matter of fact, Décamps and Lovo [10], [11] show that when market makers and traders differs in risk aversion, trades do not necessarily disclose private information and informational efficiency fails.

This suggests that a correct interpretation of the information content of a trade cannot abstract from the study of the trader risk attitude and, more generally, of the decision process underlying the trader choice. In order to better understand under which condition and to what extent trading decisions disclose investor's private information we run a series of experiments that allow to answer the following questions. What are subjects risk attitudes and their role in the use of private information? Do subjects update beliefs in a Bayesian way? What is the effect of the beliefs updating rule on the information content

of trades? We also study how these individual characteristics interact with other features of the market. Namely, to what extent and how prior public information affects an investor willingness to engage in private-information-based speculation? Does a change in the asset intrinsic risk affects the use investors make of private information? What is the effect of these factors on the speed of convergence of prices to fundamentals?

Our experiment is inspired to the framework and the findings of two important papers by Cipriani and Guarino [8] (CG henceforth) and Drehman, Oechssler and Roider [12] (DOR henceforth) that have studied the informational efficiency property of financial market through laboratory experiments. However, while their analysis is based on the underlying assumptions that subjects are risk neutral, our study is founded on the theoretical models by Décamps and Lovo [10], [11] that explicitly take into account market participants' risk attitude.

We depart from CG and DOR's experimental studies in four main features. First, our experiment is designed so that other subjects' rationality plays no role for determining a subject's optimal action. Thus, in our setting, the lack of common knowledge of participants' rationality cannot explain deviation from what the theory predicts. Second, using a "strategy method",¹ we observe subjects choices in all possible realizations of the private signals and for different levels of prior beliefs. In this way, following the spirit of Çelen and Kariv [5], we can identify situations where subjects do ignore their private information. Third, in our main treatment, 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. This set-up allows to measure the effects that the presence of intrinsic risk has on informational efficiency. It also provides a test for the hypothesis of risk-neutrality. Fourth, in order to detect non-Bayesian belief updating as well as particular attitudes towards risk, we run our experiment in two different formats so that subjects have to solve homeomorphic problems that nevertheless require different amounts of reasoning. In the format we denote, *market experiment* (ME henceforth), 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 decisions, subjects access to public and private information on the asset fundamental value. The first consists of the ex ante distribution of the asset value, the latter consists of a private signal correlated with the true value of the asset. In the format we denote *lottery experiment* (LE henceforth), subjects are asked to choose among different lotteries where the distribution function of payoffs is explicitly given. Payoffs in different lotteries reflect the value of a portfolio resulting from different trading decisions in ME. In other words, lotteries in LE are determined so that a Bayesian subject maximizing his or her expected utility faces exactly the same decision problem in the two formats

¹See Selten [18]

and hence will make the same choices in ME and LE. The observation of subjects' choices in LE allows to understand their risk attitude. The comparison of subject's choices in the two formats allows to detect deviations from the expected-utility or Bayesian-updating paradigms. In both formats the information content of the order flow is directly observable.

Our observation in the LE shows that contrary to what implicitly assumed in most experimental papers on market efficiency, subjects are not risk neutral. We show that the behavior of about two third of subjects in LE is compatible with CARA and or CRR utility functions with subject risk attitude ranging from high degrees of risk aversion to risk loving. The fraction of risk neutral subject is negligible. This has clear implications on how the information content of the order flow varies with the strength of common prior belief. We show that the stronger are prior beliefs regarding the asset actual fundamental value, the smaller the proportion of subjects who make use of private information to determine their trade. Thus, the information impounded in the order flow shrinks when the market has strong prior belief on what the asset value is. We find that the presence of a non-learnable shock on the asset fundamentals 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 [10], [11]: when the market is sufficiently convinced about a stock positive or negative perspective, little of the 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 what observed in LE. After running a control experiment we find little evidence that this difference is due to the distinct "framings" of the two formats. More precisely in comparison with LE, the ME presents for strong prior belief, an increase of strategies that consist of following the private signal when this confirms the prior belief and no-trading when the private signal and prior beliefs are mutually contradicting. For weak prior belief, we observe an increase on no trading decisions whatever the private signal. 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.

In the LE and ME we observe each subject's contingent trading strategy for different levels of prior public belief and different realizations of the private signal. Thus, while in our experiment subjects do 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 sequence of trading prices that we use to measure market informational efficiency. We find that the absence of risk neutral behavior sharply slows down the pricing convergence 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 it reduces efficiency when the market has no precise orientation. Our simulations also generate sequences of trades that allow us to compare our results with what observed by CG and DOR. In comparison with these papers, our simulations generate higher frequency of no-trade decisions and lower frequency of trades opposite to the private signal. We think that this difference is due the distinct link between trading prices and asset fundamentals. Our simulations are based on the behavior of subjects facing prices that reflect the objective distribution of the asset fundamentals. In GC and DOR sequential trade experiments, prices updating rule assumes subjects behave in a risk-neutral way, assumption that is not consistent with subjects actual behavior. As a consequence in CG and DOR prices tend to over react to the order flow. This increases the amount of trades contrary to the private signal and reduces the percentage of no-trade decisions. Our simulation shows that ex-post trade is not a reliable statistics to measure risk neutral behavior. In fact, while just 3% of subjects display risk neutral strategy, four out of ten ex-post simulated trades are according to the signal.

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. Section 3 presents the experimental design. Section 4 presents the results of the experiment and compares it with findings in CG and DOR. Section 5 concludes.

2 Theoretical framework

In this section we first describe the theoretical setting borrowed from Décamps and Lovo [10]. Second, we illustrate the model with some numerical examples that help presenting the main predictions we tested in our experiment.

2.1 The model structure

We consider a discrete time sequential trade model à la Glosten and Milgrom [14]: 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 everybody. 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, 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, given that private signal precision is bounded, the closer is the prior π_t to 1 (or to 0) the smaller will be the change in belief $|\pi_t^s - \pi_t|$ induced by the private signal. For this reason we will adopt the following convention. We will say that a prior belief π_t is *stronger* the larger is $|\pi_t - 0.5|$. Also, we will call *positive prior*, a public belief π_t that is larger than 0.5 and *negative prior*, a public belief π_t that is lower than 0.5.

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 contemplates risk-neutrality, risk-aversion and risk-loving. The variable m and x represent the trader's initial monetary wealth and 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 satisfied at any time t . Private information enters prices when market makers can surmise 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 [2] 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 trading price eventually converges² to \tilde{V} . Décamps and Lovo [10] and [11] show that if market makers and traders differ in their risk aversion, and if agents' set of actions is discrete³, 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 [2] and Décamps and Lovo [10], [11] 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 to be set equal to the expected asset value, conditional on the information available at time t , i.e., $P(Q) = \mathbb{E}[\tilde{v}|H_t]$.⁴

Let us now assume a 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 0 amount of the risky asset and 12 units of money. The problem that such a trader 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* * (\tilde{v} - *trading price*) + 12. Tables 1, 2 and 3 only differ for the probabilities of getting the payoffs in each column.

Table 1 represents the problem faced by a trader who received no private signal. By definition of risk aversion, 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-lover trader will typically strictly prefer selling to the

²The same feature holds for instance in Glosten and Milgrom [14] and Easley and O'Hara [13]

³If agents were able to trade a continuum of quantities, risk aversion alone would not be enough to generate market inefficiency. See for instance Glosten [15] and Vives [21].

⁴This 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 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

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 \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 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 $s = h$. 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 strictly prefer selling to the other alternatives and this choice will not be affected by the sign of the private signal. Hence, a trader that is either sufficiently risk averse or sufficiently risk-lover 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 lack of risk neutrality and, on the other hand, the fact that a private signal does not affect posterior belief sharply. Considering that when prior belief is extreme (that is when π_t is sufficiently close to 0 or to 1) a partially informative signal will 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 embody the private signal irrespectively of the strength of prior beliefs.⁵

The impact that risk attitude and prior belief have 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 π . The first column of each table reports the different level g of *signal imbalance* defined as the difference between the number of positive and negative signals that the market has inferred from previous trades. The one-to-one mapping between signal imbalance g and public belief π will make easier the exposition of the results. 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.⁶

Several comments are in order. Note first that traders strategy 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 $-g$.⁷ Second, very risk-averse trader ($\gamma > 0.078$ and $\alpha < -0.84$ in Table 5) always choose strategy N-N whatever the public belief. Third, strategy S-B is optimal for all 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), such traders will submit non-informative orders. More precisely, risk-averse traders will prefer not to trade ignoring their private signal. Risk lover traders, (i.e. $\gamma < -0.25$ and $\alpha > 4.8$ in Table 5) will

⁵In fact risk-neutrality implies ex-ante indifference among 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.

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

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

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 lead to several 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 π one can estimate the trader's risk attitude. In particular a risk-neutral trader will choose strategy S-B for all levels of π .*

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 larger the frequency of non-informative orders and the lower the information content of order flow.*

It is interesting to link non-informative strategies with what the herding literature classifies as herding or contrarian behavior.⁸ 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) whenever $\pi_t \geq \pi^*$ (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. By contrast 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 should not be observed.*

⁸The literature on rational herding starts with the seminal papers by Bikhchandani, Hirshleifer and Welch [3], Welch [22] and Banerjee [4]. Herding with endogenous prices has been recently studied by a series of papers including Avery and Zemsky [2], Lee [17], Cipriani and Guarino [7], Chari and Kehoe [5], Décamps and Lovo [10] and [11]. See for instance Hirshleifer and Teoh [16] for a survey on herd behavior in capital markets and Chamley [6] for an extensive study on rational herding.

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 actual behavior for a given risk attitude will differ from the one described above. In order to disentangle 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 were stated in the same form of the menu of lotteries in tables 2 and 3. By explicitly providing subject with the distribution function of payoffs, we ensure that the belief updating rule plays no role in their decision. In ME, subjects were informed of the prior π , of the precision of their private signal and asked to declare their preferred trading position contingent to the realization of the private signal. In other words, while in LE posterior probability were explicitly provided, in ME subjects were in possession of all elements necessary to compute posterior probability. The two formats have been designed so that a rational Bayesian expected utility maximizer would have found them perfectly equivalent.

Implication 5 *The behaviors of a Bayesian expected utility maximizer in format LE and ME are identical.*

3 Experiment Design

We perform our experiment under two different but to some extent equivalent formats: ME and LE. Each subject participated to both formats. Our main treatment matches the theoretical setup 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 *no-unlearnable risk treatment* (NUR treatment henceforth) and the *simplified market experiment* (SME henceforth) that will be detailed in sections 4.3 and 4.2.1 respectively.

Below, we describe precisely the Market Experiment (ME), the Lottery Experiment (LE), the subjects' payoff and the implementation.

Market Experiment. The Market Experiment consists of a series of 17 questions or "rounds". 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 . As described in Section 2, the fundamental value of asset τ is a random variable $\tilde{v}_\tau = \tilde{V}_\tau + \tilde{\epsilon}$ with $\tilde{V}_\tau \in \{4, 12\}$. The trading price for asset τ was fixed at $P_\tau = \pi_\tau 12 + (1 - \pi_\tau) 4$. Both π_τ and P_τ are announced to subjects in round τ (see a screen shot in Figure 7). Moreover, in each

round each subject receives a private signal $\tilde{s} \in \{h, l\}$ with precision $p = 0.65$. Before being informed of the private signal and after observing π_τ and P_τ , 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 given by the probability π_τ and the corresponding trading price P_τ . For 17 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 $+8$.⁹ ¹⁰

Lottery Experiment. The 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 ME but with a different formulation. Instead of asking a subject the position to take into a given financial asset, he or she 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” implemented in ME, at each stage, each subject was proposed two menus 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 reflecting 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 including two menus.¹¹ Overall each subject had to choose 34 lotteries among 34 menus. It was never mentioned to 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 equivalent.

Implementation. The experiment was run in HEC School of Management and Toulouse University. We recruited 228 subjects from undergraduate finance classes in HEC and Toulouse University. The subjects had no previous experience in financial market experiments. In each session between 16 and 43 subjects participated as decision makers. The main treatment involved 135 subjects (5 cohorts) while the two control treatments involved 93 subjects (2 cohorts for the NUR treatment and 3 cohorts for SME). At the beginning of a session we gave written instruction that were 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

⁹Only from -7 to $+7$ for subjects in cohorts 1 and 2, that is, 15 assets.

¹⁰See for example Table 5 for a correspondence between the numbering of the asset, the corresponding π_τ and g .

¹¹Only 15 rounds for cohorts 1 and 2.

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 took 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 was paid to each subject. Subjects were also rewarded with bonus points valid to improve their mark in the Financial Markets course. Subjects' payoffs were determined on the basis of the gain on only one round for the market experiment and one for the lottery experiment. The payoff-relevant rounds were randomly selected at the end of the experiment.¹² 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 of 4,147 in the main treatment, 1,428 in the NUR treatment and 850 in the SME.

4 Experimental Results

We describe in this section the main results of the experiment. Sections 4.1 and 4.2 examine results obtained in the settings LE and ME. Section 4.4 is devoted to the analysis of market informational efficiency. Section 4.5 compares our results with those of CG and DOR.

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 only depends 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 subjects' behavior conforms to the expected utility assumption. This format also allows to measure subjects' risk attitude.

Before proceeding with the detailed analysis of the data, we stress the main results of the lottery experiment for the main treatment. First, we find that only 3% 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 15, 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

¹²See the experiment instruction for a precise description on the algorithm determining a subject's payoffs.

of strategy S-B for $g = 0$. The proportion of strategy N-N increases with $|g|$. Strategy B-B decreases with g while strategy S-S increase with g . We observe a pick of strategy N-B for weakly negative priors (that is $g \in [-4, -1]$) and, symmetrically, a pick of strategy S-N for weakly positive prior. The fraction of non-informative strategies (i.e., sum of N-N, B-B and S-S) as a function of g display a U shape. That is, the information content of the order flow shrinks 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 problem a subject faces for a given level of public belief π is the same faced when public belief is $1 - \pi$. This feature is apparent in LE as the lottery in the menus corresponding to prior π and to prior $1 - \pi$ are the same but 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 π with the one chosen for prior $1 - \pi$. The score gives us the proportion of a subject's choices that respects the symmetry rule. 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 a score larger than 0.57. Overall these data suggest that subjects' behavior is consistent with Implication 1 and does not contradict this basic test of rationality. Nevertheless, there is a small fraction of subjects (6.36%) that clearly behave inconsistently and display a symmetry score of less than 0.33.

4.1.2 Risk attitude

In order to infer subjects' risk attitude, we compare each subject 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 better matches the subject's observed behavior. Then to each subject we assign a matching score between observed behavior and closest theoretical behavior. The closer the matching score is to 1 the better the subject behavior can be explained with a CARA or a CRR utility function.

The median matching score is 0.83 with 77.1% of subjects having a matching score greater than 0.66. Thus, the behavior of about two third of subjects is reasonably well explained by a CARA or a CRR utility function.

Within subjects with an overall matching score greater or equal to 0.66, the distribution of risk aversions is reported in Table 6 for CARA and CRR utility functions. About 46.5% display sufficiently strong risk aversion to induce N-N in at least two third of the situations. The percentage of risk lovers is

18%, while 32% of subjects in this sub-sample display intermediate levels of risk aversion. Surprisingly only 3% of the participants to the experiment are close to behave in a risk neutral way.¹³ This finding is in sharp contradiction with the assumption that 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 (i.e. N-N, S-S, and B-B) changes with the signal imbalance. For the sake of clarity, frequencies are reported by class of signal imbalance (or, equivalently, by strength of belief):

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

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 52.67% of all trades for neutral prior belief and this percentage falls to 16.19% and 20.62% for strong negative and strong positive prior belief, 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.¹⁴ Corresponding percentages are reported in Table 8.

Strategies S-S (B-B) are related to contrarian behavior when they occur for positive (negative) prior. We observe both buy and sell contrarian behavior. The frequency of strategies S-S for neutral priors is

¹³ Actually, only one subject has chosen the “risk-neutral” strategy S-B whatever the signal imbalance g .

¹⁴ Strategy “Other” includes B-N, B-S, and N-S.




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

7.63%, and is much larger for strong positive prior (20.18%). Similarly, the frequency of strategies B-B rises from 2.29% for neutral priors to 14.86% for strong negative priors. According to Implication 4, contrarian behavior can be ascribed to the fraction (about one fifth) of subjects that are risk lover. 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 risk aversion.

Strategies S-S (B-B) are related to herd behavior when related to negative (positive) prior. While our theory gives no justification for herding strategies, a small fraction of subjects engaged nevertheless in herd behavior. Namely sell herding (S-S for strong negative prior belief) amount to 5.99% and buy herding (B-B for strong positive prior belief) to 8.20%.

4.2 Market Experiment

Implication 5 of the theory suggests subjects should choose exactly the same conditional strategies in LE and in ME. This prediction is clearly rejected by our data. Overall only in 43.59% of the observations, subjects' answers are the same for the LE and the ME. For only 27.48% of subjects, the answers in LE and ME were consistent in at least 3/4 of the questions. In most of these cases subjects preferred strategy N-N for all levels of prior belief in both formats.

Figure 2 and the corresponding Table 10, summarize the distribution of subjects strategies.




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

The only common pattern with LE is the symmetry of 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 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.¹⁵ Fifth, herd behavior increases in ME.

The effect of the ME format on the information content of the order flow is ambiguous. As illustrated in Table 7, informative strategies rises from 28.26% of all choice 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 information content of the order flow is lower in ME compared to LE. However for strong prior belief, in ME private information will be better signaled through subjects' strategies. It results that 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.

¹⁵With the exception of $g = 0$.

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 might 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 indeed asked to take trading decision in “virtual financial assets”. This framing could trigger heuristic behaviors in ME that are absent in LE.¹⁶ Another possible explanation for the difference of subjects’ behavior in the two formats is that while in LE probabilities are explicitly provided, in ME subjects have to interpret public and private information when forming their decision. Subjects who do not conform to Bayes rule should 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 have not to interpret private and public information. Namely, similarly to ME, subjects had the opportunity to trade a financial asset at a given price. However instead of providing subjects with prior belief and private signals, we directly supplied them with the posterior probability that $\tilde{V} = 12$. Namely, these probabilities were computed as follows. For a given asset τ we considered the prior belief corresponding to the proposed trading price and update 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 subjects’ behavior in SME is closer to what we observed in LE, than to what we observed in ME. Namely, similarly to what happens in LE, the information content of the order flow decreases as the prior belief becomes stronger. Also, strategies S-B (N-N) are more (resp. less) frequent for neutral prior, as observed in LE. In one respect however behavior in SME is closer to the one in ME rather 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 subjects trading decisions.

In order to analyze the hypothesis of non-Bayesian updating we proceed as follows. We focus on those subjects for which there exist a CARA or a CRR utility function that explains at least 66% of their choices in LE. For this subset of subjects, we find that about two third of subjects’ choices in LE match choices in ME. Given 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 s . Within this set of posterior beliefs we focus on the one closest to the Bayesian posterior belief and we denote it $\hat{\pi}_i(g, s)$.

¹⁶Some example of heuristic behavior are: buying (selling) when the prior information is strong and positive (resp. negative); always trading according to the signal; buying when the price is low and sell when the price is high, etc.




Figure 3: Distribution of subjects strategies in Simplified Market Experiment main treatment

Repeating the process for both realizations of the private signals in all levels of g we obtain two point-wise functions $\{\hat{\pi}_i(g, l), \hat{\pi}_i(g, h)\}_{g=-8, -7, \dots, +7, +8}$ mapping prior belief and private signals into posterior beliefs. The difference between these functions and Bayesian posterior belief provides an estimation of the subject's bias in updating beliefs.

Using this method we can explain about 87% of subjects choices in ME while for the remaining 13% of choices there exist no belief that coupled with the subject utility function explains the choice in ME.

Figure 4 represents the average bias in posterior beliefs for subjects that are not Bayesian i.e. subjects for which $\hat{\pi}_i(g, l)$ is different from the Bayesian belief in at least 1/3 of the cases.

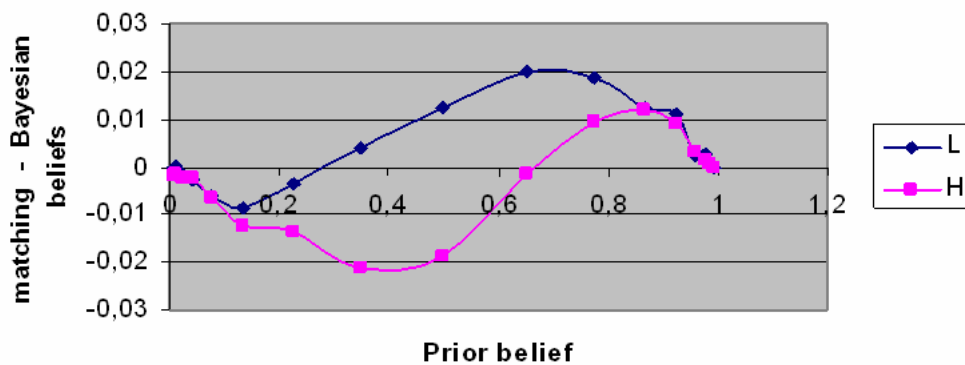


Figure 4: Bias in belief updating rule

For strong public belief, we find that on average subjects display confirmation bias and tend to interpret private signals in a way that either confirms prior belief or does not challenge it. Namely Figure 4 shows that, when prior beliefs are above 0.7 (below 0.3), 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 strong positive public belief. On the one hand the impact of a signal *high* is reinforced by the public belief and traders are more inclined to buy when they receive such a signal. On the other hand 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. Overall the information content of the order flow is more important than in LE. The argument is symmetric for strong negative prior beliefs.

For weak public belief the analysis is different. In that case, we find evidence of underconfidence, namely that subjects undervalue their private signal. As shown in Figure 4, for prior belief around 0.5, behaviors of subjects with signal *low* (*high*) in ME can be explained with 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. It follows that the information content of the order flow for weak public belief is lower in ME than in LE.

4.3 Effect of intrinsic uncertainty.

We now describe the results of a NUR treatment in which the intrinsic risk component $\tilde{\epsilon}$ is absent, i.e., $\underline{V} = 4$, $\overline{V} = 12$, $\epsilon = 0$. By comparing 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.

The main features of the distribution of subjects strategies obtained in the experiments are overall the same than in the main treatment. 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 and 6 synthesize these observations. There are however two noticeable differences between the two treatments. First, the frequency of informative orders for strong negative and strong positive prior belief is higher for the control treatment. For strong negative prior, this frequency rises from the 16.19% of the main treatment to 34.52%, and from 20.62% to 36.31% for strong positive prior. 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 important in the NUR treatment than in the main treatment. Specifically, sell herding (S-S for strong negative prior belief) in the NUR treatment amount to 13.69% and




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

buy herding (B-B for strong positive prior belief) to 12.50%. In the main treatment these percentages fall to 5.99% and 8.20%. A possible explanation could be subject's tendency of rounding probabilities of the lotteries and/or ignoring events whose probability is sufficiently small. In the control treatment, the event $\tilde{V} = V$ is considered by traders as virtually sure for strong beliefs. 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 $\tilde{\epsilon}$ makes any trade intrinsically risky even when \tilde{V} is known. Section 4.4, devoted to the market informational efficiency, will go further in the analysis of the effect of intrinsic uncertainty.

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 increasing with the strength of public belief. In this section we try to measure the actual impact of these figures on the price dynamics and 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 trader's private information. In




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

our experiment, however, we do not observe trading histories as subjects do 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 this purpose, we exploit our observations of contingent trading strategies at different levels of public belief. In order to generate virtual trading histories we assume virtual subjects randomly come to the market to trade once and behave as real subjects did in the actual experiment.¹⁷ 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.

¹⁷We 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 $\pi_0 = 0.5$. In each trading round, first, one among the 9 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 correct 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.

These frequencies are those observed in the experiment. After observing a given action the public belief will change according to the Bayesian probability that this order comes from someone who received a signal l or a signal h . The trading price is updated accordingly. We will denote this pricing rule $R1$.¹⁸

We simulated about 5,000 trading histories per treatment and format, each trading history lasting a maximum of 20 trading rounds. Figure 10 reports the evolution of the average pricing error in the main treatment. The pricing error in ME is consistently higher when compared to LE. After 20 trades the average (median) pricing error is 30% (25.14%) in LE and 34% (34.6%) in ME. Figure 11 reports the distribution of pricing errors at the 20th 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%, while in 8.34% (8.68%) of histories the pricing error is larger than 70% suggesting that in both format information cascades in the “wrong direction” are not rare. Figures 12 and 13 reports 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 at the 20th 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 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 ϵ 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, thus initially prices and public belief tend to stagnate. On the other hand if public belief reaches some strength non-Bayesian behavior tends to increase the information flow.

4.5 Ex post decisions

In this section we compare our findings with those of CG and DOR. To this purpose we focus on the virtual trading histories based on subjects’ behavior in ME, NUR treatment. This is, in fact, the format and treatment that is the closest to the flexible price treatments of CG and DOR.¹⁹ As CG and DOR only observe subjects’ decisions corresponding to the realized private signals, we will focus on the equivalent

¹⁸This is different from the price updating rule adopted in CG and DOR. See next section for the details.

¹⁹We compare our finding to the flexible price treatment in CG and treatment $P50 - 66$ in DOR.

metric generated in our simulations, i.e. virtual traders decision corresponding to the realization of their private signals. These data are reported in Table 4.

Table 4: Comparison with CG and DOR

	ME, $\epsilon = 0$	CG	DOR
Trade according to the signal	39.5%	61 – 65%	59%
No Trade	59.5%	22 – 25%	24%
Trade opposite to the signal	1%	13 – 14%	17%

It is worth noting that our simulations are based on a population in which virtually no subject consistently followed strategy S-B. Nevertheless, 39.5% of simulated ex post trades are according to the signal. This suggests that a relative high frequency of ex post trades according to the signal provides little or no proof that agents have a risk neutral behavior.

When comparing the distribution of our simulated ex-post trades with those reported by CG and DOR two differences are striking. First, the frequency of no-trades in our simulation is 59.5%. In comparison no-trades in CG and DOR range between 22% and 25%. Second, trading opposite to the signal is virtually absent in our simulation while it represents between 13% to 17% of trading decisions in CG and DOR. These differences can be imputed to the specific pricing rules adopted in theirs and ours experiments. While in our experiment prices always reflects objective probability on the asset fundamentals, this relation is generally not satisfied for the prices proposed in CG and DOR. Namely, the price updating rule in CG and DOR systematically interprets a buy decision, a sell decision and a no-trade decision into signal h , signal l and no-signal, respectively. We will refer to this pricing rule as R2. Note that R2 provides prices that are consistent with the actual information impounded by the trading history only if all previous subjects followed strategy S-B. However the proportion of subjects' behavior that is consistent with such behavior in CG and DOR is at most 65%. As a consequence, a rational subject participating to CG and DOR's experiments and anticipating previous subjects do not follow necessarily strategy S-B, would perceive a difference between his or her estimation of the expected value of the asset and the proposed trading price.²⁰ This confirms what has been pointed out by CG and DOR: prices in R2 tend to over-react to the order flow. This induces subjects to sell after a large increase in the asset and to buy after a fall in price. This behavior is correlated to trades opposite to the signal. To summarize, the

²⁰In treatments denoted AS in DOR, the price is set at its full information level and hence reflects objective probabilities. Unfortunately when this treatment was used, subjects were not given the option of not trading. Hence, it is difficult to compare their result in the AS treatment with ours.

reason why in our simulation, in comparison with CG and DOR, no-trades are more frequent and trades opposite to the signal are rare is the absence of mispricing implied by R1.

5 Conclusion

We report results of an experiment that simulates trading in financial market. We adopt two formats for our experiment: the Lottery Experiment and the Market Experiment. This allows 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” behavior are not so if one takes into account subjects’ risk attitude. The design of LE allows to measure subjects’ risk attitude. We find that CARA and CRR utility functions explain the behavior of about two third of subject in LE. While risk neutral subjects are rare, most subjects display risk aversion and some subjects display risk loving. This has 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 predicted by the theory subjects behave differently in ME and LE. This discrepancy can be ascribed to non-Bayesian belief updating and only partially to framing effects. More specifically, by considering that subjects’ utility functions are those consistent with their behaviors in LE, we find that confirmation bias and underconfidence seems to fit well subjects behavior in ME. Its effect on market efficiency is ambiguous. Non-Bayesian updating reduces (improves) the information flow when subject prior belief is weak (strong).

References

- [1] **Anderson, Lisa R. and Holt, Charles A.** “Information Cascades in the Laboratory.” *American Economic Review*, 1997, 87(5), pp. 847–62.
- [2] **Avery, Christopher and Zemsky, Peter.** “Multidimensional Uncertainty and Herd Behavior in Financial Markets.” *American Economic Review*, 1998, 88(4), pp. 724–48.
- [3] **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.
- [4] **Banerjee, Abhijit V.** “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, 1992, 107(3), pp. 797-817.
- [5] **Çelen, Boğaçhan and Kariv, Shachar.** “Distinguishing Informational Cascades from Herd Behavior in the Laboratory.” *American Economic Review*, 2004, 94(3), pp. 484–98.
- [6] **Chamley, Christophe P.** “Rational herds: Economic models of social learning. Cambridge: Cambridge University Press, 2004.
- [7] **Cipriani, Marco and Guarino, Antonio.** “Herd Behavior and Contagion in Financial Markets.” Unpublished Paper, 2001.
- [8] **Cipriani, Marco and Guarino, Antonio.** “Herd Behavior in a Laboratory Financial Market.” *American Economic Review*, 2005, 95(5), pp. 1427–1443.
- [9] **Chari, V. V. and Kehoe, Patrick J.** “Financial Crisis as Herds: Overturning the critiques.” *Journal of Economic Theory*, 2004, 119(1), pp. 128–150.
- [10] **Décamps, 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.
- [11] **Décamps, 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.
- [12] **Drehman, Mathias, Oechssler, Jorg and Roider Andreas.** “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment.” *American Economic Review*, 2005, 95(5), pp. 1403–1425.
- [13] **Easley, David, and O’Hara, Maureen.** “Time and the Process of Security Price Adjustment, *Journal of Finance*, 57, pp. 577–605.

- [14] **Glosten, Larry, and Milgrom, Paul.** “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents”, *Journal of Financial Economics*, 1985, 14, pp. 71–100.
- [15] **Glosten, Larry.** “Insider Trading, Liquidity and the Role of the Monopolist Specialist.” *Journal of Business*, 1989, 62, pp. 211–235.
- [16] **Hirshleifer, David and Teoh, Siew H.** “Herd Behaviour and Cascading in capital Markets: a Review and Synthesis.” *European financial Management*, 2003, 9(1), pp. 25–66.
- [17] **Lee, In Ho.** “Market Crashes and Informational Avalanches.” *Review of Economic Studies*, 1998, 65(4), pp. 741–59.
- [18] **Selten, Reinhard** “Die strategiemethode zur erforschung des eigescharankt rationalen verhaltens im rahmen eines oligopolexperiment”. In *Beitrage zur experimentellenwirtschaftsforschung*, 1967, H. Sauer mann, editor, Tübingen: J.C.B. Mohr. pp. 136–168.
- [19] **Shefrin, Hersh.** “A Behavioral approach to Asset Pricing,” 2005, Elsevier Inc.
- [20] **Smith, Lones and Sorensen, Peter.** “Pathological Outcomes of Observational Learning.” *Econometrica*, 2000, 68(2), pp. 371–98.
- [21] **Vives, Xavier.** “The speed of information revelation in a financial market mechanism.” *Journal of Economic Theory*, 67, pp. 178-204.
- [22] **Welch, Ivo.** “Sequential Sales, learning , and Cascades.” *Journal of Finance*, 1992, 47(2), pp. 695–732.

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

CRR: $U(x) = \frac{x^\alpha}{\alpha}$		$\alpha < -0.85$	$\alpha = -0.46$	$-0.034 < \alpha < 4.7$	$\alpha = 4.98$	$\alpha = 5.91$
CARA: $U(x) = -\gamma e^{-\gamma x}$		$\gamma > 0.078$	$\gamma = 0.02$	$-0.25 < \gamma < 0.003$	$\gamma = -0.26$	$\gamma = -0.3$
g	π					
-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 $U(x) = \frac{x^\alpha}{\alpha}$	Number of subjects	Average matching score
High risk averse	$\gamma > 0.078$	$\alpha < -0.84$	45	96.21
Medium risk averse	$0.005 < \gamma < 0.078$	$-0.84 < \alpha < -0.065$	36	81.99
Close to risk neutral	$-0.25 < \gamma < 0.005$	$-0.065 < \alpha < 4.8$	3	73.33
Risk lover	$\gamma < -0.25$	$\alpha > 4.8$	17	87.19
Not classed			30	56.44

Table 7: Informative orders (in %) in the main treatment ($\epsilon = 1$) and in the NUR control treatments ($\epsilon = 0$)

Signal imbalance g	Main treatment:		NUR treatment:	
	ME	LE	ME	LE
$g \in [-8, -5]$	41.29	16.19	69.64	34.52
$g \in [-4, -1]$	40.96	33.02	74.40	51.79
$g = 0$	26.15	52.67	35.71	54.76
$g \in [+1, +4]$	46.35	34.35	76.79	55.36
$g \in [+5, +8]$	45.09	20.62	68.45	36.31
All	42.35	28.26	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
$g \in [-8, -5]$	14.86	0.44	1.11	5.10	62.97	1.33	4.43	3.77	5.99
$g \in [-4, -1]$	16.22	1.15	0.76	19.27	46.76	0.95	6.11	4.77	4.01
$g = 0$	2.29	0.00	0.00	5.34	37.40	0.00	37.40	9.92	7.63
$g \in [+1, +4]$	2.86	0.38	0.57	4.20	43.89	0.57	9.54	19.08	18.89
$g \in [+5, +8]$	8.20	0.44	0.67	4.88	51.00	1.55	4.88	8.20	20.18

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

OK

Figure 8: Screen layout in a Lotterie 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

OK

Figure 9: Screen layout in a Simplified Market Experiment




Figure 10: Pricing error in the main treatment




Figure 11: Distribution of pricing error in the main treatment




Figure 12: Pricing error in the NUR treatment




Figure 13: Distribution of pricing error in the NUR treatment