



Working Paper 08-41
Statistics and Econometrics Series 10
September 2008

Departamento de Estadística
Universidad Carlos III de Madrid
Calle Madrid, 126
28903 Getafe (Spain)
Fax (34) 91 624-98-49

**AGGREGATION AND DISSEMINATION OF INFORMATION IN
EXPERIMENTAL ASSET MARKETS IN THE PRESENCE OF A
MANIPULATOR**

HELENA VEIGA* AND MARC VORSATZ**

Abstract

We study with the help of a laboratory experiment the conditions under which an uninformed manipulator - a robot trader that unconditionally buys several shares of a common value asset in the beginning of a trading period and unwinds this position later on - is able to induce higher asset prices. We find that the average contract price is significantly higher in the presence of the manipulator if, and only if, the asset takes the lowest possible value and insiders have perfect information about the true value of the asset. It is also evidenced that the robot trader makes trading gains; i.e., independently on whether the informed traders have perfect or partial information, it earns always more than the average trader. Finally, not only uninformed subjects suffer from the presence of the robot trader, but also some of the imperfectly informed insiders have lower payoffs once the robot trader is added as a market participant.

Keywords: Asset Market, Experiment, Price Manipulation, Rational Expectations.

* Veiga, Helena, Departamento de Estadística, Universidad Carlos III de Madrid, C/ Madrid 126, 28903 Getafe (Madrid), e-mail: mhveiga@est-econ.uc3m.es .

** Vorsatz, Marc, Fundación de Estudios de Economía Aplicada (FEDEA), Calle Jorge Juan 46, 28001 Madrid, e-mail: mvorsatz@fedea.es.

Aggregation and Dissemination of Information in Experimental Asset Markets in the Presence of a Manipulator

Helena Veiga* Marc Vorsatz†

August 1, 2008

Abstract

We study with the help of a laboratory experiment the conditions under which an uninformed manipulator – a robot trader that unconditionally buys several shares of a common value asset in the beginning of a trading period and unwinds this position later on – is able to induce higher asset prices. We find that the average contract price is significantly higher in the presence of the manipulator if, and only if, the asset takes the lowest possible value and insiders have perfect information about the true value of the asset. It is also evidenced that the robot trader makes trading gains; *i.e.*, independently on whether the informed traders have perfect or partial information, it earns always more than the average trader. Finally, not only uninformed subjects suffer from the presence of the robot trader, but also some of the imperfectly informed insiders have lower payoffs once the robot trader is added as a market participant.

Keywords: Asset Market, Experiment, Price Manipulation, Rational Expectations.

JEL-Classifications: C90, G12, G14.

1 Introduction

The efficient market hypothesis of Hayek (1945), operationalized first by Muth (1961) in form of the rational expectations equilibrium, predicts that the price of a financial security summarizes all information about the asset's intrinsic value held by traders. The capacity to aggregate information through prices makes it, in principle, possible to use security markets to predict the likelihood of future events (see Wolfers and Zitzewitz, 2004), but it has also

*Department of Statistics, Universidad Carlos III Madrid, Calle Madrid 126, 28903 Getafe, Spain. E-mail: mhveiga@est-econ.uc3m.es. Financial support from the Spanish Ministry of Education and Science, through the research project SEJ2006-03919, is gratefully acknowledged.

†Corresponding author. Fundación de Estudios de Economía Aplicada (FEDEA), Calle Jorge Juan 46, 28001 Madrid, Spain. Email: mvorsatz@fedea.es. Financial support from the Ramón y Cajal program of the Spanish Ministry of Education and Science is gratefully acknowledged. This work has been initiated while the author was employed by the Economics Department at Maastricht University.

been argued recently that prediction markets may suffer from manipulations because some markets participants can have incentives to incur into trading losses in order to distort, indirectly, policies. Two newspaper articles by Pearlstein (2003) and Wyden and Dorgan (2003) reported precisely that a prediction market with the aim to assess the probabilities of geopolitical events was shut down because of the fear that terrorists would be willing to distort the process of information aggregation.

The capacity of asset markets to aggregate disperse information has been analyzed extensively with the help of laboratory experiments (see, among others, Forsythe et al., 1982; Plott and Sunder, 1982, 1988; Forsythe and Lundholm, 1990; Sunder, 1992), but only little effort has been made to assess their manipulability. In the study of Hanson et al. (2006), some traders got an additional payoff depending on the median contract price. The authors found that manipulators place, as expected, higher bids. The accuracy of the market, however, remained unaffected because the other traders offset this effect. Veiga and Vorsatz (2008) studied a short-selling constrained asset market in which subjects traded a common value asset that took either a high or a low value. Information was distributed asymmetrically in the sense that only three out of twelve subjects knew the actual value of the asset when the market was open for trading. It was found that a robot trader that bought ten shares of the asset in the beginning of the trading period and sold them afterwards again was able to induce significantly higher prices whenever the actual value of the asset was low. This occurred because some of the uninformed traders took the initial buy offers to be informative, revised their beliefs about the actual value of the asset upwards, and drove the price even higher.¹ This study provided therefore some experimental evidence for the recent theoretical finding of Chakraborty and Yilmaz (2008) that strategic traders have incentives to mimic insiders by moving prices initially away from the fundamentals before reversing the position later on at

¹Following Nöth et al. (1999), we say that some subjects fell into an *information trap*; that is, these traders had misaligned beliefs about the state (the true value of the asset) and the traders who recognized them did not have the resources to adopt their behavior (due to short-selling constraints). Information traps are to be distinguished from *information mirages* and *information cascades*. According to Camerer and Weigelt (1991), an information mirage is obtained if some of the traders see information that could be there but is not. The term information cascades, on the other hand, refers to a different decision environment. Consider a situation in which every individual receives some private information about the state and, subsequently, individuals have to announce publicly their guess about the actual state. Then, individuals may discard their private information following the decisions of the former players (see, among others, Bannerjee (1992) and Bikhchandani et al. (1992)).

more favorable prices.²

The objective of this paper is to go one step further and analyze the conditions on the distribution of private information such that manipulations occur. In particular, we address the question whether manipulations happen primarily if private information is perfect or whether similar results can also be obtained when the insiders hold partial information about the true value of the asset (but there is no aggregate uncertainty). A priori there is no straightforward answer to this question. Since all traders are aware about the distribution of private information, the initial purchases of the robot trader may, on the one hand, be considered less informative by the other traders if information is imperfect. This suggests that manipulations are more unlikely in this case. But, on the other hand, now also partially informed traders and not only uninformed subjects may be deceived by the manipulator.

We propose a three-by-two between subjects design to investigate this tradeoff. The three baseline treatments, which are inspired by the study of Plott and Sunder (1988), are as follows: twelve subjects trade a common value asset that takes the values 125, 375, and 525 ECU (experimental currency units) with equal probability for ten periods in a short-selling constrained and computerized double auction market. Every round, each trader receives four shares of the asset and an interest free loan of 25.000 ECU that has to be returned once the market closes. The three baseline treatments differ only in the allocation of private information to subjects. In the first case, six subjects are imperfectly informed about the true value of the asset before trading starts in such a way that there is no aggregate uncertainty, whereas the other six subjects remain uninformed; in the second case, all participants are imperfectly informed; and, in the last case, four subjects have perfect information about the true value of the asset and the remaining eight subjects stay uninformed. The three manipulation treatments differ from their associated baseline treatments since a robot trader is added as a potential manipulator. The uninformed robot trader buys unconditionally ten shares of the asset in the beginning of a trading period and sells them afterwards unconditionally together with its initial endowment of four shares.

In our statistical analysis, we concentrate on the price of the asset and the payoffs of

²Further theoretical and empirical literature on market manipulations includes, among others, Allen and Gale (1992), Benabou and Laroque (1992), Kumar and Seppi (1992), Bagnoli and Lipman (1996), Camerer (1998), Strumpf and Rhode (2003), Hanson et al. (2004), Chakraborty and Yilmaz (2004), Ottaviani and Sørensen (2007), Hanson and Oprea (2007), and Goldstein and Gumbel (2008). See Veiga and Vorsatz (2008) for a short summary.

different groups of traders. With respect to the price of the asset, we find that the market disseminates information better than it aggregates it and that the aggregation of information becomes better if there are both more imperfectly informed subjects and the value of the asset is high (Result 1). Moreover, it is evidenced that the robot trader never influences the average last contract price, but it affects the average contract price positively if, and only if, information is perfect and the value of the asset is 125 ECU (Result 2). This finding has two implications. First, the manipulability of the market by means of the suggested trading pattern is closely linked to the distribution of information since the initial purchases of the robot are only believed to be informative if signals are perfect. Second, any successful manipulation is of temporary nature that vanishes until the market closes.

Regarding payoffs, we may conclude that the robot trader manages to be quite profitable. It earns always more than the representative uninformed subject, its payoff is never statistically lower than the one of the representative informed trader, and it earns significantly more than the average trader if all subjects are imperfectly informed (Result 3). Consequently, manipulations are individually rational. We also find, as expected, that the payoff of the informed (uninformed) subjects is highest (lowest) when private information is perfect (Result 4). Finally, we observe that the payoff of the informed (uninformed) subjects as a group tends to increase (decrease) thanks to the presence of the robot trader, and there is evidence that some of the informed traders suffer from the intended manipulations: if all subjects are imperfectly informed, those traders who receive the message that the value of the asset is not 375 ECU earn significantly less in the presence of the robot trader (Result 5).

The remainder of the paper proceeds as follows. In the next section, we describe our experimental design and procedures. Afterwards, we introduce the hypotheses that are tested with the help of the experiment. In Section 4, we present our results together with the corresponding statistical analysis. Finally, we conclude. The graphical representations of the price paths and the experimental instructions are relegated to the Appendices.

2 Experimental Design and Procedures

2.1 Setting

We consider a three-by-two between subject design. In all three baseline treatments, twelve subjects trade a common value asset, which takes the values 125, 375, and 525 ECU (experi-

mental currency units) with equal probability, for ten rounds in an electronic double auction market. We fixed one series of values similar to the one applied by Plott and Sunder (1988) in their market number 8 before conducting the experiment.

Every round of a session, subjects receive four shares of the asset and an interest free loan of 25.000 ECU that has to be returned at the end of the trading period. Using these endowments the asset is traded for five minutes in an electronic double auction market by submitting bid and ask prices. A trade takes place whenever a subject accepts a standing buy or sell offer. It is not possible to trade multiple units of the asset simultaneously and short sales are not permitted.

It is common knowledge that no subject has information about the asset's value before the market closes in the first round of the experiment and that some subjects are privately informed before the market opens from round two on. The first round has only been included in our design to familiarize subjects with the trading platform and, therefore, it is excluded from the statistical analysis later on.

In continuation, we are more specific on how private information is distributed. In treatment *B6*, six randomly selected subjects receive some private information regarding the true value of the asset. For example, if the actual value of the asset is 125 ECU, three subjects get to know that the value is not 375 ECU and three subjects learn that the value is not 525 ECU. Formally, we can describe the assignment of information in this treatment as follows: if the state space is denoted by $S = \{x, y, z\}$ and the actually selected state is x , then three traders get to know that the state is not y and three traders learn that it is not z . Hence, private information is imperfect but there is no aggregate uncertainty. The probability to be privately informed is the same for all subjects in all rounds. Subjects are fully aware about how information is allocated.

Treatment *B12* differs from *B6* only because every subject gets partially informed. Information is again uniformly distributed meaning that if the value of the asset is x , six subjects learn that it is not y and six subjects get to know that it is not z . Finally, in treatment *BP*, four randomly determined traders are perfectly informed (*i.e.*, they learn the actual value of the asset), while the other eight subjects stay uninformed.

The manipulation treatments, abbreviated by *M6*, *M12*, and *MP*, differ from their associated baseline treatment due to the introduction of a computer program as an additional

uninformed trader.³ This robot trader is active from round two on and once the market is open for 25 seconds, it places additional random bids in the following way: 20 ECU are added to the highest bid price to come up with a hypothetical price the robot would pay at most. Then, a random bid between the highest bid price and the hypothetical price calculated before is placed. So, eventually one share of the asset is bought. After a random delay time between 4 and 8 seconds, this procedure is repeated until 10 of the 48 outstanding shares are bought. Since it may happen that the robot has a standing buy offer when it is about to place a new bid, the old buy offer is deleted after calculating the new bid price but before placing the new offer. In this way, it is insured that the robot has at any point in time at most one standing buy offer. Finally, all shares (the 10 shares bought plus the endowment of 4 shares) are sold in a similar way. The program stops automatically 50 seconds before the end of a trading period. Table 1 below summarizes the experimental design.

Distribution of Information	Environment	
	Baseline	Manipulation
6 imperfectly informed / 6 uninformed	<i>B6</i>	<i>M6</i>
12 imperfectly informed / 0 uninformed	<i>B12</i>	<i>M12</i>
4 perfectly informed / 8 uninformed	<i>BP</i>	<i>MP</i>

Table 1: Experimental design.

2.2 Implementation

We conducted the experiment, which was programmed within the z-Tree toolbox (see Fischbacher, 2007), in the computer laboratory at Maastricht University between June and December 2007. Since all students from the Faculty of Economics and Business Administration have an e-mail account associated with their student-ID, we promoted the experiment mainly via electronic newsletters and gave students the opportunity to register online for their preferred session. In total, 432 unexperienced undergraduates participated in one of the 36 experimental sessions (six sessions per treatment). No student took part in more than one session and within a given session, subjects did not know each other.

At the beginning of a session, students met outside the laboratory. We prepared cards

³In our experimental instructions, we informed subjects that an uninformed robot trader participated in the market (*i.e.*, the robot was not programmed conditional on the asset's value), but the exact program structure was not revealed. See Appendix B for a detailed description.

with the numbers from one to twelve and let each student draw one card. If more than twelve students showed up for a particular session, we offered three Euros in case somebody was willing to leave. If an insufficient number of students decided to leave, we put additional empty cards into the deck and determined the participating students randomly. Left-out students received a compensation of 3 Euros. Finally, we reminded everybody that any kind of communication inside the laboratory would lead to an immediate cancellation of the session.

Students then entered the laboratory and took seat in front of the computer corresponding to the number of their card. The computers were distributed in such a way that subjects could not see each other and, next to each computer, we placed the instructions, an official payment receipt, and a set of control questions (see Appendix B). Subjects could study the instructions at their own pace and eventual occurring doubts were privately clarified. The experiment started once everybody answered all control questions correctly.

At the end of a session, subjects were paid privately one-by-one. In addition to the three Euro show-up fee, we offered 90 Euro-cents for every 1000 ECU obtained in the experiment. As a result, the average payment of the 90 minutes session was equal to 15.8 Euro in the baseline treatments. The average payoff of the manipulation treatments was slightly lower due to the partial gains of the robot trader.

3 Hypotheses

In this section, we present our experimental hypotheses. We distinguish between hypotheses that regard the price of the asset and hypotheses related to payoffs.

3.1 Prices

Following Table 1, we divide our analysis with respect to the price of the asset into two dimensions. The first dimension focuses on price variations between treatments when the distribution of private information (six imperfectly informed, twelve imperfectly informed, or four perfectly informed subjects) changes but the environment (baseline or manipulation) is held fixed, whereas the second dimension regards price variations between treatments when the environment changes but the distribution of information is maintained. The former dimension therefore addresses the question whether the market can aggregate information

as well as disseminate information and whether the aggregation of information depends on the number of imperfectly informed traders. The latter is concerned with the manipulability of the market.⁴

Regarding the effect of the distribution of private information on prices, the null hypothesis is derived from the rational expectations equilibrium. Given that there is no aggregate uncertainty in the market, it predicts that the market price converges to the true value of the asset independently of the distribution of private information. The corresponding alternative hypothesis is based on experimental evidence that the aggregation of information by the market price is more difficult than the dissemination of information whenever markets are incomplete (see Plott and Sunder, 1988). Hence, in the baseline environment, prices should be closest to the rational expectations equilibrium in treatment *BP*. Moreover, more informed subjects should induce a better and faster aggregation of information and, therefore, we hypothesize that prices are closer to the rational expectations equilibrium in *B12* than in *B6*. The former statements should also be true for the manipulation environment because the manipulator is not expected to be dominant enough to overturn this relation.

HYPOTHESIS 1: For a given environment, the market disseminates information better than it aggregates it and the aggregation of information becomes better as the number of imperfectly informed subjects increases.

Regarding the manipulability of the market, the null hypothesis states that the manipulation attempt of the robot trader fails. This implies that given any distribution of information, the last contract price – calculated conditionally on the true value of the asset – should be the same for both environments. The alternative hypothesis, on the other hand, states that the manipulation attempt is successful and that some of the traders fall into an information trap yielding higher prices in the manipulation environment.⁵ At this point it is important to note

⁴We only discuss between-treatment comparisons of prices at this point. The additional within-treatment tests presented later on reveal that the market price is positively correlated with the value of the asset. These tests confirm therefore the very basic idea that markets are indeed able to incorporate some information in prices and that the robot trader is not strong enough to eliminate this property. Consequently, the within-treatment comparisons of prices constitute a lower bound to test the proper functioning of the market.

⁵Strictly speaking, the robot trader cannot manipulate the market when the true value of the asset is 525 ECU. But it may act as an additional competitor on the buying side of the market, which could lead to a better aggregation of information if prices in the baseline environment fall short of the rational expectations equilibrium.

that not only uninformed traders but also some of the imperfectly informed subjects may be affected by the presence of the robot trader. To clarify this point, suppose, for instance, that true value of the asset is low. Then, the insiders with the information that the value of the asset is not 375 ECU are particularly prone to fall into an information trap, because the initial purchases of the manipulator may be interpreted as a signal for a high value. The same rationale applies to subjects with the information that the asset is not worth 125 ECU when its actual value is 375 ECU.⁶

HYPOTHESIS 2: For a given distribution of information, prices are higher in the manipulation treatment.

3.2 Payoffs

The competing hypotheses with respect to payoffs are divided into two main parts: first, we analyze payoff differences between distinct groups of traders (uninformed, informed, and robot) in a given treatment. Afterwards, we check for payoff differences across treatments for a given group of traders. To be consistent with Table 1, the latter analysis distinguishes again whether the environment or the distribution of information is held fixed.

Regarding the within-treatment tests, the null hypothesis of the rational expectation equilibrium implies that insiders do not earn more than uninformed traders. However, if information is not aggregated instantaneously, as it is usually the case in experiments, the informational advantage of the insiders should be reflected by higher payoffs in comparison to the uninformed traders. Therefore, the main question becomes how the payoff of the robot in a given manipulation treatment relates to the ones of the other two groups of traders.

Since the robot trader is itself uninformed, the group of uninformed traders forms a natural peer-group for evaluating the success of the manipulator in monetary terms. Under the null hypothesis, the noise trading does not affect prices and, therefore, the robot trader should not earn more than the representative uninformed trader. The alternative hypothesis states the opposite. Also note that we are silent on the payoff comparison between the robot and the group of all insiders because it is very likely that the uninformed robot trader earns less than the representative informed trader even if the manipulation is successful.

⁶We come back to this point when discussing how the payoff of different groups of traders is affected by the manipulator.

HYPOTHESIS 3: In treatments M6 and MP, the payoff of the robot trader is higher than the one of the average uninformed trader.

Now, we focus on the comparative statics with respect to the distribution of information. We start by considering the baseline environment. According to the rational expectations equilibrium both the informed and the uninformed subjects should earn the same in all treatments (the fourth null hypothesis). However, if information is valuable, the payoff of the insiders may change with the presence/absence of uninformed traders. Naturally their payoff is lowest in the absence of uninformed traders (treatment *B12*), because then there is no informational rent to be earned. Moreover, since insiders can still misjudge the value of the asset when information is imperfect, the payoff of the informed (uninformed) subjects should be higher (lower) in treatment *BP* than in treatment *B6*. Again, there is no particular reason why this relationship should be overturned in the manipulation environment.

HYPOTHESIS 4: For a given environment, the payoff of the informed (uninformed) subjects is highest (lowest) when private information is perfect.

The more interesting question is how the payoffs of insiders and uninformed traders change due to the presence of the robot trader. Under the null hypothesis, the noise trading is not successful and, given any distribution of information, the payoffs of uninformed and informed traders should not be lower in the manipulation than in the corresponding baseline treatment. Hence, the manipulator will fully account the burden of its failure.

The alternative hypothesis is explicit about who suffers and who gains in case the manipulation is successful. As it has already partly been expressed in Hypothesis 3 and in the discussion before Hypothesis 2, not only the uninformed but also some of the imperfectly informed traders may get affected negatively. In particular, if the value of the asset is 125 (375) ECU, the traders with the signal that the value of the asset is not 375 (125) ECU may be misled by the manipulator. At the same time, perfectly informed insiders and traders with the signal that the value of the asset is not 525 ECU, who should be on the selling side of the market if the manipulation works out, would be able to earn some extra money thanks to the higher prices. This discussion also shows that is not possible to come up with a clear alternative hypothesis on how the payoff of the insiders as a group should be affected by the

presence of a successful manipulator if information is imperfect.

HYPOTHESIS 5: *Uninformed traders and subjects with the information that the value of the asset is not 125 (375) ECU earn more in a baseline than in the associated manipulation treatment. Perfectly informed insiders and traders with the signal that the value of the asset is not 525 ECU earn more in a manipulation than in the associated baseline treatment.*

4 Results

4.1 Average Price Paths

In this section, we present the results of our experiments. Figure 1 shows the average sequences of contract prices. In order to make the data as visible as possible, it is divided into six panels such that every column corresponds to a different value of the asset (high, middle, and low) and every row belongs to a different environment (baseline and manipulation). The top left panel, for instance, corresponds to the average price paths of the three baseline treatments when the fundamental value of the asset is 525 ECU. In this panel, the continuous line corresponds to treatment *B12*, the dashed line to treatment *B6*, and the dotted line to treatment *BP*.⁷

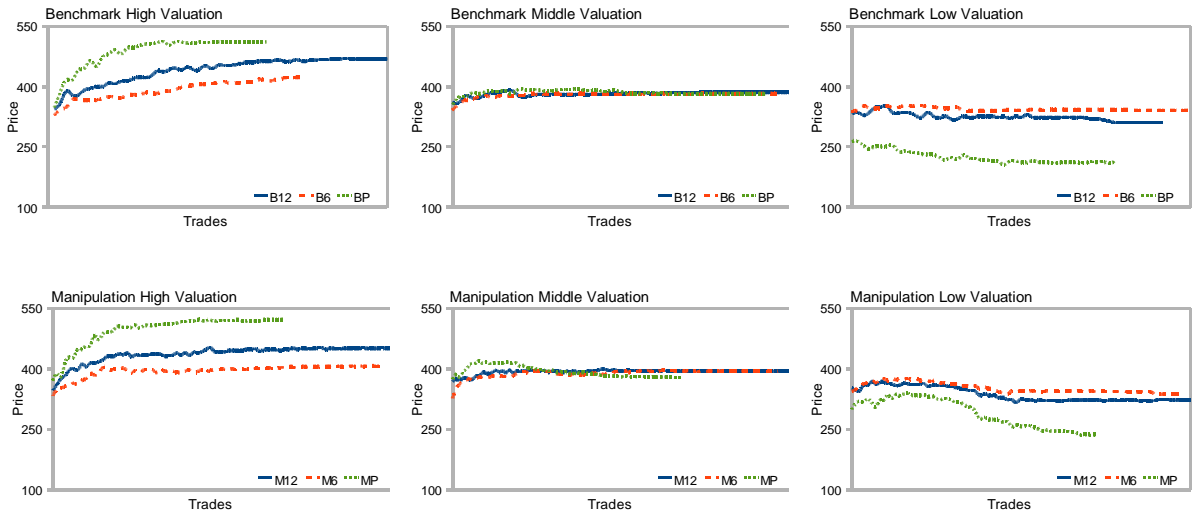


Figure 1: Average price paths.

⁷Appendix A presents the series of contract prices for all the 36 sessions. The value of the asset for a given round is indicated by a straight line in all panels.

If the value of the asset is 525 (125) ECU, different distributions of information can be clearly ranked with respect to their capacity to summarize information: independently of the environment, convergence to the rational expectations equilibrium becomes better as we move from six imperfectly informed to twelve imperfectly informed to four fully informed subjects. It is also worth noticing that increasing the number of partially informed traders has only a minor impact, while going from imperfect to perfect private information improves the convergence to the rational expectations equilibrium substantially. In fact, in these two states, the average last contract price is only close to the equilibrium prediction if private information is perfect and the value of the asset is 525 ECU. Finally, since all average price paths are almost straight lines close to the equilibrium prediction when the value of the asset is 375 ECU, there is substantial evidence that the market disseminates information better than it aggregates it and that the aggregation of information is to some extent positively correlated to the number of imperfectly informed traders (Hypothesis 1).

So far, we have informally described how different distributions of information affect average prices over time, but the main research objective is to investigate whether the robot trader is able to alter prices. Regarding this question, it can be observed that the robot trader hardly influences the average price path if the value of the asset is either 525 or 375 ECU. If the asset is worth 125 ECU, the same observation can be made in case information is distributed imperfectly. Things are only different when treatment *BP* is compared with treatment *MP*: in treatment *BP*, the average price path starts close to 250 ECU and decreases slowly to about 220 ECU, where it stays from there on. The first average contract price in treatment *MP* is already substantially higher and even increases for a considerable time up to values close to 300 ECU before falling back to about 240 ECU. This is an indication that the robot trader is able to manipulate prices if, and only if, information is perfect and the value of the asset is low.

4.2 Statistical Analysis: Prices

In our statistical analysis, we take into account the last contract price from each round. We proceed by calculating, for each session, the average last contract price over rounds conditional on the value of the asset. For each treatment and for each possible value of the asset, we obtain therefore six independent observations (one per session). The entries in Table 2 correspond to

the average independent observation. They coincide with the last points of the average price paths depicted in Figure 1. In continuation, we run Wilcoxon–signed rank tests (for within treatment comparisons) and Mann–Whitney U tests (for between treatments comparisons) on the independent observations. All reported p –values are one–sided.

		High	Middle	Low
Baseline	<i>B6</i>	423	384	346
	<i>B12</i>	469	386	311
	<i>BP</i>	512	384	213
Manipulation	<i>M6</i>	407	394	338
	<i>M12</i>	456	394	323
	<i>MP</i>	524	379	237

Table 2: Average last contract prices.

Before focusing on the main hypotheses, we check whether prices and values are positively correlated within a given treatment. In treatment *M6*, there is no significant difference between prices when the value of the asset is 525 and 375 ECU, respectively ($p=0.1042$). This result is not caused by particularly high prices when the value of the asset is 375 ECU, as a possible manipulation the robot trader would suggest, but rather by the poor aggregation of information when the value of the asset is high. In fact, Table 2 reveals that the average independent observation only amounts to 407 ECU if the value of the asset is 525 ECU.⁸ All other within treatment tests lead to the intuitive conclusion that higher values lead to higher prices ($p < 0.0300$, throughout).

4.2.1 The Effect of the Distribution of Information

We analyze whether the distribution of information influences the final contract price in a given environment. Suppose first that the value of the asset is 525 ECU. We find that the average last contract price when four subjects are perfectly informed is in both environments significantly higher than the one when all subjects are imperfectly informed ($p = 0.0386$ for the baseline and $p = 0.0025$ for the manipulation environment) and when six traders receive

⁸A detailed look at the data in Appendix A shows that information is very poorly aggregated in session 3 of treatment *M6* whenever the value of the asset is high. However, this is the only session where a lower value of the asset leads to a higher price.

partial information ($p = 0.0025$ for both the baseline and the manipulation environment). This result is rather intuitive. Since perfectly informed insiders are willing to buy the asset at any price below 525 ECU, one would expect that competition drives the price up to the rational expectations equilibrium. In fact, the average last contract price is 512 ECU in treatment *BP* and 524 ECU in treatment *MP*.⁹ On the other hand, if information is imperfect, traders always bear some risk and may thus be reluctant to buy the asset at very high prices. As a result, not all private information is aggregated by the market. Finally, we also find that the last contract price is positively correlated with the number of imperfectly informed traders ($p = 0.0388$ for the baseline and $p = 0.0066$ for the manipulation environment).

We now focus on the case when the value of the asset is 375 ECU. If the price of the asset is below 375 ECU, some of the insiders will always be able to make riskless gains, independently of the distribution of information. Consequently, the price of the asset should not be below the rational expectations equilibrium. But the price may very well end up above 375 ECU due to existing short-selling constraints. Wilcoxon signed-rank tests provide evidence that the average last contract prices are not statistically significant different from the rational expectations equilibrium in treatments *BP* ($p=0.1807$) and *MP* ($p=0.2309$). Consequently, in this state, information is very well disseminated. Regarding treatments *B6* and *B12*, the statistical analysis reveals that both treatments generate prices that are significantly higher than the equilibrium prediction ($p = 0.0296$ for *B6* and $p = 0.0180$ for *B12*) but not significantly different between each other ($p = 0.4051$). Exactly the same results are found in the manipulation environment. The last contract price is significantly higher than 375 ECU in *M6* ($p = 0.0296$) and *M12* ($p = 0.0180$) but there is no statistical difference between the two treatments ($p = 0.4051$). Again, the aggregation of information turns out to be more complicated than its dissemination.

Finally, we analyze the case when the value of the asset is 125 ECU. We observe in Table 2 that in the baseline environment, the average last contract price is the lowest in treatment *BP*. The statistical analysis confirms this, because the p -values for the comparison with treatment *B6* (*B12*) is 0.0041 (0.0153). Also observe that the still considerable difference to the rational expectations equilibrium in treatment *BP* could be caused by the absence of short-selling opportunities. Moreover, there is statistical evidence that the average last contract price in

⁹The slight difference in treatment *BP* with respect to the equilibrium prediction stems from the fact that the price in round 3 of session 4 is only 350.

treatment $B6$ is not different from the one in treatment $B12$ ($p = 0.1149$). The very same results are found in the manipulation environment. The p -values are 0.2352 when comparing treatment $M6$ with treatment $M12$, 0.0065 when comparing $M6$ with MP , and 0.0868 when comparing $M12$ with MP .

Table 3 and Result 1 summarize the results found so far.

Value	Baseline	Manipulation
High	$BP > B12 > B6$	$MP > M12 > M6$
Middle	$BP = 375; 375 < B12 = B6$	$MP = 375; 375 < M12 = M6$
Low	$BP < B12 = B6$	$MP < M6 = M12$

Table 3: Price differences across treatments.

Result 1. *In both environments, the market disseminates information better than it aggregates information. If the value of the asset is high (not high), the aggregation of information is better (not worse) the more imperfectly informed traders there are.*

4.2.2 The Effect of the Manipulator

Our next objective is to assess the impact of the robot trader on prices. Figure 1 suggests that the presence of the robot trader affects prices positively if, and only if, private information is perfect and the value of the asset is low. But Table 2 only reveals a small difference between treatment BP and treatment MP (213 vs. 237 ECU), which does not turn out to be statistically significant ($p = 0.3445$). Moreover, since the p -values of all other possible comparisons between the two environments for a fixed distribution of information lie between 0.1679 and 0.4681, we can conclude that the manipulator does not affect the last contract price significantly.

A closer look at Figure 1 suggests above all that the convergence paths and not the final contract prices differ across environments if the value of the asset is low and information is perfect. In order to test for this, we consider from now on average instead of final contract prices. The average independent observations are presented in Table 4.

We observe that if the asset has a low value, the average contract in treatment MP is higher than in treatment BP (303 vs. 234 ECU). This difference is statistically significant at the 5%

		High	Middle	Low
Baseline	<i>B6</i>	385	379	353
	<i>B12</i>	429	381	334
	<i>BP</i>	476	385	234
Manipulation	<i>M6</i>	392	385	360
	<i>M12</i>	435	392	350
	<i>MP</i>	488	382	303

Table 4: Average contract prices.

significance level ($p = 0.0464$). For all other possible comparisons, the differences between baseline and manipulation treatments are still rather negligible. In fact, no difference is significant at any conventional significance level.¹⁰ The following result summarizes our main findings with respect to the manipulability of the market.

Result 2. *The presence of the robot trader does not affect the last contract price. It has a positive impact on the average contract price if, and only if, private information is perfect and the value of the asset is low.*

4.2.3 Interpretations and Implications

According to Result 1 the market disseminates information better than it aggregates information. The aggregation of information is also positively correlated to the number of partially informed traders if the value of the asset is 525 ECU. Since the aggregation of information does not depend on the number of imperfectly informed subjects in the other two states, we can confidently reject the first null hypothesis in favor of Hypothesis 1.

Result 2 offers two main implications with respect to the manipulability of the market. First, a necessary condition is that information is perfect. Otherwise, subjects do not update their beliefs regarding the true value of the asset upon the initial trades of the robot. Second, the manipulability of the market also depends on the number of informed traders and the market liquidity. In our earlier study (see Veiga and Vorsatz, 2008), where subjects were endowed with two instead of four shares and three instead of four subjects were informed, it

¹⁰The lowest p -value (0.0737) corresponds to the comparison between the treatments *M12* and *B12* (392 vs. 381 ECU) when the value of the asset is 375 ECU. This result nicely captures the initial increase in the average price path in treatment *M12*. Nevertheless, it is impossible to conclude that the manipulation is successful in this case, because the price increase is not very persistent and the robot trader loses, on average, money.

was found that the robot was able to induce a higher final contract prices if the value of the asset was low. Here, we only find manipulations in the average contract price. Observe that this difference should not be attributed to the fact that in the current study there are three instead of two states. Information is very well disseminated in both environments if the value is not 125 ECU and, therefore, the results are not caused by a potentially more complicated experimental setup.

4.3 Statistical Analysis: Payoffs

Next, we compare the payoffs of different groups of traders (informed, uninformed, and robot) for different distributions of information and different environments. The data is prepared as follows: we calculate first, for each session, the average per round payoff of the representative trader in a given group. For example, in treatment *B6*, we sum up the per round payoffs of all imperfectly informed traders over the whole session and divide it by the product of the number of imperfectly informed traders (six) and the number of rounds per session (nine). Consequently, for each representative trader, we have six independent observations per treatment. The entries in the following table are then obtained by averaging over all independent observations. Note that the average payoff per round is equal to $\frac{(125+375+525) \cdot 3}{9} \cdot 4 = 1366.67$ ECU. This is the payoff any subject can insure herself/himself by not trading.

		Robot	Uninformed	Insider (all types)
Baseline	<i>B6</i>	–	1317	1416
	<i>B12</i>	–	–	1367
	<i>BP</i>	–	1274	1552
Manipulation	<i>M6</i>	1465	1230	1487
	<i>M12</i>	1502	–	1355
	<i>MP</i>	1391	1238	1618

Table 5: Average payoff per round (overall).

Since insiders earn more than uninformed subjects, information is valuable. We also observe that the robot trader turns out to be very successful. It earns not only more than the representative trader of its natural peer-group, the group of uninformed traders, but its payoffs is also higher than the average payoff of 1367 ECU. In fact, the payoff of the robot

trader is higher in *M12* than in *M6* and, in comparison to the representative informed trader, the robot performs better in *M12*, similar in *M6*, and worse in *MP*. Finally, in treatments with both informed and uninformed subjects, the insiders seem to profit from the presence of the robot trader, while the uninformed traders tend to have lower payoffs.

4.3.1 Within Treatment Comparisons

We start our analysis by comparing the average payoff of the group of all insiders, independently on whether they are partially or perfectly informed, with the one of the uninformed subjects. We find that insiders earn significantly more than uninformed subjects, the only exception is treatment *B6* ($p = 0.0711$). In this treatment payoffs are not significantly different from each other because three imperfectly informed subjects, who received in rounds 3 and 9 of session 2 the information that the value of the asset is not 375 ECU, bought simultaneously a lot of shares. In this way, they created an information trap by believing that the value of the asset was 525 ECU although, effectively, it was 125 ECU. As a consequence, the payoff of the group of all informed traders was slightly lower than the one of the uninformed traders.

The performance of the robot trader is evaluated next. In treatment *M6*, the average payoff of the robot trader is significantly higher than the one of the typical uninformed subject ($p = 0.0468$). The underlying reason of this result is that the robot trader earns significantly more than the uninformed subjects whenever the value of the asset is low ($p=0.0206$) and that payoffs are not significantly different from each other in the other two states ($p = 0.3412$, in both cases). Comparing the earnings of the robot trader with the one of the typical uninformed subject in treatment *MP*, it turns out that the difference is not significant ($p = 0.1473$) even though the manipulator earns substantially more (1391 vs. 1238 ECU). This result is due to the fact that the robot trader is very profitable in sessions 1–3, in which manipulations are often successful, but performs rather poor in sessions 4–6, where the manipulations fail.

Finally, we compare the payoff of the robot trader with the one of the representative insider. In treatment *M6*, payoffs are not significantly different from each other ($p = 0.3375$), which is already suggested by the very similar per round payoffs (1465 vs. 1487 ECU). In treatment *M12*, the average per round payoff of the robot trader (1502 ECU) is higher than the one of the average insider (1355 ECU), but the difference is not significant ($p = 0.0711$). The typical insider gains substantially more than the robot trader in treatment *MP* (1681

vs. 1391 ECU). This difference is again not significant ($p = 0.2009$), because the average per round payoff of the robot trader is actually higher than the one of the representative insider in sessions 1–3. For this to happen two conditions should be met: first, if the value of the asset is 375 or 525 ECU, the dissemination of information has to be rather fast so that the informed traders cannot profit too much from the noise trading of the robot. Second, if the value of the asset is 125 ECU and the informed traders are on the selling side of the market, the initial purchases of the robot have to keep up prices sufficiently long so that the robot trader can unwind its position and still earn more than the average informed trader. The latter is consistent with the hump-shaped average price path in the lower right panel of Figure 1 and our earlier finding that the average contract price in treatment *MP* is rather high if the value of the asset is low.

Result 3. *Insiders earn always more than uninformed subjects and with exception of treatment B6, the difference is significant. The average per round payoff of the robot trader is never significantly lower than the one of either the uninformed traders or the insiders and in treatment M6, it earns significantly more than the representative uninformed subject.*

4.3.2 Between Treatments Comparisons

Regarding Hypothesis 4, we analyze how the payoff of a given representative trader is affected by changes in the distribution of private information. In treatment *B12*, the representative informed trader earns 1367 ECU, which is, as expected, the minimum for this group of traders. Moreover, perfect private information should allow the insiders to exploit their informational advantage at maximum. This intuition is confirmed by the statistical analysis. The representative informed trader earns significantly more in *BP* than in *B12* ($p = 0.0180$) and *B6* ($p = 0.0101$). The difference of 46 ECU between treatment *B6* and *B12* is, however, not significant at a 5% significance level ($p=0.0711$). Similar results are found for the manipulation environment. The insiders have significantly higher payoffs in *MP* than in *M12* ($p = 0.0026$) and *M6* ($p = 0.0153$). They also earn significantly more in *M6* than in *M12* ($p = 0.0026$).

We also test whether the payoffs of the representative uninformed subject and the robot trader change with the distribution of information. Since all p -values are higher than 0.1149, we conclude that their payoff is not significantly affected by this treatment variable.

Result 4. *The average payoff per round of the representative informed trader is highest*

(lowest) if private information is perfect (if all subjects are imperfectly informed). The average per round payoffs of the representative uninformed subject and the robot trader do not depend on the distribution of information.

We study next how the presence of the robot trader affects the payoff of the informed and uninformed subjects. Table 5 shows that independently of the distribution of information, insiders tend to earn more in the presence of the manipulator (given that also uninformed subjects participate in the market), whereas the uninformed traders earn less. Statistically however, only one difference is significant: if information is imperfect, the uninformed subjects suffer from the presence of the robot trader ($p = 0.0464$). If, on the other hand, information is distributed perfectly, the difference in earnings across environments for this group of traders is not significant ($p = 0.1149$). Finally and regarding the representative informed trader, the p -value is 0.0641 when comparing payoffs in treatments $M6$ and $B6$, 0.0711 for $M12$ and $B12$, and 0.1490 for MP and BP .

Our final tests regard the question whether imperfectly informed subjects may also be deluded by the manipulator. Table 6 clearly shows that this is not the case when there are six imperfectly informed traders because independently of the type of information received, insiders earn more in $M6$ than in $B6$. This also implies that all potential gains of the robot trader are on behalf of the uninformed traders. On the other hand, if there are twelve imperfectly informed subjects, we observe that all changes the robot trader provokes are in the direction suggested by Hypothesis 5. Most importantly, the traders who receive the signal that the value of the asset is not 375 ECU earn significantly less in the manipulation environment ($p = 0.0464$). This result is caused by the fact that subjects with this information earn significantly less if the true value of the asset is 125 ECU ($p = 0.0227$). In fact, in this state their average per round payoff is 339 ECU in $B12$ and only 5 ECU in $M12$. This result, together with the earlier finding that the robot trader earns more than the average subject in treatment $M12$, can therefore be taken as a hint that also imperfectly informed subjects may fall into information traps. All other payoff differences between treatment $M12$ and treatment $B12$ are, as it can be seen in Table 6, not significant.¹¹

¹¹We want to make a final comment on Table 6. Among the imperfectly informed subjects, the traders with the information that the asset is not 525 ECU profit more from the presence of the robot trader, while those with the information that the value of the asset is not 375 ECU profit less (suffer more). This finding is in line with Hypothesis 5.

<i>Group of Traders</i>	<i>M6 vs. B6</i>	<i>M12 vs. B12</i>	<i>MP vs. BP</i>
Insider (all)	71 (0.0641)	-12 (0.0711)	66 (0.1490)
Insider (not 525 ECU)	96 (0.1893)	118 (0.1893)	—
Insider (not 375 ECU)	52 (0.3445)	-157 (0.0464)	—
Insider (not 125 ECU)	63 (0.3445)	-10 (0.2875)	—
Uninformed	-87 (0.0464)	—	-36 (0.1149)

Table 6: Payoff differences across environments (with p -values).

Result 5. *If there are twelve (six) imperfectly informed subjects, the traders with the information that the asset is not worth 375 ECU (the uninformed traders) suffer significantly from the presence of the robot trader.*

4.3.3 Interpretations and Implications

We have seen in Result 3 that the robot trader earns always more than the average trader, but, nevertheless, we cannot reject the third null hypothesis in favor Hypothesis 3. This is because the payoff difference with respect to the uninformed traders is only significant in *M6*. Result 3 suggests nevertheless that it is individually rational for a potential manipulator to follow the trading strategy suggested here (and in the theoretical work of Chakraborty and Yilmaz (2008)).

We also found that the average payoff of the representative informed subject is higher if private information is perfect. On the other hand, the average per round payoffs of the representative uninformed subject and the robot trader do not depend on the distribution of information and, therefore, we cannot reject the fourth null hypothesis in favor of Hypothesis 4. Nevertheless it should be observed that the signs of the changes are consistent with Hypothesis 4.

Finally, and most importantly, even though the uninformed subjects suffer most from the presence of the manipulator, partially informed subject can get affected as well. Traders with the signal that the value of the asset is not 375 ECU earn significantly less in treatment *M12* than in treatment *B12*. This is a new result which complements our earlier finding that uninformed traders can be deluded by a robot if information is perfect. Consequently, the fifth null hypothesis can be partly rejected in favor of Hypothesis 5.

5 Conclusion

In this paper, we studied experimentally the conditions on the distribution of private information under which a simple asset market can be manipulated by a robot trader. Our main results with respect to the price of the asset revealed that: (i) the short-selling constrained double auction market disseminates information better than it aggregate information, (ii) the aggregation of information improves sometimes when the number of imperfect subjects increases, and (iii) the intended manipulation of the robot trader is only successful when private information is perfect and the value of the asset is low. Since the latter impact is only significant for the average contract, this result differs from our previous findings, where we evidenced that the last contract price was affected by the presence of the robot trader when insiders had perfect information and the value of the asset was low (see Veiga and Vorsatz, 2008). However, the weaker effect of the manipulator in comparison to our former work can be due to the facts that, in the present work, the market is more liquid (traders are endowed with four instead of two shares) and that there are more perfectly informed traders in the market (four instead of three).

Regarding payoffs we found that (i) the computer program is profitable in monetary terms, especially in the treatment when all traders are partially informed, (ii) the uninformed traders (insiders) have the lowest (highest) payoff if private information is perfect, and (iii) the uninformed traders and the subjects with the signal that the value of the asset is not 375 earn less in the presence of the robot trader. The first result clearly shows that manipulations are to be expected in this kind of market (they are individually rational), whereas the latter establishes for the first time that also insiders may be negatively affected by the manipulator.

We regard this research as a further step in order to understand how markets aggregate and disseminate information in the presence of manipulators. In our future research, we are going to assess the robustness of our results with respect to the distribution of payoffs. In case the asset is privately and not commonly valued and the market is still manipulable in a similar way as presented here, an efficiency loss occurs. To quantify this efficiency loss is certainly an important task.

References

- Allen, F. and D. Gale (1992). Stock Price Manipulation. *Review of Financial Studies* 5, 503–529.
- Bagnoli, M. and B. Lipman (1996). Stock Price Manipulation Through Takeover Bids. *RAND Journal of Economics* 27, 124–147.
- Bannerjee, A. (1992). A Simple Model of Herd Behavior. *Quarterly Journal of Economics* 107, 797–817.
- Benabou, R. and G. Laroque (1992). Using Privileged Information to Manipulate Markets: Insiders, Gurus and Credibility. *Quarterly Journal of Economics* 107, 921–958.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades. *Journal of Political Economics* 100, 992–1026.
- Camerer, C. (1998). Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting. *Journal of Political Economy* 106, 457–482.
- Camerer, C. and K. Weigelt (1991). Information Mirages in Experimental Asset Markets. *Journal of Business* 64, 463–493.
- Chakraborty, A. and B. Yilmaz (2004). Informed Manipulation. *Journal of Economic Theory* 114, 132–152.
- Chakraborty, A. and B. Yilmaz (2008). Microstructure Bluffing with Nested Information. *American Economic Review* 98, 280–284.
- Fischbacher, U. (2007). Z-Tree – Zurich Toolbox for Readymade Economic Experiments. *Experimental Economics* 10, 171–178.
- Forsythe, R. and R. Lundholm (1990). Information Aggregation in an Experimental Market. *Econometrica* 58, 309–347.
- Forsythe, R., T. Palfrey, and C. Plott (1982). Asset Valuation in an Experimental Market. *Econometrica* 50, 537–567.

- Goldstein, I. and A. Gmbel (2008). Manipulation and the Allocational Role of Prices. *Review of Economic Studies* 75, 133–164.
- Hanson, J., C. Schmidt, and M. Strobel (2004). Manipulation in Political Stock Markets-Preconditions and Evidence. *Applied Economic Letters* 11, 459–463.
- Hanson, R. and R. Oprea (2007). A Manipulator Can Aid Prediction Market Accuracy. *Economica*, forthcoming.
- Hanson, R., R. Oprea, and D. Porter (2006). Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization* 60, 449–459.
- Hayek, F. (1945). The Use of Knowledge in Society. *American Economic Review* 35, 519–530.
- Kumar, D. and D. Seppi (1992). Futures Manipulation with Cash Settlement. *Journal of Finance* 47, 1485–1502.
- Muth, J. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica* 29, 315–335.
- Nth, M., C. Camerer, C. Plott, and M. Weber (1999). Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs. *Mimeo*.
- Ottaviani, M. and P. Srensen (2007). Outcome Manipulation in Corporate Prediction Markets. *Journal of the European Economic Association* 5, 554–563.
- Pearlstein, S. (2003). Misplacing Trust in the Markets. *Washington Post* July 30.
- Plott, C. and S. Sunder (1982). Efficiency of Experimental Security Markets with Insider Information: An Application to Rational-Expectations Models. *Journal of Political Economy* 90, 663–698.
- Plott, C. and S. Sunder (1988). Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets. *Econometrica* 56, 1085–1118.
- Strumpf, K. and P. Rhode (2003). Historical Presidential Betting Markets. *Journal of Economic Perspectives* 18, 127–141.

Sunder, S. (1992). Market for Information: Experimental Evidence. *Econometrica* 60, 667–695.

Veiga, H. and M. Vorsatz (2008). Price Manipulation in an Experimental Asset Market. *European Economic Review*, <http://dx.doi.org/10.1016/j.euroecorev.2008.05.004>.

Wolfers, J. and E. Zitzewitz (2004). Prediction Markets. *Journal of Economic Perspectives* 18, 107–126.

Wyden, R. and B. Dorgan (2003). Call for Immediate Halt to Tax-funded “Terror Market” scheme. *Press Release July 28*.

Appendix A: Price Paths

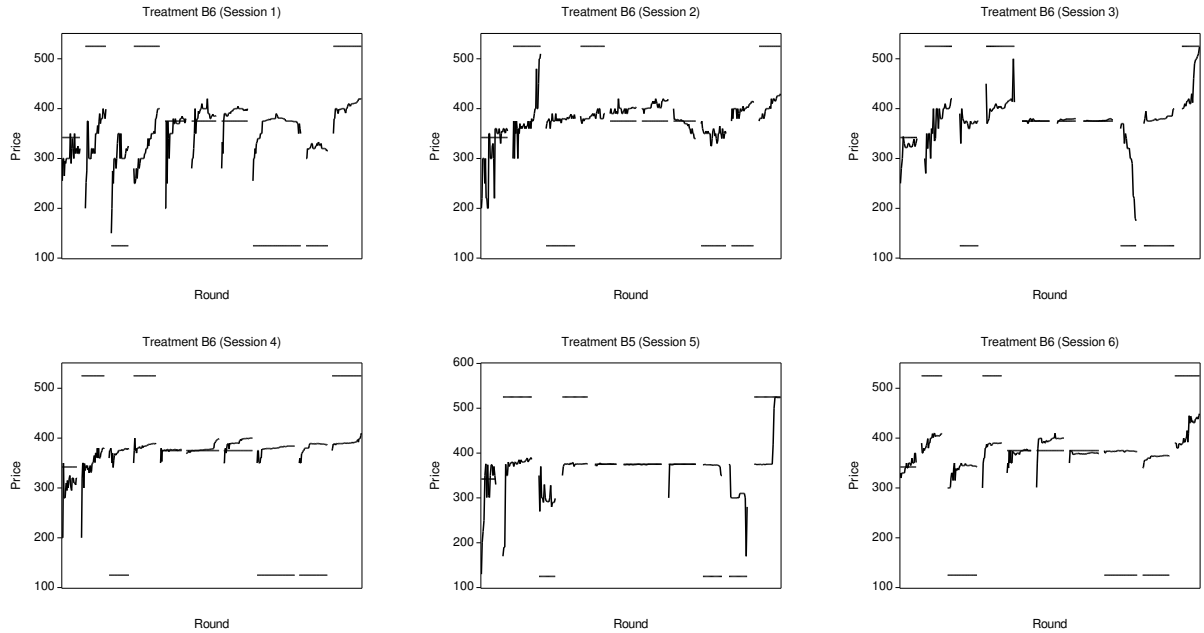


Figure 2: Price development in B6.




Figure 3: Price development in M6.

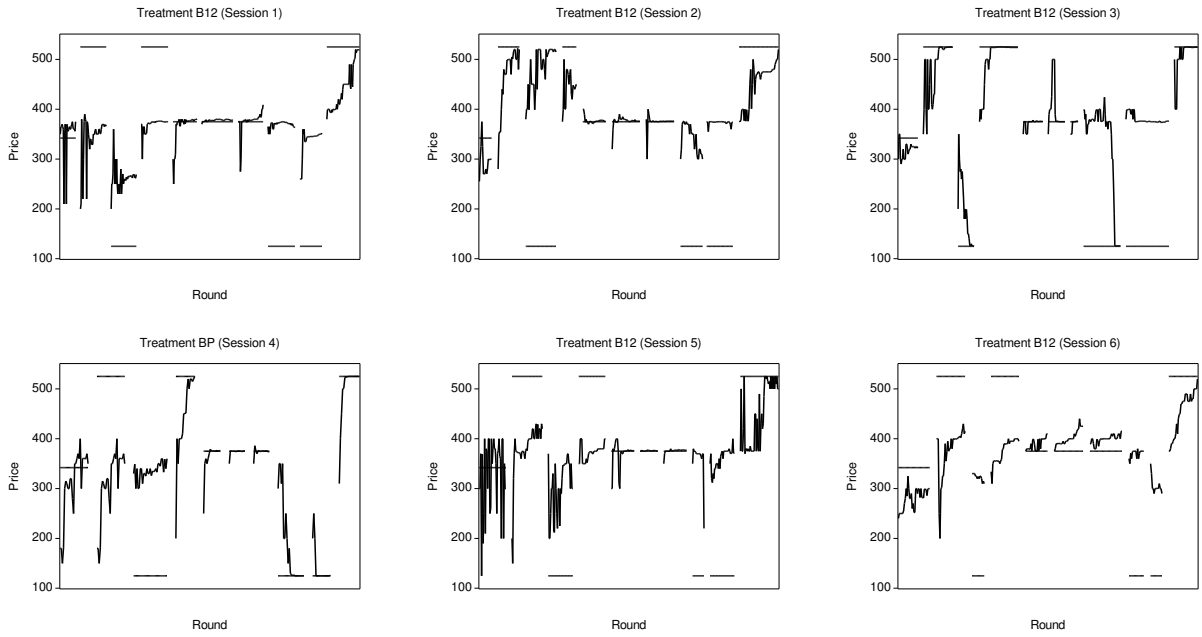


Figure 4: Price development in B12.

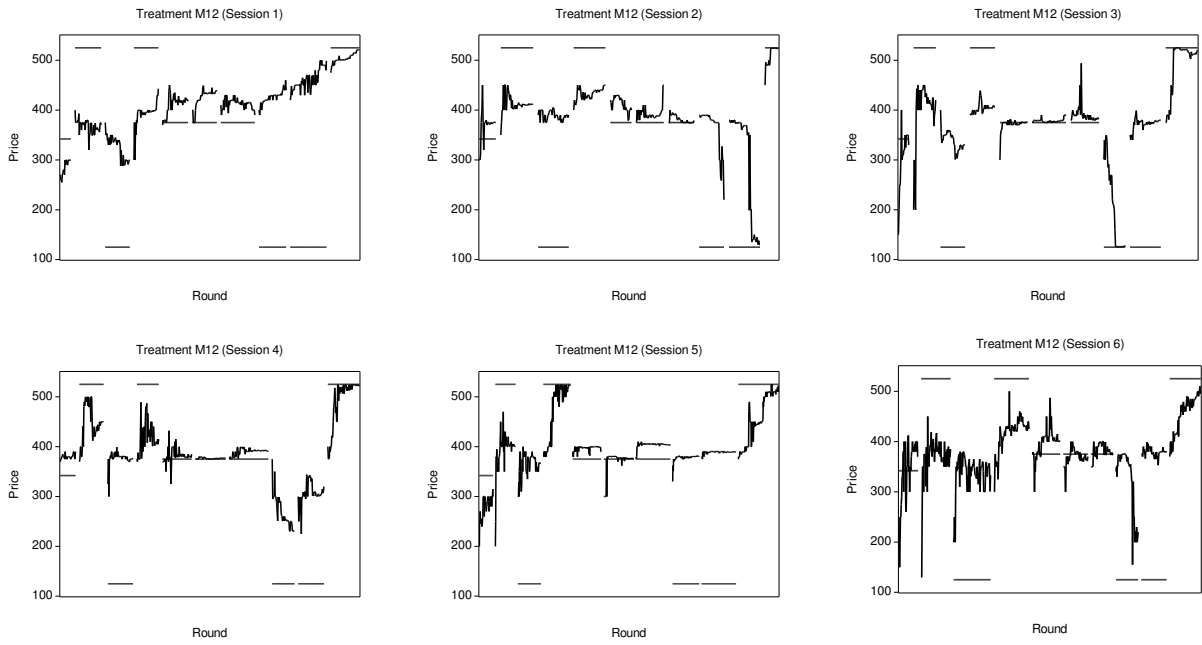


Figure 5: Price development in M12.




Figure 6: Price development in BP.

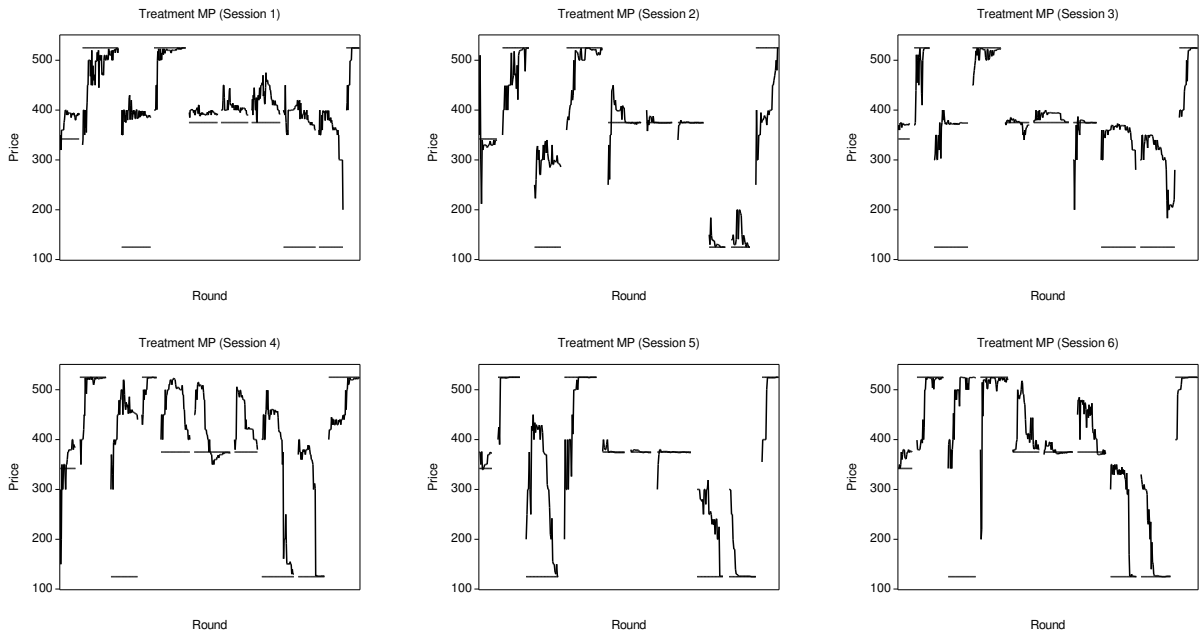


Figure 7: Price development in MP.

Appendix B: Instructions (M12)

Welcome

Dear participant, thank you for taking part in this experiment. It will last about 90 minutes. If you read the following instructions carefully, you can – depending on your decisions – earn some more money in addition to the 3 Euro show-up fee, which you can keep in any case. The entire of money which you earn with your decisions will be added up and paid to you in cash at the end of the experiment. These instructions are solely for your private information.

We will not speak of Euros during the experiment, but rather of ECU (Experimental Currency Units). Your whole income will first be calculated in ECU. At the end of the experiment, the total amount you have earned will be converted to Euro at the following rate:

$$\mathbf{1000\ ECU = 90\ Eurocents.}$$

In order to ensure that the experiment takes place in an optimal setting, we would like to ask you to abide to the following rules during the whole experiment:

- do not communicate with your fellow students!
- do not forget to switch off your mobile phone!
- read the instructions carefully. If something is not well explained or you have any question now or at any time during the experiment, then ask one of the experimenters. Do, however, not ask out loud, raise your hand instead. We will clarify questions privately.
- you may take notes on this instruction sheet if you wish.
- after the experiment, remain seated till we paid you off.
- if you do not obey the rules, the data becomes useless for us. Therefore we will have to exclude you from this experiment and you will not receive any compensation.

Environment

You will learn how the experiment will be conducted later, first we introduce you to the basic situation. You will find control questions at the end of the description that help you to understand it better.

The Asset Market

In this experiment, you have the possibility to trade a financial asset in a stock exchange market for a total of ten rounds. Your final income (in Euro) will be determined by the sum of your per-round payoffs (in ECU). Every round, before the market opens for trade, the actual liquidation value of the asset is determined. It is either 125, 375, or 525 ECU and all values are equally likely to occur. The liquidation value is also common; that is, it is the same for all traders. Once the market closes, you receive the liquidation value for every share of the asset in your portfolio. For example, if the actual liquidation value of the asset is 375 ECU and you have a total of five shares of the asset, then you receive 1.875 ECU.

In the beginning of every round, every trader is endowed with four shares of the asset and 25.000 ECU. Yet, the 25.000 ECU are an interest free loan from a bank; that is, you will have to pay them back at the end of the very same round.

Information Structure

No trader is informed about the actual liquidation value before the trading stops in the first round of the experiment. From round two on, all traders receive some information about the liquidation value before the market opens. In particular, everybody gets to know one of the two values the asset **does not take**. Six traders learn one value the asset does not have and six traders the other. As an example consider a situation when the actual liquidation value is 125. In this case, six traders learn that it is not 375 and six traders learn that it is not 525.

The Computer Trader

In addition to the twelve human traders there is one computer program that participates in the market. **Most importantly, the computer trader does not receive any information with respect to the liquidation value of the asset before the market closes.**

Control Questions

Please answer the following control questions before you continue reading the instructions. Once you have written down all your answers, please raise your hand so that one of the experimenters can check them.

1. How many traders participate in the same market?
2. How many shares of the asset do you receive in the beginning of round 5?
3. Do you have to return your monetary endowment of 25.000 ECU at the end of the very same round?
4. How many traders receive information with respect to the liquidation value in round 1 of the experiment?
5. How many traders receive information with respect to the liquidation value in round 9 of the experiment?
6. Does the computer trader ever receive information with respect to the liquidation value of the asset?
7. Suppose you learn in round 5 that the liquidation value is not 525 ECU. How many traders receive the same information as you?
8. Imagine you learn in the end of round 6 that the liquidation value of the asset is 375.
 - (a) If you bought one share of the asset at a price of 300 ECU, what was be your monetary gain (in ECU) from the trade?
 - (b) If you bought one share of the asset at a price of 450 ECU, what was your monetary gain (in ECU) from the trade?
9. Imagine you learn in the end of round 6 that the liquidation value is 375.
 - (a) If you sold one share of the asset at a price of 300 ECU, what was be your monetary gain (in ECU) from the trade?
 - (b) If you sold one share of the asset at a price of 450 ECU, what was your monetary gain (in ECU) from the trade?

The Experiment

In the next step, we will now go over a brief instruction period so that you get used to the computer interfaces.

The Trading Mechanism

After the determination of the liquidation value, a stock exchange market opens for 300 seconds. On the top of the corresponding computer screen you can identify the current trading round, how long the market remains to be open, and the total amount of ECU you have gained so far. In our example, we are in the second out of ten trading rounds, the market remains to be open for one second, and the trader has earned 1500 ECU so far.

round		2 out of 10		remaining time [sec]: 1		total payoff: Your total payoff so far in ECU is 1500.			
Information: The actual liquidation value of the asset in ECU is not 375				Asset Market Ask Price: <input type="text"/> SUBMIT DELETE					
Inventory: Number of shares: 5 Amount of ECU: 24730		Available: Number of shares: 5 Amount of ECU: 24260		Own Ask if ""		Ask Price 360 400			
Traded Prices in ECU (most recent on top) 524 524 400 524 270 471 524				Own Bid if ""				Bid Price 240 230	
Own Trades: Someone sold to you at price: 270				<input type="text" value="270"/> SUBMIT DELETE					

The screen is further divided into two main parts, the boxes on left hand side and the boxes on the right hand side. The boxes on the left hand side provide different pieces of information whereas the boxes on the right hand side are needed to trade the asset. At first, we introduce the purposes of the boxes on the left hand side.

1. The box on the top is entitled **Information**. As it had been said before, in round one

of the experiment no trader has information about the liquidation value of the asset. From round two on, you will get to know one of the two values the asset does not take. You will see each of the two possible values the asset does not have with the same probability. In our example, the trader learns that the liquidation value is not **375**.

2. The box in the middle gives you an overview about your portfolio and your cash account. In the left part of this box, the **Inventory**, you find how many shares of the asset you possess (in our example you possess 5) and how many ECU you have in your cash account (in our example you have 24.730 ECU). The right part of the box, which is denominated **Available**, has the following aim:

Any sell offer you make is binding. So, if you want to sell one share of the asset, then you must be able to deliver it at any time in the future once the offer is accepted. To insure this, we reduce the number of shares available to you by one whenever you place a new sell offer. As the number of available shares is not allowed to be negative, you can have in total at most as many standing sell offers as shares in the inventory. In our example, the number of available shares is equal to the number of shares in the inventory because the individual does not have any open sell offers.

A similar approach applies to the available ECU. If you want to buy one share of the asset, then you must be able to pay for it in the future. To insure this, we reduce the available ECU by the amount you are willing to pay whenever you place a new buy offer. As the available ECU is not allowed to be negative, the total value of your buy offers cannot exceed the amount of ECU in the inventory. In our example, you see on the right hand side of the screen that the trader has two standing buy offers (the ones marked with a star at the prices of 230 and 240 ECU). Hence, a total of 470 ECU is subtracted from the ECU in the inventory to come up with 24.260 ECU this individual can still use for additional buy offers.

3. The box at the bottom is called **Own Trades**. This box contains a list of your own trades during a round. The most recent trade is on the top of the list. In our example, the individual made so far one trade in this round. S/he bought one share of the asset at a price of 270 ECU.

The boxes on the right hand side of the screen are denoted **Asset Market**. We are going to explain next how the asset is bought and sold using these boxes.

1. If you want to sell one share of the asset, enter the minimum amount of ECU you want to obtain in the field denominated **Ask Price**. You have to confirm your decision by pressing the button **Submit**. Your offer appears immediately in the column **Ask Prices** where all open sell offers are collected. The open sell offers are ordered with the lowest ask price being on the top of the list. You can easily identify your own open sell

offers; they are marked with the symbol *. We want to remind you that any additional sell offer decreases the amount of available shares by one. You are allowed to withdraw a sell offer that has not found a buyer. To do so, you only have to select the sell offer you want to eliminate from the list and to click on the button **Delete**. As a consequence, the amount of available shares rises by one again.

2. If you want to buy one share of the asset, enter the maximum amount of ECU you are willing to pay in the field denominated **Bid Price**. You have to confirm your decision by pressing the button **Submit**. Your offer appears immediately in the column **Bid Prices** where all open buy offers are collected. The open buy offers are ordered with the highest bid price being on the top of the list. You can easily identify your open buy offers; they are marked with the symbol *. We remind you that any additional buy offer decreases the amount of available ECU by the value of your bid. You are allowed to withdraw a buy offer that has not found a seller. To do so, you only have to select the buy offer you want to eliminate from the list and to click on the button **Delete**. As a consequence, the available ECU goes up again.
3. **When and how does a trade take place?** A trade is possible if the **highest bid price** is at least as high as the **lowest ask price**. In this situation, one bidder is willing to pay for the asset at least as much as one seller asks for it. These situations are recognized by the system and trade takes place automatically. The traded price will be equal to the one proposed by the first of the two parties. One simple example clarifies this: Suppose that in a certain situation the lowest ask price is 170 ECU and that the highest bid price is 169 ECU. Then, no trade is possible. If another bidder is willing to pay 177 ECU for the share, the only thing s/he needs to do is to enter a bid of 177 ECU into the system following the procedure of point (2) above. The system recognizes that a trade is possible; that is, the seller receives 170 ECU from the buyer's inventory (because the seller was first in the market) and the buyer receives one share of the asset from the seller's inventory.

An important box on the right hand side is called **Traded Prices**. In it, you find a list of all prices at which a trade took place. The most recent trade price is on the top of the list. In our example, the most recent price is 524 ECU.

Round Summary

Once the market closes, the asset is liquidated. In the corresponding computer screen you find different pieces of information. On the left hand side of the screen you find (a) your particular information with respect to the actual liquidation value of the asset, (b) how many shares and how many ECU you have in your inventory, and (c) a history of your trades.

On the right hand side you find the summary statistics of this round. We inform you about the

round	2 out of 10	remaining time [sec]: 8	total payoff	Your total payoff so far in ECU is 3855
<p>Your information was:</p> <p>Every trader knew the number of informed individuals.</p> <p>You got the following information about the value: 375</p>				
<p>Inventory:</p> <p>Number of shares: 5 Amount of ECU: 24730</p>		<p>This round's results:</p> <p>The actual liquidation value of the asset in ECU is: 525</p> <p>The total liquidation value of your portfolio in ECU is: 2625</p> <p>Therefore, your total amount of ECU is equal to: 27355</p> <p>Your final payoff of ECU after subtracting costs is: 2355</p>		
<p>Own Trades:</p> <p>Someone sold to you at price: 270</p>				
		OK		

actual liquidation value of the asset (in our example it is 525 ECU). This value is multiplied with the number of shares in your inventory to determine the liquidation value of your portfolio (since the trader has five shares in her/his portfolio, the value of the portfolio is in our example equal to 2.625 ECU). Afterwards, we add the amount of ECU in your inventory to it (the sum is in our example equal to 27.355 ECU). Finally, we subtract the 25.000 ECU that have been given to you in the beginning of the round as an interest free loan. As a result, we obtain the final payoff of the round (in our example it is 2.355 ECU). This amount is added to your earlier payoffs. For example, in the top right corner of the screen you see now that your total current payoff is 3.855 ECU (1.500 ECU from the first plus 2.355 ECU from the second round).

Click on the button **OK** to proceed to the next trading round. Note that every trader starts again with 25.000 ECU from an interest free loan and four shares of the asset. At the end of the last round, you will get a short electronic questionnaire regarding your personal background. This data will only be used for statistical purposes.

FINAL QUESTION: Suppose you want to sell the asset at a price of 400. In which field do you have to enter the price, in the ASK or in the BID cell? Once you have written down the answer, please raise your hand.