

Affecting Policy by Manipulating Prediction Markets: Experimental Evidence¹

Cary Deck
University of Arkansas
and Economic Science Institute
Chapman University

Shengle Lin
Economic Science Institute
Chapman University

David Porter
Economic Science Institute
Chapman University

September 2010

Documented results indicate prediction markets effectively aggregate information and form accurate predictions. This has led to a proliferation of markets predicting everything from the results of elections to a company's sales to movie box office receipts. Recent research suggests prediction markets are robust to manipulation attacks and resulting market outcomes improve forecast accuracy. However, we present evidence from the lab indicating that well funded, single minded manipulators can in fact destroy a prediction market's ability to aggregate information. Our results clearly indicate that the usefulness of prediction markets as inputs to decision making may be limited.

Keywords: Information Aggregation, Prediction Markets, Manipulation

JEL: C9, D8, G1

¹ Support from the Economic Science Institute at Chapman University and the Information Technology Research Institute in the Walton College of Business at the University of Arkansas is gratefully acknowledged.

Introduction

Information needed by decision makers is often decentralized. The Office of Homeland Security does not know when or where a terrorist will strike. A company executive has to forecast sales in a given period to schedule production, but lacks information about localized market conditions in different sales regions. A motion picture company has to conjecture how much box office revenue a movie is likely to generate when deciding whether or not to finance the project, but it does not know viewer preferences. The ability to collect and aggregate information can be tremendously valuable, but the process can be difficult and time consuming.

Recently there has been a dramatic increase in the use of markets for the purpose of aggregating information and forming forecasts.² Chen and Plott (2002) set up internal prediction markets at Hewlett-Packard to forecast sales. They report that these markets “performed better than traditional methods employed inside Hewlett-Packard” (p.1). Pennock, et al. (2001) study the Hollywood Stock Exchange (HSX), where people trade futures for a movie’s box office receipts, and find that the market provides highly accurate box office forecasts. Perhaps the most well know prediction market is the Iowa Political Stock Market, which has systematically outperformed opinion polls in predicting election outcomes (Berg, et al. 2003). These success stories are not surprising from a theoretical standpoint. Under the assumptions of rational expectations and low transaction costs, asset markets should aggregate all of the disparate information of the market participants as better informed traders would arbitrage the market (Hayek 1945, Muth 1961). This information aggregation occurs in standard stock markets as well. By examining stock prices, Maloney and Mulherin (2003) describe how Wall Street quickly discovered the company responsible for the Challenger Shuttle explosion as opposed to the months long period required for the official investigation.

The apparent forecasting power of prediction markets suggests that decision makers should use these tools to implement policy (see Hahn and Tetlock 2005 for such an appeal). However, this raises a critical question about the manipulability of prediction markets. If an actor’s salary is based upon the HSX prediction of the box office sales, then the actor has a large incentive to push up the futures price, which is denominated in nonconvertible play money at HSX. Had liability for the Challenger disaster been determined by stock price movement, involved companies would have had tremendous incentives to manipulate their stock prices. This concern was explicitly raised with the Policy Analysis Markets (PAM) that were proposed after the terrorist attacks of September 11, 2001 (Pearlstein 2003 and Wyden and Dorgan 2003).

Despite the potential for deteriorated or even misleading predictions due to traders attempting to manipulate market outcomes, relatively little work has directly examined this issue. The work that has been done suggests that manipulation may not be very effective, but those studies have considered agents with relatively weak manipulation incentives. In this study, we use laboratory methods to examine information aggregation in prediction markets when some agents are motivated exclusively to mislead market observers. As an analogy to the PAM markets, if a terrorist’s only motivation is to cause

² See Wolfers and Zitzewitz (2004) for a general discussion of prediction markets.

significant loss, his own financial profit or loss does not enter his decision process. Ultimately, we find that these single objective manipulators do in fact destroy the predictive value of the prediction markets. This research suggests that prediction markets may only be effective in certain circumstances where traders do not have strong incentives to manipulate outcomes.

Manipulating Prediction Markets

Can prediction markets be manipulated? Wolfers and Zitowitz (2004, p. 119) assert that “The profit motive has usually proven sufficient to ensure that attempts at manipulating these [prediction] markets were unsuccessful.” Failed attempts at manipulating markets include political candidates betting on themselves (Wolfers and Leigh 2002) and bettors placing large wagers at horse races (Camerer 1998). Hansen, et al. (2004) did successfully manipulate election prediction markets, but the effects were short lived. In fact, Rhode and Strumph (2009, p. 37) provide an extensive discussion of attempts to manipulate political markets and conclude that “In almost every speculative attack, prices experienced measurable initial changes. However, these movements were quickly reversed and prices returned close to their previous levels.”

Prediction markets have also been created in the laboratory and subjected to attempted manipulation, again with little success. Plott and Sunder (1988) created markets where subjects traded assets that paid one of three dividend amounts. The amount was randomly determined and half of the traders were informed of one non-realized amount and the other half were informed of the other non-realized amount. In this way the market had perfect information as to what dividend the asset would pay. Plott and Sunder (1988) report the markets successfully identify the actual dividend payment, a result confirmed in the Replication Treatment of Hanson, et al. (2006). Hanson, et al. (2006) also conducted series of Manipulation Treatment experiments where some traders received *additional* payment based upon the median transaction price and thus had an incentive to push prices upwards. They conclude that other subjects account for this bias, which was in a known direction, making manipulation ineffective.

Oprea, et al. (2007) also consider the manipulation of prediction markets. Their information structure follows previous experimental work by Anderson and Holt (1997) and Hung and Plott (2001). Here traders are informed that assets are worth either 0 or 100. Each trader receives a signal of the true value that is accurate with a two thirds chance and inaccurate with a one third chance. Thus, the market as a whole receives multiple independent signals of the true value and the price should reflect the Bayesian expected value of the asset. An innovation of Oprea, et al. (2007) is the introduction of market observers who make predictions about the realized value. These observers receive no private information and cannot trade in the market. The observers make a binary prediction regarding the true value of the assets (0 or 100). If and only if an observer makes an accurate prediction does he then receive a fixed payoff. In some sessions manipulators are introduced. The manipulators are given a common target value and each receives a bonus based upon the percentage of observers predicting that the asset value equals the target. The target is independent of the actual value and thus contains no information. The manipulators’ bonuses are *added* to their regular trading profits. The experimental results indicated that, absent manipulators, prices are correlated with the Bayesian price prediction, but

that prices do not converge to this theoretical benchmark. Still, observers' forecasts are improved (relative to random guessing) with functioning prediction markets. The manipulators attempt to manipulate prices, but have no effect on information aggregation and do not reduce the accuracy of observer forecasts.

Taken together, the evidence suggests prediction markets are quite robust to attempted manipulation. Hanson and Oprea (2004) even show how the increased liquidity provided by manipulators may actually serve to enhance prediction market accuracy. One limitation of the above tests for manipulability of prediction markets was that manipulators suffered the financial losses associated with the activity.³ This is true in any market, but in some cases the relative value of manipulating the market may dominate the financial loss associated with attempting to do so. Further, manipulators may have an informational advantage over other traders. With the PAM markets there was concern that terrorists would use the markets to reduce the perceived likelihood of a planned attack or time attacks to occur when the markets suggested they would not. Almost paradoxically, if these perfectly informed traders are active in the market (or if people expect this kind of behavior on the part of terrorist), then the market should be able to incorporate this information, stymieing attempts at manipulation.

Experimental Design

Our experimental design is similar to that of Oprea, et al. (2007) in that traders receive a signal as to what the true asset value will be (0 or 100) and can trade assets via a double auction market. Forecasters can watch the market and make predictions as to what outcome will be realized. However, our design differs in some potentially important ways. First, rather than having a fixed number of assets as in Oprea, et al. (2007), our traders can create shares of the security by simply selling shares they do not have and setting aside enough money to cover the share's maximum potential value of 100. Our share creation is similar to naturally occurring predictions markets like intrade.com and is used to test the conjecture that the fixed number of shares in Oprea, et a. (2007) drove those prices above 50. Second, rather than making binary predictions, our forecasters have a range of investment opportunities to measure the intensity of their beliefs. Previous work cannot distinguish between a forecaster who thinks the chance a particular event will occur is 51% and a forecaster that believes the likelihood the event will occur is 99%, but the policy implications differ greatly. Finally, our manipulators are given perfect information as to the true asset value and are *not* paid by their own final asset holdings, but paid *only* based upon the amount invested *incorrectly* by the forecasters. This gives our manipulators maximal incentive to disrupt the market's ability to aggregate information. In this sense, our experiment more directly captures the confrontational interplay between manipulators and those attempting to use prediction market outcomes to set policy.

We conducted experiments in three treatments: Baseline, Liquidity and Manipulation. The Liquidity treatment differs from Baseline treatment by the amount of cash holdings in the market; the

³ As pointed out by Hanson, et al. (2006), their design, as well as that of Oprea, et al. (2007), creates a public goods problem among manipulators. Each manipulator would individually prefer to not engage in costly manipulation and instead free-ride on the attempted market manipulations of others.

Manipulation treatment differs from Baseline treatment by introducing manipulators in the market. The focus of the experiment is on the ability of forecasters to make accurate predictions. Therefore, we conducted one session per treatment each with multiple independent forecasters. This important design choice controls the information and experience of forecasters within a treatment. In no case did forecasters receive any feedback regarding the decisions of other forecasters observing the prediction markets.

Baseline Treatment

In the Baseline treatment there were two types of subjects: Traders and Forecasters. The eight traders were each endowed with 300 experimental currency units (ECU) at the start of each of the 20 periods.⁴Subjects were informed that one of two possible events would occur: Black or White. If the event was Black, asset shares were worth 100 and if the event was White, shares were worth 0. A trader's earnings each period were (initial cash of 300 + price received for shares sold – price paid for shares purchased + ending number of shares x share value). A trader who was a net creator of shares would have a negative ending number of shares and the sum of ending share balances across all traders in the market was necessarily 0. Thus, the total wealth among buyers was always 2400 (= 300 endowment x 8 traders), regardless of the state being drawn. Subjects in the role of traders were paid their cumulative earnings in cash at the end of the experiment at the rate 200 ECU = \$1.

At the start of each period, each trader observed a private independent signal about which event would occur that period. Signals were either Dark Gray or Light Gray. If the event was Black, then there was a two-thirds chance that a signal would be Dark Gray and a one-third chance that a signal would be Light Gray. If the event was White then there was a two-thirds chance that a signal would be Light Gray and a one-third chance that a signal would be Dark Gray. With no signal the expected value of a share is 50. With one Dark Gray signal the expected value is 67 and with a Light Gray signal the expected value is 33. Figure 1 shows a sample screen for a Trader. The signal is displayed on the bottom right portion of the screen along with a reminder of terminal share values. Also visible are the Trader's holdings in terms of shares and cash. The large graph on the left of the screen shows current bids and asks as well all previous transactions in the market period. Subjects could enter bids and asks and accept standing bids or asks in the section labeled "Orders" at the top right of the screen.

⁴ Neither shares nor cash carried forward from period to period in any treatment.

Figure 1. Screen Shot for Subjects in the Role of a Trader



If the information of the 8 traders is completely aggregated then expected value of a share when a total of n Dark Gray signals are observed is $\frac{\frac{1}{3} \frac{2^{(8-n)}}{3} + \frac{2^n}{3}}{\frac{1}{3} \frac{2^{(8-n)}}{3} + \frac{1^n}{3} \frac{2^{(8-n)}}{3}} \times 100$. Table 1 gives the

number of Dark Gray signals, the expected price, and realized event for each period. For example, the top row of Table 1 indicates that in period 1 of the Baseline treatment, the event was Black, but only three private Traders observed a Dark Gray signal. Therefore, the Bayesian expected price taking all signals into account would be 20. The last three columns of Table 1 pertain to the Manipulation treatment and are described below.

Table 1. Market Information

Period in Baseline and Liquidity Treatments	Truth	Number of Dark Gray Signals	Bayesian Expected Price	Period in Manipulation Treatment	Manipulator Present
1	Black	3	20.0	11	M3
2	Black	6	94.1	12	None
3	White	2	5.9	13	M1
4	Black	6	94.1	14	M2
5	White	1	1.5	15	M4
6	White	1	1.5	16	None
7	Black	4	50.0	17	M2
8	White	3	20.0	18	M3
9	Black	7	98.5	19	M1
10	Black	5	80.0	20	M4
11	White	1	1.5	1	M3
12	White	5	80.0	2	M4
13	White	2	5.9	3	None
14	Black	5	80.0	4	M2
15	Black	6	94.1	5	M1
16	White	2	5.9	6	M2
17	White	2	5.9	7	M3
18	White	3	20	8	M4
19	White	3	20	9	None
20	Black	5	80.0	10	M1

In the Manipulator treatment, the Events and signals were reversed. For example, in the first market period in the Manipulator Treatment the event was Black and there were 7 Dark Gray signals. M_i denotes a market period in which Manipulator i was active in the Manipulator treatment. None indicates that no Manipulator was active in a given market period in the Manipulator Treatment.

Forecasters were given a budget of 100 ECU every period and could decide how much to invest in Black, with the remainder automatically invested in White. The investment options of Forecasters are shown in Table 2. For example, suppose a forecaster invested 85 in Black. If the event was Black, then forecaster would earn 71 ECU. If the event were White, then the forecaster would earn 15. The payoff is linear in the event for investments below 50 and concave in the event for investments above 50. Given the structure of the payoff function, Forecaster investments in Black should be monotonically increasing with the forecaster's belief that the event will actually be Black. That is, this novel design allows forecasters to express the intensity of their beliefs. The specific investment that a forecaster with a given belief will make depends upon the forecaster's risk attitude. Table 2 also shows the optimal investment for risk neutral forecasters. For example, a forecaster who believes the likelihood that the event will be Black is between 0.77 and 0.83 should invest 85 in Black. At the conclusion of the experiment, Forecasters were paid based upon their cumulative earnings at the rate 33 ECU = \$1.

Before any market signals arrive, the best Forecaster estimate of the likelihood that Black occurs is 50%. Therefore, the default investment is set at 50 each period. As information arrives in the market

with each bid, offer or transaction, Forecasters should use the information shown on their screen in the same graph observed by Traders, see Figure 1, to update their beliefs about the likelihood that Black occurs. This dynamic has been ignored in previous experimental research. To explore how forecasts are affected by the data stream, a Forecaster's ability to update his or her investment had an unknown termination. Specifically, there was an 80% chance that a Forecaster could update his or her investment throughout the market period and a 20% chance that he or she could not. Conditional on not having the full period to update one's investment, the investment stopped r seconds *before the end of the period* with r distributed uniformly over $[0,120]$. Subjects were only informed that there was some chance they would not be able to update investments for the entire period and therefore should keep updating their investments continually. As described later in the paper, this rule was common information among all participants.

Table 2. Investment Opportunities of Forecasters

Amount Invested in Black	Payoff if Event is Black	Payoff if Event is White	Implied Likelihood that Event is Black for Risk Neutral Forecaster
100	72.7	0	[0.97,1.00]
95	72.5	5	[0.91,0.96]
90	72	10	[0.84,0.90]
85	71	15	[0.77,0.83]
80	69.5	20	[0.72,0.76]
75	67.5	25	[0.67,0.71]
70	65	30	[0.63,0.66]
65	62	35	[0.59,0.62]
60	58.5	40	[0.56,0.58]
55	54.5	45	[0.53,0.55]
50	50	50	[0.48,0.52]
45	45	54.5	[0.45,0.47]
40	40	58.5	[0.42,0.45]
35	35	62	[0.38,0.41]
30	30	65	[0.34,0.37]
25	25	67.5	[0.29,0.33]
20	20	69.5	[0.24,0.28]
15	15	71	[0.17,0.23]
10	10	72	[0.10,0.16]
5	5	72.5	[0.04,0.09]
0	0	72.7	[0.00,0.03]

Manipulation Treatment

The Manipulation treatment consisted of the Baseline environment with the addition of subjects in the role of Manipulators, referred to as Target Traders in the experiment (following the terminology

of Oprea, et al. 2007). Instead of receiving a signal before the market period began, Manipulators were informed of the actual event that would be realized. Manipulators were endowed with 1200ECU and could participate in the market like regular Traders. The Manipulator's endowment is 4 times that of a regular trader resulting in the Manipulator having one third of the money in the market. However, Manipulators were not paid in any way for their market earnings. Instead, they were paid solely based upon the average amount that Forecasters invested in the wrong event. There were four subjects acting as Manipulators. As opposed to previous studies, only one Manipulator was active in any given period. Each Manipulator was only active for 4 periods and there were 4 periods in which no Manipulator was active. While inactive, Manipulators were still informed of the realized event, able to observe the market, and informed of the average Forecaster investment.

Our design choices mean that subjects in the role of Manipulators had very strong incentives to mislead Forecasters, a feature further encouraged by the exchange rate of 8ECU = \$1 that was applied to Manipulators' cumulative earnings. The extreme financial resources associated with having a third of the market liquidity and the strong incentives to manipulate the market are designed to stress the robustness of previous results that prediction markets are difficult to manipulate.

Traders and Forecasters were aware of the possibility that Manipulators could be active in a given period, but did not know if and when such subjects were active. Table 1 indicates which markets did not have an active Manipulator present. For reasons discussed below, it was necessary to relabel the events and signals and reorder the market periods for the Manipulator treatments. Thus, row 1 of Table 1 pertains to period 11 of the Manipulation Treatment. In this period the event was White, there were 3 Light Gray signals, the Bayesian expected price was 80, and Manipulator 3 was active. Forecaster investment opportunities were as shown in Table 2.

Liquidity Treatment

The Liquidity treatment was identical to the Baseline treatment, except that each of the 8 traders was endowed with 450ECU at the start of each period and the cumulative earnings of Traders in the Liquidity treatment were converted into cash at the rate 300 ECU = \$1. Therefore, the liquidity in the system is the same as in the Manipulation treatment, but the real wealth is the same as in the Baseline treatment. Including this treatment allows us to separate a change in forecaster accuracy due to increased liquidity from a change due to the actions of manipulators while maintaining the same starting liquidity and distribution of signals for Traders in the Baseline and Manipulator treatments. The signals, expected prices and realized events are those shown in Table 1 for the Baseline treatment. Trader endowments were unknown to the Forecasters in every treatment, but exchange rates were public information. Thus, there was no difference between the Baseline and Liquidity treatments from the perspective of the Forecasters, whose own exchange rates were unchanged across all treatments. Forecaster investment opportunities were those shown in Table 2. There were no Manipulators in the Liquidity treatment.

Experimental Procedures

The experiments were conducted at the Economic Science Institute at Chapman University over the course of three distinct days in a single week. One set of subjects was recruited to participate in a two day experiment, where the lab session each day would last two hours. These subjects participated in the Baseline treatment on the first day. A second set of subjects was recruited for another two day experiment, with two hour sessions each day. These subjects completed the Liquidity treatment on the first day. Both sets of two day experiment subjects returned to the lab at the same time as a third group of subjects recruited for a one day, two hour experiment reported to the lab. The Manipulation treatment was completed using subjects from all three groups. In all cases, traders were drawn from the subjects who were completing the first day of their respective experiments. Manipulators, only present on the last day, were drawn from the experienced Traders returning for their second day.⁵ In the Baseline and Liquidity treatments, the 10 Forecasters were Inexperienced. For the Manipulation treatment, there were 30 Forecasters: 10 drawn from the experienced Forecasters returning for the second day, 10 from the experienced traders returning for the second day, and 10 from the inexperienced subjects arriving for the first day. Because some Forecasters in the Manipulation treatment were experienced as Traders or Forecasters, it was necessary to reorder market periods and reliable the signals and events so that these subjects could not rely on memory to make investments while maintaining comparability across treatments. This design feature allows us to compare how different experiences affect a Forecaster's ability to glean information from prediction markets. Other than as described here, no subjects had prior experience with prediction market experiments. Table 3 summarizes our experimental design.

At the start of each session, the directions were read aloud to everyone so that everyone heard what subjects in every role were told. The directions were identical for the Baseline and Liquidity treatments, but differed for the Manipulation treatment, which was the only version that discussed the possibility of Manipulators being active in the market. After the directions were completed, subjects completed a role specific quiz online. This process lasted approximately 30 minutes, with the remaining 90 minutes used to complete the 20 separate three minute trading periods and allow subjects time to review results between periods. The appendix contains a copy of the instructions. All subjects received a participation fee of \$10 per day. The participation payment(s) and all salient earnings were paid on the last day. The average payment was \$33 per day.

⁵ The Traders from the first two treatments were ranked according to their earnings on the first day. The first, third, fifth, and seventh traders were selected to be Manipulators. This procedure was designed to give manipulation its best chance of success given that previous research had found prediction markets were robust to manipulation, while still allocating some of the top performing traders to the role of Forecaster.

Table 3. Experimental Design

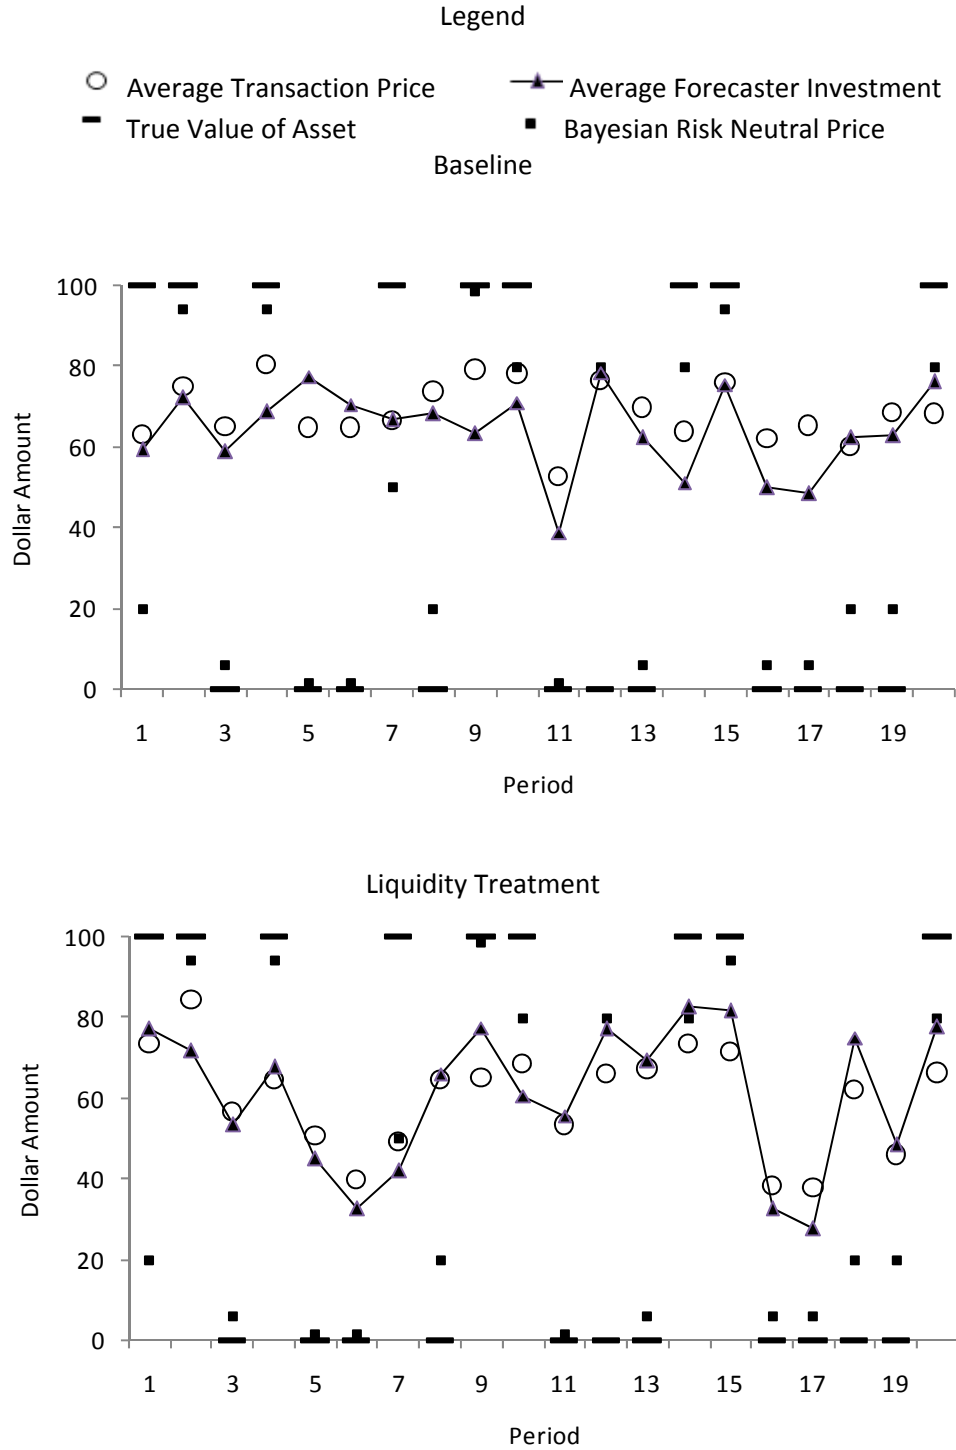
Treatment	Baseline	Manipulation	Liquidity
Number of Forecasters	10	30	10
Forecaster Experience	None	10 experienced as Forecasters 10 experienced as Traders 10 with no experience	None
Number of Traders	8	8	8
Experience of Traders	None	None	None
Number of Manipulators	0	4, with either 0 or 1 manipulator present in a period	0
Experience of Manipulators	-	Experienced as Traders	-
Periods	20	20	20
Total Cash (ECU)	2400	3600 when Manipulator is active 2400 when Manipulator is inactive	3600

Experimental Results

Before looking at the impact that manipulators have on the market we consider performance in the Baseline and Liquidity treatments. Our two main findings are that: 1) markets are aggregating information consistent with previous experimental studies and 2) forecasters are using this information to make their predictions.

Figure 2 plots the average transaction price (large circles), average forecaster investment (small triangles), the true value of the assets (thick dashes) and the Bayesian risk neutral price (solid squares) for each market period in the Baseline and Liquidity treatments, respectively. Average prices are positively and significantly correlated with the Bayesian predictions. For Baseline $\rho = 0.73$ (p-value < 0.001) and for Liquidity $\rho = 0.66$ (p-value = 0.002). Two additional features of the data warrant attention. First, we do not find support for the claim by Hanson and Oprea (2004) that increased liquidity improves information aggregation. A second feature is that transaction prices tend to be less extreme than predicted, and tend to be too high, often above 50 even when the Bayesian expected price is quite low. . This is the same pattern reported by Oprea, et al. (2007) and suggests that the fixed number of shares in their design is not the cause for this behavioral phenomenon.

Figure 2. Average Transaction Price and Average Forecaster Investment by Period



Without market information, Forecasters should do no better at predicting the outcome than random guessing. However, our Forecasters are able to extract information from the markets and make better predictions. The appropriate litmus test for Forecaster performance is the Bayesian price since

this reflects the best available information. In the Baseline treatment, investments are positively and significantly correlated with Bayesian Prices ($\rho = 0.45$, p -value = 0.043). For the Liquidity treatment the results are even stronger ($\rho = 0.64$, p -value = 0.002).

To determine what information Forecasters are using, we estimated a random effects model with clustered standard errors to reflect the fact that our data includes repeated observations from the same Forecasters. The data for each period are broken up into 15 second intervals and the dependent variable is the change in a Forecaster's investment.⁶ We regress the change in a forecaster's investment on the change in the moving average of the 5 most recent transaction prices and the change in excess bid (number of bids minus number of asks) during the time interval. Intuitively, these two variables capture the new information that the Traders are discovering as Traders only have four actions: bid, ask, accept a standing bid, and accept a standing ask. An increase in bidding relative to asking or higher transaction prices suggests that traders are discovering that the current price is too low. Falling prices or increasing offers to sell relative to the number of bids suggests that Traders have discovered that prices are too high. Table 3 reports the estimation results for the Baseline and Liquidity treatments, estimated separately. It is clear from the estimation results that as prices increase, forecasters increase their investment and as the number of excess bids increases, forecasters increase their investments. Both of these influences are consistent with Forecasters assuming that the market is aggregating information. Looking separately at the top 5 and bottom 5 earning Forecasters in these two treatments (see Table 3), top 5 earners in both liquidity treatment and baseline treatment are more sensitive to changes in moving average price than bottom 5 earners, as reflected in the larger regression coefficient on prices. While everyone is using price to inform their investments, the top earners are also using excess bids.

The results in Table 3 reflect the direction of the change in Forecasters' estimates, but Forecaster behavior could also be influenced by risk attitude, optimism, or the price dispersion in the market. In all three cases, these variables could influence how aggressively, measured as an absolute deviation from 50, a Forecaster invests. Table 4 reports similar regression analysis to that reported in Table 3, where the dependent variable is the aggressiveness of the investment and the dependent variables include risk attitude index, an optimism index, the standard deviation of prices in the interval as well as the absolute value of the difference between the moving average of the last five transactions and 50 and the absolute value of the number of excess bids.⁷The same pattern shown in Table 3 holds in Table 4. More extreme prices and a more extreme number of excess bids leads to more extreme investments. An increase in the standard deviation of prices is associated with less extreme investments, but the risk and optimism measurements have no predictive power.

⁶ Only Forecasters who were still able to update their investment at the end of the 15 second window are included in the analysis. The results are qualitatively similar if the window is changed to 30 seconds.

⁷ The risk attitude and optimism index were jointly estimated for the subjects based upon responses the subjects gave to a battery of questions approximately 4 months before this experiment. John Dickhaut and Nat Wilcox had developed these measurement tools and collected information for a large subsample of the subject pool at Chapman University. Recruitment for this experiment was limited to those people in this subsample so that we could exploit this information.

Table 3. Forecaster Investment in Baseline and Liquidity Treatments

	Coefficient	Robust Standard Error	z-statistic	p-value
Baseline - All Forecasters (Observations = 836)				
Constant	0.034	0.817	0.04	0.967
Moving Average of 5 Most Recent Transaction Prices	0.553	0.043	12.77	<0.001
Number of Excess Bids	0.516	.108	4.79	<0.001
Liquidity – All Forecasters (Observations = 1408)				
Constant	-0.054	0.864	-0.06	0.950
Moving Average of 5 Most Recent Transaction Prices	0.393	0.061	6.47	<0.001
Number of Excess Bids	0.876	0.162	5.38	<0.001
Top Five Earners – Baseline and Liquidity (Observations =1125)				
Constant	-0.032	0.6870	-0.05	0.963
Moving Average of 5 Most Recent Transaction Prices	0.567	0.038	14.77	<0.001
Number of Excess Bids	1.012	0.099	10.26	<0.001
Bottom Five Earners - Baseline and Liquidity (Observations = 1120)				
Constant	0.390	0.992	0.39	0.694
Moving Average of 5 Most Recent Transaction Prices	0.419	0.057	7.30	<0.001
Number of Excess Bids	0.211	0.142	1.48	0.139

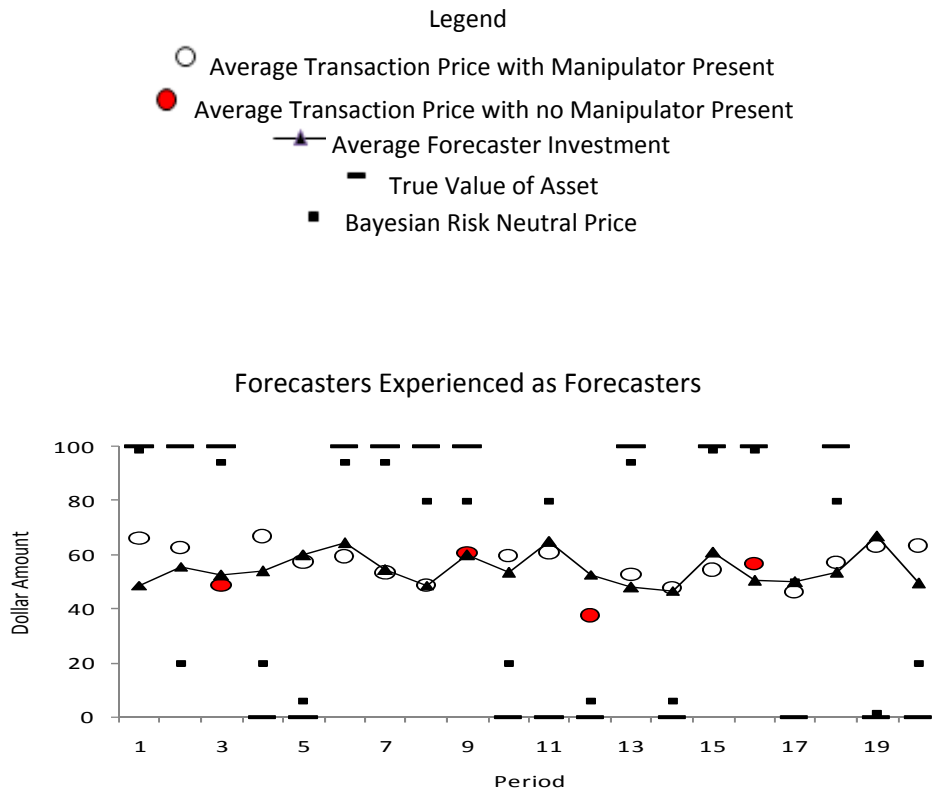
Table 4. Aggressiveness of Forecaster Investment in Baseline and Liquidity Treatments

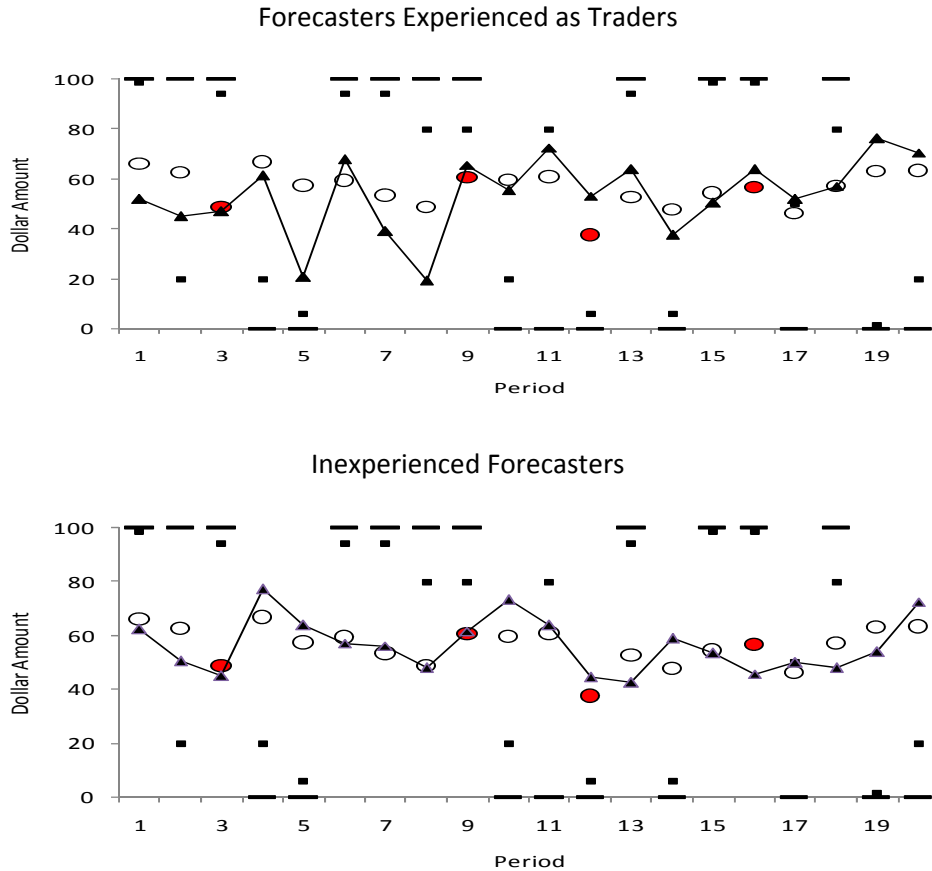
	Coefficient	Robust Standard Error	z-statistic	p-value
Constant	12.56	2.98	4.22	0.000
Standard Deviation of Prices Average of Five Most Recent Transaction Prices – 50	-0.09	0.05	-1.81	0.071
Excess Bids	0.26	0.03	10.27	0.000
Risk Index	0.67	0.08	8.85	0.000
Optimism Index	1.52	4.80	0.32	0.751
	0.62	5.21	0.12	0.905

Having established that in the absence of manipulators, markets aggregate information which is utilized by Forecasters, consistent with previous experimental work, we now turn to the impact that manipulators have on market and Forecaster performance. Contrary to previous studies, we find that the markets fail to aggregate information. The correlation between average price and the Bayesian price is indistinguishable from zero: $\rho = 0.00$ (p-value > 0.999).

Figure 3 plots the average transaction price (large circles), the true value of the assets (thick dashes) and the Bayesian risk neutral price (solid squares) for each market period for the Manipulation treatment. A shaded circle indicates a period in which no Manipulator was active. The small triangles denote the average Forecaster investment for forecasters experienced as forecasters in the top panel of Figure 3, forecasters experienced as traders in the middle panel of Figure 3, and novice forecasters in bottom panel of Figure 3. As compared to the other two treatments, prices in the Manipulation treatment tend to be lower and less variable between periods. It is also clear from Figure 3 that forecasters are reacting differently depending on their past experience. Inexperienced Forecasters follow transaction prices closely, while experienced Forecasters make predictions that are counter to price movements. Table 5 provides econometric support for these conclusions. Using the same random effect model with standard clustering errors, the econometric results also indicate that Forecasters who had previous Trader experience followed market prices and reacted to the number of excess bids. This pattern is somewhat intuitive; if markets are correctly aggregating information then investments should follow prices, as in the Baseline and Liquidity treatments. The question a Forecaster faces here is whether or not the markets are aggregating information. Novice Forecasters seem to believe the markets while experienced Forecasters suspect that manipulators are taking prices in the “wrong” direction.

Figure 3. Mean Price and Average Investment by Period: Manipulation Treatment





As with the no manipulation treatments, separate analysis was conducted for the highest and lowest earning Forecasters in this treatment. Based upon the results shown in Table 5, those earning the most profit relied upon excess bids and ignored prices, while the lowest earners simply followed market prices and mis-interpreted excess bids. Additional analysis was conducted to look at how aggressive forecasters were in their investments taking into account their own risk attitude and optimism, in addition to the variation in market price, see Table 6. The qualitative results of the base model in Table 5 are unchanged. Risk and optimism do not have a significant influence, but in this case an increase in price variation leads to more aggressive forecaster investment in aggregate.

Given that the markets fail to aggregate the private Trader information, it is not surprising that Forecaster investments are not correlated with the Bayesian Risk Neutral predictions. The correlation of average investment and the Bayesian Price for Forecasters previously experienced as Forecasters is $\rho = -0.03$ ($p\text{-value} = 0.902$) and for Forecasters previously experienced as Traders it is $\rho = 0.08$ ($p\text{-value} = 0.725$). For inexperienced Forecasters the correlation is $\rho = -0.37$ ($p\text{-value} = 0.106$) providing weak evidence that these people are making bad decisions. One explanation for the poor market performance could be that Manipulators in their attempts to mislead the Forecasters are actually revealing their own perfect information, but this is not the case. Investments of experienced Forecasters are uncorrelated with the event ($\rho = -0.09$, $p\text{-value} = 0.707$ for forecasters with forecasting

experience and $\rho = -.12$, p-value = 0.625 for forecasters with trading experience). Inexperienced Forecasters make investments that are significantly negatively correlated with the truth ($\rho = -0.51$, p-value = 0.023). Therefore, information markets with manipulators do not help forecasters make better predictions and in fact cause inexperienced forecasters to make bad predictions.

Table 5. Forecaster Investment in Manipulation Treatment Based on Prior Experience

	Coefficient	Robust Standard Error	z-statistic	p-value
Novice (Observations = 1589)				
Constant	0.085	0.811	0.10	0.917
Moving Average of 5 Most Recent Transaction Prices	0.131	0.066	1.97	0.049
Number of Excess Bids	-0.866	0.135	-0.64	0.520
Forecaster Experience (Observations = 1598)				
Constant	-0.202	0.501	-0.40	0.687
Moving Average of 5 Most Recent Transaction Prices	-0.098	0.041	-2.40	0.017
Number of Excess Bids	0.054	0.082	0.66	0.511
Trader Experience (Observations = 1595)				
Constant	0.134	0.617	0.22	0.828
Moving Average of 5 Most Recent Transaction Prices	0.377	0.052	7.30	<0.001
Number of Excess Bids	-0.548	0.102	-5.38	<0.001
Top Five Earners (Observations = 792)				
Constant	0.360	0.961	0.37	0.708
Moving Average of 5 Most Recent Transaction Prices	0.111	0.080	1.39	0.165
Number of Excess Bids	0.647	0.161	4.01	<0.001
Bottom Five Earners (Observations = 800)				
Constant	-0.287	1.347	-0.21	0.831
Moving Average of 5 Most Recent Transaction Prices	0.443	0.110	4.02	<0.001
Number of Excess Bids	-0.498	0.219	-2.27	0.023

Table 6. Aggressiveness of Forecaster Investment in Manipulation Treatment

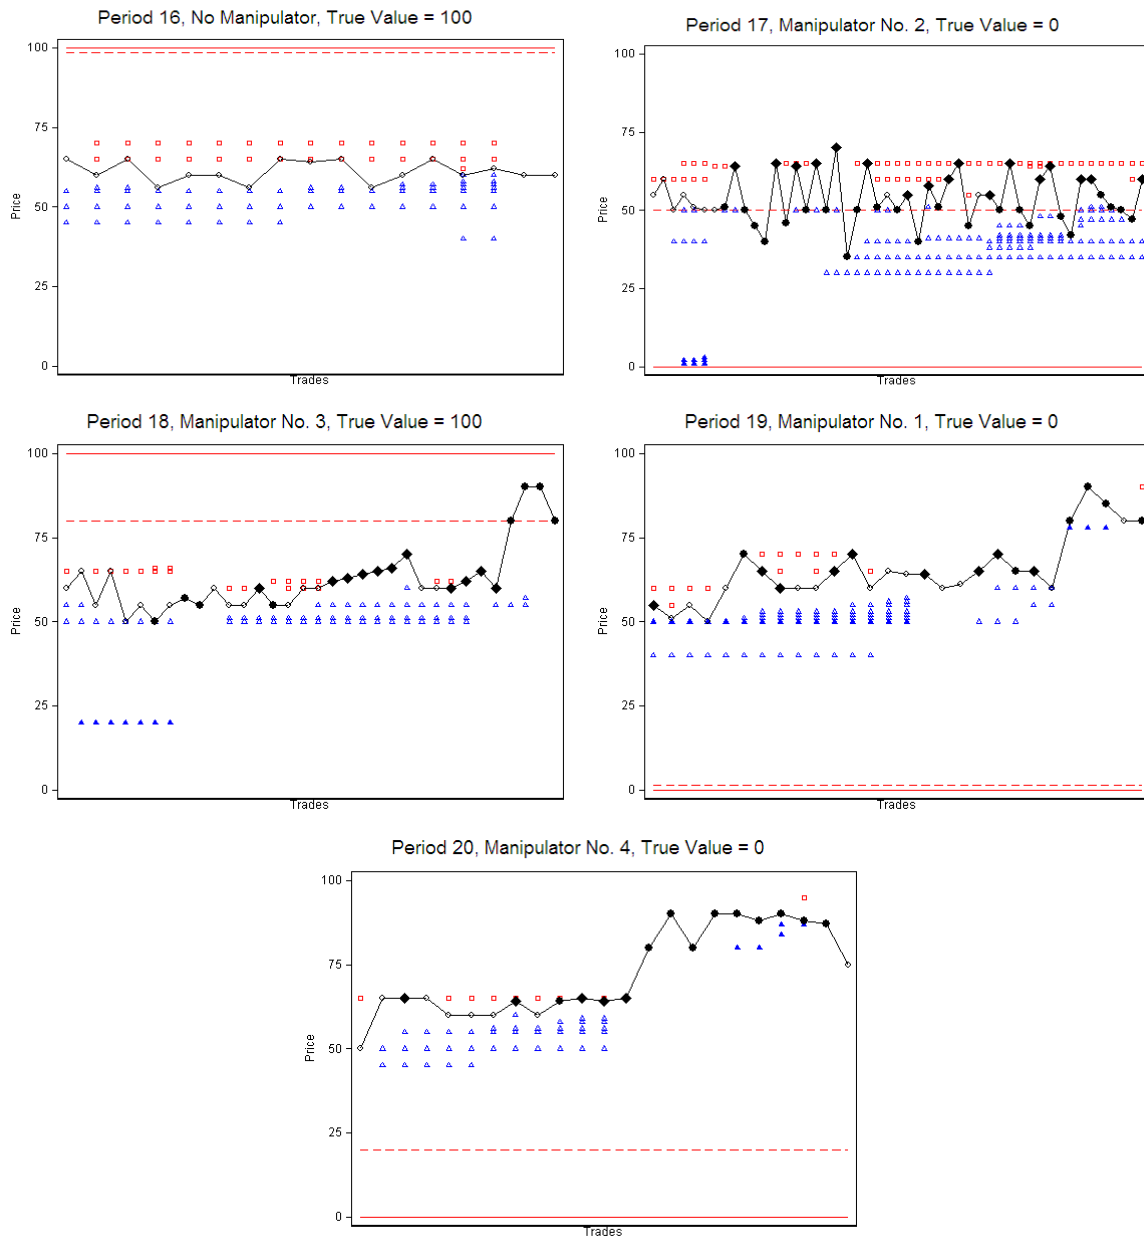
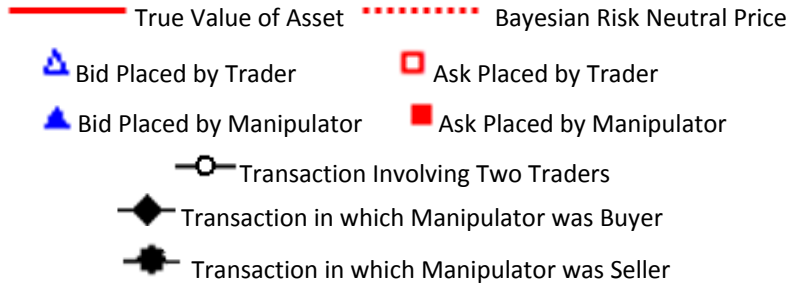
	Coefficient	Robust Standard Error	z-statistic	p-value
Constant	8.55	3.46	2.47	0.013
Standard Deviation of Prices Average of Five Most Recent Transaction Prices – 50	0.20	0.04	4.66	0.000
Excess Bids	0.11	0.02	5.05	0.000
Risk Index	0.34	0.06	5.35	0.000
Optimism Index	0.65	8.03	0.08	0.935
	-1.38	7.08	-0.19	0.846

Given that manipulators are successfully preventing the market from providing useful information to Forecasters, it is worth considering what strategies the manipulators are using. Averaging across the markets in which Manipulators were active, Manipulators were involved in 57% of the trades. The range varied from a low of 25% of trades to a high of 82% of the trades. Manipulator trading does increase trade volume as compared to markets without manipulators. A market by market comparison between the Liquidity and Manipulator treatments reveals greater trade volume when manipulators are present in every case. Closer inspection of the data indicates that in about half of the manipulated markets, Manipulators primarily accepted Trader bids or asks while in the remaining markets Manipulators were primarily posting bids or asks that Traders were accepting. This variation by Manipulators helps explain why it is so difficult for Forecasters to consistently predict the event.

Figure 4 shows the standing bids, asks, and contract prices for each trade in last 5 periods with solid markers indicating activity by the manipulators. In period 16 there was no manipulator. In this period there are relatively few trades, but prices are on the correct side of 50. Each of the four manipulators was active exactly once in the last four periods. There is no obvious pattern to manipulator behavior with two pushing prices away from the truth, one pushing prices towards the truth and one with flat prices. In period 17 the manipulator actively picks off standing bids and asks creating flat prices. This manipulator accounts for 12% of the bids placed in the market in Period 17, but none of the asks. In period 18, the manipulator is involved in 14 of the last 16 trades initially as a seller and then as a buyer ultimately pushing prices to the Bayesian prediction. Here the manipulator accounts for 20% of the bids and 8% of the asks (all of the manipulator asks are masked in Figure 4 because in each instance a Trader immediately accepted it resulting in a transaction). Periods 19 and 20 look very similar to period 18 despite the fact that the signals and the event are reversed and different manipulators are involved. In these two periods the manipulator accounted for 33% of the bids and 3% of the asks on average, consistent with a strategy of trying to raise market prices. Summarizing over all 16 periods where manipulators are present, manipulators pushed prices away from the true value in 8 periods, kept prices flat in 6 periods, and drove prices closer to true value in 2 periods.

Figure 4. Bids, Asks, and transactions in Periods 16-20 of the Manipulation Treatment

Legend



Policy Implications

The experimental results clearly show that Manipulators were able reduce the predictive power of information markets and create a situation where the Forecasters were unable to make good decisions. Manipulators do this by actively trading in the markets, which increase the trade volume. The silver lining of this is that the active trading by manipulators provides a means for identifying when manipulators are likely to be active in a market. Table 7 provides the estimation of a random effect probit model where the dependent variable is one if there is a manipulator present in the market and 0 otherwise. Based upon this estimation, manipulators increase trade volume while making prices less extreme and less varied within a trading period. Based upon the information in Table 7, manipulators do not impact the number of excess bids or asks.

Table 7. Probit Estimation for Presence of a Manipulator

	Coefficient	Standard Error	z-statistic	p-value
Constant	-1.971	1.109	-1.78	0.076
Average Price – 50	-0.043	0.025	-1.70	0.088
Excess Bids	0.0009	0.0008	1.17	0.244
Trade Volume	0.162	0.051	3.20	0.001
Standard Deviation of Prices	-0.172	0.060	-2.89	0.004

N=60, Pseudo R²=0.497, |Average Price – 50| is the absolute value of a period's average price minus 50, |Excess Bids| is the absolute value of the sum of excess bids in a period.

Table 8 reports random effects probit estimation for predicting the event when there are no manipulators aggregating across all treatments and when manipulators are present. The results in Table 8 confirm that forecasters should use price to inform their decisions in the absence of manipulators. This affirms the behavior reported above showing that forecasters in the Baseline and Liquidity treatments relied upon prices. Unfortunately, the results in Table 8 also indicate that when manipulators are present prices and excess bids cannot be used to accurately predict the event. This suggests that when forecasters observe increases in trade volume, they should be reticent and incorporate information from other sources of intelligence. Subjects in the experiment had no such ability and thus had no chance to outperform random guessing.

Table 8. Probit Estimation for Predicting the Event

When No Manipulator				
	Coefficient	Standard Error	z-statistic	p-value
Average Price – 50	0.044	0.020	2.18	0.029
Excess Bids	0.0002756	0.00015	1.80	0.071
Constant	-3.018	1.321	-2.28	0.022
When Manipulator is present				
	Coefficient	Standard Error	z-statistic	p-value
Average Price – 50	-0.188	0.116	-1.62	0.105
Excess Bids	0.0007	0.0005	1.37	0.171
Constant	10.99	6.838	1.61	0.108

Note: in the two regressions above, there are 44 periods of non-manipulation, and 16 periods of manipulation. State is either 0 or 100, |Excess Bids| is the sum of excess demand for a period, and average is the average price for a period.

Discussion

Previous research has demonstrated that prediction markets can be quite successful in the absence of manipulation, a result we replicate in this study. Our results also demonstrate that “high” prices observed by other researchers are not the result of a fixed number of available shares. Success has led to calls for markets to be relied upon in increasingly complex and important policy decisions by firms and government agencies. The recent proliferation in the use of prediction markets necessitates the need to scrutinize the boundaries in which such markets can function well. Upon inspection of previous stress tests of prediction markets, success has been limited to relatively tame environments with small incentives for manipulation. However, in some situations, such as determining how to spend budget against a terrorist attack or how to set monetary policy (Sumner 2006), manipulators may have very strong incentives to manipulate markets thereby encouraging particular policy actions.

We find clear evidence that highly incentivized manipulators can destroy the predictive power of a prediction market. That is we have identified a case where manipulators do cause human forecasters to make predictions that are no better than random guessing would generate showing that prediction markets can be manipulated. In fact, the manipulators are so effective that the econometrician observing the market cannot predict the truth more accurately than a coin flip. This finding demonstrates that policy makers should not indiscriminately rely upon market predictions. A secondary question is if one can identify when manipulators are present in a market. Our results show that manipulators increase trade volume, make prices less extreme, and reduce price variation. Further, our results show that the effects of introducing manipulators is due to more than the large influx of liquidity in the market. Our results may help identify when policy makers should rely upon prediction markets. Clearly, further research is needed to identify how specific environmental and intuitional factors affect the ability of prediction markets to aggregate information in the presence of attempted manipulation and the strategies employed by manipulators.

References

- Anderson, Lisa R., and Charles A. Holt, 1997, Information cascades in the laboratory, *American Economic Review* 87, 847-862.
- Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz, 2008, Results from a dozen years of elections futures markets research, *Handbook of Experimental Economics Results*, ed: Charles A. Plott and Vernon L. Smith, Elsevier: Amsterdam, 742-751.
- Camerer, Colin, 1998, Can asset markets be manipulated? A field experiment with racetrack betting, *Journal of Political Economy* 106, 457-82.
- Chen, Kay-Yut, and Charlie A. Plott, 2002, Information aggregation mechanisms, concept, design and field implementation, *Social Science Working Paper no. 1131*, California Institute of Technology.
- Hahn, Robert, and Paul Tetlock, 2005, Using information markets to improve public decision making, *Harvard Journal of Law & Public Policy* 29, 213-289.
- Hansen, Jan, Carsten Schmidt, and Martin Strobel 2004, Manipulation in political stock markets - Preconditions and evidence, *Applied Economics Letters* 11, 459-463.
- Hanson, Robin and Ryan Oprea, 2004, Manipulators increase information market accuracy, *Working Paper*, George Mason University.
- Hanson, Robin, Ryan Oprea and David Porter, 2006, Information aggregation and manipulation in an experimental market, *Journal of Economic Behavior and Organization* 60, 449-459.
- Hayek, Friedrich A., 1945, The use of knowledge in society, *American Economic Review* 35, 519-30.
- Hung, Angela A., and Charles R. Plott, 2001, Information cascades: replication and an extension to majority rule and conformity rewarding institutions, *American Economic Review* 91, 1508-1520.
- Maloney, Michael T., and J Harold Mulherin, 2003, The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash, *Journal of Corporate Finance* 9, 453-479.
- Muth, J.F., 1961, Rational expectations and the theory of price movements, *Econometrica* 29, 315-335.
- Oprea, Ryan, David Porter, Chris Hibbert, Robin Hanson, and Dorina Tila, 2007, Can manipulators mislead market observers? *ESI Working Paper*, Chapman University.
- Pearlstein, S., 2003, Misplacing trust in the markets. *Washington Post*, July 30.
- Pennock, David M., Steve Lawrence, C. Lee Giles, and Finn Arup Nielsen, 2001, The power of play: Efficiency and forecast accuracy in web market games, *Science* 291, 987-988.
- Plott, C., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56, 1085-1118.

Rhode, Paul, and Kolemanstrumpf, 2009, Manipulating political stock markets: A field experiment and a century of observational data, NBER Working Paper.

Rhode, Paul W., and Kolemanstrumpf, 2004, Historical presidential betting markets, *Journal of Economic Perspectives* 18, 127-142.

Sumner, S., 2006, Let a thousand models bloom: the advantages of making the FOMC a truly 'open market, *Contributions to Macroeconomics* 6, Article 8.

Wolfers, Justin, and Andrew Leigh, 2002, Three tools for forecasting federal elections: Lessons from 2001, *Australian Journal of Political Science* 37(2) 223-240.

Wolfers, Justin, and Eric Zitzewitz, 2004., Prediction markets, *The Journal of Economic Perspectives* 18(2), 107-126.

Wyden, R. and B. Dorgan 2003. Wyden, Dorgan call for immediate halt to tax-funded "Terror Market scheme, Press Release, July 28.

Instructions

This is an experiment in market decision making. You will be paid in cash for your participation at the end of the experiment. Different participants may earn different amounts. What you earn depends on your decisions and the decisions of others.

The experiment will take place through the computer terminals at which you are seated. If you have any questions during the instructions, raise your hand and a monitor will come by to answer your question.

If any difficulties arise after the experiment has begun, raise your hand, and someone will assist you.

In this experiment there will be **Traders**, **Target Traders** and **Forecasters**. On the sheet of paper located at your workstation, you will find the role you will be taking. More details are coming about what each kind of participant does.

There will be several market rounds in this experiment. Each market round is separate, although everyone will maintain the same role in all market rounds. Your payoff will be the cumulative sum of your earnings in each market round. On your sheet of paper you will find your **Exchange Rate** which converts your experimental earnings into US Dollars.

There are two possible events that could occur each round called “White” and “Black.” Both events are equally likely. This means that there is a 50% chance that the event is White and a 50% chance that it is Black. You can think of this like flipping a coin; if it is heads the event is White and if it lands on tails the event is Black.

Target Traders learn which event occurred at the start of each round, but **Traders** and **Forecasters** only learn which event occurred at the end of the round. Events each round are independent, so the event that occurred in a previous round will not affect the event that will occur in any other round.

Before each round begins, all **Traders** will receive their own private information about which event might occur. The signal will be either “light gray” or “dark gray.” If the event is going to be White, there is a two-thirds chance that the signal will be light gray and a one-third chance it will be dark gray. However, if the event is going to be Black, there is a two-thirds chance that the signal will be dark gray and a one-third chance it will be light gray.

Each **Trader** receives their own independent signal so no one will know what signal anyone else observed. Even though each **Trader** observes an independent signal, there is only one event each round and it is common to everyone in the market.

You can visualize this process as follows. We have two bowls, each with three balls. The white bowl has two light gray balls and one dark gray ball, while the black bowl has two dark gray balls and one light gray ball. If the event is White, the computer will draw a ball from the white bowl, show a **Trader** the ball, put the ball back in the bowl, draw a ball for the next **Trader** from the same bowl, and so on. The computer will do the same thing using the black bowl if the event is Black. A **Trader** will see what ball was drawn, but not which bowl was used. A **Target Trader**, sees which bowl was used, but not which balls were drawn and shown to **Traders**.

Traders

Traders can trade “Shares” that work like lottery tickets. The share is worth 100 if the event is Black and worth 0 if the event is White.

Each round, **Traders** will start with some “Cash” and can buy or sell shares if they wish. If a **Trader** buys a share, she pays the price of the share and earns 0 if the event is White and 100 if it is Black. If she sells a share that she bought previously during the round, she will receive the price but will not earn 100 if the event is Black or 0 if it is White.

As a **Target Trader**, you will also start with some “Cash” and can trade shares like a **Trader**, but the value of the shares and your cash holdings do not affect your earnings. Your earnings are based solely upon the choices of the **Forecasters**. More information is coming about **Forecasters**, but first we review how the market for shares works.

Traders

If a Trader does not have a share, but still wishes to sell, she can create a share. Creating a share is identical to selling an existing share except that the seller will have to pay the buyer 100 if the event is Black and 0 if the event is White. In order to make sure the seller creating a share can cover the value at the end of the round, the computer automatically takes 100 from the **Trader** and puts it into a reserve account to cover payment if the event is Black. So to create a share a **Trader** must have 100 in cash after adding the selling price, which the seller receives. The reserves are given back to the trader if the event is white.

Now we will describe a Trader's screen

At the beginning of each round, Traders will be given some Cash shown in **Your Holdings** section.

Cash: the available cash in a Trader's account.

Shares: number of shares the Trader owns (or has created if negative)

Your Holdings

Cash	807
Shares	0
Cumulative Earnings 0	

At the beginning of each round Traders will be provided **Information** on the potential value of a share

Your clue: independent signal to a Trader on the potential value of a share

Shares are worth either 100 (Black) or 0 (White)

Information

Your clue
Light gray

Shared Message
If the event is Black shares are worth 100. If the event is White shares are worth 0.



During every round, Traders can buy or sell shares from one another by making offers to buy or to sell.

The existing offers are shown on the **Market Graph** to the left.

Your Holdings

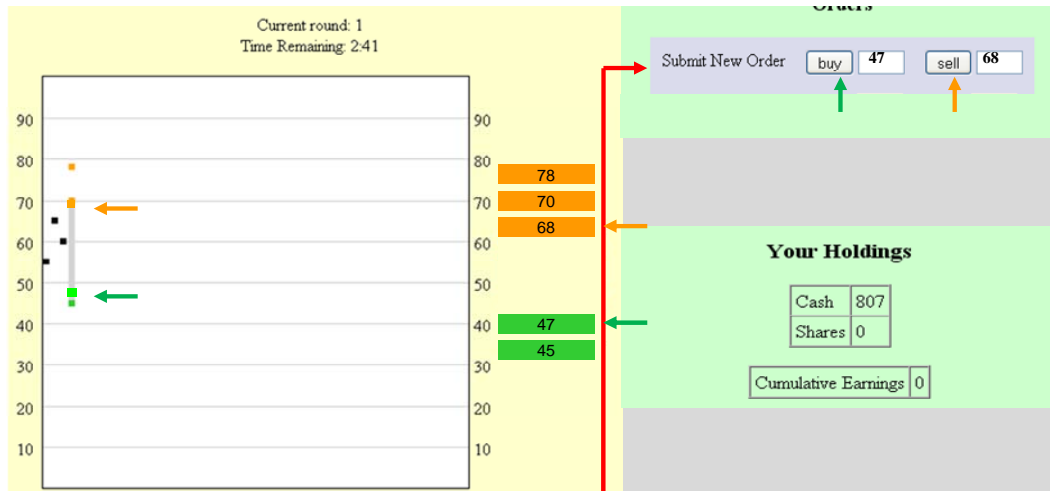
Cash	807
Shares	0
Cumulative Earnings 0	

On top of the graph, the **Current round** is shown.

Below that, the **Time Remaining** for the trading round is shown. Each round lasts several minutes. The vertical axis lists the **Price** for the offers.

Every time someone makes an offer to **buy** a share, a **GREEN** dot will appear on the graph to the left. Every time someone makes an offer to **sell**, an **ORANGE** dot will appear on the graph to the left. Once a **trade** is actually made, the trade will be shown as a **BLACK** dot in the graph.

Offers are also listed on the **Market Book** to the right of the graph.

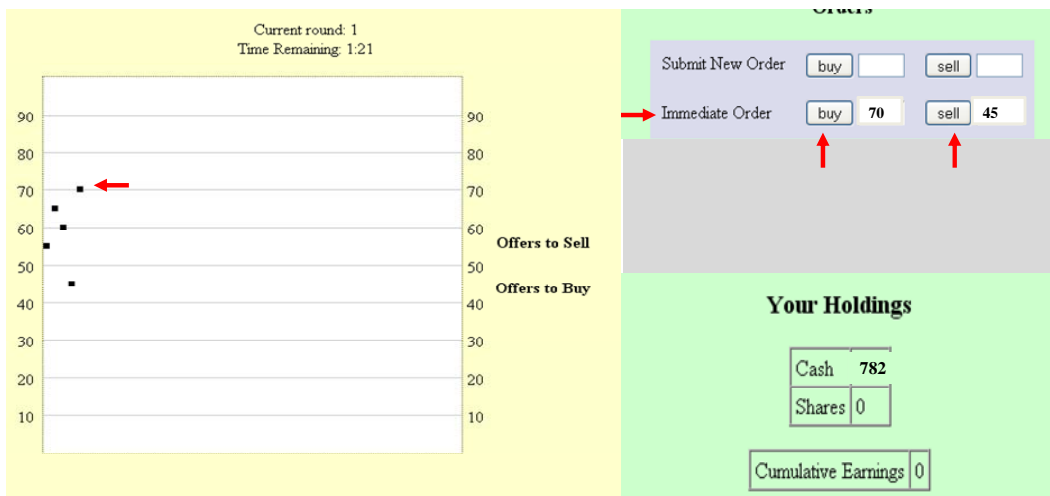


The top right section of the screen is the **Orders** box.

To enter a **New Order** to buy or to sell, Traders type in the price at which they would like to buy, or sell, in the appropriate **Submit New Order** box and click the **Buy** or **Sell** button to submit the order. Once the order is entered, the offer will be updated on both the **Market Book** and the **Market Graph**.

Suppose a Trader wants to place an order to buy, it must be higher than the current best offer to buy, which is now 45. Say, the Trader wants to buy at 47, she types in 47 and clicks buy.

Suppose a Trader wants to place an order to sell, it must be lower than the current best offer to sell, which is now 70. Say, the Trader wants to sell at 68, she types in 68 and clicks sell.



To accept an existing offer from another participant, a Trader can click the Buy or Sell button in the **Immediate Order** section above. The Immediate Order section shows the best prices to buy, or sell, that are currently available on the market.

By clicking on the **Sell** button, a Trader **sells** at the listed price.

The current best offer to buy is 45, if a Trader clicks Sell, she sells a share at the price of 45 immediately. Her shares go to -1 (she is short a share).

Her cash holdings will initially increase by 45, but her cash will then decrease by 100 as money is put in reserves to cover the share if the event is Black (pays 100). The net change in her cash is $+45(\text{price}) - 100(\text{in reserves}) = -55$. If the event turns out to be White, the 100 in reserves will be given back to her.

By clicking on the **Buy** button, a Trader **buys** at the listed price.

The current best offer to sell is 70. If a Trader clicks Buy, she immediately buys a share at a price of 70 (and 100 in reserves is returned to cash).



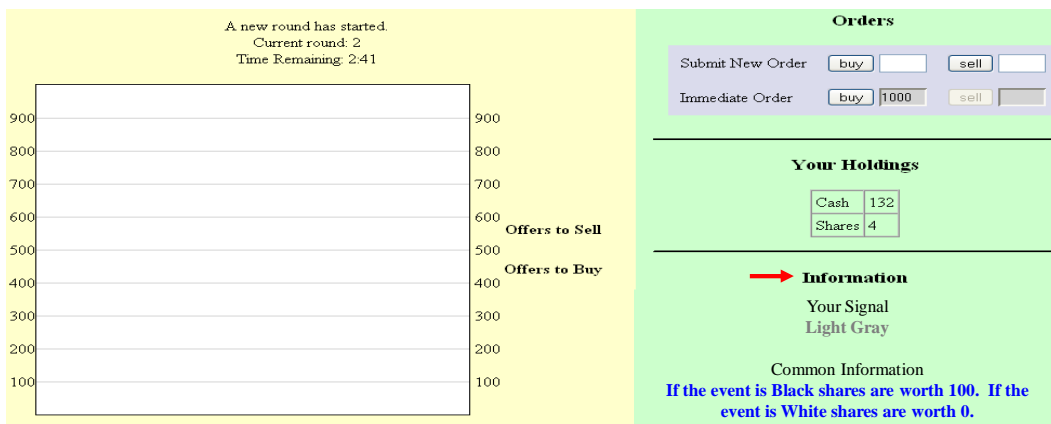
Whenever a Trader enters new offers to buy, or sell, she will have those offers appear as **Buttons** below the order box.

A Trader's outstanding offers to buy cannot exceed her cash holding; her outstanding offers to sell cannot exceed her ability to meet the reserves holdings.

Therefore, a Trader may have to delete offers under "**Cancel Orders**".

By clicking on these buttons, a Trader can take them off of the market.

Suppose a Trader clicks on the bid button 25, she will remove it from the market.



The **Information** section will provide Traders updates on the following:

- Her signal (it is light gray in this example);
- Common information (this tells her how the shares payoff);

The share earnings each round will be added to the cash account of the holder.

A Traders earnings will accumulate each period.

Her cash and shares do not carry over to the next round.

Forecasters

Some subjects are in the role of **Forecasters**. These subjects cannot trade in the market nor do they own any shares. Further, they do not receive any signal regarding the event that will occur (White or Black). However, these subjects can observe the market and see any bids, asks, or contracts that occur.

Forecasters have a budget of 100 each period to invest between Black and White. The more they invest in the event that actually occurs, the more money they earn. Their investment does not affect how much **Traders** earn.

Each round, Forecasters decide how much of their budget to invest in Black. Any money not invested in Black is automatically invested in White by the computer. So the amount invested in Black + amount invested in White = 100. Forecaster investments must be in increments of 5.

A Forecaster's payoff will be determined by the amount invested in the actual event, according to the accompanying table. This table is also provided to you at your workstation if you are a Forecaster.

Amount Invested in		Payoff if Event is	
BLACK	WHITE (=100 - Investment in Black)	Black	White
100	0	72.7	0
95	5	72.5	5
90	10	72	10
85	15	71	15
80	20	69.5	20
75	25	67.5	25
70	30	65	30
65	35	62	35
60	40	58.5	40
55	45	54.5	45
50	50	50	50
45	55	45	54.5
40	60	40	58.5
35	65	35	62
30	70	30	65
25	75	25	67.5
20	80	20	69.5
15	85	15	71
10	90	10	72
5	95	5	72.5
0	100	0	72.7

For example, if a Forecaster invests 75 in Black and the event is Black, he will earn 67.5. However, if the event is White he would earn 25 (as he had invested 100-75 = 25 in White).

Thus, the more confident a Forecaster is that the event will be Black, the more he should invest in Black. Similarly, the more confident he is that the event will be White, the more he should invest in White (by investing less in Black).

Amount Invested in		Payoff if Event is	
BLACK	WHITE (=100 - Investment in Black)	Black	White
100	0	72.7	0
95	5	72.5	5
90	10	72	10
85	15	71	15
80	20	69.5	20
75	25	67.5	25
70	30	65	30
65	35	62	35
60	40	58.5	40
55	45	54.5	45
50	50	50	50
45	55	45	54.5
40	60	40	58.5
35	65	35	62
30	70	30	65
25	75	25	67.5
20	80	20	69.5
15	85	15	71
10	90	10	72
5	95	5	72.5
0	100	0	72.7

Current round: 1
Time Remaining: 2:26

Offers to Sell	80
Offers to Sell	75
Offers to Buy	50
Offers to Buy	25

Your Earnings

Cumulative Earnings: 0

Amount to be invested in Black:

0 20 40 60 80 100

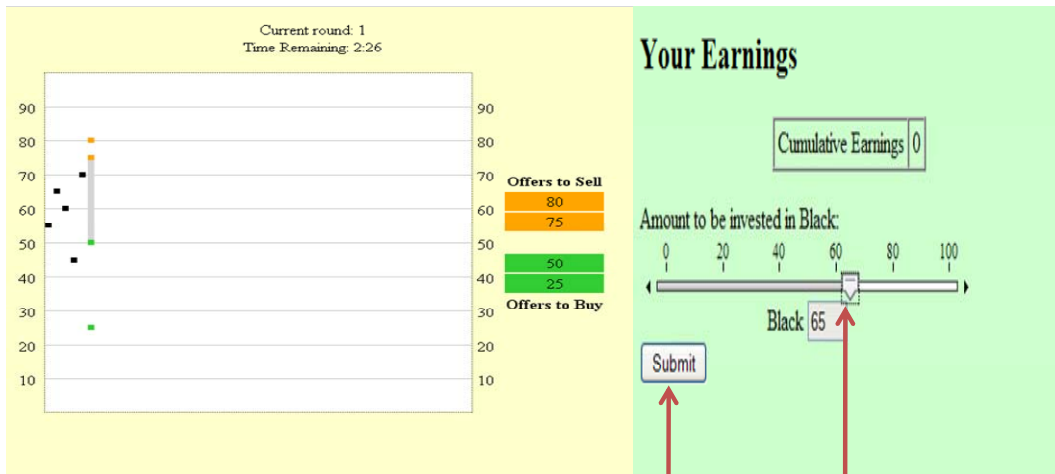
Black: 50

Submit

Now we will describe a Forecaster's screen.

At the start of each round the investment in Black will be 50 since Black and White are equally likely.

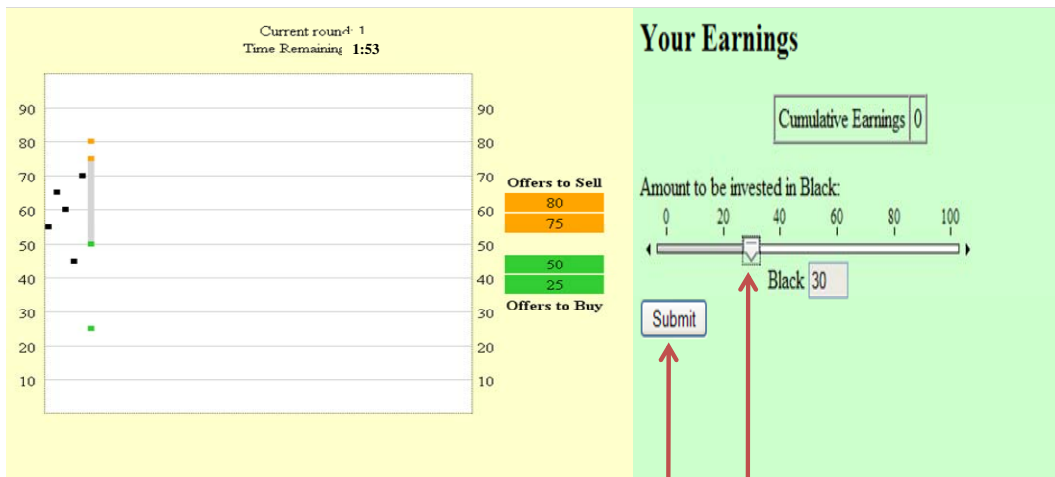
A Forecaster can change his investment as many times as he wants during a round, but only the last investment in a round will be used to determine his payoff. The amount of time available for investing is determined randomly. A Forecaster will not know how much time he has and it may be less than the time remaining in the market. Thus, Forecasters should keep their investment updated.



If during the round a Forecaster wants to change his investment, he can just move the Slider to the desired investment level and press Submit.

For example, if the Forecaster wanted to invest 65, he can just move the slider to 65 and then press submit.

His investment level would be 65 unless he changes it before his time runs out and his slider is disabled for the round.



The Forecaster can change his investment amount again, for example to 30.

Target Traders

A **Target Trader's** earnings are based upon the average amount that **Forecasters** invest in the *wrong* event. The more money **Forecasters** invest in the event that *does not* occur, the more a **Target Trader** earns. The more **Forecasters** invest in the event that does occur, the less a **Target Trader** earns. This is the only source of **Target Trader** earnings.

Target Traders are given an endowment of cash and can trade shares in the market. All of the rules describing how **Traders** can create shares, hold reserves, place and cancel offers, and accept trades apply to **Target Traders** as well. It is just that **Target Traders** are not paid based upon their cash or shares, They are only paid based upon the average amount invested in the wrong event by **Forecasters**.

There is no way for a **Trader** or a **Forecaster** to determine if an offer or an acceptance was made by a **Trader** or a **Target Trader** or if there are any **Target Traders** active in a given market round.

Summary for Traders and Target Traders

1. Traders and Target Traders will be given an initial amount of Cash and can create shares. For Traders, every share is worth 0 if the event is White and 100 if the event is Black. The event is equally likely to be White or Black. Thus, the average value is 50.
2. Each period, Traders will receive a signal about the event. If the event is White there is a two-thirds chance the signal will be light gray and a one-third chance the signal will be dark gray. If the event is Black there is a two-thirds chance the signal will be dark gray and a one-third chance it will be light gray.
3. Target Traders will know the actual event that will occur.
4. Traders and Target Traders can submit offers to BUY shares and offers to SELL shares. If a Trader or a Target Trader creates a share by selling a share when she does not have one, she will have to set 100 in a reserve account to make sure she can cover the payment to the share's owner if event is Black.
5. Trades can make trades by buying at the current lowest offer to sell or selling at the current highest offer to buy.
6. Target Trader earnings are only based upon the amount **Forecasters** invest in the event that does not occur. Traders earnings are only based upon the value of their shares and cash.

Summary for Forecasters

1. Forecasters will be given 100 to invest in Black and White. The more that is invested in the event that occurs, the more the Forecaster will earn. The amount the Forecaster will earn is shown in the accompanying table.
2. A Target Trader may or may not be present in a market round.
3. Forecasters do not receive a signal about the event, but can observe the market.
4. The amount of time a Forecaster has to adjust his investment is determined randomly and may be shorter than the time remaining in the market.

A practice round will be given to familiarize everyone with the program.

A short quiz follows. When you are ready to take the review quiz follow the appropriate link.

Traders press [Here](#) to go to their quiz.

Forecasters press [Here](#) to go to their quiz.

Target Traders press [Here](#) to go to their quiz.