

Can Manipulators Mislead Market Observers? □

□

Ryan Oprea
UC Santa Cruz

David Porter
Chapman University

Chris Hibbert
Google

Robin Hanson
GMU

Dorina Tila
GMU

Can Manipulators Mislead Market Observers?

Abstract

We study experimental markets where privately informed traders exchange simple assets, and where uninformed third parties are asked to forecast the values of these assets, guided only by market prices. Although prices only partially aggregate information, they significantly improve the forecasts of third parties. In a second treatment, a portion of traders are given preferences over the forecasts made by observers. Although we find evidence that these traders attempt to manipulate prices in order to influence the beliefs of observers, we find no evidence that observers make less accurate forecasts as a result.

1 Introduction

An important property of prices in asset markets is their ability to summarize dispersed information. Indeed, prices in secondary financial markets are important forecasting tools for primary investors and serve a vital role in guiding the allocation of capital. This property of prices has motivated practitioners in firms, governments and academia to establish markets for the express purpose of forecasting uncertain future events. These "prediction markets" have been used to forecast presidential elections, sporting outcomes, project completion dates, geopolitical events and futures sales of soon to be released products. These markets appear to be very good predictors of future events (Wolfers and Zitzewitz , 2004; Forsythe et. al. , 1992) and therefore may hold great promise as forecasting tools for decision makers in business and government.

Recently, critics have raised two concerns calling into question the usefulness of asset markets as decision making tools. First, there is some debate about the interpretability of prices as estimators of probabilities of future events. Although recently this debate has centered on theoretical concerns about aggregation, there exists long standing experimental evidence suggesting that asset market prices are often imperfectly efficient and therefore serve as imperfect probability estimates (Forsythe and Lundholm , 1990; Hanson et. al. , 2006). Wealth constraints, informational thinness, insufficient time for proper convergence, insufficient trader experience and even trader irrationality can result in imperfect information aggregation. In such contexts it is unclear whether and how successfully decision makers can use prediction markets as forecasting tools.

Second, some observers worry that asset markets may be vulnerable to price manipulation. If traders have preferences over decisions made based on asset prices, they may be willing to take losses in the market in order to influence prices and thereby alter observers' beliefs. Indeed it was, in part, such concerns that

lead Congress to shut down the Policy Analysis Market (Tucker and Meirowitz , 2004), a prediction market established by the defense department to aggregate intelligence data for policy makers.

We report an experiment designed to assess these concerns. In our experiment, traders with noisy signals regarding the value of an asset trade that asset in a double auction. A separate set of uninformed subjects observe the price series and, afterwards, attempt to forecast the asset's value. Although prices in these markets are imperfectly efficient (and are, in some cases, severely biased), observers manage to use them to significantly improve their forecasts.

In a second treatment, half of the population of traders are given preferences over observers' forecasts. These traders attempt to manipulate prices and in some cases succeed in mildly altering average contract prices. This manipulation, however, has no negative impact on the accuracy of observer forecasts. Our findings suggest that even if prediction markets are imperfectly efficient and subject to manipulation they, nevertheless, may be useful forecasting tools.

It is worth noting that, although our experiment was motivated by prediction markets, the problems we examine are relevant to naturally occurring asset markets as well. Our results suggest that primary investors observing secondary prices may be capable of anticipating and correcting for some biases in the price formation process when making forecasts. Thus even inefficient prices may serve as useful guides for capital allocation decisions. Moreover, our experiment suggests that it is difficult for traders to affect capital flows by manipulating asset market prices.

The aggregative properties of prices were first noted by Hayek (1945) and were formally examined by Muth (1961). In 1988, implementation and field experimentation with prediction markets began in the Iowa Electronic Market, a prediction market established to forecast the outcomes of political elections. (Wolfers and Zitzewitz , 2004) provides a useful overview of progress in the field.

Plott and Sunder (1988) conducted pioneering experiments studying information aggregation in asset markets. Though they reported aggregation in some of their markets, this aggregation was imperfect and severely sensitive to theoretically benign institutional features of the market in which trade took place. Follow up work by Forsythe and Lundholm (1990) examining factors affecting aggregation concluded, in part, that trader experience has a substantial effect on informational efficiency of prices. Recently Manski (2004) has argued that theoretically, even absent such factors, prediction market prices may do a poor job at summarizing dispersed information. Wolfers and Zitzewitz (2005) and Gjerstad (2005) have countered that, assuming moderate levels of risk aversion, prediction market prices can credibly be interpreted as probability estimates.

Prior evidence on the vulnerability of prediction markets to manipulation is mixed. A field experiment by Camerer (1998) failed to produce evidence of successful manipulation in racetrack betting markets. However, in a second field experiment Hansen et. al. (2004) successfully influenced prices in the Iowa

Electronic Market. Hanson and Oprea (2004) present a microstructure model in which manipulators create liquidity by trading and can therefore have net positive effects on price accuracy. Hanson et. al. (2006) conduct a laboratory experiment in which a subset of traders are given incentives to directly raise the level of the price. Though these traders do in fact attempt to manipulate prices, their incentives are common knowledge and they are countered by the rest of the market. Strumpf and Rhode (2003) provide evidence that 19th century efforts to manipulate betting markets also failed to negatively affect efficiency.

The remainder of this paper will be organized as follows. In section 2, we describe our experimental design. In section 3 we present our empirical questions and present corresponding results in section 4. We conclude the paper in section 5.

2 Experimental Design

In a baseline treatment, a subset of subjects (traders) trade a common value asset in a standard double auction after receiving noisy signals about the asset's value. A set of uninformed subjects (observers) watch the market and, using trading data, attempt to forecast the value. In a second treatment, half of the trading subjects are paid a bonus depending on how well the observers' forecast matches a private target. These traders therefore may have incentives to use the market to affect observer beliefs.

2.1 Information Environment and Trader Profits

Our experiment consist of a series of asset markets with 8 traders, each endowed with 200 in currency and 2 shares of a binary option (an all or nothing contract) with a common value, $v \in \{0, 100\}$. Traders in these markets are allowed to buy and sell shares in a standard double auction ¹, though the true value of these shares is uncertain. It is common knowledge, *ex ante*, that assets will be worthless ($v = 0$) or valuable ($v = 100$) with equal probability. Once the market closes, Trader i 's total earnings are:

$$\pi_i = C - \sum_{j=1}^{J_i} B_{ij} + \sum_{k=1}^{K_i} S_{ik} + V(N + J_i - K_i) \quad (1)$$

C = Endowed Cash (200)

V = Realized Value of the Security (0,100)

N = Endowed Shares (2)

J_i = Number of units trader i Buys in the Market

K_i = Number of units trader i Sells in the Market

B_{ij} = Price of Buy Contract j purchased by trader i

S_{ik} = Price of Sell Contract k sold by trader i

¹Note that traders are not allowed to short sell in these markets

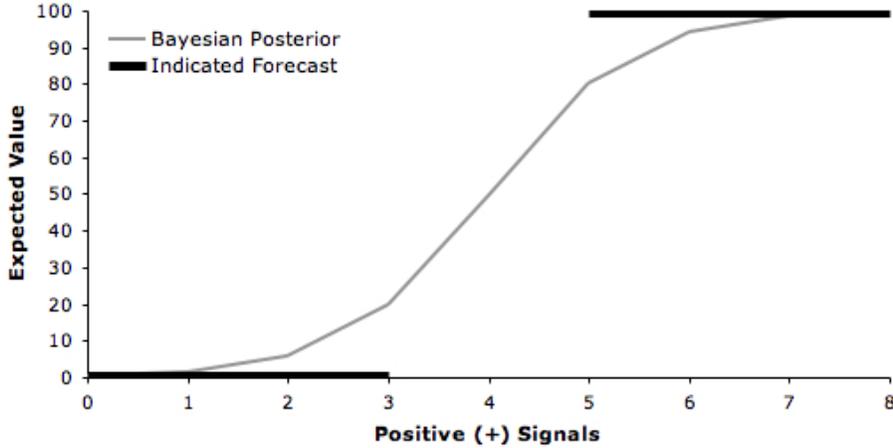


Figure 1: The expected value a fully informed Bayesian would assign to assets as a function of the number of positive signals distributed to the market. Dark black lines show the forecast a fully informed Bayesian would make.

Prior to a trade each trader, i , receives a private noisy signal $s_i \in \{-, +\}$ with replacement regarding the value of the asset. The market, overall is therefore equipped with a vector of signals, $\vec{s} = (s_1, s_2, \dots, s_8)$. The distribution from which signals are drawn depends on the true state, 0 or 100, of the security. If the true state is 0, each signal is drawn uniformly and independently from $\{-, -, +\}$; if the true state is 100, each signal is drawn similarly from $\{+, +, -\}$. Urn designs such as this one are common in information aggregation experiments. Our information environment is particularly similar to ones implemented in Anderson and Holt [1997] and Hung and Plott [2001].

Notice that, prior to receiving a signal, a rational Bayesian trader will assign an expected value of 50 to each share of the security. After receiving a signal, a Bayesian trader will revise this value to $\frac{1}{3} \times 100$ if the signal is - and about $\frac{2}{3} \times 100$ if the signal is +.

In a perfectly revealing rational expectations equilibrium, the market price should aggregate all of the information contained in the market, serving as a sufficient statistic for the 8 signals held by market participants. Such a price will be the expected value a Bayesian trader would assign to the asset if she had access to all 8 signals. This Bayesian price will therefore equal:

$$Ev(n) = \frac{\frac{1}{3} \frac{n}{3} 2^{8-n}}{\frac{1}{3} \frac{n}{3} 2^{8-n} + \frac{2}{3} \frac{n}{3} 1^{8-n}} \times 100 \quad (2)$$

where n is the number of + signals in the market. Bayesian prices are charted in Figure 1.

In this study we are interested in how information influences third party decision making through the

market price. It will therefore be useful to have a metric describing the degree to which the information behind market activity indicates one value realization versus the other. We construct a variable called signal strength, which indicates how consistent underlying information is but is independent of the direction of information. When $n = 4$, \vec{s} is uninformative in the sense that the prior estimate of the value and the posterior coincide. The further n is from 4, the stronger is the indication that the asset will take one value versus another. Denoting $x = (n - 4)$, $prob(100|x) = prob(0|-x)$, since the distributions are symmetric. We then define *signal strength* as $m = |x|$. A signal strength of 0 means that observation of signals would not change the beliefs of a Bayesian after viewing \vec{s} . A signal strength of 4 means that such a Bayesian would be virtually certain of the asset's value.

2.2 Observer Incentives

In addition to 8 traders inside the market, there are 5 *observers* outside the market who observe trade in the market and, after the market closes, make a forecast $f \in \{0, 100\}$. Observers have no money or shares and are unable to trade. Instead they earn money if they make the decision $f = v$, correctly forecasting the value of the asset based on their observation of market prices. Observer i 's earnings are

$$\pi_o = \begin{cases} 250 & \text{if } f = v \\ 0 & \text{if } f \neq v. \end{cases} \quad (3)$$

Observers are induced with the same priors as traders. However, unlike traders, observers are provided with no private information and must, instead, rely entirely on data from the market to guide their decisions.

In order to judge how well observers are informed by prices, it will be useful to have a benchmark to compare their forecasts. A natural benchmark is what we will call the *indicated value*. The indicated value, v^I , is simply the prediction in $\{0, 100\}$ a risk neutral Bayesian would make after observing all 8 draws. To be precise, the indicated value is:

$$v^I = \begin{cases} 100 & \text{if } prob(v = 100|\vec{s}) > 0.5 \\ 0 & \text{if } prob(v = 100|\vec{s}) < 0.5 \end{cases} \quad (4)$$

2.3 Manipulation Incentives

The environment described in sections 2.1 and 2.2 comprises our *Baseline treatment*. In it, a cohort of privately informed traders are allowed to use private information to trade in a standard double auction. A group of observers observe trade and, after the market closes, attempt to forecast the true value of assets based on the time series of prices they have observed.

In a second treatment, called *Forecast Preferences*, we provide half of the traders (4 traders) in the market with preferences over observers forecasts, f . Before the market opens these preference traders are assigned

the same random *target* t , drawn uniformly from $\{0, 100\}$. *Preference traders* are provided incentives to influence observers to make decisions, f , which approach this target. After trade, preference trader total earnings are equal to their earnings as a trader, π_t , plus a bonus that depends on the average forecasts of observers. Earnings for preference traders are:

$$\pi_p = \pi_t + 200 - 2\left|t - \frac{1}{5} \sum_{j=1}^5 F_j\right| \quad (5)$$

where F_j is the forecast of observer j . Preference traders, therefore, have some incentives to influence decision making and to influence it in a direction which is, *ex ante*, uncorrelated with the true value of the asset.

2.4 Discussion of Design

Our experimental design is optimized to study how asset market prices perform as forecasting tools when markets suffer from two pathologies. First, we introduce bias into the price by mirroring the distribution of assets and currency used in Hanson et. al. (2006). The limited number of shares distributed to subjects (combined with no short selling provision) limits the degree to which traders can influence prices downwards while ample currency makes bidding prices up relatively easy. The result in Hanson et. al. (2006) was a persistent bias in prices relative to perfect aggregation. Replicating this environment allows us to study how well observers foresee and adjust for these biases in prices when forecasting.

Second, our Forecast Preferences treatment is designed to give traders a strong opportunity and incentives to obscure information in the market price through manipulation. A key feature of the implementation of the treatment is that the target observer forecast is known neither to non-preference traders nor to observers. As a result non-preference traders cannot anticipate and automatically react to manipulation attempts as they do in Hanson et. al. (2006). Moreover, without knowing the incentives of preference traders, observers can expect noisier information coming from a successfully manipulated market, rendering it a less reliable tool. This feature of the design allows us to study whether unpredictable and unobserved manipulation campaigns might short circuit the use of and reliability of prices for the purpose of forecasting.

2.5 Experiments and Experimental Procedures

Each experimental session consisted of 16 periods,² with each period constituting a distinct market. In each period a new asset value was drawn, trader accounts were reset so that they had 200 in cash and 2 shares of the asset, traders were assigned new signals, preference traders received new targets and observers made new decisions once the market closed.

Subjects were recruited from the undergraduate population at George Mason University and given a \$5 show-up payment upon arrival. The subjects went through a detailed instruction period with written and oral

²Subjects were not told how many periods would be conducted, but were recruited for 2 hours.

Period	+ Signals	Indicated Value	Signal Strength	Preference Target, t	Actual Value
1	6	100	2	0	100
2	1	0	3	0	0
3	2	0	2	0	0
4	3	0	1	100	0
5	5	100	1	0	100
6	0	0	4	100	0
7	3	0	1	0	0
8	7	100	3	0	100
9	3	0	1	100	100
10	1	0	3	100	0
11	4	NA	0	0	100
12	2	0	2	100	0
13	6	100	2	100	100
14	4	NA	0	100	100
15	5	100	1	100	100
16	2	0	2	0	0

Table 1: Parameter design by period, providing the number of positive signals distributed, the forecast of a fully-informed risk-neutral Bayesian, the coherence of signals distributed to the market, the observer forecast that preference traders would prefer and the actual value of the asset. The indicated value is undefined (we write "NA") when only 4 signals are distributed to the market since, in this case, it is unclear what a risk neutral Bayesian would forecast.

instructions. Care was taken to demonstrate the information and incentive structure in the experiment.³ We used visual props to demonstrate the distribution from which the values and signals were drawn. In addition, two detailed practice periods, in which payoffs were not calculated in an individual's total earnings, were used to provide experience. The value of the security for each round was set to be either 0 or 100 according to a uniform random draw. At the start of each round, each trader is given a signal about the asset's value. The traders are informed that the signals were independently drawn, and with a 2/3 probability a + would be drawn when the value was 100, and with a 2/3 probability a - would be drawn when the value was 0.⁴

Between the two treatments we kept all elements in the design identical, including the value realizations and the signals to traders each period, except that, in the Forecast Preferences treatment, half of all traders were provided with a bonus based on the average of observers decisions each period. Table 1 shows, for each period in an experimental session, the information draws (number of + signals drawn), the indicated value, the signal strength, the preference traders' target (known only to preference traders), and the actual realization of the security value.⁵

³Instructions for the Baseline treatment can be found at <http://ices2.gmu.edu/dorina/No.Manipulation.Instructions/website11.htm> and for the Forecast Preferences treatment at <http://ices2.gmu.edu/dorina/Manipulation.Instructions/start.treatment1.htm>

⁴During the experiment subjects use e-cash, which was converted to US=dollars at a rate of 190 e to 1 for traders and 110 e = 1 for judges. The exchange rate used was known to the subject prior to the start of the experiment. Subjects average earnings per session were \$32.

⁵Traders, preference traders and observers interacted on networked computers using Zocalo, an open-source market software package (available at <http://zocalo.sourceforge.net>). The screen design and wording used on the screens can be found in the instructions.

3 Experimental Questions

Previous experimental studies on information aggregation in asset markets have focused directly on the relationship between trader information and the market price. Typically such experiments were designed to examine the degree to which various functions of the price series (typically the closing price) reflect or aggregate information provided to market participants. Although not our main research questions, we are interested in the degree to which the price transmits information underlying the market.

Question 1 *Do prices aggregate distributed information?*

A common concern about prediction markets is that traders with preferences over the beliefs of market observers may attempt to manipulate prices. By taking positions that either contradict or overstate their beliefs, a group of traders can conceivably cause observers to take on false beliefs about the direction or quality of information distributed to traders. Our Forecast Preferences treatment allows us to examine whether traders with preferences over the beliefs of observers attempt to manipulate these beliefs through the market price.

Question 2 *Do preference traders attempt to manipulate market prices?*

A direct extension of this question is whether, following Hanson et. al. (2006) manipulation attempts significantly interfere with information aggregation in prices.

Question 3 *Does manipulation damage the aggregative property of prices?*

Our main question is whether and how well observers use market prices to inform their forecasts. Prices in experimental prediction markets are often imperfectly efficient and cannot be straightforwardly interpreted as estimates of expected value. It is therefore unclear whether observers will make use of prices at all. Even if they do use prices as forecasting tools, it is difficult to know exactly how *ex ante*. We suspect that if observers do use prices to inform their forecasts, relatively simple functions of the price are unlikely to capture all of the information a decision maker is likely to glean from a price series. Including market observers as subjects endogenizes the problem allowing us to empirically pose this fourth question:

Question 4 *Do observers use information from markets to enhance their decision-making? If so, how much does the market improve forecasting?*⁶

It will be difficult to fully assess the damage done by price manipulation by looking directly at the effects of extra-market incentives on the market price. Even if traders who care about observer forecasts can affect

⁶Note that since observers have no private information, absent a market they cannot systematically make correct forecasts. Any deviation from random forecasting is evidence that markets improve forecasting

the market price through trade, it is not necessarily true that this will constitute manipulation in the sense intended here. Because it is difficult to discern how observers interpret market prices (as discussed above), it is difficult to know what manipulations of the market price will actually be effective in altering observer beliefs. Furthermore, it is unclear that would-be manipulators will correctly make this assessment and succeed in manipulating the market price. Our experimental design endogenizes the relationship between the trades of potential manipulators and the forecasts of observers and therefore allows us to pose a final question:

Question 5 *Do manipulators negatively impact the ability of observers to use prediction markets as a forecasting tool?*

4 Experimental Results

We report experimental results in three sections. In section 4.1 we provide evidence that preference traders do attempt to manipulate prices and succeed in affecting average contract prices. In section 4.2 we show that although prices in our markets correlate with the informationally efficient price, information aggregation is decidedly imperfect. Mirroring Hanson et. al. (2006) (in which subjects have similar levels of currency and shares of assets), our environment results in upwardly biased prices which imperfectly adjust to variation in underlying information. Price manipulation, however does not ultimately affect the quality of aggregation. In section 4.3 we provide evidence that, although prices are imperfectly efficient, observers use them to significantly improve the quality of their forecasts.

4.1 Price Manipulation

We begin by examining whether preference traders attempt to influence observer forecasts by manipulating the market price (Question 2). In order to test for price manipulation we compare price offers and acceptances made by preference traders with offers made by a similar pool of subjects from the Baseline treatment. Intuitively, if preference traders wish to convince observers that $v=100$, we suspect they will mimic the behavior of traders with positive expectations of value and make relatively high bids to buy. Likewise, preference traders wishing to convince observers that $v = 0$, they will make relatively low asks to sell.

To test for manipulation, we compare the bids, asks and acceptances of preference traders with those of a pool of subjects with identical information (and history of information) from the Baseline treatment. We estimate the following mixed effects model which is essentially a means test of offers across treatments, controlling for subject and session level effects.

$$price_{ijk} = \alpha_1 + \alpha_2 \times pref_i + u_i + \eta_j + \epsilon_{ijk} \tag{6}$$

Variable	Offers: t=0	Offers: t=100	Contracts: t=0	Contracts: t=100
<i>Intercept</i>	67.77177 (0.000)	55.68884 (0.000)	61.53712 (0.000)	60.53099 (0.000)
<i>pref</i>	-7.579531 (0.011)	7.621655 (0.045)	-1.853006 (0.344)	6.785543 (0.001)

Table 2: Mixed effects estimates from model (6) on four subsets of the data. Offers include bids in t=100 rounds and asks in t=0 rounds (and price acceptances in both). p-values are included below estimates in parentheses.

This model is estimated on two subsets of the data. The first is bids (and acceptances) in rounds in which t=100 and the second is asks (and acceptances) from rounds in which t=0. Thus $price_{ik}$ is bids in one regression and asks is the other. Indices are i for subject, j for session and k indexes trades within that session. $pref_i$ is an indicator variable that takes a value of 1 if session i is under the Forecast Preferences treatment and 0 otherwise, u_i is a random effect on subjects assumed to be distributed $N(0, \sigma_u)$, η_j is a random effect on sessions assumed to be distributed $N(0, \sigma_\eta)$, and ϵ_{ijk} is a disturbance term assumed to be distributed $N(0, \sigma_\epsilon)$. Estimates are presented in Table 2 (under the Offers column).

Under the hypothesis that forecast preferences have no effect on offers, the coefficient on $pref$ will be insignificantly different from zero. However a positive coefficient when t=100 and a negative one when t=0 will support the alternative hypothesis that traders respond to forecast preferences by attempting to manipulate prices. Indeed this is exactly what our estimates indicate. The coefficient on $pref$ is 7.62 when t=100 and -7.57 when t=0 and is statistically significant in both cases. This results in our first finding:

Finding 1 *Subjects with incentives over observer beliefs attempt to manipulate price. When t=100, bids are higher under the Forecast Preferences treatment than under the Baseline treatment. When t=0, asks are lower under Forecast Preferences than in the Baseline.*

A related question is whether these efforts to manipulate actually alter contract prices. In order to investigate this, we estimate (6) using only prices from the set of completed contracts. Once again we produce separate estimates for t=0 and t=100 rounds. Results are presented in Table 2 (under the Contracts columns). When t=100, manipulation efforts succeed in raising prices by 6.78 (p=0.001). Manipulation appears to have no corresponding effect on prices in rounds in which t=0 (p=0.344).

Finding 2 *Attempts to raise the price through manipulation succeed in raising prices by nearly 7 points; attempts to lower prices fail.*

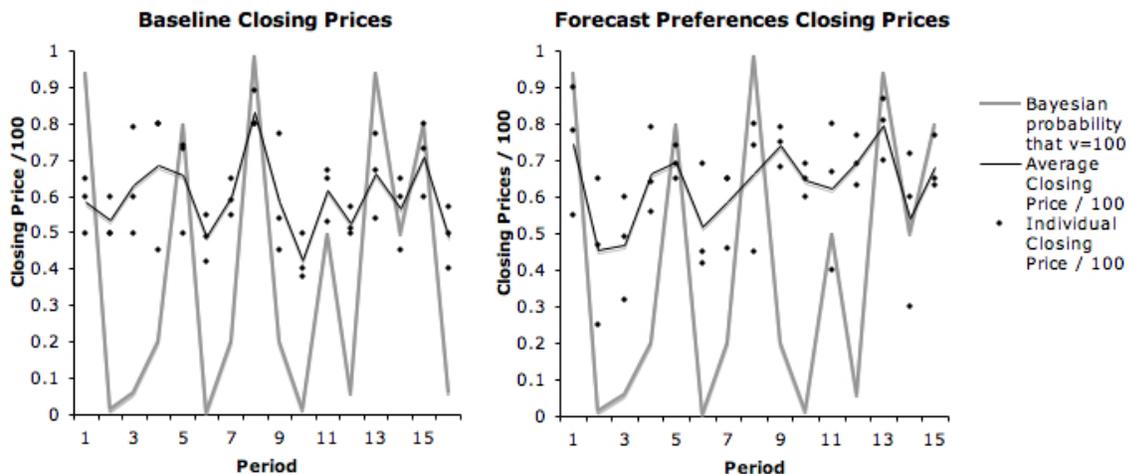


Figure 2: Observed closing prices (dots) and the average of these prices across sessions (black line) with fully informed expected values (gray line).

4.2 Information Aggregation in Closing Prices

When a market fully aggregates information, its contract prices approach the expected value of assets, given all of the information in the market. We are interested therefore (Question 1) in estimating aggregation by looking at the relationship between the prices we observe in our markets and the expected value a fully informed Bayesian would assign to the asset. Figure 2 shows the fully informed expected value of assets and the average closing prices for each period, broken up by treatment. Though it is clear that closing prices are strongly correlated with information, it is also clear that they fail to adjust fully to changes in probabilities from period to period. Hanson et. al. (2006) report similar qualitative but incomplete correlation between expected value and price in an experimental prediction market with similar currency and asset distributions but a very different information environment.

We formally support these observations by estimating the following random effects regression:

$$price_{jt} - 50 = \alpha_1 + \alpha_2 \text{pref}_j + \beta_1 (Ev(n)_j - 50) + \beta_2 \text{pref}_j (Ev(n)_j - 50) + u_j + \epsilon_{jt}. \quad (7)$$

The dependent variable, $price_{jt} - 50$ is the price amount of the t 'th contract in treatment j , relative to the uninformed price estimate of 50. The variable pref_j is an indicator variable for the Forecast Preferences treatment, $Ev(n)$ is defined as in (2) and u_j is a random effect on session and ϵ_{jt} is an error term (both assumed to be normally distributed). If the market aggregates perfectly in the Baseline treatment, the intercept α_1 will be indistinguishable from zero and the coefficient on expected value, β_1 , will be equal to 1.

Variable	Q1	Q2	Q3	Q4	F10
<i>Coefficient</i>	10.20671 (0.008)	11.01617 (0.002)	11.20923 (0.000)	12.30581 (0.000)	13.03491 (0.001)
<i>pref</i>	2.324719 (0.669)	1.206104 (0.812)	2.00208 (0.652)	1.849452 (0.643)	1.598616 (0.781)
$Ev(n) - 50$.0431189 (0.099)	0.0147719 (0.511)	0.0854649 (0.000)	0.1636089 (0.000)	0.1955946 (0.000)
$pref \times (Ev(n) - 50)$	-0.0107977 (0.756)	0.0147755 (0.624)	-0.0124539 (0.669)	-0.0296694 (0.32)	-0.0397304 (0.372)

Table 3: Results from random effects model (7) estimating the relationship between the expected valuation of a fully informed Bayesian to the price. Estimates are given for the four consecutive quarters of contracts (Q1-Q4) and for the final ten percent of contracts (F10). p-values are included below estimates in parentheses.

We estimate this model on 5 subsamples of the data including the first, second, third and fourth quarters of trades (Q1, Q2, Q3 and Q4) and the final 10% of trades. Results for each of these sub-samples are shown in Table 3.

A first observation is that the value of the coefficient on $Ev(n)_i$ increases over the four quarters of trade and takes its highest value during the final 10 percent of trades, indicating that the market price increasingly incorporates distributed information over time. During the final half of trade (Q3, Q4, F10) the slope term on expected value is significantly greater than zero, indicating a relationship between the efficient price and the realized price. Moreover, in all phases of trade, the intercept term is significantly greater than zero indicating a bias in the price. This is clear from Figure 2 in which prices rarely fall below 50, even when much of the information in the market is negative.

By the end of trade (F10) there is a significant relationship between the full-information expected value and the market price (β_2 is significant), though at 0.19 the coefficient still lies significantly below 1. Further, the intercept coefficient, α_1 , is estimated at 13, meaning that, even when the market is equipped with only negative signals, the price fails to predict a low enough value realization (a price near 50, taken as a probability estimate, simply reiterates the markets priors).⁷

We summarize these observations in our third finding:

Finding 3 *Over the course of trade, prices increasingly correlate with underlying information. However, even towards the end of trade they fail to fully aggregate and therefore cannot be straightforwardly interpreted as probability estimates. These markets have particular trouble aggregating negative information (information indicating a low value) resulting in prices which are upwardly biased relative to fully aggregated prices.*

We can also use the estimates from (7) to assess the effects of manipulation on aggregation (Question

⁷The upward bias in price we observe may be due, in part, asymmetries in subject endowments; each subject could only sell two units, but had enough cash to buy about three to four units at the observed prices. Similar bias was observed in (Hanson et. al. , 2006)

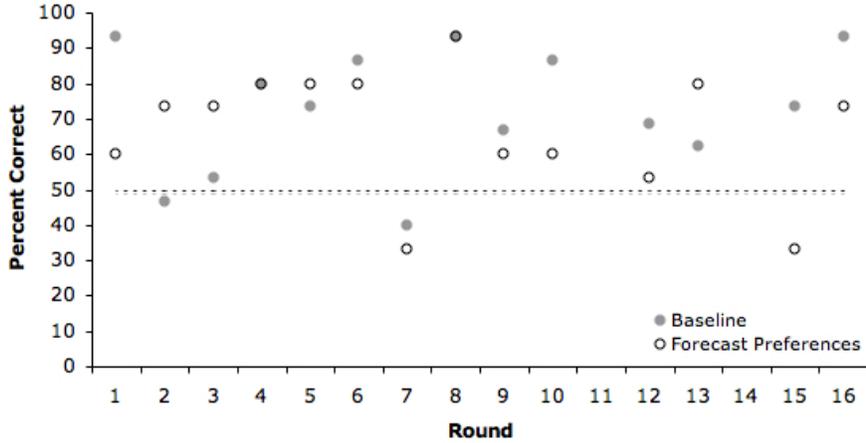


Figure 3: Proportion of forecasts matching the indicated value in each period and for each treatment.

3). If the interaction terms or indicator variables for the Forecast Preferences treatment are significantly different from zero, we have evidence that the relationship between price and expected value is affected by manipulation. In fact, these interaction terms are never significantly different from zero, indicating that aggregation is not weakened by preference traders.

Finding 4 *Manipulation has no effect on aggregation.*

4.3 Observer Forecasts

After observing the full price series in a market, do observers, in fact, tend to make decisions that reflect signals in the market (Question 4)? Figure 3 charts the proportion of observers' predictions that match the indicated value in each period for each treatment. It is clear that (i) in most periods the indicated value is predicted more than 50% of the time (observations are generally above the dotted line) and (ii) there is little difference between data in the Baseline treatment and the Forecast Preferences treatment. This figure therefore provides qualitative answers to both Questions 4 and 5: Observers do use information from prediction markets to enhance their decision-making and preference traders do not affect observers decision-making abilities.

To answer these first questions more formally, we estimate the following random effects probit model (without intercept):

$$\text{prob}(f_{ij} = v_j^I | \vec{s}) = \text{prob}(\eta_i + \alpha_1 b_i + \alpha_2 \text{pref}_i + m_j(\beta_1 b_i + \beta_2 \text{pref}_i) > \epsilon_{ij}) \quad (8)$$

where f_{ij} is the forecast by observer i in period j , v_j^I is the indicated value, v^I (the forecast a risk neutral

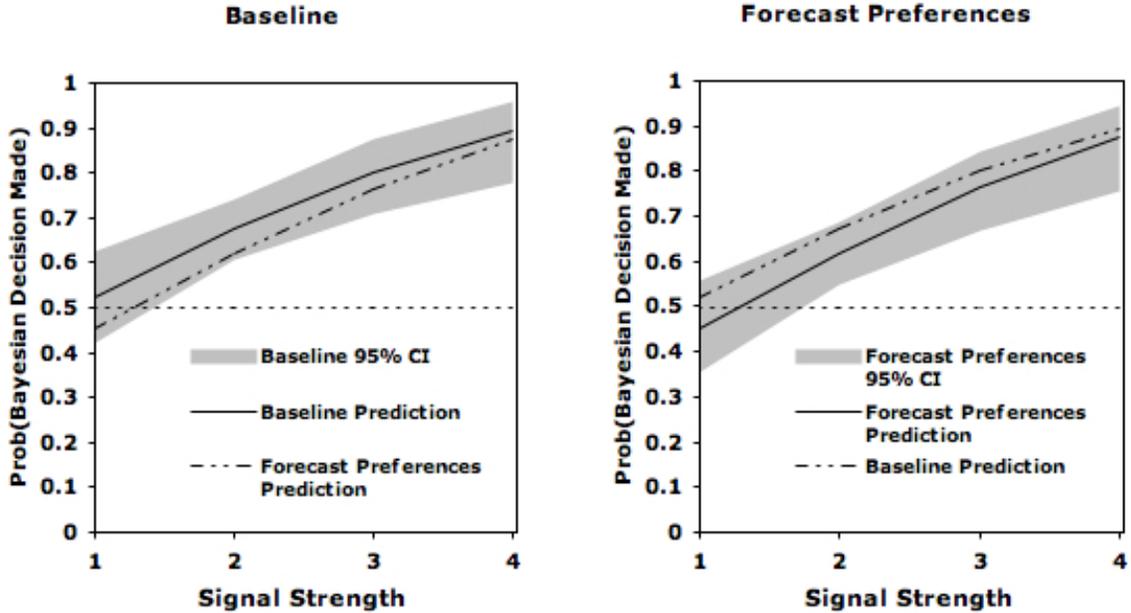


Figure 4: Predictions from probit model (8), showing, for each treatment, the estimated probability the observer forecasts the indicated value as a function of signal strength. Gray shaded areas show the 95% confidence interval. The dotted line on the Baseline (Forecast Preferences) panel shows the Forecast Preferences (Baseline) estimates for easier comparison.

Bayesian would make, defined in 7) in period j given \vec{s} , b is an indicator variable taking a value of 1 under the Baseline treatment, pref_i is an indicator variable taking a value of 1 under the Forecast Preferences treatment, m_j is the signals strength in period j , η_i is a normally distributed random effect drawn by subject and ϵ_{ij} is an error term assumed to be distributed normally with mean 0 and variance 1.

Table 4 provides the estimates from this model and Figure 4 shows the resulting predicted probability of the indicated decision being made at each level of signal strength in the Baseline and Preference treatments. The gray areas around these estimates give 95 % confidence intervals on the predictions.

The coefficient on signal strength in the baseline treatment, β_1 , is significant and positive, indicating that observers forecast the indicated value more frequently when the information in the market is more consistent. From Figure 4, it is clear (with 95% confidence) that, when the signal strength is greater than 1, the probability that observers forecast the indicated value is greater than .5 (the confidence interval does not contain 0.5). However, the same is not true when the signal strength is 1. Since 0.5 is the prior distribution over states, there is no evidence in this case that observers gather usable information from the prediction market. We sum up these observations as a fifth finding.

Finding 5 *Asset market prices significantly improve the accuracy of observer forecasts, but only when the*

Variables	Coefficients (p-values)	Marginal Effects
<i>pref</i>	-0.5381658 (0.013)	-0.197
<i>b</i>	-0.3386443 (0.123)	-0.125
<i>pref</i> × <i>m</i>	0.4208571 (0.000)	0.156
<i>b</i> × <i>m</i>	0.3957784 (0.000)	0.146

Table 4: Probit estimates from model (8).

signal strength is sufficiently high. Observers are more likely to forecast the indicated value at higher signal strengths.

We finally examine whether the manipulation efforts reported in findings 1 and 2 have a negative effect on the accuracy of observer forecasts.

In order to do this we use a Chow-type test suggested in Greene (2002) testing the hypothesis that the intercept and slope term on signal strength jointly differ across treatments. We first estimate the following random effects probit model⁸: $prob(F_{ij} = I_j | \vec{s}) = prob(\eta_i + \alpha + s_i \beta > \epsilon_{ij})$ once with all of our data and once each for the Baseline and Forecast Preferences treatments. Under the hypothesis that there is no difference across treatments, twice the difference between the sum of the likelihoods in the individual treatment estimates and the likelihood from the pooled estimate will be chi-square distributed with 1 degree of freedom. Summing our likelihoods in this way, we arrive at a test statistic of 1.5 (p=0.220), evincing no significant difference between the treatments. A simple (though extremely conservative) alternative test of such a difference across the treatments can be accomplished visually using the predictions in Figure 4. The dashed lines in each panel of Figure 4 show predictions from the other treatment. Predictions from the preference treatment fall, at each level of signal strength, within the 95 % confidence interval of the baseline prediction. Thus, our data fails to produce evidence that manipulation significantly alters observer forecasts. This provides us with a sixth finding.

Finding 6 *Manipulation does not reduce the accuracy of observer forecasts.*

5 Discussion

We conducted laboratory experiments studying how well asset markets function as forecasting tools. Uninformed third party observers were asked to forecast a stochastic event after watching trade in a market

⁸Variables have the same interpretation as corresponding variables in (8)

populated by partially informed traders. In one treatment, selling constraints lead to severe biases in the price formation process. Still, observers tended to correctly forecast the event using the market data. In a second treatment, half of the traders in the market were given preferences over the forecasts of observers. These subjects attempted to manipulate prices by submitting more extreme orders than normal traders. These manipulation attempts, however, failed to negatively impact the accuracy of observer forecasts.

Two main conclusions can be drawn from our study. First, inefficient asset prices do not necessarily imply equally inefficient capital allocation decisions. Our subjects seem to exhibit rational expectations regarding pathologies in the price formation process and at least partially filter these pathologies when forming forecasts. The result is forecasts that are superior to those that would be made by simply interpreting prices as efficient probability estimates. Second, our results suggest that price manipulation is a largely ineffective strategy for altering the beliefs of investors and other decision makers. Several features of our experimental design were included to favor successful manipulation. In each market, half of the traders were given coordinated incentives to manipulate and these incentives were sufficient to significantly alter the orders these subjects submitted. Moreover, other traders and observers were kept in the dark about the direction in which manipulation incentives ran, making it difficult to counteract or discount their effects on prices. Still, the manipulation effort we observe failed to reduce the accuracy of observers' forecasts.

Both of these conclusions merit further study and robustness tests. Experiments designed with even stronger incentives to manipulate seem especially appropriate. While it seems likely that this would lead to more intensive manipulation attempts, it is unclear whether this would significantly alter the aggregation and forecasting conclusions reported here. Further research also seems warranted on the correspondence between secondary market prices and the efficiency of capital allocation. While our results suggest that rational expectations can limit the degree to which inefficient prices mislead primary investors, it is unclear how far this reasoning can be taken. Experiments varying the efficiency of market prices (perhaps by manipulating share endowments) would provide a clearer picture of the relationship between price efficiency and forecast efficiency. It also seems fruitful to study forecasting in markets which are inefficient for less predictable reasons than order constraints. One obvious candidate is markets which are subject to endogenous price bubbles ala (Smith, Suchanek and Williams, 1988).

References

- Allen, F., Gale, D., 1992. Stock-Price manipulation. *The Review of Financial Studies* 5, 503-529.
- Anderson, L. and Holt, C. Information cascades in the laboratory. *American Economic Review*, December 1997, 87(5), pp. 847-62.

- Camerer, C., 1998. Can asset markets be manipulated? A field experiment with race-track betting. *Journal of Political Economy* 106, 457-482.
- Chakrabortya, A., Ylmaz, B. , 2004. Manipulation in market order models. *Journal of Financial Markets* 7, 187-206.
- Forsythe, R., Nelson, F., Neumann, G., Wright, J., 1992. Anatomy of an experimental political stock market. *American Economic Review* 82, 1142-1161.
- Forsythe, R., Lundholm, R., 1990. Information aggregation in an experimental market. *Econometrica* 58, 309-347
- Gjerstad, S. , 2005. Risk aversion, beliefs and prediction market equilibrium. mimeo University of Arizona.
- Greene, W., 2002. *Econometric Analysis*, 5th ed. (Prentice Hall, New York, NY).
- Hansen, J., Schmidt, C. , Strobel, M. , 2004. Manipulation in political stock markets – preconditions and evidence. *Applied Economics Letters* 11, 459-463.
- Hanson, R., Oprea, R., 2004. Manipulators increase information market accuracy, George Mason University Working Paper.
- Hanson, R., Oprea, R. and Porter, D., 2006. Information aggregation and manipulation in an experimental market *Journal of Economic Behavior and Organization* 60, 449-459.
- Hayek, F., 1945. The use of knowledge in society. *American Economic Review* 35, 519-530.
- Kumar, P., Seppi, D. , 1992. Futures manipulation with 'cash settlement'. *Journal of Finance* 47, 1485-1502.
- Manski, C. , 2004. Interpreting the predictions of prediction markets. NBER Working Paper 10359.
- Muth, J., 1961. Rational expectations and the theory of price movements. *Econometrica* 29, 315-335.
- 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.
- Plott, C., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56, 1085-1118.
- Smith, V., Suchanek, S., and Williams, A., Bubbles Crashes and Endogenous Expectations in Experimental Spot Asset Markets, *Econometrica*, 56(5), 1988
- Strumpf, K., Rhode, P., 2003. Historical presidential betting markets. *Journal of Economic Perspectives*, 18 127-141.

Tucker, J. Meirowitz A., 2004. Learning from terrorist markets. *Perspectives on Politics*, 2 331-6.

Wolfer, J., Zitzewitz, E., 2004. Prediction markets. *Journal of Economic Perspectives*, 18 107-126.

Wolfer, J., Zitzewitz, E., 2005. Interpreting prediction market prices as probabilities. NBER Working Paper 12200.