Security Design and Information Aggregation in Markets*

Yiling Chen Anthony M. Kwasnica

Abstract

It has been well-recognized that markets can aggregate less-than-perfect information across market participants. With two differently designed securities, this work examines the impact of security design on the information aggregation ability of markets in laboratory experiments. Results show that markets with one security aggregate information significantly better than markets with the other security, implying that information aggregation ability of markets is affected by the security design. Behavior of individual participants is then investigated to understand the observed market behavior.

JEL Classification: C92; C91; D80

Keywords: Security design; Information aggregation; Information market; Price convergence.

1 Introduction

Although the information aggregation ability of markets has been supported by a long history of theory and evidence [10, 12, 13, 20], it is only recently, with the advent of the Internet and other information and communication technologies, that information aggregation becomes the main purpose and function of certain markets. Such markets are called *information markets* or *prediction markets*. They are designed specially for information aggregation and prediction. To achieve these goals, information markets associate payoffs of securities with outcomes of well-defined future events, and provide marketplaces to trade the securities. For example, in an information market to predict whether the Democratic Party will win most votes in the next presidential election, the

*The authors sincerely thank the eBusiness Research Center (eBRC) and the Laboratory for Economics Management and Auctions (LEMA) at The Pennsylvania State University for their support on this work.

security pays a certain amount of money (e.g. $1) per share to its shareholders if and only if the Democratic Party gets most votes in the election. Otherwise, it pays nothing. Thus, before the election, the security price may reflect the expectations of market participants about how likely Democratic Party is to win the election.

Predictions provided by information markets have proved to be more accurate than or as accurate as polls, individual expert opinions, or aggregated expert opinions in many domains, including politics [2, 3, 8, 9], business [4], and entertainment [18, 5]. However, as information aggregation becomes the primary function of information markets, it is important to understand what factors or conditions facilitate or prevent information aggregation.

Past studies in this direction are represented by a series of experimental works that utilized controlled human-subjects experiments. Plott and Sunder [19, 20] found that when participants had homogeneous preferences, which means that the value of the security given the state of the nature is the same for all participants, information was successfully aggregated in double auctions. Forsythe and Lundholm [7] further showed that if participants had heterogeneous preferences, experience of participants is a necessary condition for information aggregation. Kagel and Levin [15] experimented with sealed bid auctions where information was not aggregated in their settings. Guarnaschelli, Kwasnica, and Plott [11] used the same information structure as that of Kagel and Levin [15], but paired it with double auctions, and obtained somewhat more successful information aggregation. These studies partly revealed how individual preferences, structure of information, experience of participants, and market mechanism influenced the information aggregation ability of markets.

The research reported below is an attempt to examine the effect of another factor, design of securities, on information aggregation using laboratory experiments. Since information markets can have as novel securities as one can imagine, it is essential to know whether different designs of securities affect markets' ability to aggregate information. The information market model proposed by Feigenbaum et al. [6] touches the relationship between security design and information aggregation theoretically. According to it, when the security traded in an information market satisfies certain conditions, information aggregation is guaranteed regardless of the structure of information. Otherwise, there exists some information structures such that information aggregation can not be achieved. Sensing that the theoretically model inevitably makes simplified assumptions, such as non-strategic behavior of market participants, we take an empirical approach to investigate to what extent information aggregation is affected by security design in our experimental markets.

Our main result is that security design does affect the information aggregation ability of markets. Experimenting with two differently designed security, we observed significantly better information aggregation with one security than with the other. Evidence suggests that individual behavior does not conform to the assumptions made by Feigenbaum et al.'s model [6], which indicates room for improvements. We briefly discuss potential changes to the model that may produce a more realistic model.

The paper is organized as follows. Section 2 contains the detailed design of the experimental markets created for this study. Section 3 outlines Feigenbaum et al.'s model and its projections for our experimental markets. Experimental results and analysis are reported in Section 4. Section 5 is a summary of our conclusions.

2 Design of Markets

Our market design closely follows the model of Feiganbaum et al. [6], which will be introduced in Section 3.3. We design two securities such that markets with one security should be able to aggregate information according to the theoretical model of Feiganbaum et al., while markets with the other security should not.

The markets were conducted as a series of trading periods. All periods in one experimental session were identical except for the information that participants received. There was one security traded in each market. The security paid a dividend of either 150 experimental dollars or 50 experimental dollars at the end of a period, contingent on which of the two states of nature, "Good" or "Bad", was realized. Each market consist of five participants. Each participant received a clue or signal about the nature of the state at the beginning of a period. The clue was either a "G" or a "B", and was provided by the computer through a random draw. It was equally likely for a participant to get a "G" clue or a "B" clue. Each participant knew his/her own private clue and the fact that other participants had independently drawn clues.

Two different securities that we examined were a majority security and a parity security. If a market traded a majority security, the state of nature was "Good" if and only if the majority of participants (i.e., three or more in a market with five subjects) got a "G" clue. Otherwise, the state of nature was "Bad." When a parity security was traded, the state of nature was "Good" if and only if there were odd number of "G" clues in the market (i.e., one, three, or five "G" clues). For both securities, the dividend was 150 experimental dollars when the state was "Good", and 50 experimental dollars when the state was "Bad".

Trading periods were designed as modified Sharpley-Shubik market games [21]. Each period consisted of at least 2 and up to 10 trading rounds each lasting a maximum of 2 minutes. [1] At the beginning of each period, each participant was given 1 unit of the security and 50 units of cash in experimental dollars. At each round, each participant was asked to enter a bid into the computer indicating the price at which the participant wanted to buy or sell a unit of the security. The bid was required to be between 50 and 150. A round ended when all traders had submitted bids or 2 minutes had elapsed. The market price for the round was then calculated by taking the average of all bids in a round. At the market price, total demand equaled total supply and the market cleared. All transactions happened at the market price. The net quantity traded for a participant was given by:

$$\text{Quantity Traded} = (\text{Bid of the participant - Market Price})/\text{Market Price}.$$

If the bid of the participant was higher than the market price, the participant would buy securities. If the bid of the participant was lower than the market price, the participant would sell securities. In general, the further a bid of a participant was from the market price, the more the participant bought or sold. Negative security or cash holdings were allowed. The trading period then proceeded to the next round where each participant could submit a new bid. After a minimum of 2 rounds, the trading period ended if any one of the following happened: (1) no traders submitted a bid in the last round, (2) everyone's bids are the same, (3) the market price did not change for two consecutive rounds, or (4) ten rounds were completed. At the end of a trading period, the true state of nature and true dividend were announced. Profit of a participant was the sum of his/her cash inventory and dividend from security holdings. [2]

[1] In a Sharpley-Shubik market game, a trading period ends when the market price or traders' bids do not change from round to round. Feigenbaum et al.'s model predicts that a trading period will end in the second or third round. Since conducting experiments that have infinite number of rounds was not practical, we set the maximum number of rounds in a trading period to be 10. Since 10 rounds is far greater than the predicted number of rounds for any theory, it was our feeling that it was more than sufficient to capture any convergence activity.

[2] In a Shapley-Shubik market game, each participant is asked to put up a bid b_i and a quantity q_i in each round,

Eleven sessions were conducted in the Laboratory for Economic Management and Auctions at The Pennsylvania State University. Five of the sessions used the majority security. The other six sessions used the parity security. The experiments are implemented using the z-Tree (Zurich Toolbox for Readymade Economic Experiments) software [22]. Each session had five Penn State undergraduates or graduates as subjects. All subjects did not have previous experience with this particular experimental setting, but some of them were familiar with other market institutions. All subjects were given the opportunity to familiarize themselves with the experimental procedure by participating in a practice trading period. Each experimental session consisted of eight trading periods and lasted for about an hour and a half. A trading period was essentially a separate instance of market, with the clues of participants reseted.

All experiments were conducted using experimental dollars. At the end of each session, the profit of each participant was converted to U.S. dollars at the conversion rate of 80 experimental dollars per 1 U.S. dollar. This amount plus a $7 show-up fee were paid in private to the participant. A copy of the instructions is available from the authors upon request.

3 Theoretical Models

Two models that represent two extremes are the fully revealing rational expectation equilibrium (FRRE) model and private information (PI) model. The FRRE model assumes that market participants behave as if they know all available information in the market. At FRRE equilibrium, information of market participants is fully aggregated into the equilibrium price. PI model makes the assumption that market participants only use their private information in trading. Market price at PI equilibrium does not reflect the aggregated information. According to both models, security design does not affect the information aggregation of markets. Both models are static in the sense that they only conjecture the equilibrium state of the market. As static models are usually subject to the criticism of not explaining how equilibrium is reached, Feigenbaum et al. [6] propose a dynamic process and establish the connection between the two extremes. In the Feigenbaum

where b_i is the amount of money that the participant want to spend on buying securities and q_i is the unit of security that the participant want to sell. Market price clears demand and supply of securities. Our modified Shapely-Shubik game is equivalent to always setting q_i as 1. In other words, participants are required to put 1 unit of the security for sale in each round.

et al.'s model, market participants at the beginning of the market use their private information in forming their expectation of the dividend, which is the same as projected by the PI model. But market participants also learn from prices of previous rounds, and revise their expectations in later rounds. If the security design satisfies certain conditions, the market will eventually reach the FRRE equilibrium. Otherwise, there are situations where the FRRE equilibrium can not be reached. In the following of this section, we briefly introduce the FRRE and PI models, and then visit the Feigenbam et. al's model. We examine the implications of these models for the outcomes in our experimental setting.

3.1 Rational Expectations

In a market that participants have the same or symmetric information, a competitive equilibrium is a price and net trade, at which a market clears. In other words, a competitive equilibrium is reached when demand equals supply. The market price at a competitive equilibrium reflects the preferences and budget constrains of market participants. Rational expectation equilibrium [17] extends the concept of competitive equilibrium to situations where participants may have asymmetric information. With asymmetric information, future utilities of items traded are uncertain and participants may have different information about future state. Market activities may reflect the information of participants in addition to their preferences and budget constrains. Thus, market prices potentially provide informational feedback for market participants.

At a FRRE equilibrium, the equilibrium price, reveals all available information. Hence, all participants in equilibrium behave as if they know the pooled information of all participants in the market. Allen [1] and Jordan [14] have shown that FRRE equilibrium generically exist.

In our experimental setting, the FRRE model implies that market participants know the total number of "G" clues in the market. Therefore, they know for certain what the dividend will be at the end of the period. The bid of each participant and hence the market price according to this model would be equal to the dividend of the security. Table 1 lists the equilibrium price predictions of the RE model when markets have different numbers of "G" clues. RE model is an extreme since it assumes that individual behavior of market participants is based on all information in the market.

Table 1: Equilibrium Price Predictions – FRRE Model

Security	Number of "G" Clues in a Trading Period					
	0	1	2	3	4	5
Majority	50	50	50	150	150	150
Parity	50	150	50	150	50	150

3.2 Private Information

The private information (PI) model, also called the prior information model, is propsed by Plott and Sunder [20], while a similar model is used by Lintner [16] in analyzing US securities market. As opposed to the FRRE model, the PI model assumes that market participants do not condition expectations upon market price. In stead, market participants apply Bayes rule to update the likelihood of the future outcome upon receiving their private prior information, and act according to the derived probability.

According to the PI model, if a risk neutral trader gets clue "G" in a market with a majority security in our experiments, his/her expectation of the security payoff v is

$$E(v|G) = 150 \times P(v = 150|G) + 50 \times P(v = 50|G) = 118.75.$$

If the trader gets a clue "B", his/her expectation of the security payoff is

$$E(v|B) = 150 \times P(v = 150|B) + 50 \times P(v = 50|B) = 81.25.$$

Assume that market participants truthfully bid their expectations. If the market price clears the market, it will be the average of expected security payoffs of all participants.

With a parity security, a participant's expectation of the security payoff is always 100 in our experiments regardless of his/her private information, because the probability that the dividend is 150 is always 0.5, i.e.

$$P(v = 150|G) = P(\text{There are 0, 2, or 4 "G" clues in the rest of the market}) = 0.5, \text{and}$$

Table 2: Equilibrium Price Predictions – PI Model

Security	Number of “G” Clues in a Trading Period					
	0	1	2	3	4	5
Majority	81.25	88.75	96.25	103.75	111.25	118.75
Parity	100	100	100	100	100	100

$$P(v = 150|B) = P(\text{There are 1, or 3 “G” clues in the rest of the market}) = 0.5.$$

Being the average of all bids, the market price for a parity security should always be 100 according to the PI model.

Table 2 shows the price predictions of PI model if market participants bid their expectations.

3.3 Dynamic market model

Both the FRRE and PI models are static models that provide an equilibrium prediction, Feigenbaum et al.'s model is a dyanmic model that connects the two static models.

The dynamic market model of Feigenbaum et al. [6] takes a computational approach. It treats a market as a computational device. In a market with n traders, each trader holds a piece of private information, whose value is either 0 or 1. Let $x \in \{0,1\}^n$ denote the information vector $(x_1, x_2, ...x_n)$. Then, x is the input to the market, and equilibrium price p is the output. The key component of the device is the security F whose payoff is a Boolean function of the private information, $f(x) : \{0,1\}^n \rightarrow \{0,1\}$. The market mechanism is modified Sharpley-Shubik games as described in Section 2.

Feigenbaum et al. studied under what conditions the market can aggregate the information of market traders and correctly compute the value of the function $f(x)$. With the assumption that in each round market participants truthfully bid their expectations conditional on their private information and market price of previous rounds, the main result of Feigenbaum et al. is that

> If the function f is a weighted threshold function, then the market will take at most n rounds of trading to reach an equilibrium, where the equilibrium price of F is equal to $f(x)$. If f is not a weighted threshold function, then there exists a prior probability

Table 3: Price Predictions – Feigenbaum et al.'s Model

Security	Round	Number of "G" Clues in a Trading Period					
		0	1	2	3	4	5
Majority	1	81.25	88.75	96.25	103.75	111.25	118.75
	2 and all other	50	50	50	150	150	150
Parity	1	100	100	100	100	100	100
	2 and all other	100	100	100	100	100	100

distribution of x for which the price of the security F does not converge to the value of $f(x)$.

A function $f : \{0,1\}^n \rightarrow \{0,1\}$ is a weighted threshold function if and only if there are real constants $w_1, w_2, ..., w_n$ such that $f(x) = 1$ iff $\sum_{i=1}^{n} w_i x_i \geq 1$.

Our majority security is based on a weighted threshold function, and our parity security is not. The implications of Feigenbaum et al.'s model in our experimental setting are that the price of the majority security will converge to its true dividend in the second round of trading, while the price of the parity security will not converge to its true dividend.

With a majority security, market participants bid their expectations of the security payoff only conditional on their private information in the first round of trading. As shown in Table 2, the correspondence between the total number of "G" clues and the market price based on private information is a one-to-one mapping. Hence, from the price of the first round, market participants can infer how many "G" clues there are in the market, and incorporate the information in forming their expectations in the second round. The market then reaches the fully revealing rational expectation equilibrium in the second trading round. With a parity security and the prior probability distribution of clues (equally likely to get a "G" clue or a "B" clue) in our experiments, market participants will not be able to get any extra information from the market price as shown in Table 2. They can only use their private information to bid in each round. Thus, the market does not aggregate information.

Table 3 shows the price predictions according to Feigenbaum et al.'s model. An equilibrium is reached within two rounds. Hence, prices in later rounds stay the same.

Table 4: Total Number of "G" Clues in Experiments

Security	Session	Trading Period							
		1	2	3	4	5	6	7	8
Majority	1	0	2	2	3	2	3	3	3
	2	3	3	3	1	3	3	3	4
	3	2	3	4	2	3	2	2	3
	4	2	3	1	2	3	2	3	2
	5	3	3	2	2	2	2	2	1
Parity	6	1	2	3	3	2	1	2	3
	7	4	0	2	2	4	1	3	3
	8	4	3	2	1	3	3	4	3
	9	3	3	1	4	5	2	4	2
	10	1	3	3	3	2	2	4	2
	11	3	3	1	3	3	3	2	2

4 Results

Table 4 summarizes the actual number of "G" clues that subjects received in each trading period of our experiments. We first investigate the aggregated behavior of our experimental markets to examine to what extend security design affects information aggregation. Then, we study individual behavior of market participants to obtain better understanding of the observed market behavior.

4.1 Aggregated Behavior

When information of market participants is fully aggregated, the market price of a security equals to the security payoff (dividend). Thus, we use how far the market price deviates from the dividend as a measure for how well information is aggregated. Results 1 and 2 show that markets with the majority security aggregate information better than markets with the parity security. Prices of the majority security are closer to the dividend on average and move toward the dividend to a larger degree than prices of the parity security.

Result 1: On average, markets with a majority security have less price errors than markets with a parity security.

Support: Let p_{ijk} be the market price for trading round k in trading period j of the experimental session i, and d_{ij} be the dividend of trading period j of experimental session i. Price error for trading round k in trading period j of experimental session i is defined as the absolute difference between p_{ijk} and d_{ij}, i.e.

$$e_{ijk} = |p_{ijk} - d_{ij}|.$$

We calculate the price errors of every round and period in an experimental session i, and use the average,

$$e_i = \frac{\sum_j \sum_k e_{ijk}}{\text{total number of rounds in session } i},$$

to represent the price error for the experimental session i. Table 5 lists the average price error e_i for all experimental sessions i from 1 to 11. Applying the non-parametric Wilcoxon-Mann-Whitney test to the data in Table 5, we test the null hypothesis that the median of the average price errors for the majority sessions is equal to that for the parity sessions, with the alternative hypothesis being that the median for the majority sessions is less than that for the parity sessions. The resulted test statistic rejects the null hypothesis at the significance level of 0.05 (W = 19, p-value=0.0276). This implies that market prices deviate from true dividends less in majority sessions than in parity sessions.

Table 5: Average Absolute Difference of Price from Dividend

Security	Session	Average Price Error (e_i)
Majority	1	39.96
	2	46.38
	3	43.39
	4	28.24
	5	47.90
Parity	6	46.03
	7	44.36
	8	55.02
	9	52.61
	10	54.98
	11	56.18

Result 2: Prices in markets with the majority security converge toward the true dividend to a larger degree than those with the parity security.

Support:

The price error of a particular trading round, the absolute difference of market price from dividend, reflects how well information is aggregated at that time. If information is fully aggregated, price error should be 0. If no individual information is aggregated and the market price only incorporates the common prior of the outcome, the price error should be around 50 (which means that the price is around 100). Loosely speaking, the closer the price error is to 0, the better the information aggregation is. Thus, we examine both the first round and the last round price errors for all majority and parity trading periods to check whether the markets move toward the direction of reducing price errors.

Dividing the interval between 0 and 100 into four subsets, [0, 25], (25, 50], (50,75], and (75, 100], we count the total number of trading periods whose price errors fall into the same subset using first round and last round market prices respectively. The frequencies of price errors for all subsets are shown in Table 6 with majority markets and parity markets listed separately.

Table 6: Frequency of Absolute Difference of Price from Dividend

Round	Period	Price Errors			
		[0,25]	(25,50]	(50, 75]	(75, 100]
First	Majority (out of 40)	6 (15%)	21 (52.5%)	9 (22.5%)	4 (10%)
	Parity (out of 48)	4 (8.33%)	17 (35.42%)	18 (37.5%)	9 (18.75%)
Last	Majority (out of 40)	17 (42.5%)	9 (22.5%)	11 (27.5%)	3 (7.5%)
	Parity (out of 48)	5 (10.42%)	19 (39.58%)	21 (43.75%)	3 (6.25%)

From Table 6, we can see that first round price errors for most majority and parity trading periods fall into the subsets (25, 50] and (50, 75]. This implies that both majority and parity markets start with a similar level of price errors. For the last round, however, 42.5% of majority trading periods have a last round price error that is less than or equal to 25, while only 10.42% of parity trading periods fall into the same range. Although the first round price errors for both majority and parity trading periods are similar, there are more majority trading periods that have less last round price errors.

We apply Chi-Square tests to examine whether the observed differences of frequencies are statistically significant. Table 7 displays the results of our four Chi-Square tests. Each table cell contains the value of the test statistic and p-value of a Chi-Square test for two rows of Table 6, indicated by the row and column names. The null hypothesis is that the frequencies of price

Table 7: Results of Chi-Square Tests on Frequencies of Price Errors

	Majority (First Round)	Parity (Last Round)
Parity (First Round)	5.059 (0.168)	3.453 (0.327)
Majority (Last Round)	10.404 (0.015)	12.619 (0.006)

* The first number of each table cell is the value of the Chi-Square statistic. Numbers in parenthesis are the p-values of the tests.

errors for the two rows are the same. For example, the first table cell tests whether the observed differences of first round price errors between majority periods and parity periods as shown in the first two rows of Table 6 are significant. From the test results and frequencies of price errors in Table 6, we can conclude in a statistical sense that (1) majority trading periods and parity trading periods have the same level of price errors at the beginning of markets; (2) majority trading periods move toward the direction of reducing price errors, while parity trading periods do not; and (3) majority trading periods have less price errors than parity trading periods at the end of the markets.

Although market prices for the majority security move toward the true dividend, they do not converge to the dividend as fast and accurately as the Feigenbaum et al.'s model predicted. We plot the actual market prices of the last period (period 8) against the predicted prices of three models for majority markets and parity markets in Figure 1 and Figure 2 respectively. Using the last period prices reduces errors caused by inexperience of traders, and was done by Plot and Sunder [20]. For majority markets, the first round price of Feigenbaum et al.'s model is the same as that of PI model, while prices in later rounds according to the Feigenbaum et al's model are the same as those of FRRE model. For parity markets, the predictions of PI model and Feigenbaum et al.'s model are the same. Upon casual inspection, the actual market prices were different from predictions of all three models most of the time, although in some cases of majority markets (e.g. sessions 1 and 4) we can see clear trends of prices slowly moving toward the FRRE prices. We do not observe any obvious trends of prices in parity markets. Result 3 supports our observation statistically.

Result 3: Actual market price is significantly different from the price predictions of all models in

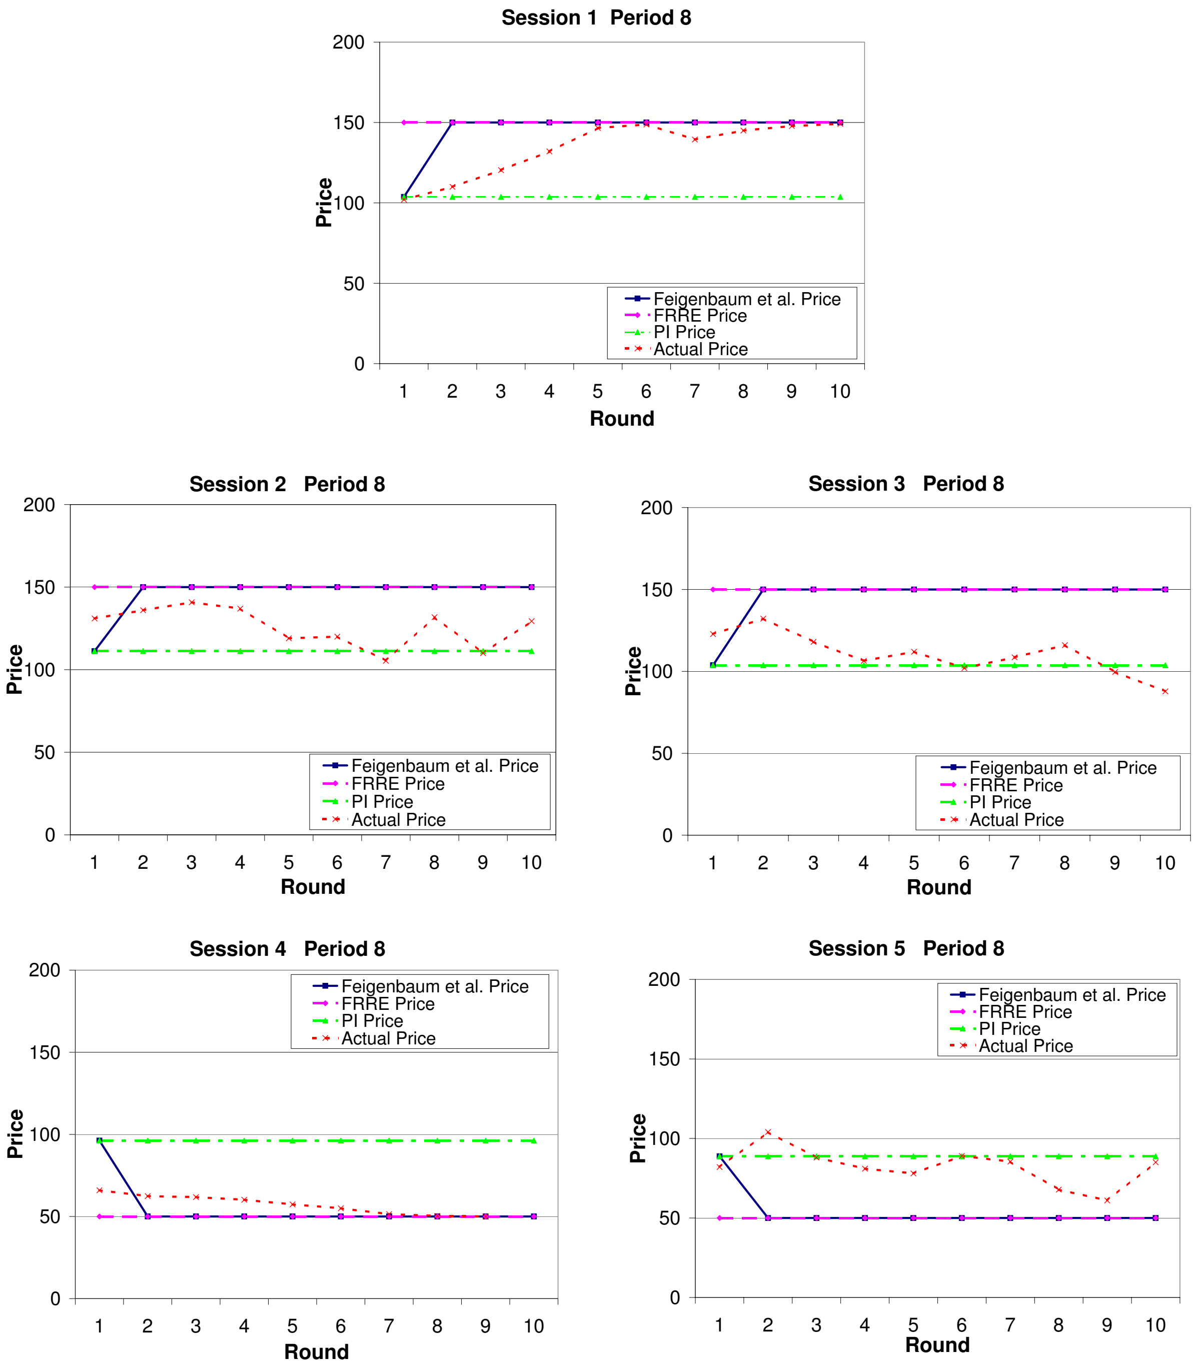

Figure 1: Price Comparison for Markets with Majority Security (Period 8)

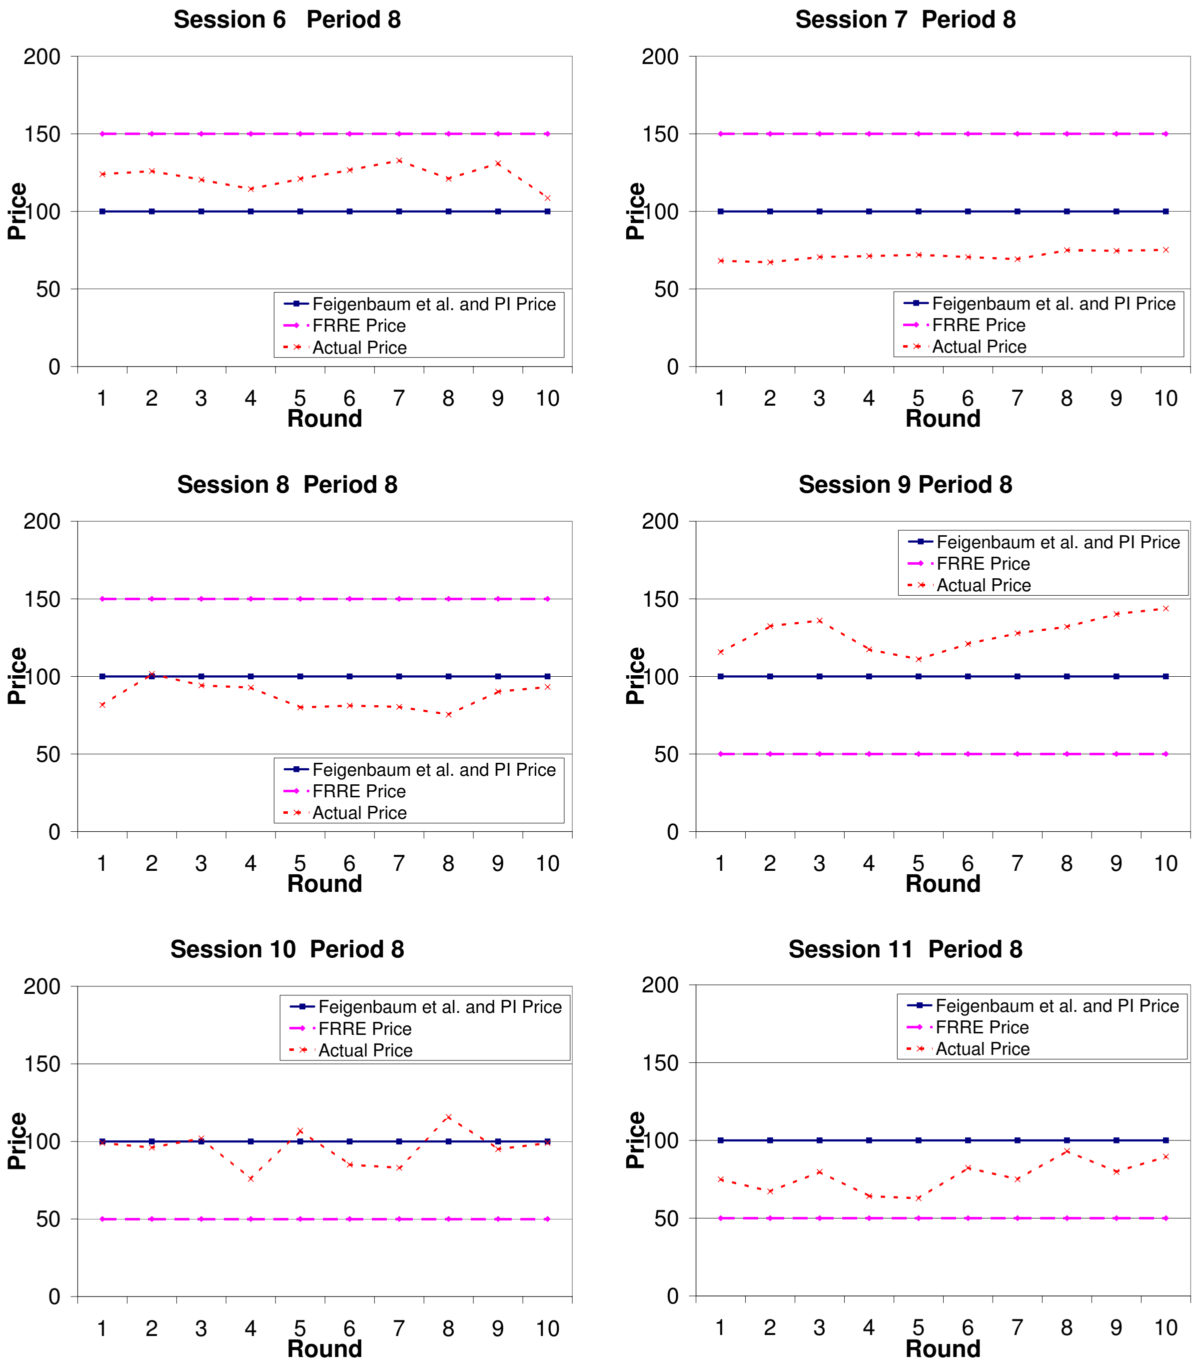

Figure 2: Price Comparison for Markets with Parity Security

both majority and parity markets.

Support:

If a model can accurately predict the market prices, the absolute difference between actual market prices and model predicted prices should roughly equal to zero. Let p_{ijk} be the market price for trading round k in trading period j of the experimental session i and m^l_{ijk} be the price predicted by model l for the same trading round, the average absolute difference between market prices and model l predicted prices for trading period j of the experimental session i is calculated as:

$$f^l_{ij} = \frac{\sum_k |p_{ijk} - m^l_{ijk}|}{\text{total number of rounds in period} j}. \tag{1}$$

For each experimental session i and each theoretical model l, we test the null hypothesis that the mean of f^l_{ij} (j across 8 trading periods) is equal to zero using a one-sample t test. Our test statistics indicate that, for all eleven experimental sessions (five majority and six parity) and all three theoretical models, the mean of f^l_{ij} is statistically different from zero at the significance level of 0.05. This indicates that Feigenbaum et al.'s model can not be claimed as accurate in an absolute sense, neither can the FRRE and PI models be. Such result is typical in experimental studies, since simplified assumptions of theories are usually not completely conformed in laboratories. Plott and Sunnder [20] and Guarnaschelli, Kwasnica, and Plott [11] all found that market behavior relative to predictions of theoretical models, including FRRE and PI, differed substantially.

4.2 Individual Behavior

The theoretical results of Feigenbaum et al.'s model, and the other models to some extent, rely on two critical assumptions about individual behavior. First, it is assumed that market participants are rational in the sense that they can correctly calculate conditional expectations and make inferences based on market prices. All participants know that all participants are rational, and know that all participants know this, and so forth ad infinitum. Second, the model assumes that participants are not strategic; they truthfully bid their conditional expectations in each rounds. Thus, in a market with the majority security, a participant bids 118.75 in the first round if his/her private clue is "G", and bids 81.25 if his/her private clue is "B". From the first round market price, rational participants can infer how many "G" clues and "B" clues are there in the markets, and get to know the true dividend. In the second round, participants will then bid the true dividend, knowing the total number of "G" clues in the market. Hence the market price converges to the true dividend in the second round. On the other hand, in a market with the parity security, participants' conditional expectations of the security value are always 100 regardless of their private clues. The market price does not reveal private information. The market hence does not converge to the true dividend.

However, the results of Feigenbaum et al.'s model, are not robust when these assumptions are violated. Even if there is only one participant who is not fully rational or does not bid truthfully, the market may fail to converge. The following two results are centered on observed individual behavior in our experimental markets. They show that these assumptions are substantially violated.

Result 4: For both majority and parity markets, most market participants do not bid the conditional expectations of the security value in first rounds of trading.

Support:

In a market with the majority security, a market participant's conditional expectation of the security value is 81.25 if he/she receives a "B" clue, and 118.75 is he/she has a "G" clue. In a market with the parity security, a market participant's conditional expectation of the security value is 100 regardless of the clue that the participant gets. We examine whether market participants actually bid their conditional expectations of the security value in first rounds. For majority markets, we put some buffer around the expected security value – a participant is considered to bid his/her conditional expectation if his/her bid falls into the interval [77.5, 85] with a "B" clue, or [115, 122.5] with a "G" clue.

In our experiments, there are 200 first round bids for majority markets, and 240 for parity markets. Only 13 out of 200 (6.5%) first round bids in majority markets are consistent with the buffered conditional expectation of the security value. 37 out of 240 (15.42%) first round bids in parity markets have the value 100. Most first rounds bids in both majority and parity markets deviate from the market participants' conditional expectations of the security value. These conclude that individual behavior assumptions of Feigenbaum et al.'s model are substantially violated, which might account for why the model is not fully supported by our experimental data.

Result 5: For majority markets, most market participants do not infer the direction of security value and bid accordingly in first rounds.

Support:

In majority markets, an individual may not correctly (or do not want to) calculate the conditional expectation of the security value, but he/she may know that the expected security value is lower than 100 if his/her first round clue is “B”, and higher than 100 if his/her first round clue is “G”, and bid toward the expected direction. In order to capture such possible bidding behavior of market participants, we use the following exclusive categories to classify first round bids.

1. Bid 100 in the first round.

   In majority markets, biding 100 means that a participant ignores his/her private clue and only relies on the common prior. In parity markets, bidding 100 is the behavior based on the assumptions of Feigenbaum et al.’s model, but it can also be a result of ignoring private clues.

2. Bid lower than 100 with clue “B”.

   For majority markets, this and the next category imply that market participants use their private clue to infer the expected direction of the dividend.

3. Bid higher than 100 with clue “G”.

4. Bid lower than 100 with clue “G”.

   We consider both this category and the next one as against the theoretical assumptions on individual behavior.

5. Bid higher than 100 with clue “B”.

Table 8 presents the number and percentage of first round bids for each category. We notice that 47.5% of first round bids in majority markets get the direction right (categories 2 and 3). However, even in parity markets, there are 42.08% of first round bids, which is approximately an equal proportion as in majority markets, fall into categories 2 and 3. Since neither a “G” clue or a “B” clue alone provide useful information for the value of a parity security, there is no clear evidence to support that participants of majority markets use their information to infer the direction of dividend and bid accordingly in first rounds of trading.

Table 8: First Round Bidding Behavior of Market Participants

Behavior Category	Majority		Parity	
	Number	Percentage	Number	Percentage
1	46	23%	37	15.42%
2	60	30%	77	32.08%
3	35	17.5%	24	10%
4	38	19%	72	30%
5	21	10.5%	30	12.5%

Category 1: Bid 100 with "B" or "G"; *Category 2*: Bid lower than 100 with "B";
Category 3: Bid higher than 100 with "G"; *Category 4*: Bid lower than 100 with "G";
Category 5: Bid higher than 100 with "B";

Result 6: Market participants in general tend to bid lower than expected security value in first rounds of trading.

Support: From Table 8, we observe that there are more lower-than-100 bids than higher-than-100 bids (49% vs. 28% in majority markets and 62.08% vs. 22.5% in parity markets) in first rounds of trading. To take into accout of different clues that participants receive, we plot the empirical culmulative distributions (CDF) of first round bids for different clues and different securities in Figure 3.

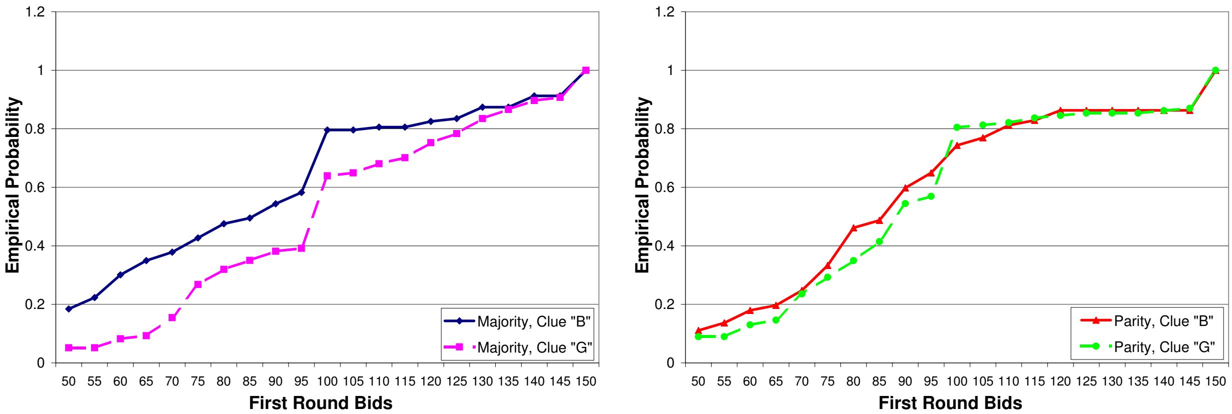

(a) Majority Markets (b) Parity Markets

Figure 3: Empirical Cumulative Probability Distribution of First Round Bids

The lower curve in Figure 3 (a) is the CDF of first rounds bids for those participants who receive a “G” clue in majority markets. More than 70% of these first round bids are lower than 118.75, the expectation of the security value conditional on a “G” clue in majority markets. For parity markets, where observing a clue “B” or a clue “G” does not help refining the prior probability of the state of nature, there are about 75% to 80% percent of first round bids that are lower than 100 as shown in Figure 3 (b). These indicate that market participants are systematically bidding low in first rounds of trading.

Bidding lower than expectations could be a kind of strategic behavior of market participants. In a Shapley-Shubik market game, an individual's bid can affect market price. Hence, if an individual intents to buy the security, he/she has the incentive to bid lower to some degree so that the market price decreases and he/she can buy the security at a lower price. On the other hand, if an individual wants to sell, he/she also has the incentive to bid lower, since it will enable them to sell more units of the security. Another possible explanation of bidding low is the risk aversion of market participants. Assuming others bid the same, bidding lower can result in less buying or more selling for a participant. In either case, it reduces the security holding of the participant, and hence is less risky.

Another observation of Figure 3 is that, while the CDF of first round bids with “G” clues and that with “B” clues are very similar for parity markets, they are clearly different for majority markets. The CDF with ”G” clues is consistently lower than the CDF with “B” clues in majority markets, which implies that first round bids in general are lower when participants receive “B” clues than when they receive “G” clues. Thus, to some degree, first round bids and hence the market price provide some information about participants' private clues in majority markets, which helps information aggregation. This observation indirectly supports our Results 1 and 2.

Result 7: For both majority and parity markets, participants adjust their bids according to observed market prices.

Support:

Results 4 to 6 reveal that in our experimental markets, participants most often do not bid their conditional expectations of security value. This deviation makes market prices less informative and inferring the total number of “G” clues from market prices hard. Based on this understanding, we start to examine whether market participants still appears to learn from prices.

Let b_t be the bid of an individual in round t, and p_t be the market price of round t, we calculate the correlation coefficients between bid adjustments, $b_t - b_{t-1}$, and the observed price-bid differences, $p_{t-1} - b_{t-1}$, for every participant of our experiments. The correlation coefficients are all positive. 45 out of 55 participants have a correlation coefficient that is greater than 0.5. The positive correlation means that market participants tend to adjust next round bids toward the direction of the observed previous round market prices. In other words, if a participant's bid is lower (higher) than the market price in the previous round, the participant will increase (decrease) his/her bid in the next round. Market participants still believe that market price contains some information of other participants and follow it. It is worth of noting that, price in a parity market may not be informative according to Feigenbaum et al.'s model, so following the price probably do not benefit market participants.

5 Conclusion

This work is an attempt to study the effect of security design on the information aggregation ability of markets in a controlled experimental market environment, in which the behavior of individuals are not well understood. The basic results of our study are as follow:

1. Security design affects the information aggregation ability of markets. Markets with the majority security on average have less price errors and move toward the dividend of the security to a larger degree than markets with the parity security.

2. The effect of security design on information aggregation is less strong than what is predicted by Feigenbaum et al.'s model. None of the theoretical models is supported in an absolute sense. Market prices in our experimental markets deviate from price predictions of theoretical models significantly.

3. Assumptions on individual behavior of theoretical models are substantially violated. Market participants do not truthfully bid their conditional expectations of the value of the security, and in general they systematically bid lower than their conditional expectations in the first rounds. Market participants demonstrate some degree of learning from prices in both majority and parity markets.

The implications of this study are that (1) security design should be one of the key factors to be considered in designing an information market for better information aggregation; and (2) theoretical models have provided useful insights into the behavior of information markets, but its explanatory power is limited due to their simplified assumptions on individual behavior. One possible direction to improve existing models of information markets is to incorporate the systematic bias of individual bids, i.e. participants tend to bid lower.

References

[1] B. Allen. Generic existence of completely revealing equilibria for economics with uncertainty when prices convey information. *Econometrica*, 49:1173–1199, 1981.

[2] J. E. Berg, R. Forsythe, F. D. Nelson, and T. A. Rietz. Results from a dozen years of election futures markets research. In C. A. Plott and V. Smith, editors, *Handbook of Experimental Economic Results*. 2001, forthcoming.

[3] J. E. Berg and T. A. Rietz. Prediction markets as decision support systems. 2003.

[4] K. Y. Chen and C. R. Plott. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. working paper No. 1131, California Institute of Technology, Division of the Humanities and Social Sciences, 2002.

[5] Y. Chen, C. H. Chu, T. Mullen, and D. M. Pennock. Information markets vs. opinion pools: An empirical comparison. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05)*, Vancouver, Canada, June 2005.

[6] J. Feigenbaum, L. Fortnow, D. M. Pennock, and R. Sami. Computation in a distributed information market. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC'03)*, San Diego, CA, June 2003.

[7] R. Forsythe and F. Lundholm. Information aggregation in an experimental market. *Econometrica*, 58:309–47, 1990.

[8] R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.

[9] R. Forsythe, T. A. Rietz, and T. W. Ross. Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39:83–110, 1999.

[10] S. J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–559, 1981.

[11] Serena Guarnaschelli, Anthony M. Kwasnica, and Charles Plott. The winner's curse in double auctions. *Information Systems Frontiers*, (5):61–75, 2003.

[12] F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.

[13] J. C. Jackwerth and M. Rubinstein. Recovering probability distribution from options prices. *Journal of Finance*, 51(5):1611–1631, 1996.

[14] J. Jordan. The generic existence of rational expectations equilibrium in the higher dimensional case. *Journal of Economic Theory*, 26:224–243, 1982.

[15] J. H. Kagel and D. Levin. The winner's curse and public information in common value auctions. *American Economic Review*, 76(894-920), 1986.

[16] John Lintner. The aggregation of investors' diverse judgments and preferences in purely competitive security markets. *Journal of Financial and Quantitative Analysis*, 4:347–400, 1969.

[17] R. E. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 28:103–124, 1972.

[18] D. M. Pennock, S. Lawrence, F. A. Nielsen, and C. L. Giles. Extracting collective probabilistic forecasts from web games. In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 174–183, San Francisco, CA, 2001.

[19] C. Plott and S. Sunder. Efficiency of experimental security markets with insider information: An application of rational expectations models. *Journal of Political Economy*, 90:663–98, 1982.

[20] C. Plott and S. Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56:1085–118, 1988.

[21] L. Shapley and M. Shubik. Trading using one commodity as a means of payment. *Journal of Political Economy*, 85:937–968, 1977.

[22] z-Tree (Zurich Toolbox for Readymade Economic Experiments). http://www.iew.unizh.ch/ztree/index.php.