Manipulation in prediction markets

Lawrence Choo Todd R. Kaplan Ro'i Zultan*

Abstract

Markets are increasingly as information aggregation mechanisms used to reduce uncertainty and predict future events. If policy makers use prediction markets to guide policy decisions, parties with vested interests in the policy may attempt to manipulate the market outcomes at a cost in order to influence policy decisions. We test manipulation in experimental markets. In the experiment, traders trade Arrow-Debreu securities corresponding to the possible states of the world. Policy makers, who observe the transaction prices, then vote on a policy, where the optimal policy depends on the true state. We find that market prices in a baseline market almost perfectly converge to reveal the true state of the world, allowing policy makers to implement the optimal policy. In contrast, information aggregation is substantially impaired if the traders suspect that manipulators operate in the market—even if this is not, in fact, the case. When there are manipulators in the market, manipulators are able to distort the prices, and draw a substantial number of votes to their preferred policy. When the existence of manipulators in the market is not common knowledge, prices do not discriminate between the true state and the state that the manipulators prefer.

Keywords: prediction markets, policy, experiment
JEL codes: C92, D53, D8, G14

*Choo: FAU-Erlangen-Nuremberg, Chair of Economic Theory, Lange Gasse 20, 90403, Nuremberg Germany. E-Mail: lawrence.cy.choo@fau.de. Tel: +49 911-5302-222. Kaplan: University of Exeter, Exeter, EX4 4PU, UK, and University of Haifa, Mount Carmel, Haifa 31905,Israel. E-mail: toddrkaplan@gmail.com, Tel: +44-4392-263237. Zultan: Ben-Gurion University of the Negev, P.O.B. 653, Beer-Sheva 84105, Israel. E-mail: zultan@bgu.ac.il, Tel:+972(0)86472306.

1. Introduction

Prediction markets, where traded assets yield payoffs based on the future realizations of uncertain events, are able to aggregate dispersed information. Predictions based on asset prices in such markets overwhelmingly outperform conventional, forecasting methods (Wolfers and Zitzewitz, 2004). This raises the possibility of using prediction markets to guide decision making. Indeed, prediction markets are increasingly being used by governments and private corporations as basis for policy decisions (Arrow et al., 2008).

In some cases, where the unknown variables are complex and there is no clear future realization, it is not possible to design a dedicated artificial prediction market. Nonetheless, prices in natural financial markets can still be used to aggregate the information existing in the market to serve policy making. Consider, for example, legislation geared towards different energy technologies. Whether traditional or alternative energy technologies are more efficient—and should therefore be supported by appropriate legislation—depends on values of myriad unknown variables. Increasing stock prices of sustainable energy technology firms may lead legislatures to believe that the state of the world is favourable to such technologies, and vote accordingly.

If policy makers listen to the market, however, parties with vested interest in the implemented policy have an incentive to try to manipulate the market prices (Hanson, 2004). In the example above, if energy companies expect stock prices to influence future legislation, they might be willing to incur market losses in the short term in order to inflate their stock prices, indirectly influencing the decision makers.

The success of prediction markets in forecasting various outcomes is generally taken to indicate that manipulation attempts are unsuccessful (Wolfers and Zitzewitz, 2004). The field evidence, however, is naturally restricted to events where there is, ex post, a clear outcome, and any manipulation would not be aimed at policy making, but instead at increasing the perceived popularity of an option. Predictions of political elections outcomes, for example, provide a situation where parties or candidates have a natural incentive to manipulate the market prices in order to create a 'bandwagon effect' (Rhode and Strumpf, 2004). Yet, opposing parties are generally equally likely to engage in manipulation attempts, which may cancel each other out. Thus, identifying actual manipulation attempts that can be studied in the field is a challenging endeavour. Camerer (1998) found that temporary placed bets on horse races had a negligible effect on closing betting odds. Hansen, Schmidt, and Strobel (2004) found one such case in political markets, where a party sent an email to its supporters asking them explicitly to affect the market in order to gain political traction, manipulation indeed influenced the market prices. Notwithstanding, if the researcher is aware of manipulation attempts, so are other traders in the market. When all the traders are fully aware of attempt to manipulate the marking, the attempt will likely prove futile since the traders can profit by gathering information of their own and being a counter-party to the manipulator (Hanson and Oprea, 2009). Even if there is a distortion, decision makers will take it into account when making decisions.

Laboratory experiments provide a controlled environment where the ability of markets to aggregate disperse information can be studied (Deck and Porter, 2013; Plott and Sunder, 1982, 1988). In an experiment, some traders can be endowed with incentives to manipulate the market. Thus, manipulation attempts can be directly measured, and the market outcomes are fully observable. Several studies looked manipulation in single asset markets, where the value of the asset depends on an unknown state of the world. In Hanson, Oprea, and Porter (2006), some traders received additional payment to their market earnings based on the median transaction price, incentivizing them to push prices up. Manipulation attempts did not have a significant effect on prices. Hanson, Oprea, and Porter (2006) found evidence that traders in the market—who know that some traders are incentivized to inflate prices—are able to counter the manipulation attempts. Similarly, in Veiga and Vorsatz (2010), a robot trader that creates artificial demand and supply was not able to manipulate prices.[1] However, in both those studies—even without manipulators—average prices did not move substantially with the true value of the asset, indicating that the market did not fully aggregate the information. The question of whether manipulators are able to impede efficient information aggregation therefore remains open.

To test directly manipulation aimed at influencing observers, Deck, Lin, and Porter (2013) introduced forecasters, who observe the market activity and make costly investments. Their experiment involved two possible states of the world, with each trader receiving an independent stochastic signal about the true state. Manipulators knew the true state with certainty, but did not receive their market earnings and were only paid based on forecast errors. Without manipulators, prices did not converge to the benchmark levels, but were informative

[1] Such manipulation did lead to higher prices if the true value was low in a different setting, where some traders have perfect information regarding the true value (Veiga and Vorsatz, 2009, 2010).

enough to improve forecasts. With manipulators, prices were completely non-informative, and forecasts made by inexperienced forecasters were even negatively correlated with the true state.

Thus, the bulk of the existing evidence suggests that, in single-asset situations—where markets were not able to fully aggregate information—non-manipulators that are aware of the existence of manipulators are able to anticipate and prevent price manipulation. This raises two important questions. First, can manipulators distort prices and influence policy decisions in markets that are otherwise able to efficiently aggregate information? Second, how does uncertainty about the existence of manipulators in the market affect the market's ability to aggregate information and guide optimal policy making? The current research aims to answer these questions.

To do so, we study a market with Arrow-Debreu securities, each corresponding to one possible state of the world.[2] These types of assets, dubbed by Wolfers and Zitzewitz (2004) 'Winner-take-all contracts', are able to efficiently aggregate disperse information even in complex situations (Choo, Kaplan, and Zultan, in press). After the end of the trading period, policy makers—who observe all transactions—vote on multiple policies, each optimal in a different state of the world. Voting for a 'safe' status quo option is also allowed, which is implemented if none of the policies receive a majority of votes. The introduction of a status quo option allows us to estimate the trust that policy makers place in the market prices, and to study how this trust varies according to market activity and the policy makers' awareness of manipulation attempts. We introduce manipulators by varying the incentives of two of the traders across market periods. With probability 0.5, these traders substantially gain from the implementation of a policy that they know is *not* socially optimal. Importantly, while in previous laboratory studies the existence of manipulators in the market was common knowledge, we compare situations with and without common knowledge.

2. Experimental design and procedure

Each period consists of a trading stage and a voting stage, with different subsets of participants active in each stage. We manipulated two independent variables in a 2×2 design. The existence of manipulators was manipulated within subjects, whereas the information regarding the existence of manipulators was ma-

[2] To continue our example above, one may think of these securities as stocks of firms specializing in different energy technologies.

nipulated between subjects. Thus, participants took part in either a *FullInfo* or a *NoInfo* treatment. Each treatment included two possible types of markets, a *baseline* and a *manipulation* market. The experimental instructions are provided in the appendix.

2.1. Market procedure

Each market involves eight traders (active in the trading stage) and four policy makers (active in the voting stage). Two of the traders—the potential manipulators—are designated as *red* ($\mathcal{R}$) traders, and the other six traders as *blue* ($\mathcal{B}$) traders. The traders are separated into two information groups of four, with the $\mathcal{R}$ traders being in the same information group. $\mathcal{B}$ traders do not know if they are grouped with the $\mathcal{R}$ traders.

2.1.1. Trading stage

Before the trading stage commences, nature selects one of three possible states of the world, X, Y, and Z, with equal probabilities. Each group of four traders is then informed that one of the other two states is not the true state.[3] For example, if nature selects state Y, one group is informed that state X is not true and the other that state Z is not true.

Traders trade three Arrow-Debreu securities x, y, and z (corresponding to the three possible states, X, Y, and Z) in three concurrent markets. Trade takes place using the continuous double auction (CDA) mechanism as follows. At the beginning of trade, each trader is endowed with 200 ECU (Experimental Currency Unit) and 5 units of each security type. During a trading period of 120 seconds, traders can place bids and asks (in the range of 0–20)—and accept open bids and asks—for each of the three securities. Short-sales are prohibited. When the markets close, each security pays a dividend of 10 ECU if it corresponds to the true state and 0 ECU otherwise,[4] which is added to the trader's capital balance at the end of trade to determine the trading stage earnings.

2.1.2. Voting stage

The policy makers observe all of the transaction prices, and proceed to the voting stage. Each policy maker casts a vote for one of three policies $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$

[3]The policy makers only know that the true state is X, Y or Z with equal probabilities.

[4]For example, if the true state is Y then security y pays a dividend of 10 ECU and the other securities 0 ECU.

TABLE 1. Implemented policy returns

	Baseline markets			Manipulation markets		
	Policy maker	$\mathcal{B}$ trader	$\mathcal{R}$ Trader	Policy maker	$\mathcal{B}$ trader	$\mathcal{R}$ Trader
True Policy	400	400	400	400	400	-400
Fake Policy	-400	-400	-400	-400	-400	1000
Neutral Policy	-400	-400	-400	-400	-400	-400
Status-quo	100	100	100	100	100	100

Note. Each row of panels A and B details the implemented policy returns to a policy maker, $\mathcal{R}$ trader and $\mathcal{B}$ trader, depending on type of policy implemented.

(corresponding to the three possible states, X, Y, and Z), or for the status quo $\mathcal{Q}$. If one of the three policies receives a majority of the votes, it is implemented. Otherwise, the status quo $\mathcal{Q}$ is implemented.[5]

2.1.3. Payoffs from the implemented policy

Implementing the status quo $\mathcal{Q}$ yields a payoff of 100 ECU each for all traders and policy makers. The $\mathcal{B}$ traders and policy makers gain 400 ECU from the implementation of the policy that corresponds to the true state, and lose 400 ECU if any of the policies that correspond to the other two states is implemented. The payoff for the $\mathcal{R}$ traders depends on the market type. In the *Baseline* markets, the $\mathcal{R}$ traders receive the same payoff as the other participants in the market. In the *Manipulation* markets, they receive a high payoff of 1,000 ECU if the implemented policy is the one that corresponds to the state they know not to be true—i.e., a policy that harms the other participants—and lose 400 ECU from the implementation of the other two policies. In the following, we refer to this policy as the *Fake* policy, to the policy that corresponds to the true state as the *True* policy, and to the remaining policy as the *Neutral* policy. For convenience, we maintain this terminology for the corresponding states of the world and assets.[6] Table 1 summarizes the returns from the implemented policy by market type and role.

[5]For example, if policies $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{Z}$ and $\mathcal{Q}$ receive 2, 1, 1 and 0 votes, respectively, then policy $\mathcal{X}$ is implemented. Alternatively, if two votes go for policies $\mathcal{X}$ and $\mathcal{Y}$ each, the status quo $\mathcal{Q}$) is implemented.

[6]For example, if the true state is Y and the $\mathcal{R}$ traders are informed that X is not true, then the True state, security, and policy are Y, y, and $\mathcal{Y}$. The Fake state, security, and policy are X, x, and $\mathcal{X}$; and the Neutral state, security, and policy are Z, z, and $\mathcal{Z}$.

2.1.4. Total payoffs

Writing π_i for the returns to individual i from the implemented policy, the payoff of each policy maker is $650 + \pi_i$. The corresponding payoff of the traders is given by

$$400 + \big[\underbrace{L_i + d(x)e^x + d(y)e^y + d(z)e^z}_{\text{Market decisions}}\big] + \pi,$$

where $L \geq 0$ is the trader's cash balance at the end of trade, e^j is his inventory of Arrow-Debreu security $j = x, y, z$ and $d(j)$ is the dividend of Arrow-Debreu security j. The difference in base payment between traders and policy makers makes up for the value of the trader's endowment (and therefore average trading stage earnings).

2.2. Treatment design and experimental procedure

The first part of the experiment was a training phase consisting of one practice and five experimental rounds, in which participants could learn the trading mechanism and information structure. Each round followed the design and procedure of the baseline markets described above, with the exception that there was no voting stage. Instead, all 12 participants participated in the role of traders, divided into two information groups of six traders each.

The main part of the experiment consisted of 14 rounds. Each round could included either a baseline or a manipulation market with equal probability. For efficient between-treatment comparisons, we pre-generated a sequence of states, which we implemented in all sessions. The two treatments differed as follows:

- *FullInfo* (5 sessions, 12 subjects per session). At each round, all subjects were informed as to whether they are participating in the baseline or manipulation market.

- *NoInfo* (4 sessions, 12 subjects per session). At each round, only the $\mathcal{R}$ traders were informed as to whether they are participating in the baseline or manipulation market (i.e., how they benefit from the implemented policy).

To facilitate comprehension, all roles ($\mathcal{R}$ traders, $\mathcal{B}$ traders and policy makers) were fixed across all 14 rounds.

The experiment was conducted in University of Exeter FEELE in 2018 and the student subjects were recruited through ORSEE (Greiner, 2015). The experiment was programmed with z-Tree (Fischbacher, 2007). At the end of each session, one round (out of five) from the training phase and two rounds (out of fourteen) from the experiment phase were randomly chosen for payment. Payoffs were converted to cash at the rate of 100 ECU equals 1 GBP, and added to a show up payment of 6 GBP.

3. Theoretical analysis

We maintain the terminology introduced above to denote the true state (and corresponding security and policy) as *True*; the state that the $\mathcal{R}$ traders know not to be true (and corresponding security and policy) as *Fake*; and the remaining state, security, and policy as *Neutral*. We denote the $\mathcal{B}$ traders who are in the same and different group as the $\mathcal{R}$ traders as $\mathcal{B}^R$ and $\mathcal{B}^B$ traders, respectively.

The complexity of the CDA mechanism makes it extremely difficult to analyse the equilibrium with standard game theoretical methods. We therefore evaluate the market's success at aggregating information by comparing market prices against the *fully revealing rational expectations equilibrium* (FRE) and *prior information equilibrium* (PIE, cf. Choo, Kaplan, and Zultan, in press). The PIE and FRE are both 'static' models that describe some form of market aggregation. They however differ on whether the true state is partially (PIE) or fully (FRE) revealed in prices.

Can prices always converge to the FRE? There is an extensive literature studying such dynamic behaviour of traders (e.g., Dubey, Geanakoplos, and Shubik, 1987; Hellwig, 1982; Ostrovsky, 2012). However, the complexity of such models are beyond the scope of our market design.[7] We therefore complement the static equilibrium analysis with a limited-rationality dynamic reasoning model (DRM) to generate predictions about the conditions under which prices converge to the FRE. This model considers a simplified discrete-time trading process, wherein supply and demand correspond to traders' beliefs, and beliefs are updated in each period based on the market clearing prices. The model offers the following two predictions:

(i) Security prices in the Baseline market can converge to the FRE.

(ii) Security prices in the Manipulation market can only converge to the PIE.

[7] Biais and Pouget, 1999 show that trade in Plott and Sunder (1988) markets, similar to the one we study, can only occur at the FRE price. Yet without trade at non FRE prices, it is difficult to see how traders can learn about the true state, and the analysis does not allow for manipulation attempts.

(iii) There is asymmetric information in the Manipulation market, where the $\mathcal{R}$ and $\mathcal{B}^R$ but not the $\mathcal{B}^B$ traders are informed about the true state.

As such, the DRM predicts that a minority of price manipulators (i.e., $\mathcal{R}$ traders in the Manipulation market) can severely impede the information aggregation properties of markets. Although the $\mathcal{B}^R$ traders are informed of the true state, they do not have enough market power to influence prices..

In the following, we first analyze the PIE and FRE in our setting, followed by the DRM analysis. The analysis assumes that policy makers are risk neutral, and use Bayes' rule whenever updating their beliefs.

3.1. Equilibrium analysis

The PIE describes the market clearing prices when traders update their beliefs about the true state given their own private information and condition their demands for securities upon such posteriors, but do not update their beliefs any further based on the observed prices. In the *baseline market*, all traders believe the True security to be True with probability 0.5, and therefore value it at 5 ECU. For each the other two securities, there is one group that values it at 5 ECU, whereas the other group values it at zero.[8] Thus, the market-clearing prices of the True, Fake and Neutral securities will be 5.00, 2.50, and 2.50 ECU, respectively.[9] Note that, although the true state is only partially revealed in prices, there is sufficient information for the policy makers to infer the true state.

In the *manipulation market*, we assume that the $\mathcal{R}$ traders trade as if they know the Fake state to be true. This creates excess demand (supply) for the Fake security at any price above (below) 5 ECU, the value assigned by the $\mathcal{B}^B$ traders. On the other hand, as only the two $\mathcal{B}^R$ traders are willing to buy the Neutral security at any positive price, excess supply will drive its prices to zero. The supply and demand for the True security remain as in the baseline market. The market-clearing prices of the True, Fake and Neutral securities are, therefore, 5.00, 5.00, and 0.00 ECU, respectively. As prices do not distinguish between the True and Fake securities, the risk-neutral policy-makers should vote for the status-quo policy.

[8]The $\mathcal{R}$ and $\mathcal{B}^R$ value the Neutral security at 5 ECU and the Fake security at 0 ECU, and vice versa for the $\mathcal{B}^B$) traders.

[9]Any price strictly between zero and 5 ECU will clear the markets for the Fake and Neutral securities. Taking the midpoint for simplicity, as we do here and in the following, does not affect the analysis.

The FRE assumes that traders continuously learn about the private information of others from market prices. This endogenous process continues until no trader can learn any more from prices. Radner (1979) shows that this equilibrium is equivalent to an outcome where all traders truthfully and openly communicate their posteriors about the true state. Accordingly, we analyze who beliefs would evolve under such open communication.

Baseline Market. If all traders openly announce (simultaneously) the states to which they assign positive probabilities, the unique conjunction of all announcements will be the true state.Thus, security prices will converge to the true values, fully informing the policy makers of the true state and correspondingoptimal policy.

Manipulation market. The FRE becomes more complicated, as we have to consider that the $\mathcal{R}$ traders have an incentive not to communicate truthfully. Given the information structure, each trader should assign a positive probability to exactly two states. Thus, the $\mathcal{R}$ traders can not credibly announce the Fake state *only*. In this case, each R trader can only (simultaneously) make one of two possible announcements: (a) Fake and True states, and (b) Fake and Neutral states.There are thus three possibilities:

- Both announce (a). The mode of all announcements is the True state and the true state is revealed.
- One announces (a) and the other (b). Again, the mode of all announcements is the True state and the true state is revealed.
- Both announce (b). The modes are the True and Fake states and the true state is not fully revealed.

Manipulation is only successful in the last possibility, where the FRE price are identical to the PIE price. Nevertheless, $\mathcal{B}^R$ and $\mathcal{R}$ traders (but not the $\mathcal{B}^B$ traders) have sufficient information to know the true state. Table 2 summarizes the FRE and PIE equilibrium prices and implemented policies.

3.2. Dynamic reasoning model

The DRM assumes that trade takes place over $t \in \{1, 2, \ldots\}$ hypothetical periods. In each t, traders proceed according to the following four stages:

TABLE 2. The FRE and PIE equilibrium prices.

	Security prices			Implemented policy
	True	Fake	Neutral	
Prior information equilibrium				
Baseline	5	2.5	2.5	True policy
Manipulation	5	5	0	Status quo
Fully revealing rational expectations equilibrium				
Baseline	10	0	0	True policy
Manipulation (unsuccesful)	10	0	0	True policy
Manipulation (succesful)	5	5	0	True policy

Stage 1 Traders observe period $t-1$ prices of all securities.

Stage 2 Traders update their beliefs about the true state.

Stage 3 Traders set their supply and demand for each security according to their updated beliefs.

Stage 4 The market clears at a price that equates supply and demand.

In the baseline market, Period 1 beliefs are set by the prior information that the trader holds. Our key assumption in analyzing manipulation markets is that $\mathcal{R}$ traders set their supply and demand of the securities *as if* they know the Fake state to be true. I.e., as if they value the Fake security at 10 ECU and the other securities at 0 ECU.

Baseline market. At period $t = 1$, the analysis is identical to that provided above for the PIE. Traders value the securities corresponding to the states that they know to be possible at 5 ECU. The resulting market clearing prices of the True, Fake and Neutral securities are 5.00, 2.50 and 2.50 ECU, respectively. This price profile uniquely identify the true state. Therefore, at period $t = 2$, all traders value the True security at 10 ECU and the other securities at 0 ECU. The resulting market clearing prices of the True, Fake and Neutral securities are hence 10.00, 0.00 and 0.00 ECU, respectively. Since traders are fully informed about the true state, there will be no further revisions to prices in period 3. That is, prices converge to the FRE.

Manipulator market. At period $t = 1$, the $\mathcal{B}^R$ and $\mathcal{B}^B$ traders behave as in the Baseline market. In contrast, the $\mathcal{R}$ traders will only demand the fake asset. The resulting market clearing prices of the true, fake and neutral assets are hence 5.00, 5.00 and 0.00 ECU, respectively, as in the PIE. These prices reveal that the Neutral state is not the true state. Therefore, the $\mathcal{R}$ and $\mathcal{B}^R$ traders—who can also rule out the Fake state—have sufficient information to deduce the true state. The symmetry between the True and Fake security prices, however, imply that the $\mathcal{B}^B$ traders are still uninformed about the true state. This symmetry persists in the next period $t = 2$. The $\mathcal{B}^B$ traders, who value both the True and the Fake securities at 5 ECU, form a majority of the market. Hence, there is excess supply (demand) above (below) the price of 5 ECU for both securities. The resulting market-clearing prices of the True, Fake and Neutral securities remain at 5.00, 5.00 and 0.00 ECU, respectively, and an equilibrium is reached. Thus, the DRM prices converge to the PIE prices in the manipulation market.

4. Results

The analysis focuses on the aggregated data, pooling over all treatments, sessions, rounds and market design.[10] This enables us to anchor our discussions around four *constellations* that differ on the treatment and market design: FullIno-Baseline, Fullinfo-Manipulation, NoInfo-Baseline and NoInfo-Manipulation.

4.1. Trading volume

The solid line of Figure 1 details the mean transaction volume of each security type over 30 second time blocks. The transaction volumes for the Fake and Neutral securities appear to steadily decline over the trading period. In contrast, no such decline is apparent for the True security.

The stacked bars on Figure 1 represent the proportion of transactions in each 30 seconds time block by trader types. For example, the RB^B (resp. $B^B B^R$) area details the the proportion of transactions where a $\mathcal{R}$ (resp. $\mathcal{B}^B$) trader is the buyer and the $\mathcal{B}^B$ (resp. $\mathcal{B}^R$) trader is the seller.

The FullInfo-Baseline constellation (first row) provides a benchmark description of trading behaviour. In the first 60 seconds of the market duration, we clearly see that traders often condition their demand for securities on their private information about the True state. For example, the Fake security is almost always purchased by the $\mathcal{B}^B$ traders—who assign a positive posterior to the Fake state given their private information—and is often sold by the $\mathcal{R}$ and $\mathcal{B}^R$ traders—who can rule out the Fake state. In comparison, purchases of the True

[10] A detailed exposition of experimental data is reported in the Appendix.

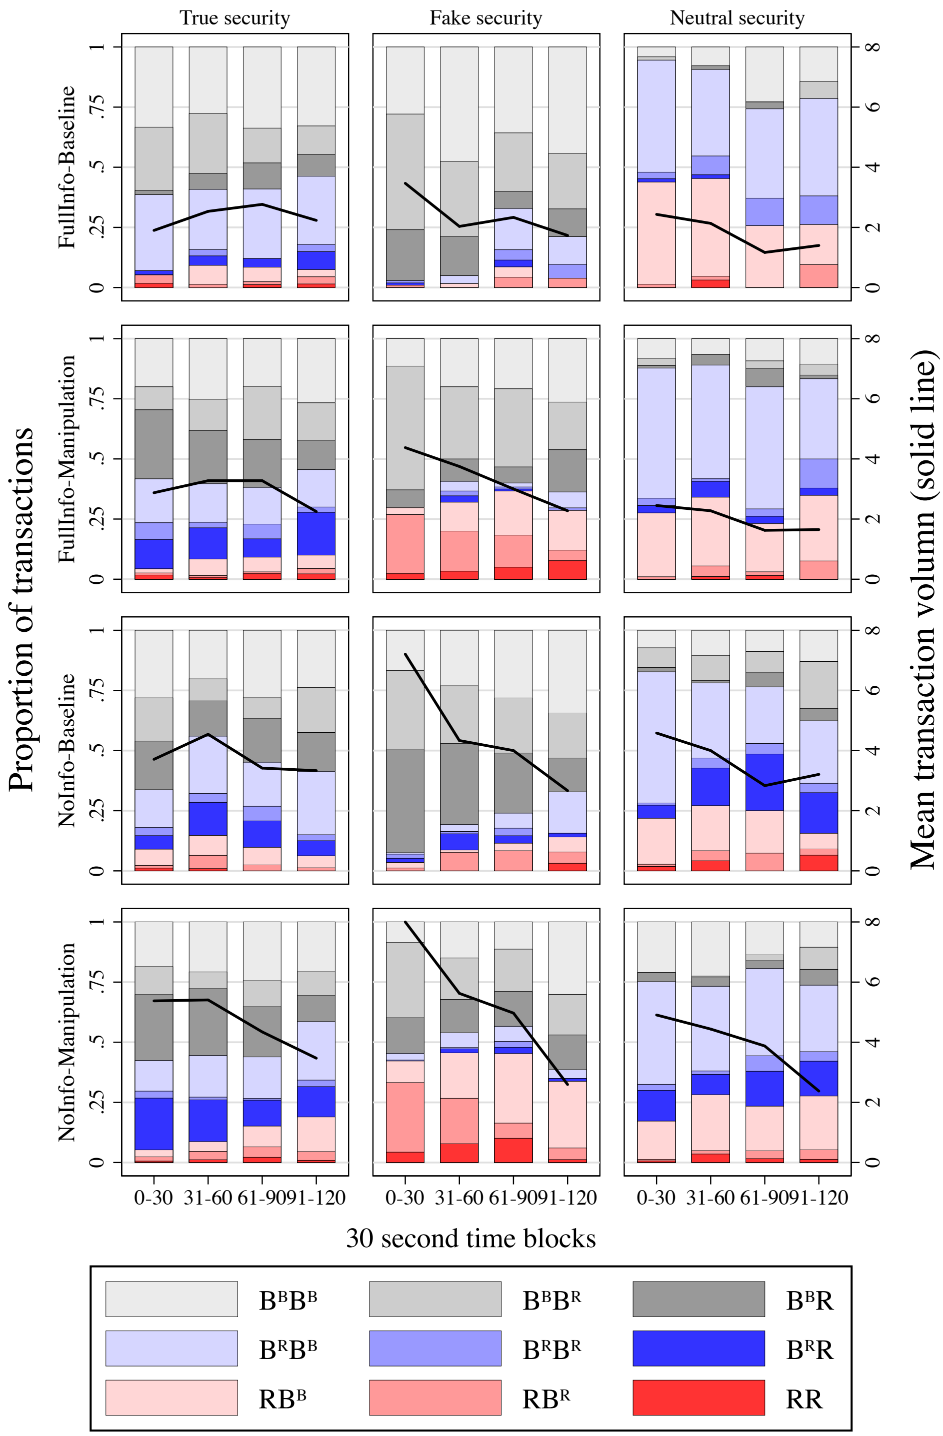

Note. The panels are organised by constellations (rows) and securities (columns). The solid line (right axis) details the mean transaction volume of each security type for each 30 seconds time block. The stacked bars detail the proportion of transactions by the trader types. For example, the area $B^B B^R$ (resp. $B^R R$) details the proportion of transactions where the buyer is a $\mathcal{B}^B$ (resp. $\mathcal{B}^R$) trader and the seller is a $\mathcal{B}^R$ (resp. $\mathcal{R}$) trader.

FIGURE 1. Mean transaction volumes and proportion of transaction by trader types.

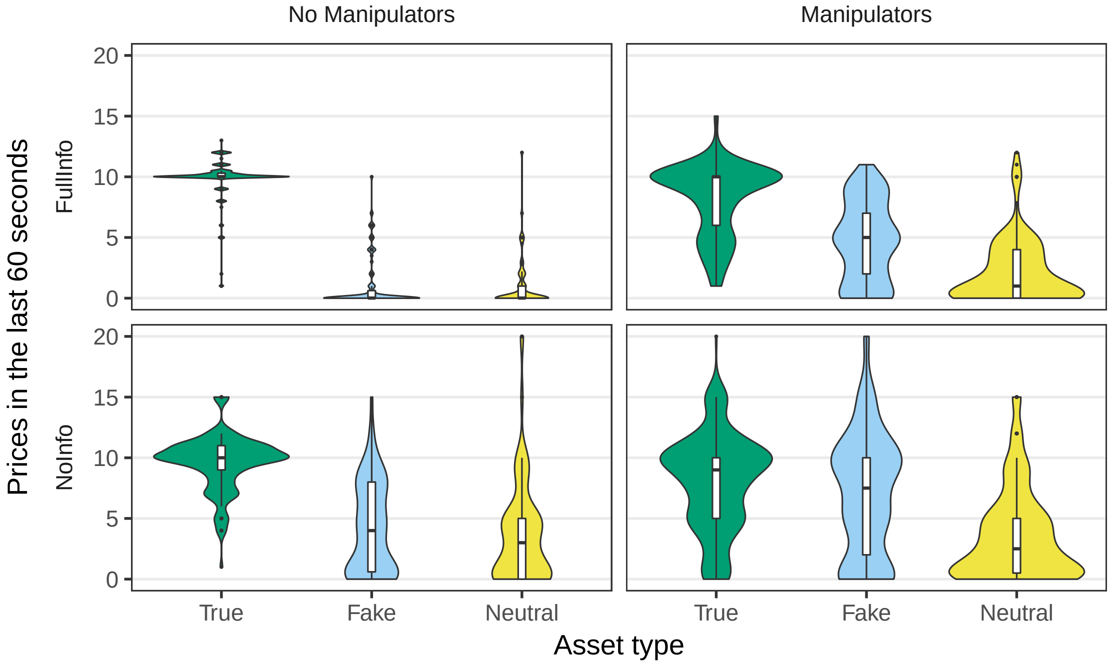

Note. Violin plots present kernel distributions of transaction prices. Box plots present the median .

FIGURE 2. Transaction prices in the last 60 second of the market.

security–to which all traders assign a positive probability—appear to be evenly distributed amongst all trader types. Interestingly, traders appear increasingly buy the securities which they know to be worthless in the last 60 seconds of trade. For example, whereas the $\mathcal{R}$ and $\mathcal{B}^R$ traders account for only 2% of Fake security purchases in the first 30 seconds, their share of such purchases increases to 21% in the last 30 seconds of trade. This perhaps suggests that some traders may have been trying to engage in speculative trade, creating demand for securities that they know to be worthless.

Turning our attention to the FullInfo-Manipulation constellation (second row), we observe that the $\mathcal{B}^R$ and $\mathcal{B}^B$ traders behave similarly to the FullInfo-Baseline constellation. In contrast, we now observe that the $\mathcal{R}$ traders account for a substantial proportion of the Fake security purchases. The $\mathcal{R}$ traders are also fairly active in purchasing the Neutral security, possibly in an attempt to obscure the True state from the policy makers by increasing noise trading.

4.2. Market Prices

We focus on prices in the second half (last 60 seconds) of trade, providing enough time for prices to converge. Figure 2 presents violin plots of transactions prices

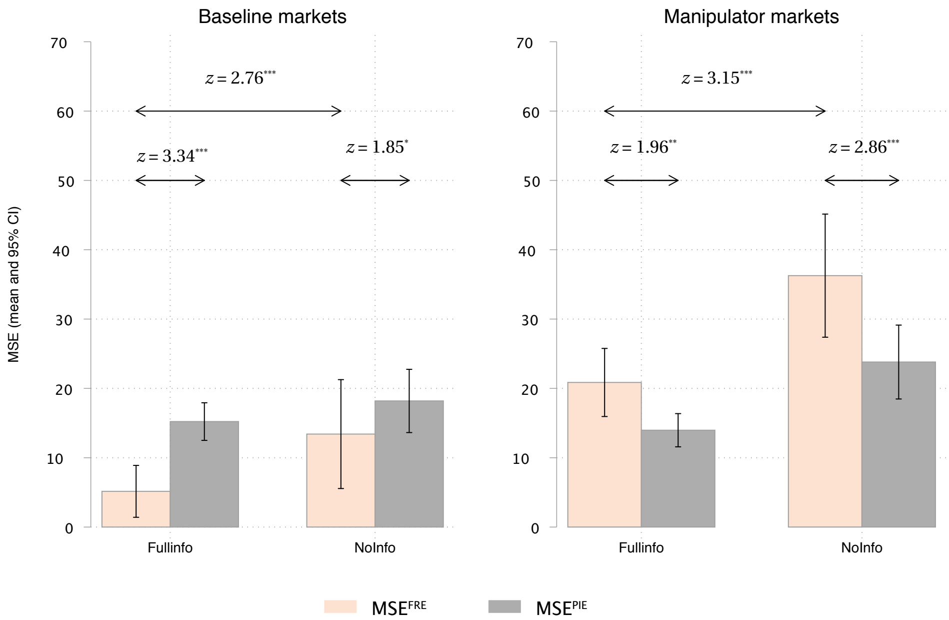

Note. z-scores for comparisons between treatments based on Mann-Whitney tests. z-scores for comparisons between market types based on Wilcoxon signed-ranks tests.

FIGURE 3. Mean square deviations from the equilibrium prices

in the last 60 seconds of trade by security type and constellation. The results in the baseline market with no manipulation and full information are striking. Prices converge almost perfectly, with both the median and mode equal exactly to the true values of 10, 0, and 0 ECU for the True, Fake, and Neutral securities, respectively. The comparison to the NoInfo treatment reveals the importance of complete information, which was maintained in all previous studies. Even when there are no manipulators in the market, the mere suspicion of manipulation is enough to impede price convergence, with many transactions closed at prices deviating from the true value of the security. Although the median transaction price for the True security still reflects its true value, the prices of the two other assets are mostly distributed around the PIE price of 2.50 ECU.

Moving to the Manipulation markets, we see that manipulators are able to influence the prices. With full information, the other traders are apparently able to partly counter the manipulation attempts, with the prices of the True and Neutral securities traded at a median price equal or close to the true value. Nonetheless, the median price of the Fake security is at 5 ECU, corresponding to the PIE price, indicating that manipulators are able to keep the prices from falling to the true value of zero. The picture is more bleak in the NoInfo treatment, where the price distributions of the True and Fake securities are mostly

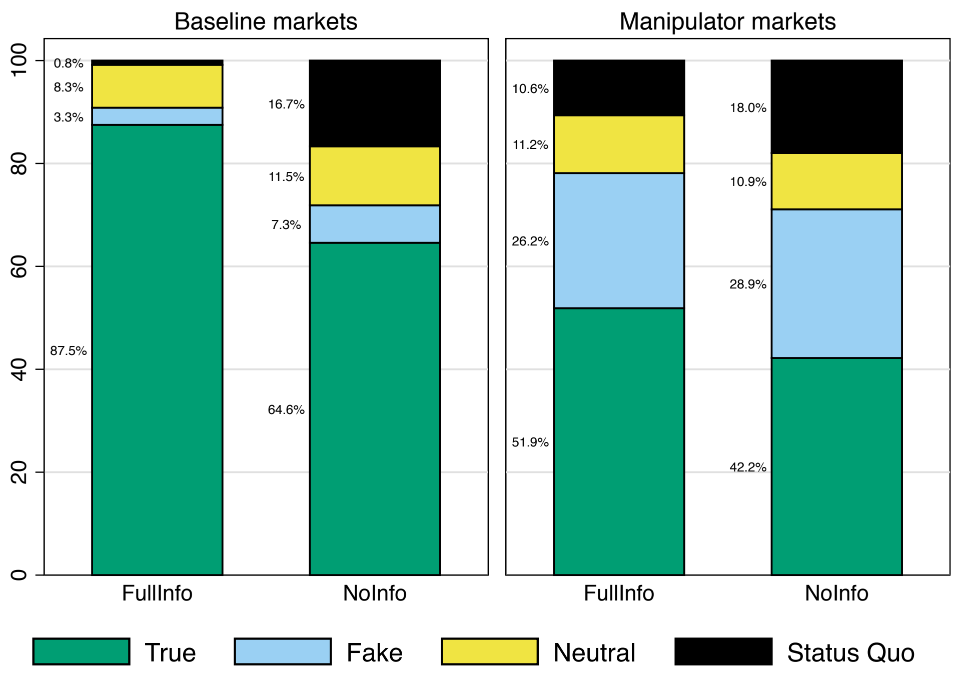

FIGURE 4. Distributions of votes.

indistinguishable. In line with the DRM predictions, prices are as in the PIE, although shifted upwards.

To test price convergence, we calculated in each constellation the mean square deviations from the equilibrium prices, separately for the FRE and for the PIE. The results are depicted in Figure 3. In the baseline market with full information, the FRE provides a significantly better fit than the PIE for the transaction prices. This difference almost disappears with no information. With manipulators, transaction prices are closer to the PIE prediction than to the FRE prediction.

4.3. Voting behavior

Figure 4 presents the distributions of votes by constellations. In the baseline markets with full information, almost 90 percent of all votes go to the True policy. With no information, confidence in the market is low, with the share of the status-quo votes increasing to one in six, compared to less than one percent in the FullInfo treatment.

We see a similar picture in the manipulator markets, where more votes go to the status quo in the NoInfo compared to the FullInfo treatment. The main finding, however, is that manipulators are fairly successful, drawing 25–30 percent of the votes to the Fake policy. Even when the policy makers perfectly know that there are manipulators in the market, they place a high degree of trust in the market prices, voting for the status quo only 10 percent of the time. Consequently, the vote share of the True policy drops to approximately half of the votes, with more than a quarter of the votes going to the Fake policy.

5. Conclusion

Motivated by advancements in the study of information aggregation in markets over the last few decades, many researchers and policy makers advocate the use of markets in guiding policy decisions. This raises the necessity of better understanding how invested parties are able to misuse the market in order to distort information and influence policy.

Deck, Lin, and Porter (2013) provided evidence that manipulators can sufficiently obscure the information as reflected in the market prices such that observers are unable to forecast the true state better than chance. We extend the market setting, allowing us to address additional issues and to draw several new conclusions. First, we study a setting were prices almost perfectly aggregate information in the baseline market with full information. We show that, with efficient markets, manipulators are still fairly successful, however prices still contain enough information to reveal the True state most of the time.

Second, we study the role of common knowledge of manipulation. In previous studies, all traders were aware of the manipulation attempts, and were able to counter them to some extent. We find that, indeed, manipulators are more successful if their existence is not commonly known, to the extent of completely eliminating the price differences between the True and the Fake securities. Perhaps more interesting is the finding that the *mere suspicion* of manipulation is enough to substantially impede information aggregation—even when there are no actual manipulation attempts. To the extent of reducing the probability of implementing the optimal policy almost as much as with (commonly known) manipulation.

Third, we disentangle two manipulator motivations, namely obscuring the True state on the one hand, and promoting the Fake state on the other hand. This brings to light the importance of common knowledge of manipulation. With common knowledge, manipulators are able to increase the noise in the market, however only without common knowledge do they succeed in increasing the price of the Fake security to significantly above the PIE price. Further-more, by introducing a status quo policy, we are able to conclude that confidence in the market remains high even with suspicion of manipulation or with actual manipulation. Of those two factors, lack of common knowledge has the larger effect on policy makers' confidence in the market. The high confidence—as reflected in the low share of votes for the status quo—results in successful manipulation. That is, manipulators are able not only to draw votes from the True policy, but also to influence policy makers into voting for the Fake policy.

Taken together, the results reveal that even markets that have the capability of efficiently aggregating disperse information into prices are susceptible to manipulation. Thus, if dedicated predictions markets are set up to inform policy, it is important to regulate participation in order to eliminate possible manipulation. Moreover, it is important that the participating traders are aware that such precautions are in place, as suspicion of manipulation per se is enough to impede information aggregation.

References

Arrow, Kenneth J et al. (2008). The promise of prediction markets. *Science* 320(5878), pp. 877–878.

Biais, Bruno and Sébastian Pouget (1999). Microstructure, incentives and convergence to equilibrium in experimental financial markets. Working Paper, Toulouse University.

Camerer, Colin F (1998). Can asset markets be manipulated? a field experiment with racetrack betting. *Journal of political economy* 106(3), pp. 457–482.

Choo, Lawrence, Todd R. Kaplan, and Ro'i Zultan (in press). Information aggregation in arrow–debreu markets: an experiment. *Experimental economics.*

Deck, Cary, Shengle Lin, and David Porter (2013). Affecting policy by manipulating prediction markets: experimental evidence. *Journal of economic behavior & organization* 85, pp. 48–62.

Deck, Cary and David Porter (2013). Prediction markets in the laboratory. *Journal of economic surveys* 27(3), pp. 589–603.

Dubey, Pradeep, John D. Geanakoplos, and Martin Shubik (1987). The revelation of information in strategic market games: a critique of rational expectations equilibrium. *Journal of mathematical economics* 16(2), pp. 105–137.

Fischbacher, Urs (2007). Z-tree: zurich toolbox for ready-made economic experiments. *Experimental economics* 10(2), pp. 171–178.

Greiner, Ben (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the economic science association* 1(1), pp. 114–125.

Hansen, Jan, Carsten Schmidt, and Martin Strobel (2004). Manipulation in political stock markets–preconditions and evidence. *Applied economics letters* 11(7), pp. 459–463.

Hanson, Robin (2004). Foul play in information markets. In: *Information markets: a new way of making decisions*. Ed. by Robert W Hahn and Paul C Tetlock. Washington, D.C.: AEI-Brookings Press. Chap. 6, pp. 126–141.

Hanson, Robin and Ryan Oprea (2009). A manipulator can aid prediction market accuracy. *Economica* 76(302), pp. 304–314.

Hanson, Robin, Ryan Oprea, and David Porter (2006). Information aggregation and manipulation in an experimental market. *Journal of economic behavior and organization* 60(4), pp. 449–459.

Hellwig, Martin F. (1982). Rational expectations equilibrium with conditioning on past prices: a mean-variance example. *Journal of economic theory* 26(2), pp. 279–312.

Ostrovsky, Michael (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica* 80(6), pp. 2595–2647.

Plott, Charles R. and Shyam Sunder (1982). Efficiency of experimental security markets with insider information: an application of rational-expectations models. *Journal of political economy* 90(4), pp. 663–698.

— (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56(5), pp. 1085–1118.

Radner, Roy (1979). Rational expectations equilibrium: generic existence and the information revealed by prices. *Econometrica* 47(3), pp. 655–678.

Rhode, Paul W and Koleman S Strumpf (2004). Historical presidential betting markets. *Journal of economic perspectives* 18(2), pp. 127–141.

Veiga, Helena and Marc Vorsatz (2009). Price manipulation in an experimental asset market. *European economic review* 53(3), pp. 327–342.

— (2010). Information aggregation in experimental asset markets in the presence of a manipulator. *Experimental economics* 13(4), pp. 379–398.

Wolfers, Justin and Eric Zitzewitz (2004). Prediction markets. *Journal of economic perspectives* 18(2), pp. 107–126.