Price manipulation in an experimental asset market

Citation for published version (APA):

Veiga, H., & Vorsatz, M. (2006). *Price manipulation in an experimental asset market*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 024 <https://doi.org/10.26481/umamet.2006024>

Document status and date:

Published: 01/01/2006

DOI:

[10.26481/umamet.2006024](https://doi.org/10.26481/umamet.2006024)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Helena Veiga, Marc Vorsatz

Price Manipulation in an Experimental Asset
Market

RM/06/024

JEL code: C90, G12, G14

Maastricht research school of **E**conomics
of **T**Echnology and **O**Rganizations

Universiteit Maastricht
Faculty of Economics and Business Administration
P.O. Box 616
NL - 6200 MD Maastricht

phone : ++31 43 388 3830
fax : ++31 43 388 4873

Price Manipulation in an Experimental Asset Market*

Helena Veiga[†] and Marc Vorsatz[‡]

Abstract

We analyze in the laboratory whether an uninformed trader is able to manipulate the price of a financial asset. To do so, we compare the results of two different experimental treatments. In the *Benchmark Treatment*, twelve subjects trade a common value asset that takes either a high or a low value. Information is distributed asymmetrically, only three out of twelve subjects know the actual value of the asset. The *Manipulation Treatment* is identical to the Benchmark Treatment apart from the fact that we introduce a computer program as an additional trader. This manipulation program buys a fixed number of shares in the beginning of a trading period and sells them afterwards again. Our results show that the last contract price is significantly higher in the Manipulation Treatment if the asset takes a low value and that there are no price differences between the two treatments if the value of the asset is high. Moreover, this simple manipulation program is, at least in some instances, profitable.

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

JEL-Classifications: C90, G12, G14.

1 Introduction

According to the efficient market hypothesis (Hayek [11]), the price of an asset summarizes all information available to market participants. One immediate implication of this hypothesis is that markets are immune to price manipulations; that is, traders are not able to influence

*We are very grateful to Martin Strobel for his suggestions and support on the experimental design and we thank all participants of the BEE-seminar at Maastricht University (in particular Arno Riedl) for their helpful comments. Marc Vorsatz acknowledges financial support from the METEOR small scale research grant 01/06. All remaining mistakes are ours.

[†]Department of Statistics, Universidad Carlos III Madrid, Calle Madrid 126, 28903 Getafe, Spain. E-mail: mhveiga@est-econ.uc3m.es

[‡]Corresponding author. Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: m.vorsatz@algec.unimaas.nl

the price of the asset without making a substantial capital loss. It is our objective to provide experimental evidence that the latter need not to be true in general. In fact, some financial markets may in deed be *manipulable* in the sense that it is not too costly for a single trader to distort prices significantly.

Muth [13] formalized the efficient market hypothesis in form of the rational expectations equilibrium. Consequently, the aim of the first experimental studies on asset pricing has been to analyze the conditions under which the price of an asset would, usually after some learning, converge to the equilibrium prediction (see, among others, Forsythe et. al [9], Plott and Sunder [15] and [16], Forsythe and Lundholm [8], and Sunder [17]). The main tenor of this literature is that markets are in the majority of the times able to aggregate disperse information and its capacity increases substantially with the existence of Arrow-Debreu state contingent claims or other certificates (i.e. futures or options) that complete the market.

Misaligned beliefs about the state are probably the most important reason why a market may fail, in a particular instance, to aggregate all available information correctly. For example, in the experimental study of Camerer and Weigelt [5] it was common knowledge to all market participants that either six out of twelve traders or no trader at all knew whether a common value asset had a high or low value. In this situation it happens occasionally that some uninformed traders interpret a buy offer above the expected value as a strong signal that there are six informed traders and that the true value of the asset is high, although, effectively, there is no private information in the market. Since these traders place then even higher bids, the final price becomes unhinged from the equilibrium prediction, the expected value. The authors refer to this type of misaligned beliefs as *information mirages*, because some market participants see information which does not exist. In a subsequent analysis, Nöth et al. [14] ran an experiment in which subjects traded a common value asset with a two-period lifespan. The dividend in the first (second) period was determined by a random choice between two (three) states. Before the opening of the market, one third of the subjects knew the dividend that would be paid at the end of period one, while the remaining two thirds of the subjects did know, on the aggregate, the two states that would not occur in period two. Using this design the authors evidence that *information traps*-some traders have misaligned beliefs and the traders that recognize the misaligned beliefs do not have incentives or resources to adopt their behavior- are the underlying reason why private information may not be aggregated

correctly through prices.

Building upon these results it is our objective to study whether a single trader is able to manipulate the price of the asset by forcing the other market participants to fall into an information trap and, if so, whether this strategy is profitable.¹ To address this question we compare the results of two different experimental series, the *Benchmark* and the *Manipulation Treatment*. In the Benchmark Treatment, we consider a common value asset that takes the values 100 and 220 with equal probability. All traders are endowed with two shares of the asset and an interest free loan that provides enough money to buy all remaining shares. Moreover, it is common knowledge that three out of twelve traders are informed about the asset's fundamental value before the opening of an electronic double auction market. Since there is no aggregate uncertainty, the equilibrium prediction of the rational expectations model is such that the price of the asset should converge in both states to the true value.

Convergence should be easily obtained in the Benchmark Treatment if the fundamental value is 220. Every informed trader has sufficient money to buy all outstanding shares and competition should therefore insure that the price of the asset rises quickly. On the other hand, information traps may occur if the fundamental value is 100 because due to the absence of short selling opportunities there is no competitive force that drives the price of the asset to its actual value. Hence, uninformed traders can only infer the fundamental value by observing that the price of the asset fails to converge to 220. This rather complicated deduction of the state may require some learning.²

Our second treatment, the Manipulation Treatment, aims consequently at analyzing whether one trader can cause an information trap if the fundamental value is 100. To study this question in a systematic way we introduce a computer program, which acts as a manipulator, as an additional uninformed trader into the Benchmark Treatment. The structure of the program is as follows: Once the electronic double auction market is open for 25 seconds, the program increases the actual highest bid price step by step until ten of the 24 outstanding

¹To our best knowledge, so far there exists only one paper that analyzes price manipulations in experimental asset markets. Hanson et al. [10] study a situation in which some traders get an additional payoff depending on the median contract price per round. They find that manipulators place, as expected, higher bids. The accuracy of the market remains however unaffected because the other traders offset this effect by setting lower ask prices.

²Actually, Sunder [17] identified some (not necessary) conditions under which a market may fail to aggregate private information. First, informed traders should have perfect information while the uninformed have none. Second, there should be enough informed traders to achieve information aggregation in at least some states. Third, the inactivity of a trader should not be observable by others. These conditions are met in our experiment.

shares are bought. Afterwards, the shares are sold again by gradually decreasing the actual lowest ask price. The simple idea behind the manipulation program is that buying the asset at prices substantially above the expected value signals to the other market participants that the fundamental value is high. If only one uninformed traders trusts this signal (assigns a very high probability to the associated belief) and places in turn even higher bids, the price of the asset will rise even more and the manipulator can possibly make a trading gain by unwinding the position before the market closes.

Our main findings regard the price of the asset in the market, the earnings of the insiders and uninformed traders, and the profits of the manipulation program. First, and with respect to prices, we find that if the fundamental value is 220, the price of the asset converges in both treatments almost always to the true value. On the other hand, if the fundamental value is 100, learning by repetition is needed in the Benchmark Treatment for the price to converge correctly. But, most importantly, our main hypothesis turns out to be true: If the fundamental value is 100, the last contract price is significantly higher in the Manipulation Treatment (Hypothesis 1). In particular, the results our empirical results reveal that the estimated price difference is in this case larger than 50. Second, and with respect to earnings, it should be noted that insiders make always higher profits because they have an informational advantage and convergence is never instantaneous. Therefore, we analyze the profit difference between the informed and uninformed traders in order to assess whether the manipulation program influences the market statistics. We find, by means of a Wilcoxon rank-sum test, that the payoff difference is the same across the two treatments if the fundamental value is 220 and that the difference is significantly greater in the Manipulation Treatment if the fundamental value is 100 (Hypothesis 2). Finally, our last result shows that although the manipulation program is not profitable on the aggregate, there are many instances where it makes trading gains (Hypothesis 3).

The topic of price manipulation in asset markets has been addressed theoretically from different angles. The seminal paper of Gale and Allen [1] shows in a simple rational expectations model that it is possible for an uniformed trader to make profits by buying and selling an asset if the investors attach a positive probability that the manipulator is informed (this kind of manipulation is thus of the very same nature as the one we implement). Benabou and Laroque [3] show in a model of strategic information transmission that an insider will

not always disclose her/his information about the fundamental value of the asset honestly. Kumar and Seppi [12] analyze a model with multiple risky assets and cash-settled futures and prove that uninformed traders can manipulate the market by sequences of trades in different markets. According to Bognoli and Lipman [2] a manipulator who acts jointly with somebody who has the possibility to alter the value of a firm can make profits by building a position, making a take-over bid, and liquidating the position again. Brunnermeier [4] establishes that if an insider is also informed about future public announcements regarding the value of the asset, then s/he has incentives to build a large position in the beginning of a trading period and to partially unwind it later on for speculative reasons. Finally, Chakraborty and Yilmaz [6] show that if there is uncertainty about the existence of insiders in the market and there is a sufficiently large number of trading periods, then insiders will manipulate in every equilibrium [buy (sell) the asset in at least one trading period although the value is low (high)].

The remainder of the paper is organized as follows: In Section 2, we discuss our experimental design and describe how a typical experimental session has been organized. In Section 3 and 4, we present the competing hypotheses and results, respectively. Finally, we conclude and indicate future research objectives. The instructions of the Manipulation Treatment, the implemented control questions, and a discussion of some additional econometric results are relegated to the Appendices.

2 Experimental Design and Procedures

We consider two different experimental treatments. In the *Benchmark Treatment*, twelve subjects trade for ten rounds a common value asset that takes the values 100 and 220 ECU (Experimental Currency Units) with equal probability. It is common knowledge that (1) no subject has information about the fundamental value before the market closes for trading in the first two rounds of the experiment (the only objective of these two rounds is to familiarize subjects with the computer interfaces, and therefore, we discard the corresponding data) and (2) from round three onwards, three subjects learn the fundamental value before the market opens for trading. We opted for a balanced design meaning that in the last eight rounds, every subject knew the fundamental value twice. Moreover, we decided to fix the fundamental value prior to collecting data. We see in Table 1 that the fundamental value was always 100 in the first and 220 in the second round. Additionally, in the last eight rounds of a session, the

fundamental value equaled five times 100 and three times 220.

Every round, subjects receive two shares of the asset and an interest free loan of 6000 ECU (this amount of money has to be given back at the end of a round). Using these endowments, subjects trade the asset for five minutes in an electronic double auction market by submit-

Session	Round									
	1	2	3	4	5	6	7	8	9	10
1	100	220	100	100	220	220	100	220	100	100
2	100	220	100	220	100	100	220	100	100	220
3	100	220	220	100	100	100	220	100	220	100
4	100	220	220	100	100	220	100	220	100	100
5	100	220	100	220	220	100	100	100	220	100
6	100	220	220	100	220	100	100	100	100	220

Table 1: Fundamental Values

ting 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 at the same time.

To see why this market may be prone to price manipulation suppose that the fundamental value of the asset is 220. Since all informed traders are sufficiently liquid to buy all remaining shares, competition should insure that the price of the asset converges rather quickly to its true value. This is different if the fundamental value of the asset is 100. Informed traders are now willing to sell their stock endowment at any price above 100 and since they cannot use their privileged situation extensively (i.e. we do not allow for short selling), the other market participants can only learn the actual value of the asset by observing that the price does not converge to 220. A market participant who is aware of this learning procedure can thus try to manipulate the market by buying a lot of shares at prices above the expected value in the beginning of a trading period. In this way, s/he signals that s/he is an informed trader and that the fundamental value is high. If at least one uninformed trader believes the signal and reacts by setting even higher bid prices, the manipulator can possibly make profits by selling the shares after some time again.

The objective of the second treatment, the *Manipulation Treatment*, is consequently to analyze whether the price of the asset can be manipulated using the argument presented above. To do so, we introduce a computer program as an additional uninformed trader (i.e. the program is not conditioned on the fundamental value) into the Benchmark Game. From round three on this manipulation program places, once the market is open for trading for 25 seconds, additional random bids in the following way: Ten ECU are added to the current highest bid price to obtain a hypothetical price the manipulator would bid at most.

After deleting all standing buy offers of the program (this is done in order to insure that the manipulation program has at most one standing buy offer), a random bid between the current highest bid price and the hypothetical price calculated before is placed. The computer program increases therefore the current highest bid price and buys eventually one additional share of the asset. After a random delay time between three and seven seconds, this procedure is repeated until the program succeeds to buy ten of the outstanding 24 shares. Finally, the manipulation program sells the shares in a similar way. The only difference to the buying algorithm above stems from the fact that the computer program decreases now the current lowest ask price by eight ECU in order to calculate the lowest hypothetical ask price. The program stops 50 seconds before the end of the trading period allowing us to analyze whether a structural break occurs if the manipulator leaves the market.

We conducted our experimental sessions in the computer laboratory at Maastricht University between February and April 2006. Since all students from the Faculty of Economics and Business Administration have an E-mail account associated to their matriculation number, we promoted the experiment mainly via electronic newsletters and gave students the opportunity to register online for their preferred session. As a result, 144 undergraduate students participated in one of the twelve experimental sessions (six on the Benchmark and six on the Manipulation Treatment). No student took part in more than one session and within a given session, subjects did not know each other.

To perform the experiment we employed the computer software Z-Tree developed by Fischbacher [7]. At the beginning of a session, students met in front of the laboratory. We prepared cards 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 register for a different session. If an insufficient number of students decided to leave, we put additional empty cards into the stack and determined the participating students randomly (students with bad luck received a compensation of three Euros). We also reminded everybody that any kind of unauthorized communication inside the laboratory would lead to an immediate cancellation of the session. Afterwards, students entered the laboratory and took seat in front of the computer with the number corresponding to their card (the computers were placed in such a way that subjects could not see each other). We placed next to each computer an envelope containing the instructions of the

experiment, an official payment receipt, and a set of control questions (see the Appendix A for the instructions and control questions corresponding to the Manipulation Treatment). Once the students finished to read the instructions, we asked them to answer the control questions. The experiment would only start if everybody answered all control questions correctly (in a typical session the whole process took roughly 30 minutes). At the end of a session, we called students one by one to step forward to the control desk for payment. We paid one Euro for every 300 ECU obtained in the experiment. As a result, the average payment in the 90 minutes session was approximately equal to 10 Euros in the Benchmark and 11 Euros in the Manipulation Treatment, with the difference being caused by the losses of the manipulator.

3 Hypotheses

Our main aim is to analyze the influence of the manipulation program on the market statistics. According to our first null hypothesis the manipulator is not successful, or, said differently, the price of the asset at the end of a trading period is, on average and independently of the fundamental value, the same for both treatments. The corresponding alternative hypothesis states instead that the computer program is able to induce misaligned beliefs about the state and cause thereby higher prices. In particular, we ask that if the fundamental value of the asset is equal to 100, the last contract price should be significantly higher in the Manipulation Treatment. Observe that we do not expect in this case that the price converges always to the true value in the Benchmark Game, because it may require some learning for the uninformed traders to be able to infer the state correctly. Moreover, as it has already been explained in the former Section, if the fundamental value is equal to 220, competition should drive the price of the asset to its true value in both treatments.

HYPOTHESIS 1: *If the fundamental value of the asset is 100, the last contract price in the Manipulation Treatment is higher than the one in the Benchmark Treatment. If the fundamental value of the asset is 220, the contract price converges in both treatments to the fundamental value.*

If the price of the asset converges instantaneously to the rational expectations equilibrium, the profits of the insiders are equal to the ones of the uninformed traders. But in experiments it takes usually some time before the price settles down, and therefore, it is rather natural

to expect that insiders can make some additional profits on behalf of the uninformed traders during the convergence process. So, if the manipulation program does not have any influence on the price development, the average payoff difference between informed and uninformed traders should be the same across the two treatments for both fundamental values. This constitutes our second null hypothesis. If, on the other hand, Hypothesis 1 turns out to be correct and the price of asset is higher in the Manipulation Treatment whenever the fundamental value is 100, it follows immediately that insiders are able to sell their shares in this treatment at higher prices. Hence, the relative position of the insiders improves with respect to the Benchmark Treatment leading to a greater payoff difference. Moreover, according to Hypothesis 1 the price of the asset is the same across the two treatments if the fundamental value is 220, and therefore, we expect in this case that the payoff difference is also the same for both treatments. Hypothesis 2 summarizes these ideas.

HYPOTHESIS 2: *If the fundamental value of the asset is 100, the average payoff difference between informed and uninformed traders in Manipulation Treatment is greater than the one in the Benchmark Treatment. If the fundamental value of the asset is 220, the average payoff difference between informed and uninformed traders is the same for both treatments.*

Finally, we analyze the profitability of the computer program. In particular, if the intended manipulation is not successful, the price of the asset in the Manipulation Treatment remains unaffected with respect to the Benchmark Treatment and the computer program will ultimately incur into losses. We test this null hypothesis against the alternative hypothesis that the computer program is able to make some trading gains; that is, it is profitable.

HYPOTHESIS 3: *The manipulation program is profitable.*

4 Results

In this Section, we present the results of our experiments. Figures 1 and 2 below show the sequences of contract prices. In every of the twelve panels, the rational expectations equilibrium for a given round -the expected value if there is no private information and the fundamental value if three subjects are informed- is indicated by a thick straight line. With respect to the Benchmark Treatment we obtain that if the fundamental value of the asset is

220, the contract price converges almost always to the fundamental value (only in session 2 convergence constitutes a slight problem). Hence, we can conclude that the available private information is in this case very well disseminated by the market. On the other hand, if the fundamental value is 100, the results are somehow different. The price of the asset is often much too high and may even converge to the high fundamental value, a clear indication of misaligned beliefs and information traps. Examples include round 5 and 6 of session 2, round 4 of session 3, round 5 and 6 of session 4, and round 6 of session 5. Since we could not find any sign that these results are caused by some kind of manipulative strategy, there is some evidence in favor of our conjecture that it is far more difficult for the uninformed traders to infer the correct state when the fundamental value is low. Nevertheless, we also see that subjects learn how to interpret the market price later on in a session, because the price drops sharply and converges nearly instantaneously to the fundamental value (i.e. round 7,9, and 10 of session 1 and 4; round 5, 8 and 10 of session 3; round 7, 8, and 10 of session 5; and rounds 6 to 9 of session 6). Again, session 2 is the only outlier.

Figure 1: Price Development in the Benchmark Treatment

Looking at the price development in the Manipulation Treatment in Figure 2 we recognize some similarities and differences. If the fundamental value of the asset is 220, the contract price converges almost always, a result we have already encountered before in the Benchmark Treatment (only in session 1 convergence is not assured). On the other hand, if the fundamental value is 100, the price of the asset does not necessarily converge even in later rounds of a session. In fact, the last contract price coincides with the fundamental value only in four

rounds (never in the sessions 4 to 6) and lies 12 times above 200. This is already a very clear indication that the computer program succeeds to manipulate prices.

Figure 2: Price Development in the Manipulation Treatment

Results on Hypothesis 1: Price of the Asset

To quantify our main conjecture, we specify an econometric model with random effects for the last contract price P . The introduction of random effects helps us to account for the heterogeneity of the groups (see Hanson et al. [10] for an application in a similar setting). Let P_{ij} be the last contract price in round i of session j (the six sessions corresponding to the Benchmark Game are indicated by $j = 1, \dots, 6$, whereas the six sessions of the Manipulation Treatment are associated to the indices $j = 7, \dots, 12$). As independent variables we consider a constant, the fundamental value of the asset V , a dummy variable M that takes the value 1 if a session corresponds to the Manipulation Treatment and 0 otherwise, and the variable $M \cdot V$ which measures the additional effect of the fundamental value on the last contract price when a round belongs to the Manipulation Treatment. Additionally, we include in our specification a variable (denoted L) whose aim is to capture the effect that convergence, in the low value case, becomes better in later rounds. This variable takes the value zero if the fundamental value of round i in session j is 220 ($L_{ij} = 0$) and k , $k = 1, \dots, 5$, if the fundamental value in round i of session j is for the k 'th time 100 ($L_{ij} = k$).³ Finally, γ_j denotes the random effect for session j and ε_{ij} is the standard idiosyncratic error for round i of session j . The random

³It turned out that the variable $M \cdot L$, which measures whether the learning pattern differs across the two treatments, is statistically insignificant. Therefore, we dropped it from our final specification.

effects and idiosyncratic errors are assumed to be distributed normally and independently with mean zero. Formally, we have then that

$$P_{ij} = (\beta_0 + \gamma_j) + \beta_1 \cdot M_j + \beta_2 \cdot V_{ij} + \beta_3 \cdot M_j \cdot V_{ij} + \beta_4 \cdot L_{ij} + \varepsilon_{ij}. \quad (1)$$

If $\beta_2 = 1$ and $\beta_i = 0$ for all $i \neq 2$, the estimated price is equal to the fundamental value and the rational expectations model explains the data perfectly. On the other hand, if β_1 and β_3 are statistically insignificant, we do not reject the null hypothesis that the manipulation program does not have any effect on the price. On the other hand, Hypothesis 1 turns out to be correct if only if the estimated price for a given round in the Manipulation Treatment is significantly higher than the corresponding estimated price in the Benchmark Treatment. A particular interesting case occurs if Hypothesis 1 is true and the estimated price in the Manipulation Treatment is the same for both fundamental values, because in this situation the manipulation is successful and the price does not reveal any information about the state.

Variable	Coefficient	Std. Error	t-Statistic	p-Value
C	119.4931	16.6868	7.1609	0.0000
M	97.9678	16.5689	5.8584	0.0000
V	0.4445	0.0868	5.1231	0.0000
$M \cdot V$	-0.4450	0.0987	-4.5068	0.0000
L	-10.5027	2.2377	-4.4737	0.0000
R ² = 0.6121; adjusted R ² = 0.5950				

Table 2: Panel-EGLS Regression Results

Table 2 reports the parameter estimates of the Efficient-GLS estimation procedure (we employed the computer software E-Views 5 for this purpose). It can be observed that all parameters are statistically different from zero (all p -values are equal to zero) and that the estimated model explains more than 59 percent of the price variation. Moreover, the parameter of the dummy variable M takes a very high value, convergence for the low fundamental value becomes better in later rounds of a session (the sign of β_3 is negative), and we do not reject the null hypothesis that $\beta_2 + \beta_3 = 0$ (the p -value of the t-statistic is almost 1). In Table 3 below, where we summarize the estimated prices, we finally see that the manipulation program has a big impact on the price when the fundamental value is 100, but for the high fundamental value there is virtually no difference between the treatments.⁴

⁴We refer to Appendix B for more details of our estimation results.

	Benchmark	Manipulation
1st Time 100	153.44	206.95
5th Time 100	111.43	164.94
Fund. Value 220	217.28	217.45

Table 3: Estimated Prices

Finally, we perform a non-parametric test on the average last contract price per session to provide further evidence in favor of Hypothesis 1. Table 4 already indicates clearly that average prices are significantly higher in the Manipulation Treatment if the fundamental value is 100. Since the p -value of the corresponding one-tailed Wilcoxon rank-sum test is smaller than 0.005 (the test statistic is 22.5), we reject the null hypothesis that prices are the same in the Manipulation Treatment and the Benchmark Treatment in favor of Hypothesis 1. On the other hand, Table 4 only reveals minor differences in the average last contract price between the two treatments if the fundamental value of the asset is 220. Since the p -value of the associated two-tailed Wilcoxon rank-sum test is bigger than 0.1 (the test statistic is 32), the difference is not statistically significant.

<i>Fundamental Value 100</i>			<i>Fundamental Value 220</i>		
Session	Benchmark	Manipulation	Session	Benchmark	Manipulation
1	107	150	1	220	206
2	176	194	2	210	218
3	125	176	3	219	219
4	140	194	4	219	219
5	140	180	5	219	219
6	106	202	6	219	218
Average	132	183	Average	218	217

Table 4: Average Last Contract Price Per Session

Results on Hypothesis 2: Earnings of the Traders

According to our second null hypothesis the payoff difference between the informed and uninformed traders is for both fundamental values the same in the Benchmark and the Manipulation Treatment. To test this hypothesis we calculate for every session the difference in the average payoff over all rounds. In this way we obtain six independent observations, one observation for every session. The values, which are presented in Table 5, reveal that if the fundamental value of the asset is equal to 100, all but one observations in the Manipulation Treatment are greater than 100. Contrarily, only one observation is greater than 100 in the

Benchmark Treatment. If the fundamental value of the asset is 220, the observations for the Manipulation Treatment show only a small variance whereas the ones for the Benchmark Treatment seem to be rather distinct from each other. The payoff difference is particularly high in the second session of the Benchmark Treatment, one of the two sessions without convergence in the high value scenario.

<i>Fundamental Value 100</i>			<i>Fundamental Value 220</i>		
Session	Benchmark	Manipulation	Session	Benchmark	Manipulation
1	63	110	1	7	63
2	74	194	2	141	23
3	64	192	3	68	47
4	137	130	4	45	23
5	98	93	5	77	21
6	23	146	6	21	32
Average	77	145	Average	38	35

Table 5: Difference of Average Payoffs

Table 5 supports Hypothesis 2, but the data is not clear-cut enough to reject the second null hypothesis straight away. Nevertheless, we obtain by means of a Wilcoxon rank-sum test that the payoff difference is significantly greater in the Manipulation Treatment if the fundamental value is 100 (the test-statistic is 25 and the corresponding p -value is smaller than 0.025) and that there is no difference between the two treatments if the fundamental value is 220 (the test-statistic is 34.5 and the corresponding p -value is greater than 0.2). Hence, we reject the second null hypothesis in favor of Hypothesis 2.

Results on Hypothesis 3: Profitability of the Manipulator

So far we have established that the manipulation program is able to distort the dissemination of information in the market. Our final analysis regards the question whether this kind of strategy is also profitable from the manipulator's points of view. We can see in Table 6, where we present the payoff of the manipulation program for every round, that large losses can occur (i.e. in round 10 of session 1, in round 6 and 9 of session 2, and in round 4 of session 3). Nevertheless, it also turns out that the manipulation program is profitable five out of 30 times when the fundamental value is 100 (these rounds are indicated by the bold font style). Since the average loss per share of the asset bought with respect to the risky investment -even in the worst state the asset pays 100 ECU, and therefore, the price paid minus 100 corresponds to a

risky investment- is 16.37%, we do not reject the third null hypothesis that the manipulation program is not profitable.

Session	Round							
	3	4	5	6	7	8	9	10
1	-51	-33	75	25	-28	-8	4	-855
2	-7	48	15	-992	4	-587	-811	42
3	22	-1066	-36	-27	18	-33	23	-991
4	194	34	-286	162	-306	55	-539	-190
5	34	120	11	-151	-42	-560	33	-474
6	84	306	-23	-48	-389	-257	-411	147

Table 6: Absolute Payoffs of the Manipulator

The reasons for the losses of the manipulation program lie presumably in its simple structure. In particular, there are two situations in which the program may fail to liquidate its position and incur into large losses. The first occurs when the price of the asset converges to 220 before the manipulation program finishes to buy ten shares of the asset (i.e. in round 4 of session 3). Since in this case often no trader is any more willing to sell her/his shares, the manipulation program does never switch to the selling algorithm in order to cash in trading gains. Consequently, the program loses the maximal amount of money possible if the fundamental value is 100 although the manipulation attempt was successful. The second occurs because the manipulation program stops fifty seconds before the end of a trading period. We opted for this procedure in order to be able to test whether there is a structural break in the price once the manipulator leaves the market. We could not find any sign of such break, but valuable time, in which a human trader would have tried to liquidate her/his speculative position, is lost. These are obviously not the only shortcomings of the manipulation program (i.e. the program could simultaneously post bid and ask prices instead of using a sequential strategy), yet it has also to be mentioned that the simple program structure is very advantageous: It allows us to associate the price distortions uniquely to the manipulation of beliefs about the fundamental value, simply because the program does not exhibit any other feature that may influence the price development. To have an idea how the two shortcomings mentioned above affect the profitability of the manipulator, we recalculate profits under the assumption that the manipulation program is able to sell all shares it possesses at the end of a trading period at the last highest bid price.

Session	Round							
	3	4	5	6	7	8	9	10
1	-5.3	-3.7	8.0	-5.7	-2.9	-0.9	0.4	-100
2	-0.9	25.3	2.2	<u>2.5</u>	1.7	<u>-9.8</u>	-70.4	4.0
3	3.1	<u>-14.7</u>	-3.1	-2.3	1.9	-2.8	2.0	-90.5
4	19.3	3.6	<u>-10.9</u>	20.3	<u>-6.6</u>	4.8	<u>-9.0</u>	<u>-2.7</u>
5	3.8	11.1	0.9	<u>-9.3</u>	-4.2	10.2	2.9	<u>-24.6</u>
6	7.9	34.6	-2.1	-4.1	<u>-7.9</u>	<u>-7.3</u>	<u>-8.3</u>	14.2

Table 7: Modified Profits (in % of the Risky Investment)

Table 7 presents the modified profits per round. As before, bold numbers indicate that the manipulation program is profitable given a fundamental value of 100. The underlined numbers represent those rounds in which the losses decline thanks to the additional assumption that the manipulation program is able to sell all shares immediately before the market closes. The losses would be lower in 13 rounds and, moreover, the manipulation program would have made trading gains instead of incurring into important losses in round 6 of session 2 and round 8 of session 5. Finally, since the average loss per share bought (or per trade made) would decrease from 16.37% to 6.19%, we conclude that there is a lot of room in order to improve the performance of the manipulator.

5 Final Remarks

In this paper, we studied the manipulability of an experimental asset market by comparing the results of two different experimental treatments. With respect to the baseline treatment we obtained that information is very well disseminated in the double auction market, only in the low value state subjects need learning by repetition in order to infer the fundamental value correctly. In the next step of our analysis we introduced a manipulator (a computer program that buys ten of shares of the asset in the beginning of a trading period and liquidates this position afterwards) as an additional uninformed trader into the baseline treatment and showed that this manipulation programs distorts prices seriously. In particular, if the fundamental value of the common value asset is low, the last contract price is now substantially higher as the one corresponding to the baseline treatment. As a consequence, insiders have also the possibility to make some additional profits on behalf of the uninformed traders.

Our results raise some serious questions with respect to the organization of markets. First,

since we consider a common value asset, efficiency is not an issue. Prices reflect only transfers between individuals and under the assumption of risk neutrality all allocations are efficient. But if price manipulations are also possible for assets whose value is distinct for different types of traders, the asset may finally end up in the hands of the wrong type of traders. The efficiency of the double auction market for private value assets in the presence of a manipulator is therefore one question that remains to be investigated in the future. Second, it is far from clear to what extent our results can be transferred to other market forms. Possible ways how to analyze the dimension of our findings are to allow subjects to short-sell the asset or to complete the market with the help of Arrow-Debreu state contingent claims, futures, options, etc. It is our first guess that these measures can reduce the manipulability of the market, but it remains doubtful whether they eliminate the problem completely. Finally, we found only in one round of a pilot session corresponding to the baseline treatment a subject that tried (successfully) to manipulate the market in the way the computer program does. One possible reason for this low level of manipulation attempts is that this strategy cannot be supported in equilibrium. Although this explanation coincides with our finding that the manipulator is not profitable, we rather believe that the manipulation strategy is too complicated for student subjects. Therefore, it would certainly be interesting to see how a human manipulator performs in a market with experienced traders.

References

- [1] F. Allen and D. Gale, *Stock Price Manipulation*, Review of Financial Studies **5** (1992), 503–529.
- [2] M. Bagnoli and B. Lipman, *Stock Price Manipulation Through Takeover Bids Mirages*, RAND Journal of Economics **27** (1996), 124–147.
- [3] R. Benabou and G. Laroque, *Using Privileged Information to Manipulate Markets: Insiders, Gurus and Credibility*, Quarterly Journal of Economics **107** (1992), 921–958.
- [4] M. Brunnermeier, *Buy on Rumors-Sell on News: A Manipulative Strategy*, Mimeo (2000).
- [5] C. Camerer and K. Weigelt, *Information Mirages in Experimental Asset Markets*, Journal of Business **64** (1991), 463–493.

- [6] A. Chakraborty and B. Yilmaz, *Informed Manipulation*, Journal of Economic Theory **114** (2004), 132–152.
- [7] U. Fischbacher, *Z-tree Tutorial Version 2.1*, Zurich University (2002).
- [8] R. Forsythe and R. Lundholm, *Information Aggregation in an Experimental Market*, Econometrica **58** (1990), 309–347.
- [9] R. Forsythe, T. Palfrey, and C. Plott, *Asset Valuation in an Experimental Market*, Econometrica **50** (1982), 537–567.
- [10] R. Hanson, R. Oprea, and D. Porter, *Information Aggregation and Manipulation in an Experimental Market*, forthcoming Journal of Economic Behavior and Organization.
- [11] F. Hayek, *The Use of Knowledge in Society*, American Economic Review **35** (1945), 519–530.
- [12] D. Kumar and D. Seppi, *Futures Manipulation with Cash Settlement*, Journal of Finance **47** (1992), 1485–1502.
- [13] J. Muth, *Rational Expectations and the Theory of Price Movements*, Econometrica **29** (1961), 315–335.
- [14] M. Nöth, C. Camerer, C. Plott, and M. Weber, *Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs*, Mimeo (1999).
- [15] C. Plott and S. Sunder, *Efficiency of Experimental Security Markets with Insider Information: An Application to Rational-Expectations Models*, Journal of Political Economy **90** (1982), 663–698.
- [16] ———, *Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets*, Econometrica **56** (1988), 1085–1118.
- [17] S. Sunder, *Market for Information: Experimental Evidence*, Econometrica **60** (1992), 667–695.

Appendix A: Instructions and Control Questions

Welcome

Thank you for taking part in this experiment. The purpose of this session is to study how people behave in stock exchange markets. You will be paid for your participation in the way as explained at the end of this document. If you have any doubts, feel free to raise your hand and your question will be answered to you in private. From now on until the end of the session, unauthorized communication of any nature with any other participant is prohibited. The experiment will be conducted through computers and all interactions between you will take place through them.

This session consists of a total of ten rounds. The first two rounds differ from the last eight rounds in the way as explained below. We will now go over a brief instruction period so that you get used to the computer interfaces.

The Asset Market

Every round, you have the possibility to trade a financial asset. Before the market opens for trade, the liquidation value of the asset is determined. It is either 100 or 220 ECU (Experimental Currency Units) and both values are equally likely to occur. Observe that the liquidation value 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 220 ECU and you have a total of five shares in your portfolio, then you receive 1100 ECU. In the beginning of a round, every trader is endowed with two shares of the asset and 6000 ECU. Yet, you should note that the 6000 ECU are an interest free loan from a bank; that is, you will have to pay them back at the end of the round.

Information Structure

No trader is informed about the actual liquidation value before the trading stops in the first two rounds of the experiment, whereas three traders (which are determined randomly) are informed about it before the trading starts in the last eight rounds of this session.

The Computer Trader

In total, there are thirteen traders participating in the market. Twelve traders are human and one trader is a computer program. The objective of the computer trader is to maximize its payoff. **Most importantly, the computer trader is never informed about the actual liquidation value of the asset before the market closes.**

The Trading Mechanism

After the determination of the liquidation value, a market opens for 300 seconds. On the top of the corresponding computer screen you can identify the number of 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 first out of ten trading rounds, the market remains to be open for 107 seconds and the trader has not made any profit so far.

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.

round
1 out of 10
remaining time [sec]: 107
total payoff
Your total payoff so far in ECU is 0.

Information:

The number of informed individuals is: 0
Every trader knows the number of informed individuals.

The actual value of the stock in ECU is: No information

Inventory:
Number of shares: 2
Amount of ECU: 6009
Available:
Number of shares: 2
Amount of ECU: 5598

Own Trades:
You bought at price: 155
Someone bought from you at price: 164

Asset Market

Ask Price:

Own Ask # ""	Ask Price
	183
	185
	189
	190

Traded Prices in ECU (most recent on top)

170
165
164
155

Own Bid # ""	Bid Price
	179
	160
	155
*	149
*	132
*	130

Bid Price:

1. The box on the top is entitled **Information**. In our example, you find the information that no trader knows the actual liquidation value of the asset and that all traders are aware of this fact. Remember that this is the case because the screen corresponds to round 1. From round 3 on, three traders are informed about the actual liquidation value. If you are randomly selected to know the liquidation value, then the corresponding value (**220** or **100**) appears in the last row of this box. If you are not selected, then you will get the message **No Information**.
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 2) and how many ECU you have in your cash account (in our example you have 6009 ECU). The right part of the box, which is entitled **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 enter

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 amount of 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 amount of available ECU by the amount you are willing to pay. As the number of 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 three open buy offers (the ones marked with a star at the prices of 130, 132, and 149 ECU). Hence, a total of 411 ECU is subtracted from the 6009 ECU in the inventory to come up with the remaining amount of 5598 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 two trades in this round. First, s/he bought a share at 155 ECU and afterwards s/he sold one share at 164 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 unit 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, because 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.
2. If you want to buy one unit 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, because they are marked with the symbol *. We remind you that any additional bid 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 amount of available ECU goes up again.
3. **When and how does a trade take place?** A trade is possible if the **lowest ask price** of all open sell offers is lower than the **highest bid price** of all open buy offers. In this situation, one bidder is willing to pay more for the asset than the seller asks for it. These situations are recognized by the system and trade takes place automatically.

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 170 ECU.

Once the market closes, the asset is liquidated. In the corresponding computer screen you find different pieces of information. In the left hand side of the screen you find (1) how many traders have been informed beforehand about the liquidation value of the asset and your particular information, (2) your inventory, and (3) a history of your trades.

22

On the right hand side you find the summary statistics of this round. We inform you about the actual liquidation value of the asset (in our example it is 220 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 two shares in her/his portfolio, the value of the portfolio in our example is equal to 440 ECU). Afterwards, we add the amount of ECU in your inventory to it (the sum is in our example equal to 6449 ECU). Finally, we subtract the 6000 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 449 ECU).

Click on the button **OK** to proceed to the next trading round. Note that every trader starts again with 6000 ECU from an interest free loan and two 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.

Payment

The points you accumulate during the course of the session will determine your payment. The exchange rate ECU/Euros is such that every 300 ECU in the experiment are equal to 1 Euro; that is, if your total payoff after 10 rounds is equal to 3000 ECU, then you get 10 Euros for your participation. *Please, do not fill in your final payment into the official receipt. This will be done by us!!!* Leave the instructions and the pen at your desk. Once you are paid, you may leave.

Control Questions

1. This question corresponds to the first figure of the instructions. Suppose that the market is still open for trading. Suppose that you bought two shares at a price of 150 ECU and sold one share at a price of 160 ECU. Moreover, you have two open sell offers (at a price of 170 ECU) and two open buy offers (at a price of 160 ECU) that have not been accepted so far by any other trader.
 - (a) How many ECU do you have in your Inventory?
 - (b) How many shares do you have in your Inventory?
 - (c) How many ECU do you have available for buying additional shares?
 - (d) How many shares do have available for additional sell offers?
2. How many traders (including the computer trader) know the liquidation value of the asset while the market is open for trading?
3. This question corresponds to the second figure of the instructions. Suppose that market has closed. During the trading activity you bought two shares (at a price of 180 ECU) and sold two shares (one at a price of 170 ECU and another at a price of 160 ECU). Moreover, you learn that the actual value of the stock is 100 ECU.
 - (a) What is the total liquidation value of your portfolio in this round?
 - (b) What is your total amount of ECU in this round?
 - (c) What is your total payoff in ECU after subtracting costs in this round?
4. Consider a different situation in which the liquidation value of the asset is either 200 or 300 ECU (the probability of the first value is 0.4 and the probability of the second value is 0.6). What is the expected liquidation value of the asset?

Appendix C: Econometric Analysis

The estimated model incorporates random effects such that

$$P_{ij} = (\beta_0 + \gamma_j) + \beta_1 \cdot M_j + \beta_2 \cdot V_{ij} + \beta_3 \cdot M_j \cdot V_{ij} + \beta_4 \cdot L_{ij} + \varepsilon_{ij}.$$

If we calculate $E(P_{ij})$ and $E(P_{ij}|\gamma_j)$ we obtain that

$$E(P_{ij}) = \beta_0 + \beta_1 \cdot M_j + \beta_2 \cdot V_{ij} + \beta_3 \cdot M_j \cdot V_{ij} + \beta_4 \cdot L_{ij}.$$

and

$$E(P_{ij}|\gamma_j) = (\beta_0 + \gamma_j) + \beta_1 \cdot M_j + \beta_2 \cdot V_{ij} + \beta_3 \cdot M_j \cdot V_{ij} + \beta_4 \cdot L_{ij}$$

The difference $E(P_{ij}|\gamma_j) - E(P_{ij}) = \gamma_j$ is called the effect of session j . In Table 8, which reports the actual effects of our estimation procedure, we observe that the extreme values correspond to the first session of the Manipulation Treatment and to the second session of the Benchmark Treatment. Looking at Figures 1 and 2 we see that these are the only sessions in which the convergence in the high value scenario fails. Consequently, these are the sessions with most noise, which is correctly captured with the help of the random effects. Moreover, the average actual effect is -0.058, and therefore, it is very close to the expected value of 0.

Session	Benchmark	Manipulation
1	-8.07	-13.65
2	12.83	3.15
3	-1.98	-2.47
4	2.84	3.48
5	2.84	3.48
6	-8.47	6.00

Table 8: Cross-Section Random Effects

The next point regards the question whether the estimated model is correctly specified in the sense that the residuals are normal i.i.d. with mean zero. If we look at the autocorrelation and the partial autocorrelation functions of the residuals in Figure 4, we observe that they present a behavior similar to the one of a white noise process (all autocorrelations have an absolute value smaller than 0.2 and are individually insignificant). Moreover, since the null hypothesis of the Ljung-Box test that the autocorrelations are jointly insignificant ($\rho_1 = \rho_2 = \dots = \rho_r = 0$) is not rejected, we conclude that the residuals do not exhibit any kind of autocorrelation.

Figure 3: Residuals

Figure 4: Autocorrelation (a) and Partial Autocorrelation Functions (b)

Finally, it remains to be analyzed if the residuals are distributed normally with mean zero. Since we do not reject the null hypothesis of the Jarque-Bera test that the residuals follow a normal distribution and the mean of the residuals is very close to zero (see Figure 5), we can finally conclude that the model is correctly specified.

Figure 5: Histogram of the Standardized Residuals