Information Aggregation and Manipulation in an Experimental Market

Robin Hanson, Ryan Oprea, David Porter*
Interdisciplinary Center for Economic Science
George Mason University

September 21, 2004

Abstract

Prediction markets are increasingly being considered as methods for gathering, summarizing and aggregating diffuse information by governments and businesses alike. Critics worry that these markets are susceptible to price manipulation by agents who wish to distort decision making. We study the effect of manipulators on an experimental market. We find that manipulators are unable to distort price accuracy. Subjects without manipulation incentives compensate for the bias in offers from manipulators by setting a different threshold at which they are willing to accept trades.

1 Introduction

Rational expectations theory predicts that, in equilibrium, asset prices will reflect all of the information held by market participants. This theorized information aggregation property of prices has lead economists to become increasingly interested in using securities markets to predict future events. Proponents claim that these "prediction markets" can aid decision makers by gathering diffuse information and beliefs into easy to interpret price statistics. Critics, however, worry that these markets are vulnerable to price manipulation. In particular, if policy makers use information from prediction markets to inform decisions, some individuals (manipulators) may be willing to take capital losses in these markets in order to distort the informativeness of prices and thereby indirectly control policy.

It was just this type of fear that lead to the shutdown of the largest planned prediction market to date. In the summer of 2003, the Defense Advanced Research Projects Agency planned to open an extensive prediction market called the Policy Analysis Market (PAM) to help decision makers predict geopolitical events. PAM was created largely to collect information across the multiple intelligence agencies in the US federal government, though its markets were open to the public. The Senate decided to shut the program down just before it was to open. PAM was derailed, in part, because of fears that terrorists would attempt to manipulate prices in order to distort the information reaching decision makers (e.g. [14], [19])

Does price manipulation really pose a threat to prediction markets? Identification of manipulation in the field is difficult and as a result empirical evidence on the phenomena is scarce and tenuous. Because incentives and information can be carefully controlled in experimental markets, laboratory study of manipulation is a promising alternative to field study for evaluating this question. This paper reports the results of experiments designed to study the effect of manipulation on prices in a futures market and suggests future research to examine the robustness of securities markets to manipulation.

Prediction markets are based on the information aggregation property of prices first noticed by Hayek [6] and formalized by Muth [13]. The best known and longest running field experiment using futures markets for the purposes of prediction is the Iowa Electronic Market. The IEM, under special permission from the Commodity Futures Trading Commission, issues stock in political candidates over a variety of elections and allows the public to trade in this stock. Forsythe, Nelson, Neumann and Wright [4] find that these markets do better than polls in predicting the outcomes of these elections.

Empirical results on price manipulation of futures markets are mixed. Hansen, Schmidt and Strobel [7] report successful attempts at manipulating prices in the IEM. However a field experiment reported by Camerer [2] and a historical study of price manipulation due to Strumpf and Rhode [17] find that efforts to manipulate price are unsuccessful. Surveying studies of price manipulation, Wolfers and Zitzewitz [18] conclude that efforts to manipulate do not "have discernible effect on prices, except during a short transition phase."

Theoretical work on price manipulation explains the logic of price manipulation but is generally not tailored to the incentives over policy outcomes that motivate concern about manipulation in prediction markets. Allen and Gale [1] give an example of a simple signalling model in which an agent can trade large volumes early on in order to mislead investors and then sell at inflated prices. Kumar and Seppi [10] provide a microstructure model in which manipulators distort futures prices in order to earn excess profits in a later spot market. Chakrabortya and Ylmaz [3] analyze a dynamic microstructure model showing that in all equilibria informed traders will attempt to manipulate price in order to make later gains. These models do not rely on exogenous preferences over prices; agents manipulate in order to gain in the market. Hanson and Oprea [8] study a microstructure model motivated by prediction markets in which manipulation ends up causing prediction market prices to be *more* accurate due to the liquidity they provide.

This paper reports the first laboratory work on price manipulation in asset markets. However, Plott and Sunder [15] conducted pioneering experimental research on information aggregation in asset markets which was extended by Plott and Sunder [16] and Forsythe and Lundholm [5]. In the classic experimental design, subjects trade a common value asset in a double auction. Subjects are differentially informed on what value the asset will not take so that, in aggregate, though not individually, participants have information sufficient to determine the true value.

Our design follows this line of research both in terms of the information structure and the institution studied. In a control treatment, we replicate a variation of the basic design studied in Plott and Sunder [16]. In our main treatment a subset of subjects, whom we call manipulators, are given a payoff opportunity outside the common value of the asset, which is based on the median contract price in the market. Thus, the manipulators have an incentive to affect contract prices that is not related to their own private information and this incentive is common knowledge. We find that although there are clear efforts to manipulate price, these efforts fail to reduce the accuracy of prices as estimates of true asset values. Non-manipulators, knowing that there are

individuals with manipulation incentives in the market, accept offers at systematically lower levels, which counteracts efforts to manipulate price.

2   Experimental Design

Our experimental design consists of two treatments. The Replication Treatment (RT), follows Plott and Sunder's [16] experimental design. Twelve subjects began each of 8 periods with an endowment of 2 shares of a common value asset and 200 units of experimental currency. Every 160 units of experimental currency could be redeemed for $1 US after the experiment. The value of the asset is unknown to subjects, but it is common knowledge that it will take on a value in experimental currency in $\{0, 40, 100\}$ with equal probability. Values were randomly drawn prior to sessions and an identical ordering was maintained across sessions. This design is summarized in Table 1.

All subjects were endowed each period with a clue about the asset value; each subject was told one of the three values that the asset would certainly not take. For example, if the realized asset value was 100, then half of all subjects were told that the value was not 0 and half were told that the value was not 40. Therefore, in each period the experimental cohort in aggregate possessed information that revealed the value of the asset though no individual subject had this much information.

In each of eight periods, subjects were given the opportunity to trade shares of the asset prior to the realization of its value in a standard electronic double auction institution with an open book. Subjects submitted bids or asks for purchase or sale of shares through computer terminals. Contracts were formed whenever a subject accepted a standing offer, and asset(s) and currency were exchanged immediately. Subjects were also allowed to submit offers for multiple unit trades. Each period lasted exactly 5 minutes.

Subjects were randomly provided with trader identification numbers which were used to assign clues. All subjects assigned as trader 1, for example, had the same sequence of clues across sessions. Subjects were also divided into two subsets. Subjects 1,3,4,7,9, and 12 were assigned to *Set 1* and subjects 2, 5, 6, 8, 10 and 11 were assigned to *Set 2*. In each period an equal number of subjects in each set had each of the two possible clues. For example, if the true state was 100, three subjects in each set were told that the state was not 0 and the other three subjects in that set were told that it was not 40. This means that there were an even number of subjects with each type of information in *Set 1* and *Set 2*.

The only difference between our Replication Treatment (RT) and Manipulation Treatment (MT) was that in the MT treatment, it was common knowledge that *Set* 1 participants would be given an additional payoff based on the median price over all contracts for the trading period.   In particular, in addition to normal earnings and losses from trades, these manipulation subjects were given a bonus based on the median price. Manipulation subjects received two units of experimental currency for every point of the median price at the end of a period. Subjects in Set 2 under MT were given an additional 100 units of earnings each period. All subjects were shown the median price throughout the period in the MT treatment. Hence, under MT, all subjects were aware that 1/2 of all subjects had the incentive to manipulate and whether or not they had this incentive but were unaware of *which* other subjects had manipulation incentives.

Subject	Subset	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6	Period 7	Period 8
1	Set 1	100	100	100	0	40	100	100	40
2	Set 2	100	40	40	0	40	0	100	0
3	Set 1	0	40	40	0	0	0	100	40
4	Set 1	100	100	100	0	0	0	0	0
5	Set 2	100	40	40	0	40	0	100	0
6	Set 2	0	40	100	0	40	100	100	0
7	Set 1	0	40	100	40	40	100	0	40
8	Set 2	100	100	100	40	0	0	0	40
9	Set 1	100	100	40	40	40	100	0	00
10	Set 2	0	100	100	40	0	100	0	40
11	Set 2	0	100	40	40	0	100	0	40
12	Set 1	0	40	40	40	0	0	100	0
True State		40	0	0	100	100	40	40	100

Table 1: Experimental Design: Listed are the subset each subject was in and what value that subject was told would not be realized in each of the eight periods.

3 Procedures

Experiments were conducted in the *Interdisciplinary Center for Economic Science* laboratory at George Mason University. We recruited 96 undergraduate subjects through an online recruiting system. Eight sessions were conducted, with four sessions in each treatment. Twelve subjects participated in each session. Recruited subjects had no experience in the experiment and were not allowed to participate in more than one of the experiments reported here. However, all recruited subjects had experience with the double auction institution. Sessions were held between September 2003 and April 2004 and each lasted approximately 2 hours.

In a typical session, subjects waited in a reception room until a sufficient number had arrived. Twelve subjects were taken into a laboratory room and seated at computers behind partitions so that they could not see each other's screens. A strict no-talking rule was enforced as soon as subjects entered the laboratory. At each session, subjects were given identical computerized instructions which gave detailed information on the environment, rules of the institution and gave subjects experience with the interface. There was a single core of instructions common to both treatments with additional instructions added in the MT treatment.[1]

Upon completing instructions, subjects were given a self grading computerized quiz which, depending on the treatment, consisted of between 6 and 10 questions. When a subject submitted an incorrect answer, the computerized quiz automatically told the subject the correct answer and explained the intuition behind it. Subjects were allowed to ask questions for clarification throughout the instructions and quiz. Typically instructions and quizzes lasted a total of 35 minutes. Finally, following [5], subjects were given paid practice guessing draws after observing their clues. Subjects were given clues drawn in the same way and with the same interpretation as those described in the previous section. They were then asked to guess which state they thought was most likely to be the true state and paid for correct guesses. Subjects were given 10 rounds of practice, during which they were paid $.50 for correct guesses and penalized $.25 for incorrect guesses.

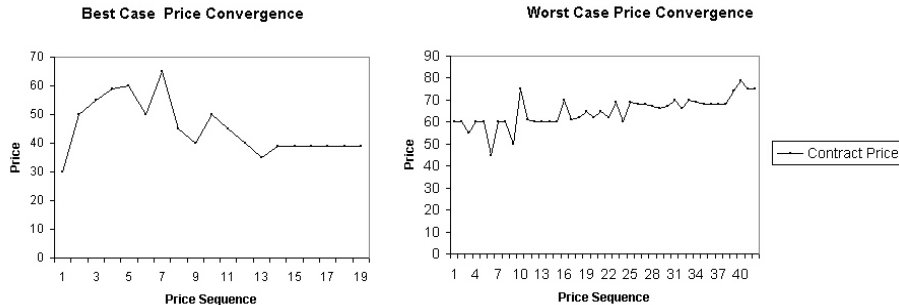[1]The instructions for the experiment can be found at http://ices3.gmu.edu/MInfoAgg/

Figure 1: Best and worst cases of price convergence. In the best case example, the true state is 40 and the last several prices are at 40. In the worst case example, the true value is 0 and the prices appear to be converging to a value around 80.

Subjects were paid $5 for showing up to the experiment on time and an average payout of $32.57 in the RT treatment and $41.82 in the MT treatment. This difference in average earnings was due to the fact that Set 2 subjects earned additional money for manipulation and that Set 1 subjects earned a flat bonus of $1 per period.

4 Results

Result 1 studies whether the subjects who were given an incentive to manipulate actually attempt to trade at elevated prices. Result 2 studies the degree of aggregation in end of period prices and the effect of manipulation on (a) the accuracy of price as a statistic for fundamental value and (b) the linear dependence of price on state. Result 3 presents evidence that non-manipulators in the MT treatment counteract the upward price bias generated by manipulators by, on average, accepting contracts with lower prices than subjects in the RT treatment.[2]

Result 1 *Manipulators submit higher bids than non-manipulators.*

Evidence. We formally establish the result by comparing the distribution of average price offers by Set 1 subjects and Set 2 subjects using a non-parametric test. Under the hypothesis that the incentives to manipulate induce inflated price offers (i.e. try to manipulate), Set 1 subjects will offer higher bids than Set 2 subjects in the MT treatment but not in the RT treatment. We conduct a pair of medians tests comparing the average offer price by each of the 24 Set 1 subjects and the 24 Set 2 subjects in the RT treatment and MT treatment. The difference is significant in MT [3] but not in RT[4]. ∎

Result 1 suggests that bid behavior is effected by the presence of the manipulation incentive. What remains to be determined is whether this behavior has a significant effect on market prices.

---

[2] Data for these experments can be found at http://ices3.gmu.edu/MInfoAggData/

[3] Pearson $\chi^2(1) = 5.333$ ($p = 0.021$) and Fisher's exact test $p = 0.021$ (one-tailed).

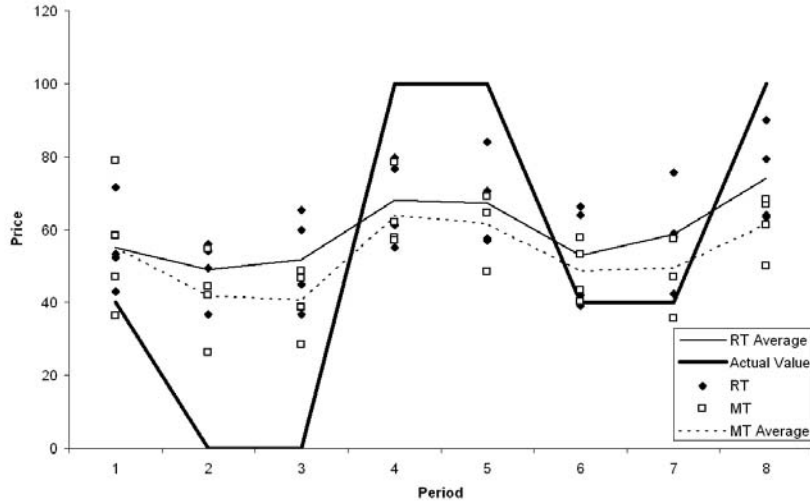[4] Pearson $\chi^2(1) = 0$ ($p = 1$) and Fisher's exact test $p = 0.613$ (one-tailed).

Figure 2: End of period prices by period.

We will formally show in the next result that, independent of the treatment, overall weak aggregation generally occurs though strong aggregation is rare. However before reporting Result 2, it is important to point out that from round to round, in this environment, aggregation is sometimes very accurate and sometimes very inaccurate. Figure 1 shows the time series of price in two different periods in the RT treatment. In the figure to the left, the true state is 40 and prices converge to this value at the end of the period. Aggregation is perfect. The figure to the right shows the opposite extreme. In this round, the true value was 0 and yet, prices appear to be converging to a value of around 80.

Result 2 *The presence of manipulators does not harm the information aggregation properties of trade; manipulation has no significant effect on the accuracy of price or the dependence of price on state.*

Evidence. We look at two measures of the correlation between price and fundamental value. The first is the linear model between the two which we call the *relationship*. The second is the squared deviation of price from fundamental value which we call the *accuracy*. Because aggregation is a process of learning, we would not necessarily expect to see aggregation or a noticeable effect of manipulation on aggregation in statistics that summarize prices over the course of an entire period of trading. Therefore in the statistical procedures reported, we focus on the properties of prices at which trade occurs at the end of each period.

   Figure 2 shows the terminal value of assets by period, the end of period prices for each session, and the average of final trading prices over sessions in RT and MT treatments in each of these periods. In neither treatment do prices converge reliably to fundamental value. But prices do qualitatively react to shifts in asset value across periods. This qualitative tracking is the price *relationship*. We statistically test whether this relationship is damaged by manipulation and whether there is an upward bias in price. That is, we are looking for differences in intercept and slope of a regression line tying price to state across treatments. We test the hypothesis that

6

Covariate	Value	t	p
Intercept	48.5879	7.897	$< .0001$
manipulation	-7.6514	-0.9770	0.3663
state	0.2008	6.2813	$< .0001$
manipulation $\times$ *state*	0.0204	0.3502	0.7275

Table 2: Results of model 2 on the relationship between end of period prices and the true state.

manipulation has an effect using the following linear mixed effects model:

$$
\begin{aligned}
price_{ij} &= \alpha + (\beta + \beta_j)manipulation_i + (\gamma + \gamma_j)state_i \\
&\quad + (\tau + \tau_j)manipulation_i \times state_i + \alpha_j + \varepsilon_{ij}
\end{aligned} \tag{1}
$$

Price is the average of the last 3 contract prices in period i of session j. [5] *Manipulation* is an indicator variable taking a value of 0 in RT observations and 1 in MT observations. *State* is the value that assets actually take at the end of period i; . Parameters $\beta_j, \gamma_j, \tau_j$ α_j are random effects for session j assumed to be distributed independently and normally with a zero mean. Finally ε_{ijt} is an error term assumed to be distributed $N(0, \sigma_\varepsilon^2)$. In perfect aggregation, the intercept α would be zero and γ would be 1. If manipulation has no effect on these measures, we would expect τ and β to be insignificant.

Results are shown in Table 2. Aggregation is far from perfect; α is high and γ is a quarter of what it would be with perfect aggregation. Thus, we can reject the hypothesis of perfect aggregation, and we can accept the hypothesis that there is a positive and significant correlation between the state and price. In addition, we cannot reject the null hypothesis that manipulation has no effect on aggregation. Because β is insignificant, we can infer that manipulation does not succeed in pushing prices up at the end of periods. Because τ is insignificant, there is no difference in the degree to which price reflects the true state.

Figure 3 shows the average price error (the squared deviation of price from fundamental value) by period under RT and MT treatments. This is the *accuracy* measurement. The differences appear negligible, an issue which can be statistically studied. To accomplish this we estimate another linear mixed effects models with random terms associated with sessions:

$$
(price_{ij} - state_{ij})^2 = \alpha + (\beta + \beta_j)manipulation_j + u_j + \varepsilon_{ij} \tag{2}
$$

As above, $manipulation_j$ is an indicator variable taking a value of 1 if the observation is under MT and 0 otherwise. The terms $price_{ij}$ and $state_{ij}$ are the average of the last 3 prices in period i of session j. β_j and u_j are random efects assumed to be normally distributed with mean zero and independent variance. Under the hypothesis that prices were less accurate in the presence of manipulation than in the absence of it, we would expect a significant and positively signed estimate for β. The results suggest that this is not the case. β is actually estimated negative ($\beta = -88.95$) but is not significant ($p = 0.739$). That is, the presence of manipulators has a tendency to *improve* accuracy but the effect is not significant. ■

Figure 2, showing average overall end of period prices by period shows just how spectacularly efforts at manipulation fail at pushing up price. Prices are, in fact, on average higher under

---

[5] We use an average because it is less susceptible to the influence of individual subjects than the final price. Regardless, all of the results reported here also hold with the final price.
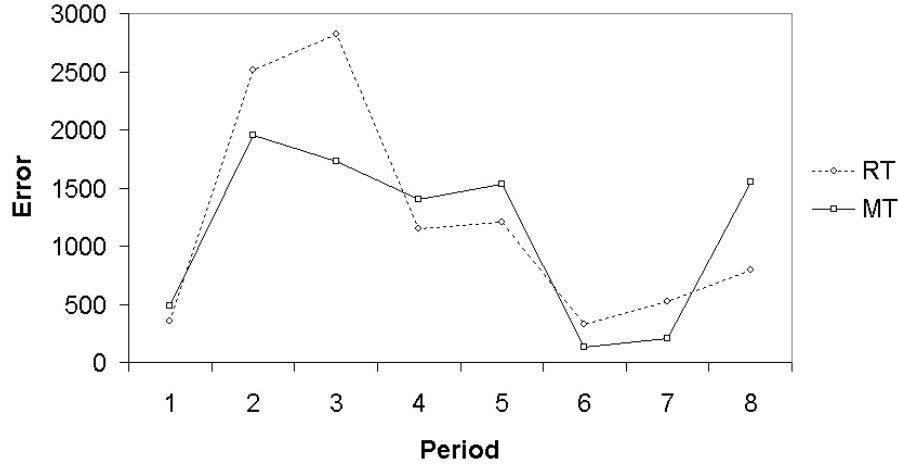
Figure 3: Average price error by treatment.

non-manipulation than manipulation. Although this difference in price levels is not significant it provides a clue as to why manipulation fails. Figure 4 shows average prices at which contracts were accepted by non-manipulators in the MT treatment and subjects in general in the RT treatment. Acceptance prices are lower under MT than RT. Subjects without manipulation incentives appear to be wary of accepting high offers because of the incentives to manipulate are present. They appear to set a lower upper threshold at which they will accept offers than subjects in the RT treatment. This observation is reported as Result 3.

Result 3 *Non manipulators tend to accept contracts at lower prices in the MT treatment than in the RT treatment.*

Evidence. We exclude manipulators from the sample studied to look for thresholds of contract acceptance when incentives to manipulate are absent. To balance samples, and control for history of information, we compare Set 2 subjects under RT with Set 2 subjects under MT. Figure 4 shows the average price at which Set 2 subjects accept offers in each period. Acceptance prices are higher in RT than in MT.

Quantitative support for this result is provided using a medians test on the average price at which each subject accepted contracts. The median average contract price accepted by Set 2 subjects under RT was 60.9 and 55 under MT and, this difference is significant[6]. Thus, in the MT treatment, subjects without incentives to manipulate on average accept prices at a lower level than in RT. ∎

---

[6]Pearson $\chi^2(1) = 7.62.82$ ($p = 0.006$) and a Fisher's exact $p = 0.009$ (two-tailed).

8
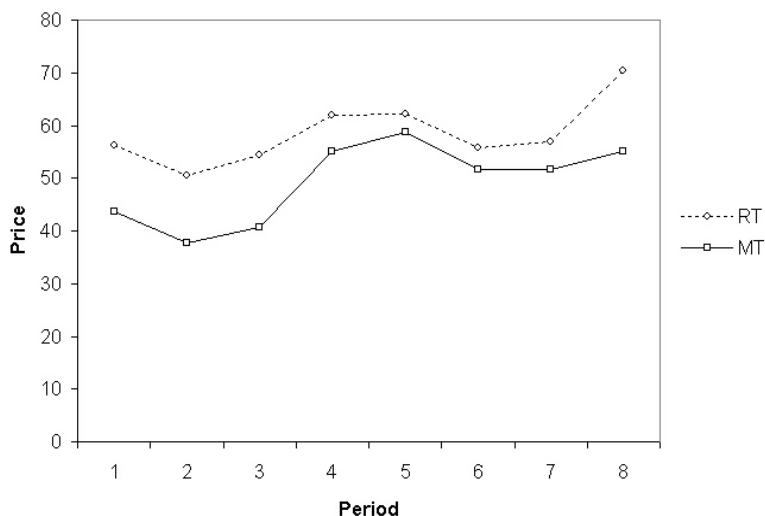
Figure 4: Average price accepted by treatment.

5   Conclusion

We replicated experiments on information aggregation in which partially informed subjects trade an asset with an unknown value. We then added a treatment in which some subjects had an incentive to manipulate market prices. We found that subjects with such incentives made systematically higher price offers than other subjects, in apparent attempts to manipulate price. However, subjects were unable to actually effect the relationship between price and asset value. We found evidence that despite the fact that manipulators make systematically higher price offers in our manipulation treatment, non-manipulators accepted systematically lower price offers. This reluctance to accept high offers by non-manipulators effectively cancelled out the distortionary effects of the efforts to manipulate price.

Our results suggest that, when agents suspect the presence of manipulators and know in what directions manipulators would like to push price, manipulation is ineffective. In a case like the one modeled by our design, this is evidence that prediction markets are safe from some kinds of efforts to manipulate. Although these results are suggestive, they are by no means conclusive. We regard this research as a first step that we hope suggests further research on the topic. Our design is specialized in that all participants know (i) that manipulators are present, (ii) how strong the incentive to manipulate is and (iii) in what direction manipulators have incentives to push price. Our design also creates a public goods problem for manipulators; no individual manipulation subject can capture all of the rents accruing from successful manipulation. All of these features of our design should be relaxed in future research before broader claims about the robustness of prediction markets to manipulation can be made.

References

[1] Allen, F., and Gale, D., 1992. Stock-Price Manipulation. The Review of Financial Studies 5, 503-529.

[2] Camerer,C., 1998. Can Asset Markets Be Manipulated? A Field Experiment with Race-track Betting. Journal of Political Economy 106, 457-482.

[3] Chakrabortya, A., Ylmaz, B. , 2004. Manipulation in Market Order Models. Journal of Financial Markets 7, 187-206.

[4] Forsythe, R., Nelson, F., Neumann, G., Wright, J.,1992. Anatomy of an Experimental Political Stock Market. American Economic Review 82, 1142-1161.

[5] Forsythe, R., Lundholm, R., 1990. Information Aggregation in an Experimental Market. Econometrica 58, 309-347

[6] Hayek, F.,1945. The Use of Knowledge in Society. American Economic Review 35, 519-530.

[7] Hansen, J., Schmidt, C. , Strobel, M. , 2001. Manipulation in Political Stock Markets – Preconditions and Evidence. Tech. rep. 2001-61.

[8] Hanson, R., Oprea, R., 2004. Manipulators Increase Information Market Accuracy, George Mason University Working Paper.

[9] Kagel, J., Roth, A., 1995. The Handbook of Experimental Economics, Princeton University Press.

[10] Kumar, P., Seppi, D. , 1992. Futures Manipulation with 'Cash Settlement'. Journal of Finance 47, 1485-1502.

[11] Lucas, R., 1972. Expectations and the Neutrality of Money. Journal of Economic Theory 4, 103-124.

[12] Milgrom, P., 1981. Rational Expectations, Information Acquisition, and Competitive Bidding. Econometrica 49, 921-943

[13] Muth, J., 1961. Rational Expectations and the Theory of Price Movements. Econometrica 29, 315-335.

[14] Pearlstein, S., 2003: "Misplacing Trust in the Markets," Washington Post, July 30.

[15] Plott, C., Sunder,S., 1982. Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models. The Journal of Political Economy 90, 663-698.

[16] Plott, C., Sunder,S., 1988. Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets. Econometrica 56, 1085-1118.

[17] Strumpf, K., Rhode, P., 2003. Historical Presidential Betting Markets. Journal of Economic Perspectives, 18(2).

[18] Wolfers, J., Zitzewitz, E., 2004. Prediction Markets. Journal of Economic Perspectives, 18(2).

[19] Wyden, R., and B. Dorgan, 2003: "Wyden, Dorgan Call For Immediate Halt to Tax- Funded 'Terror Market' Scheme," Press Release, July 28, 2003.