Manipulation and (mis)trust in prediction markets[*]

Lawrence Choo   Todd R. Kaplan   Ro'i Zultan[†]

July 13, 2021

Abstract

Markets are increasingly used as information aggregation mechanisms to predict future events. If policy makers and managers use markets to guide policy and managerial decisions, interested parties may attempt to manipulate the market in order to influence decisions. We study experimentally the willingness of managers to base decisions on market information under the shadow of manipulation. We find that when there are manipulators in the market, managers under-utilize the information revealed in prices. Furthermore, mere suspicion of manipulation erodes trust in the market, leading to the implementation of suboptimal policies—even without *actual* manipulation.

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

[*]We thank seminar participants at Samford University, University of Exeter, University of Erlangen-Nuremberg, University of Haifa, Technion, and conference participants at the Economic Science Association international meeting in Berlin, Experimental Finance meeting in Heidelberg, 9th Workshop on Theoretical and Experimental Economics in Munich, the Israel Behavioral Finance meeting, Society for the Advancement of Economic Theory (SAET) in Ischia and the Econometric Society European Meeting in Manchester for useful comments and discussion. Financial support from the International Foundation for Research in Experimental Economics is gratefully acknowledged. This research is partially supported by a National Science Foundation of China key project (no. 72033006).

[†]Choo: Southwestern University of Finance and Economics, China Center for Behavioral Economics and Finance, 555 Liutai Avenue, Chengdu, China. E-Mail: lawrence.cy.choo@gmail.com. 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.[1] Predictions based on asset prices in such markets overwhelmingly outperform conventional forecasting methods (e.g., Arrow et al., 2008; Palan, Huber, and Senninger, 2019; Wolfers and Zitzewitz, 2004). It is not surprising, then, that governments and private corporations are increasingly using prediction markets as a basis for policy decisions (e.g, Chen and Plott, 2002; Cowgill and Zitzewitz, 2015; Dianat and Siemroth, 2020; Gillen, Plott, and Shum, 2017). Moreover, trading prices in natural financial markets can be used to inform policy making if it is difficult, or even impossible, to design a dedicated artificial prediction market.[2] If, for example, there is no clear future resolution of uncertainty, or the variables of interest are unobservable and hard to measure, there is no straightforward way of fixing the redemption value of the traded assets.

Consider, for example, legislation geared towards different energy technologies. Which one of the traditional or many alternative energy technologies is most efficient—and should therefore be supported by appropriate legislation—depends on a myriad of unknown variables. Increasing stock prices of sustainable energy technology firms may lead legislators to believe that the state of the world is favorable to such technologies and vote accordingly.

If policy makers "listen" to the market, parties with a vested interest in the policy decision may have an incentive to manipulate the market prices (Hanson, 2004).[3] In the example above, if energy companies expect stock prices to guide future legislation, they might artificially inflate their own stock prices, incurring short-term market losses in order to influence the policy making in their favor. Such situations also arise naturally when private firms use predic-

[1]The information aggregation properties of markets were first formally noted by Hayek (1945) and central to the *efficient market hypothesis* (e.g, Fama, 1970). See, for example, Radner (1979), Muth (1961) and Ostrovsky (2012) for the theoretical properties of market aggregation. More recently, economists have argued that well-designed markets can be utilized as tools to gather information (e.g., Arrow et al., 2008; Plott, 2000).

[2]Futures markets can sometimes be interpreted as natural markets for information.

[3]Listening to the market can sometimes be inefficient if the market also listens to subsequent decisions. For example, central bankers may use bond markets' prices to guide their policy decisions. However, bond market participants are also reacting to the central bankers' decisions. See *The Economist* article "Can central bankers talk too much?" (24th October 2019) and theoretical treatments by Bond and Goldstein (2015) and Lieli and Nieto-Barthaburu (2020). Related topics include situations where the market participants can affect the predicted outcome (Chakraborty and Das, 2016), bluffing (Chen et al., 2010, 2007; Jian and Sami, 2012) and fraud (Blume, Luckner, and Weinhardt, 2010). See Kloker and Kranz (2017) for a typology of manipulation and fraud in markets.

tion markets. For example, when a firm runs a prediction market to forecast future sales of a new product, competitors—as well as parties within the firm whose interests do not perfectly align with the firm's—may try to manipulate the market prices in order to influence the firm's strategy.

Several theoretical and empirical studies have considered manipulation in prediction markets. Little attention has, however, been given to how managers respond to the possibility of manipulation in prediction markets. In this paper, we study manipulation and managerial decisions in an experimental asset market, where the value of the traded assets is contingent on an underlying state of the world. In each market period, each trader receives a private signal that, in itself, is not sufficient to deduce the state. Nonetheless, the combination of all signals fully reveals the true state. Managers observe all transaction prices and vote on a policy, the outcome of which depends on the true state. In some periods, a *minority* of traders stand to gain from a policy that is, on the whole, harmful to the other traders and to the managers. In one treatment, the existence of such *manipulators* (i.e., traders with incentives to deceive the managers) is commonly known. In the other treatment, traders and managers only know that there is a 50% chance that manipulators are active in the market.

When there are no manipulators in the market, we find that market prices are able to aggregate the information dispersed in the market, in the sense that the ranking of the prices reveal the true state in over 90% of the markets. Lack of common knowledge of the non-existence of manipulators leads to more volatile prices. Nonetheless, the market closing prices are still able to reveal the true state at least as often as with common knowledge. When manipulators participate in the market, they substantially affect market prices, significantly reducing price discrimination. Consequently, prices reveal the true state in only 75% (with common knowledge) and 68% (without common knowledge) of markets.

Thus, the experimental environment achieves its purpose of providing a controlled testbed to study the effect of manipulation on managerial decisions. We find that, without manipulators, the mere suspicion of manipulation undermines trust. Markets are still efficient, i.e., prices reflect the true state, and managers stand to make substantial gains by voting for the policy indicated by the highest price. Yet, they are reluctant to do so, leaving on the table 37% of the potential payoff from the policy. With manipulators, managers can still benefit in expectancy from always voting for the policy suggested by market prices. They nonetheless often choose to ignore this information or even support an alternative policy. As a result, the average payoff of the managers is even lower than what could be obtained by ignoring the market altogether. These results highlight the critical role of *trust* in the market. That is, manipulation—and mere suspicion of manipulation—hinders the ability of markets to inform optimal managerial decisions via two distinct channels. One, by influencing market prices; two, by eroding trust in the market.

The success of prediction markets in forecasting various outcomes is generally taken to indicate that manipulation attempts are unsuccessful (Wolfers and Zitzewitz, 2004). By its nature, however, manipulation is covert. Consequently, direct evidence for manipulation is difficult to obtain in the field. Exceptions include Hansen, Schmidt, and Strobel (2004), who looked at an overt manipulation attempt by a political party seeking to gain traction by calling on its supporters to inflate prices in a political prediction market and Gandal et al. (2018), who utilized a data breach to study manipulation in Bitcoin exchanges.

In contrast to field studies, laboratory experiments provide a controlled environment for studying the ability of markets to aggregate dispersed information (Deck and Porter, 2013; Plott and Sunder, 1982, 1988). In an experiment, some traders can be endowed with incentives to manipulate the market. Manipulation attempts in such a setup are directly measurable, and the market outcomes are fully observable. Several studies looked at manipulation in single asset markets, where the asset's value depends on an unknown state of the world. Manipulation was implemented by paying some traders an additional payment based on the median transaction price (Hanson, Oprea, and Porter, 2006) or by introducing bot traders that artificial demand and supply (Veiga and Vorsatz, 2010). Unlike our design, the single-asset markets did not fully aggregate information even without manipulation, and the focus was on market prices rather than subsequent managerial decisions. In contrast, we design a market environment that is successful in aggregating diverse information. Against this backdrop, we introduce manipulators and managers to study how the market informs managerial decisions, and in particular managers' willingness to trust the information revealed in such markets.

Deck, Lin, and Porter (2013) introduced forecasters, who observe the market activity without participating in it and make costly investments. Without manipulators, prices did not converge to the benchmark levels, but were informative 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. A recent study by Maciejovsky and Budescu (2020) highlighted the importance of trust. Their experiments compared the ability of group communication and markets to aggregate information (see also Maciejovsky and Budescu, 2013). Under manipulation incentives, markets outperformed the groups. Nonetheless, participants reported more trust in the groups, as did third-party observers. This highlights the need to improve the understanding of how managers incorporate information from the market into their decisions.

The current research goes beyond the existing literature in addressing several open issues. First, most existing studies used a single-asset setup, in which markets fail to aggregate information efficiently even without manipulation (cf. Corgnet, DeSantis, and Porter, 2015; Plott and Sunder, 1988). We study manipulation in markets with state-contingent assets, which aggregate information efficiently, and test whether manipulation undermines this ability. Second, previous studies focused on market behavior. Our primary focus is on the managers' response to market prices under the shadow of potential manipulation. Accordingly, we go beyond the existing results in differentiating between two possible outcomes of manipulation: obscuring the true state and promoting a false state. We do so by including a third, neutral, state, which is neither the true state nor the one favorable to the manipulators. Furthermore, we allow the managers to vote for a status-quo policy, which is not state specific. This allows us to estimate trust in the market and distinguish between ignoring the market and voting against the market. Finally, we vary whether the existence of manipulators is commonly known and test the effects of this variable on information aggregation and policy decisions.

We study a market with Arrow-Debreu securities, each corresponding to one possible state of the world.[4] These types of assets, dubbed by Wolfers and Zitzewitz (2004) "Winner-take-all contracts", can aggregate dispersed information efficiently even in complex situations (Choo, Kaplan, and Zultan, 2019). After the end of the trading period, managers—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 enables us to estimate the trust that managers place in the market prices and to study how this trust varies according to market activity and the managers' awareness of manipulation attempts.

Notably, while the existence of manipulators in the market was common

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

knowledge in previous laboratory studies, we compare situations with and without common knowledge. This comparison serves two purposes. First, it affects the ability of other traders to counteract manipulation effects actively. Second, it allows us to estimate the effect of knowledge of manipulation on the managers' behavior.[5]

This design aims to capture elements that are, in theory, critical to many information markets rather than the specifics of any one market. The aim is to create a market that successfully aggregates information in prices when there is no manipulation, yet allows manipulators to affect the market. This provides us with a testing ground against which we can study how managers react to markets that are susceptible to manipulation.

2. Experimental design and procedure

Each session included twelve participants who participated in fourteen experimental market periods. The participants were randomly allocated to roles of eight traders and four managers. Two traders—the potential manipulators—are designated as *red* ($\mathscr{R}$) traders and the other six traders as *blue* ($\mathscr{B}$) traders. To facilitate comprehension, all roles ($\mathscr{R}$ traders, $\mathscr{B}$ traders, and managers) were fixed across all periods.

Each period consisted of a trading stage and a voting stage, with different subsets of participants (traders or managers) active in each stage. We manipulated two independent variables in a 2×2 mixed between-within design. First, the existence of manipulators varied within subjects across the market periods, as a random uniform draw determined independently for each period whether the market included manipulators (*Man*) or not (*NoMan*). Second, the results of the random draw were announced only in the Common Knowledge (*CK*) sessions. In the No Common Knowledge (*NCK*) sessions, in contrast, the other traders and managers (i.e., anyone who is not a manipulator) only knew that there is a 50-50 chance that there are manipulators in the market. Table 1 summarizes the four resulting treatments. In the following subsection, we describe in detail the market procedure, followed by the detailed design.

[5]Chen et al. (2015) have shown theoretically in a sequential market with two traders that common knowledge of external incentives moderates the effect of manipulation on the market's ability to reveal private information.

Table 1: Summary of the experimental design.

Treatment	Manipulator traders?	Existence of manipulators announced?	Number of sessions
CK-NoMan	No	Yes	7
CK-Man	Yes	Yes	
NCK-NoMan	No	No	7
NCK-Man	Yes	No	

Note: each session consists of 12 participants.

2.1. Market procedure

Each market involves eight traders (active in the trading stage) and four managers (active in the voting stage). At the beginning of the market stage, the eight traders are randomly allocated into two information groups of four traders each, with the two $\mathscr{R}$ traders always placed in the same information group. The $\mathscr{B}$ traders do not know if they are grouped with the $\mathscr{R}$ traders or not.

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 information group of four traders is then informed that one of the other two states is not the true state of the world. For example, if nature selects state Y, one information group is informed that state X is not true and the other that state Z is not true. The managers do not receive any private information.[6]

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 mechanism. While not common in prediction markets, this mechanism facilitates efficient market outcomes, which is a necessary condition for the study of manipulation and trust in markets (Healy et al., 2010). The procedure is as follows. At the beginning of trade, each trader is endowed with 200 ECU (experimental currency units) and five units of each security type. During the trading duration of 120 seconds, traders can place bids and asks (in the range of 0–20 ECU)—and accept open bids and asks—for each of the three securities. Each trader can have open asks or bids for no more than one unit at any time and short sales are prohibited. When the

[6]That is, the managers only know that the true state is X, Y or Z with equal probabilities.

markets close, each security pays a dividend of 10 ECU if it corresponds to the true state of the world and 0 ECU otherwise.[7] The paid dividends are added to the traders' capital balances to determine their trading stage earnings.

2.1.2. Voting stage

The managers observe all of the transaction prices (but not open bids and asks) in the trading stage and proceed to the voting stage. Each manager casts a vote for one of three policies $\mathscr{X}$, $\mathscr{Y}$, and $\mathscr{Z}$ (corresponding to the three possible states, X, Y, and Z), or for the status quo $\mathscr{Q}$. The policy ($\mathscr{X}$, $\mathscr{Y}$, $\mathscr{Z}$ or $\mathscr{Q}$) that receives the most votes is implemented. In case of a tie, the status quo $\mathscr{Q}$ is implemented by default.[8] This rule is a simple collaborative decision procedure allowing aggregation of information and beliefs, which is easy to follow and reflects the fact that managerial decisions, and in particular those based on information markets, are made collaboratively. On the other hand, this mechanism—as any voting mechanism—has multiple equilibria and is open to strategic manipulation. This, however, requires several implausible assumptions regarding beliefs. Specifically, a manager has a strategic incentive to vote against her preferred outcome only if she believes that the other managers split their votes between two options that she herself finds suboptimal. An alternative procedure, such as a random dictator rule, is not manipulable but is not a common procedure and is less intuitive.

2.1.3. Payoffs from the implemented policy

Independent of the true state, implementing the status quo $\mathscr{Q}$ yields a payoff of 100 ECU for each trader and manager. Participants' payoffs from implementing any of the three policies $\mathscr{X}$, $\mathscr{Y}$, or $\mathscr{Z}$ depend on the state of the world, their role, and the market type.

The $\mathscr{B}$ traders and managers gain 400 ECU from the implementation of the policy that corresponds to the true state and lose 400 ECU from the implementation of any of the policies that correspond to the other two states.[9] Note that

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

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

[9]Suppose that the true state is X, the $\mathscr{B}$ traders and managers receive 400 ECU if policy $\mathscr{X}$ is implemented, -400 ECU if policies $\mathscr{Y}$ or $\mathscr{Z}$ is implemented and 100 ECU if the status quo $\mathscr{Q}$ is implemented.

Table 2: Implemented policy payoffs.

Implemented policy	NoMan markets manager	NoMan markets $\mathscr{B}$ trader	NoMan markets $\mathscr{R}$ Trader	Man markets manager	Man markets $\mathscr{B}$ trader	Man markets $\mathscr{R}$ Trader
True Policy	400	400	400	400	400	−400
Fake Policy	−400	−400	−400	−400	−400	1,000
Neutral Policy	−400	−400	−400	−400	−400	−400
Status quo	100	100	100	100	100	100

Note: The true policy corresponds to the true state. The fake policy corresponds to the state that the $\mathscr{R}$ traders can rule out.

in the absence of the trading stage, voting for the status quo maximizes the expected payoff for a manager unless she assigns a probability of at least 0.625 to one of the three states.[10] This design distinguishes between three managerial responses to the market information: *following* the market (voting for the policy associated with the highest-priced asset), *going against* the market (voting for a different policy), and *ignoring* the market (voting for the status quo). Thus, the inclusion of the status-quo policy allows us to disentangle and identify active trust and mistrust in the market.

The payoff for the $\mathscr{R}$ traders depends on the market type. In the *NoMan* markets, the $\mathscr{R}$ traders receive the same payoff as the other participants in the market. In the *Man* 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 either of the other two policies. This payoff structure incentivizes the $\mathscr{R}$ traders to manipulate prices in the *Man* markets in order to influence the managers' beliefs and consequently the implemented policy.

Henceforth, we refer to true state of the world as the *True* state, the state that the $\mathscr{R}$ traders know not to be true as the *Fake* state, and to the remaining state as the *Neutral* state. For convenience, we maintain this terminology for the corresponding policies and securities.[11] Table 2 summarizes the payoffs from the implemented outcome by market type and role.

[10] Recall that managers do not receive any private information about the true state. This implies that in the absence of the trading stage, the manager should assign equal posteriors to each possible state of the world.

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

2.1.4. Total payoffs

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

$$400 + [\underbrace{L_i + d(x)e_i^x + d(y)e_i^y + d(z)e_i^z}_{\text{Trading stage earnings}}] + \pi_i,$$

where $L_i \geq 0$ is the trader's cash balance at the end of the trading stage, e_i^j is her inventory of security $j = x, y, z$ at the end of trading stage and $d(j) \in \{0, 10\}$ is the dividend of security j. The difference in base payment between traders and managers makes up for the value of the trader's endowment and, therefore, their 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 periods, in which participants could learn the trading mechanism and information structure.[12] Each period followed the design and procedure of the *NoMan* market described above, with the exception that there was no voting stage. Instead, there were no managers, and all twelve participants participated in the role of traders, divided into two information groups of six traders each.

The main part of the experiment consisted of one practice period and fourteen experimental periods. Each period included either a *NoMan* or a *Man* market design with equal probabilities. At the end of each period, subjects received feedback about their own payoff, the true state, the implemented policy, and the market type (*NoMan* versus *Man*). For efficient between-treatment comparisons, we pre-generated a sequence of states and market types, which we implemented in all sessions. We ran seven sessions for each of the *Common Knowledge* (*CK*) and *No Common Knowledge* (*NCK*) treatments.

CK: At the beginning of each period, all participants were informed about whether they are participating in a *NoMan* or a *Man* market.

NCK: At the beginning of each period, only the $\mathscr{R}$ traders were informed about whether they are participating in a *NoMan* or a *Man* market.

[12] See the online appendix for the detailed instructions.

The experiment was conducted at the University of Exeter FEELE laboratory in 2018 and 2019. 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 period (out of five) from the training phase and two periods (out of fourteen) from the experimental 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 5 GBP. The average final payoffs for managers, $\mathscr{B}$ traders, and $\mathscr{R}$ traders were (standard deviations in parentheses): 21.18 (5.46), 23.27 (4.58) and 21.27 (6.31) pounds, respectively, in the *CK* sessions; and 21.93 (4.83), 21.24 (4.77) and 21.76 (6.83) pounds, respectively, in the *NCK* sessions.

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 $\mathscr{R}$ 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 $\mathscr{B}$ traders who are in the same and different information groups as the $\mathscr{R}$ traders as $\mathscr{B}^1$ and $\mathscr{B}^2$ traders, respectively.[13]

We evaluate the market's success at aggregating information by comparing market prices against two benchmark models: the *rational expectations equilibrium* (Radner, 1979, henceforth REE) and the *prior information equilibrium* (Choo, Kaplan, and Zultan, 2019; Plott and Sunder, 1982, 1988, henceforth PIE). The REE and PIE are both static models, which differ with respect to whether beliefs are exogenous or endogenous to the market activity. Given these beliefs, both models assume that the standard principles of supply and demand determine the market prices. Namely, excess demand equals zero.

The REE and PIE are competitive equilibrium models that do not address the strategic incentives faced by manipulators.[14] We, therefore, complement the static equilibrium analysis with a dynamic myopic reasoning model (henceforth, MRM) and consider the strategic possibilities of the manipulators within this framework. The benchmark model considers a simplified discrete-time

[13] Recall that traders do not know whether they are $\mathscr{B}^1$ or $\mathscr{B}^2$ traders.

[14] It is possible to combine rational expectations equilibrium with strategic aspects when decisions affect traders' utility by altering the values of the securities, as in, e.g., Lieli and Nieto-Barthaburu (2020). This is not the case in the current analysis, where traders have incentives that are external to the market.

trading process, wherein supply and demand correspond to traders' beliefs, and beliefs are updated in each period based on the market clearing prices. The MRM generates predictions not only for price convergence, but also for the spread of information in the market. The model was supported by the results of Choo, Kaplan, and Zultan (2019), who found that the traders who learn the true state first according to the model's predictions indeed buy more of the valuable security and are instrumental in price convergence to equilibrium.

In the following subsections, we first analyze the PIE and REE in our setting, followed by the MRM analysis. But first, we remark on how supply and demand determine the market prices. The standard assumption in analyzing experimental markets is that prices converge to the highest valuation in the market (e.g., Plott and Sunder, 1982, 1988). The rationale is that, while the number of securities in the market constrains supply when short sales are prohibited, the high liquidity implies that demand is (in practice) unlimited. In the continuous double auction, however, each trader is limited to active bids or asks of only one unit of any security at a time. The number of units demanded or supplied at any specific time is thus determined by the number of traders willing to buy or sell, respectively, at a given price—regardless of the total number of securities or liquidity in the market. We shall refer to this "local" supply and demand as the *short-run* supply and demand. We use *long-run* supply and demand to refer to the overall supply and demand in the market, which is independent of the trading mechanism.[15]

The equilibrium analysis follows the literature by considering the long-run demand and supply. The dynamic analysis, in contrast, acknowledges that prices temporarily follow the short-run supply and demand, and allows traders to update their beliefs based on these short-run prices. After beliefs stabilze, all traders except for those with the highest valuation will eventually sell their complete inventory, at which point the prices will converge towards the highest valuations.

[15]To illustrate the difference between short-run and long-run supply and demand, consider beliefs at the beginning of trade. All traders value the True security at 5.00 in expectancy, so that excess supply and demand is zero at a price of 5.00. Each of the other two securities has four traders valuing it at 5.00 with the other group valuing it at zero. Due to the high liquidity, there is excess (long-run) demand at any price below 5.00. The short-run demand, however, is limited to the number of traders willing and able to buy—which at any price $0 \leq p \leq 5$ is four, the same as the number of traders willing and able to sell. Thus, any positive price up to 5.00 clears the market.

3.1. Equilibrium analysis

In the REE, beliefs are Bayesian-rational given the market prices. Generically, REE prices reflect the aggregate information held by all traders (Radner, 1979).[16] In our setting, this means that the true state of the world is fully revealed, and the REE prices exactly match the true values of the securities. Let us call the REE in which all traders are informed about the true state of the world and the securities are traded at their true values the *fully revealing equilibrium* (henceforth FRE).

The PIE model, in contrast, assumes that traders form beliefs based on the exogenously provided information only and do not condition expectations on observed prices (*Unsophisticated equilibrium* in the language of Radner, 1979). The PIE model was shown to provide a better fit to observed prices than the REE model in single-asset markets (Corgnet et al., 2018; Corgnet, DeSantis, and Porter, 2015; Plott and Sunder, 1988) and with inexperienced traders (Choo, Kaplan, and Zultan, 2019). The PIE describes the market clearing prices when traders update their beliefs about the true state given their 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 *NoMan* markets, all traders believe the True security to be true with probability 0.5 and therefore value it at 5.00 ECU. For each of the other two securities, one group values it at 5.00 ECU, whereas the other group values it at zero.[17] Given the constrained supply of units in the market, all of the prices will eventually converge in the long run to the highest valuations in the market, which are 5.00 ECU for all three securities.

3.2. Dynamic myopic reasoning model

The MRM assumes that trade takes place over $t \in \{1, 2, \ldots\}$ hypothetical periods (the analysis here follows Choo, Kaplan, and Zultan, 2019). In each period t, traders proceed according to the following three stages:

Stage 1 Traders place bids and asks for each security according to their beliefs.

Stage 2 The market clears at a price that equates the short-run supply and de-

[16]There is an extensive literature studying how REE prices can result from dynamic behavior of traders (e.g., Dubey, Geanakoplos, and Shubik, 1987; Hellwig, 1982; Ostrovsky, 2012).

[17]The $\mathscr{R}$ and $\mathscr{B}^1$ value the Neutral security at 5.00 ECU and the Fake security at 0.00 ECU, and vice versa for the $\mathscr{B}^2$ traders.

mand.[18]

Stage 3 Traders observe the prices and update their beliefs about the true state.

These stages repeat until there is no further belief updating. At this point trade continues until the limited supply is exhausted, and prices converge to the highest valuation in the market.

No Manipulation (NoMan) markets. In the *NoMan* markets, period $t = 1$ beliefs are set by the prior information that the trader holds. All eight traders believe the True security to be true with probability 0.5, and therefore value it at 5.00 ECU. For each of the other two securities, one group values it at 5.00 ECU, whereas the other group values it at zero.[19] Thus, the short-run market-clearing prices of the True, Fake, and Neutral securities will be 5.00, 2.50, and 2.50 ECU, respectively.[20] This price profile uniquely identifies the true state of the world. Therefore, at period $t = 2$, all traders value the True security at 10.00 ECU and the other securities at zero 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 $t \geq 3$. That is, prices converge to the FRE.

Manipulation (Man) markets. To account for the external incentives that the manipulators face, we allow traders to misrepresent their supply and demand and consider the possible dynamics. We first describe price development if the manipulators artificially inflate their demand for the Fake security. We next argue that this is the only plausible dynamic in the markets with manipulators

Assume that at period $t = 1$, the $\mathscr{B}^1$ and $\mathscr{B}^2$ traders set supply and demand based on their private information, whereas the $\mathscr{R}$ traders demand the Fake security at any price up to 10.00 ECU and set their demand of the other securities to zero. The resulting short-run market clearing prices of the True, Fake and Neutral securities are 5.00, 5.00, and 0.00 ECU, respectively. These prices reveal that the Neutral state is not the true state of the world. Therefore, the

[18]As noted above, we assume that traders update their beliefs as transactions occur, i.e., according to the short-run prices). If there is no further belief updating, prices will eventually converge to the long-run prices.

[19]The $\mathscr{R}$ and $\mathscr{B}^1$ value the Neutral security at 5.00 ECU and the Fake security at 0.00 ECU, and vice versa for the $\mathscr{B}^2$ traders.

[20]Any price strictly between zero and 5.00 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.

$\mathscr{R}$ and $\mathscr{B}^1$ 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 $\mathscr{B}^2$ traders are still uninformed about the true state.

This symmetry persists in the next period $t = 2$. The $\mathscr{B}^2$ traders, who value both the True and the Fake securities at 5.00 ECU, form a majority of the market. Hence, there is excess supply (resp. demand) above (resp. below) the price of 5.00 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 there is no more belief updating. Note that there is full symmetry between the True and Fake securities, with four traders trading both based on a value of 5.00, two traders demanding each security at prices up to 10.00, and two traders supplying each security at any price. As the $\mathscr{B}^1$ and $\mathscr{R}$ traders buy all of the True and Fake units, respectively, the long-run prices of these securities will converge towards 10.00 ECU, while the price of the Neutral security remains at zero.

We argue that any dynamic other than the one described above is implausible. To see this, consider first the manipulators' options at $t = 1$, when they cannot distinguish between the True and the Neutral securities. The price of the True security is set at 5.00, its valuation by all the other six traders. Similarly, the same six traders supply the Neutral security at 5.00 ECU, with the four $\mathscr{B}^2$ traders supplying it at any price. The short-run market-clearing price must therefore be strictly below 5.00 regardless of the manipulators' trading strategy. Given these two observations, $t = 1$ prices necessarily discriminate between the True and Neutral states. At $t = 2$, then all traders value the Neutral security at zero, and the two $\mathscr{B}^1$ traders are informed of the True state. It follows that the manipulators can only manipulate the price of the Fake security (even if they could differentiate between the True and Neutral securities at $t = 1$). Furthermore, they cannot push its price above that of the True state. Finally, there is no way for the other traders to differentiate between the True and Fake securities, as the $\mathscr{B}^2$ traders are uninformed, and the $\mathscr{B}^1$ traders already demand the True security and supply the Fake.

We can draw the following conclusions from the MRM analysis. First, manipulators have the power to obscure the true state. Second, manipulators are not able to promote the Fake over the True state. The MRM further predicts that the $\mathscr{B}^1$ traders will be informed and buy the True security, while the manipulator $\mathscr{R}$ traders buy the Fake security. Table 3 summarizes the REE, PIE, and MRM predictions and implemented policies.

Table 3: Theoretical closing-price predictions.

	Security prices			Implemented policy
	True	Fake	Neutral	
Myopic Reasoning Model				
NoMan	10	0	0	True policy
Man	10	10	0	Status quo
Rational Expectations Equilibrium				
	10	0	0	True policy
Prior Information Equilibrium				
	5	5	5	Status quo

4. Results

We commence with the analysis of security prices in the trading stage. After establishing the effects of manipulators and common knowledge on market prices and dynamics, we proceed to look at voting behavior in the voting stage. Finally, we combine the market and voting data to reveal the effects of manipulators and common knowledge on the managers' strategies, and more specifically on the level of trust and mistrust in the market when casting a vote.[21]

4.1. Trading stage

We define the *market closing price* to be the average price over the last five transactions of a security.[22] Figure 1 presents violin plots of the market closing prices by security type and treatment.[23] The findings in the *CK-NoMan* treatment, where it is common knowledge that there are no manipulators, are striking. Prices converge almost perfectly to the FRE prices. Both the median and mode of the True security are equal to the true value of 10.00. The modal prices of the Fake and Neutral securities are zero, with the median prices not far above zero,

[21] We find no systematic or significant effects of experience on prices or on voting behavior. We therefore report results aggregated over all periods. See the online appendix for detailed transaction prices and votes by session and period.

[22] Our interest in market closing prices is consistent with the theoretical analysis of information aggregation as a dynamic process. We take the last five (or, in case there are less than five trades in the market, all) transactions to "smooth out" the volatility in prices. All of the results are robust to using the last ten transactions or the transactions taking place in the last 60 seconds of trade.

[23] Our dataset is comprised of 14 sessions with 14 periods in each session, totalling in 196 markets. There was at least one transaction for each security type in 98.9% of all markets. The market closing prices involve around 48%, 36%, 44% and 33% of all transactions in the *CK-NoMan*, *NCK-NoMan*, *CK-Man* and *NCK-Man* treatments, respectively.

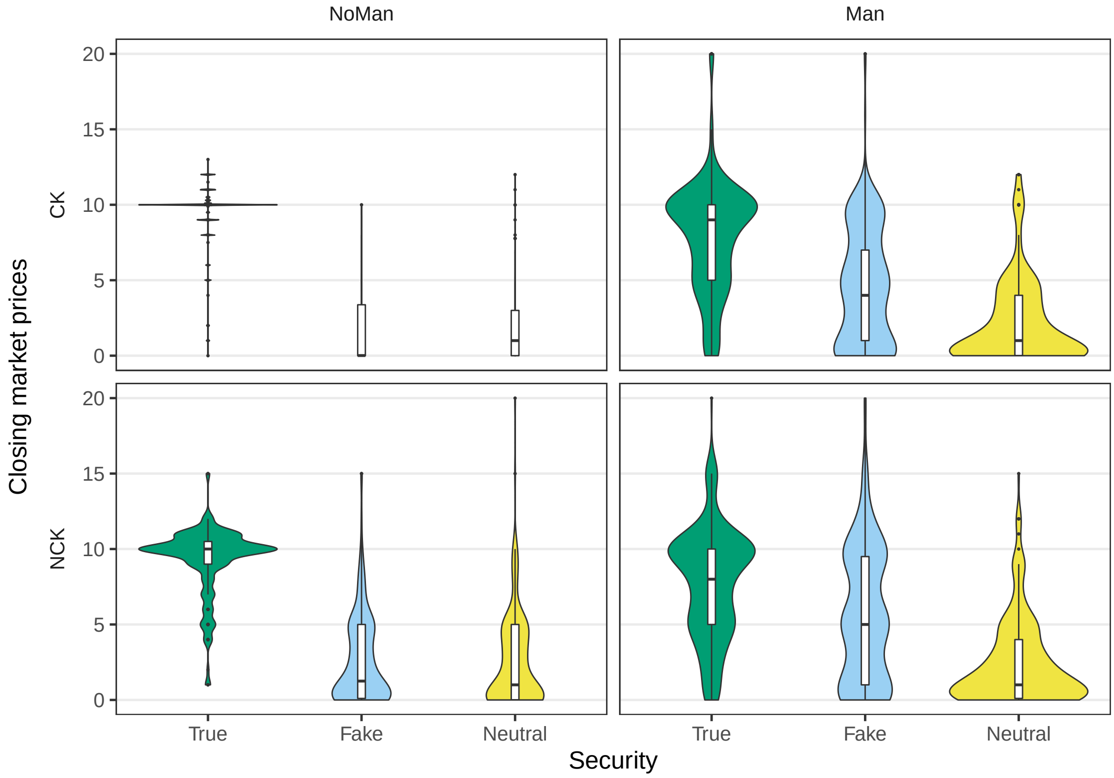

Notes. Market closing prices are defined as the average price in the last five transactions for that security. The violin plots (shaded areas) present kernel distributions of market closing prices. The boxplots (candlesticks) present the median, interquartile region, and outliers.

Figure 1: Market closing prices.

at 0.46 and 1.20, respectively. Thus, we state our first result:

Result 1. *When it is common knowledge that there are no manipulators in the market, Arrow-Debreu markets are successful at aggregating information about the true state of the world into prices.*

The comparison with the *NCK-NoMan* treatment suggests that—even when there are no manipulators in the market—lack of common knowledge leads to suspicion of manipulation and reluctance to learn from prices, as many market closing prices are substantially above or below the true values of the securities. Although the median transaction price for the True security still reflects its true value, and the median prices of the two other securities are not far above zero, many transactions take place at prices further away from the FRE prices. Without common knowledge, only 28.6% of market closing prices are within 1 ECU of the FRE prices, compared to 45.2% with common knowledge. This comparison, however, is not significant ($z = 0.868, p = .386$, clustered Wilcoxon rank-sum test, Rosner, Glynn, and Ting Lee, 2003). In fact, the higher volatility does not affect the ability of market prices to reveal the true state, as the True security is almost always priced above the other two securities regardless of common knowledge (92.9% in *CK* and 95.2% in *NCK*). Our next result summarizes:

Result 2. *The mere suspicion of manipulators—even when there are none in the market—somewhat increases price volatility. Nonetheless, this does not undermine the market's efficacy in revealing the underlying state.*

Moving to the *Man* markets, we see that manipulators have a substantial influence on prices. The median price of the True security is now below the true value of 10.00 ECU, while the prices of the Fake security vary around the PIE price of 5.00 ECU. This pattern is more pronounced in the *NCK-Man* treatment compared to the *CK-Man* treatment. Whereas a Kolmogorov-Smirnov test for the equality of the distributions of the True and Fake securities' prices (averaged within each treatment over all periods) is significant in *CK-Man* ($p = .008$), the distributions do not differ significantly in *NCK-Man* ($p = .575$). This result suggests that, when there is common knowledge that manipulators are active in the market, non-manipulator traders are able to counter the manipulation attempts, albeit only to a small extent.[24] Our third result summarizes the effect of manipulation in the market:

Result 3. *A minority of manipulators are able to harm the information aggregation properties of Arrow-Debreu markets substantially. Common knowledge of manipulation attempts somewhat mitigates the effect of manipulators on market prices.*

Results 1-3 are confirmed by testing price convergence against the theoretical equilibrium predictions. To do so, define for each market the variables M^F, M^P, and M^M as the mean square deviations of the market closing prices of each security from the FRE, PIE, and MRM prices, respectively (cf. Table 3). Note that the MRM predictions coincide with the FRE in the no-manipulation markets and are, therefore, included only in the manipulation markets. Figure 2 plots the means and 95% confidence intervals of M^F, M^P and M^M (based on robust standard errors clustered on sessions). The z-scores are based on clustered Wilcoxon rank-sum (*CK* vs. *NCK*) and signed-rank (pairwise comparisons between M^F, M^P, and M^M) tests (Rosner, Glynn, and Ting Lee, 2003, 2006).[25]

[24] The manipulators are able to gain from manipulating the markets, obtaining mean payoffs of 605 ECU and 667 ECU in the *CK-Man* and *NCK-Man* treatments, resepctively, compared to 250 ECU obtainable by doing nothing, assuming that the True policy is then implemented.

[25] The M^F and M^P values are computed for each of the 196 markets. This resulted in 42 observations each in *CK-NoMan* and *NCK-NoMan*, and 56 observation each in *CK-Man* and *NCK-Man*. The tests cluster on sessions.

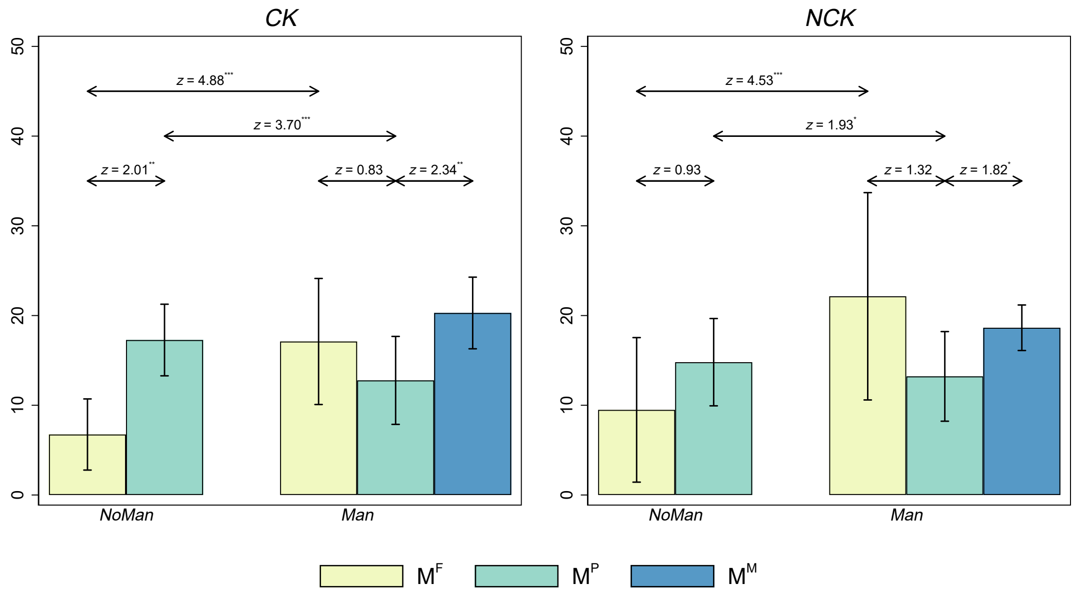

Notes. The M^F, M^P, and M^M are the mean square deviations of closing market prices from the FRE, PIE, and MRM prices, respectively. The z-scores are based on clustered Wilcoxon signed-rank and rank-sum tests. Confidence intervals are based on robust standard errors clustered on sessions. ***, **, and * indicate significance at $p < .001$, $p < .05$, and $p < .10$, respectively.

Figure 2: Mean and 95% CI of M^F, M^P, and M^M by treatment.

In the *CK-NoMan* treatment, closing market prices are significantly closer to the FRE than to the PIE predictions, indicating successful information aggregation. The picture somewhat changes in the *NCK-NoMan* treatment, where the traders are not informed that there are no active manipulators in the market. The FRE provides a better fit in 34 of 42 (81.0%) markets in *CK-NoMan*, but in only 28 of 42 (66.7%) markets in *NCK-NoMan*. The mean deviation of prices from the FRE prices in *NCK-NoMan* is larger than in *CK-NoMan*, and not significantly smaller from the deviations from the PIE prices.[26] Manipulators in the *CK-Man* and *NCK-Man* treatments are successful at impeding information aggregation in prices. In comparison with the no-manipulation treatments, the M^F is significantly higher, whereas the M^P is significantly lower, and is now lower than the corresponding M^F.

The PIE appears to provide a better fit than the dynamic MRM predictions, suggesting that manipulators obfuscate information in the market rather than successfully manipulate prices in favor of the Fake security. This conclusion,

[26]The difference between M^F and M^P in *NCK-NoMan* remain nonsignificant in higher-powered regression analysis with markets as the unit of observation and fixed effects for sessions.

however, relies on the auxiliary assumption that prices converge to the long-run prices—as the deviations from the short-run prices predicted by the MRM are slightly *smaller* than the deviations from the PIE.[27]

4.2. Voting stage

Figure 3 presents the distributions of votes (panel A) and implemented policies (panel B) by treatments. Statistical tests reported below are based on a multinomial logistic regression of the policy voted for based on treatment and standard errors clustered on sessions. In the *CK-NoMan* treatment, managers exhibit high trust in the market, voting for one of the policies, $\mathscr{X}$, $\mathscr{Y}$ or $\mathscr{Z}$, in over 97% of the time. As prices fully reveal the true state, the managers learn from the market, with close to 90% of the votes cast for the True policy, which is consequently implemented in 93% of all markets.

Common knowledge of active manipulators substantially effects policy making, with the True policy implemented in around three-quarters of all markets in the *NCK-NoMan* treatment, a decrease of 16.1 percentage points compared to *CK-NoMan* ($p = .065$). The difference is mostly due to an increase of 10.7 percentage points in status quo votes ($p = .045$), but also a nonsignificant increase of 5.4 percentage points in votes cast for the Fake and Neutral policies. The effect of (lack of) common knowledge on policy making could be attributed to the higher variance of market prices evident in Figure 1, but may also arise from the erosion of trust in the market. We explore this issue in Section 4.3. The following result summarizes the results regarding policy making without manipulators:

Result 4. *When managers know that the market is free of manipulators, they trust the market and implement the True policy with high probability. Conversely, uncertainty regarding the existence of manipulators substantially impedes policy decisions—even when there are no manipulators in the market.*

As we saw, manipulators were successful in undermining the information aggregation properties of the market. This result carries over to the voting stage, with only around half of the votes cast for the True policy in the *Man* compared to the *NoMan* treatments ($p < .001$, separately by common knowledge or combined), and an increase of roughly 15–20 percentage points in votes cast for the

[27]The online appendix provides additional details on the trading behavior in the different treatments. Manipulators ($\mathscr{R}$ traders in the *Man* markets) indeed create artificial demand for the Fake security, despite valuing it at zero. They are nonetheless unable to push prices up to the same level as those of the True security.

A. Votes.

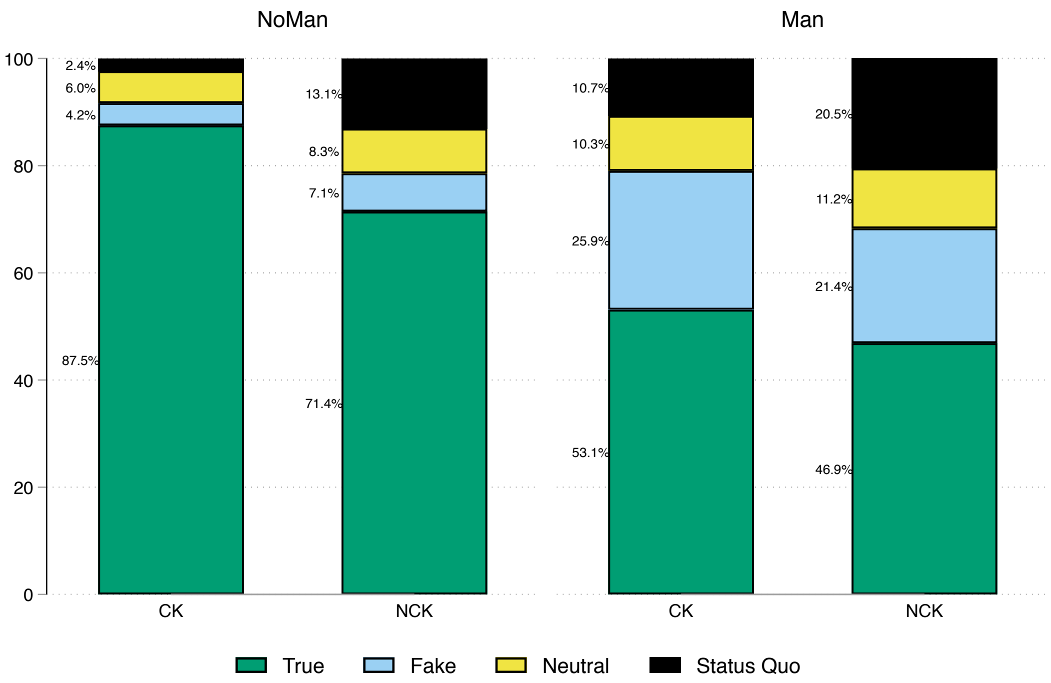

B. Implemented policies.

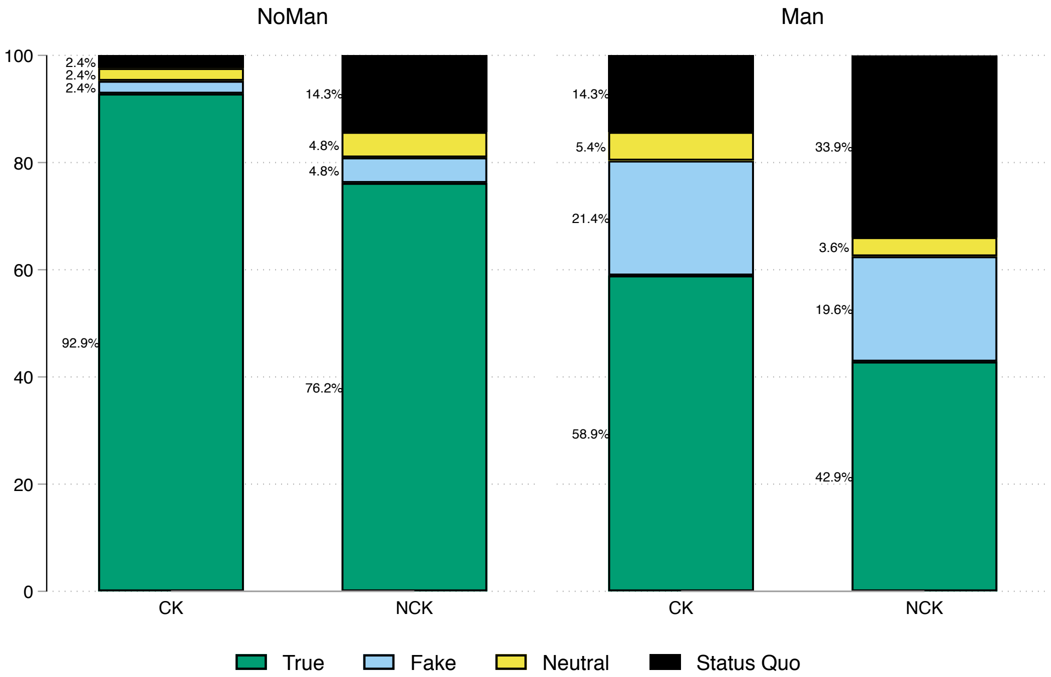

Figure 3: Distributions of votes and implemented policies.

Fake policy ($p = .002$ with common knowledge and $p = .027$ without). Furthermore, in *Man* (but not in *NoMan*), there were significantly more votes for the Fake policy than for the Neutral policy ($p = .002$ in *Man*, $p = .516$ in *NoMan*), which indicates that managers were "tricked" by manipulators rather than simply increased random voting.[28]

Lack of common knowledge appears to lower trust in the market, almost doubling the share of status quo votes in the *NCK-Man* treatment compared to the *CK-Man* treatment, albeit not significantly ($p = .171$). Accordingly, the status quo policy was implemented in one-third of the markets in the *NCK-Man* treatment compared to one in seven markets in the *CK-Man* treatment ($p = .103$).

Once more, we may ask to what extent the effects of manipulators and common knowledge on voting behavior are mediated by the level of information aggregation in market prices, and to what extent are these effects due to engendered mistrust in the market. To address these questions, we now turn to an analysis of the voting strategies. First, we state the result concerning the effect of manipulators on policy:

Result 5. *Manipulators are successful in manipulating around 25% of the votes. Uncertainty about the existence of manipulators leads to less trust in the market, as reflected in more votes cast for the status quo policy in the no common knowledge of manipulators (*NCK*) markets.*

4.3. Voting strategies

In analyzing the voting strategies, we consider whether managers vote in line with the observed market prices. To do so, write S_1 and S_2 for the securities with the highest and second-highest market closing prices, respectively.[29] We categorize all votes into three categories accordingly:

Following the market. Voting for the policy corresponding to S_1 security.

Opposing the market. Voting for a policy that is not associated with the S_1 security.

[28] To see if managers are attempting to learn how to respond to manipulation from past experiences, we regressed whether or not managers follow the market on whether the True asset was priced highest in the previous period (including random effects for managers and clustering standard errors on sessions). The coefficients in the manipulation markets are small, not systematic in sign, and far from being significant.

[29] If the market closing prices of two securities are equal, we break the tie according to the average prices in the last ten (rather than five) transactions.

Table 4: Shares of managers who follow, oppose and ignore the market.

	CK-NoMan	*NCK-NoMan*	*CK-Man*	*NCK-Man*
Follow the market	89.3%	75.6%	58.9%	58.0%
	(4.4%)	(6.4%)	(6.8%)	(6.3%)
Oppose the market	8.3%	11.3%	30.4%	21.4%
	(3.6%)	(3.1%)	(5.0%)	(5.8%)
Ignore the market	2.4%	13.1%	10.7%	20.5%
	(1.2%)	(5.2%)	(3.1%)	(6.4%)
n	168	168	224	224

Note. Robust standard errors clustered on sessions based on a multinomial logistic regression in parentheses.

Ignoring the market. Voting for the status quo.

Table 4 reports the proportion of managers in each treatment who follow, oppose or ignore the market. All statistical tests reported in this section are based on a multinomial logistic regression predicting the vote category based on the treatment with standard errors clustered on sessions.

Managers must choose a voting strategy based on the observed market activity and on the trust they place in the market. Trust, in turn, is influenced by the observed market activity and by the managers' prior information regarding manipulation. Given the infinite trading profiles in continuous time, the full strategy space is non-tractable. To estimate the extent to which managers can extract information from the market, we therefore consider the expected payoffs obtained if all managers follow a *simple heuristic* based on the ability of market closing prices to differentiate between the securities.

Let P_1 and P_2 be the corresponding market closing prices of securities S_1 and S_2, respectively. The heuristic, henceforth α*-strategy*, takes one parameter, $0 \leq \alpha \leq 1$, which can be interpreted as the degree to which the voter trusts the market. An α-strategy dictates following the market if and only if $P_2/P_1 \leq \alpha$, and otherwise ignore the market and vote for the status quo.[30] Note that $\alpha = 0$ implies always ignoring the market, for the status quo payoff of 100 ECU. As α increases, the manager trusts the market more and is willing to follow the market for lower price differentiations. At the upper end, we have full trust in the market at $\alpha = 1$, for which the manager always follows the market.[31]

[30]For example, if the closing prices of securities x, y and z are 10, 4, and 2 ECU, respectively, then an α-strategy dictates voting for policy $\mathscr{X}$ for any $\alpha \geq 0.4$, and vote for the status quo $\mathscr{Q}$ otherwise.

[31]To confirm that the family of α-strategies plausibly approximates actual voting strategies, we use a logistic regression to predict whether the managers vote for the policy corresponding

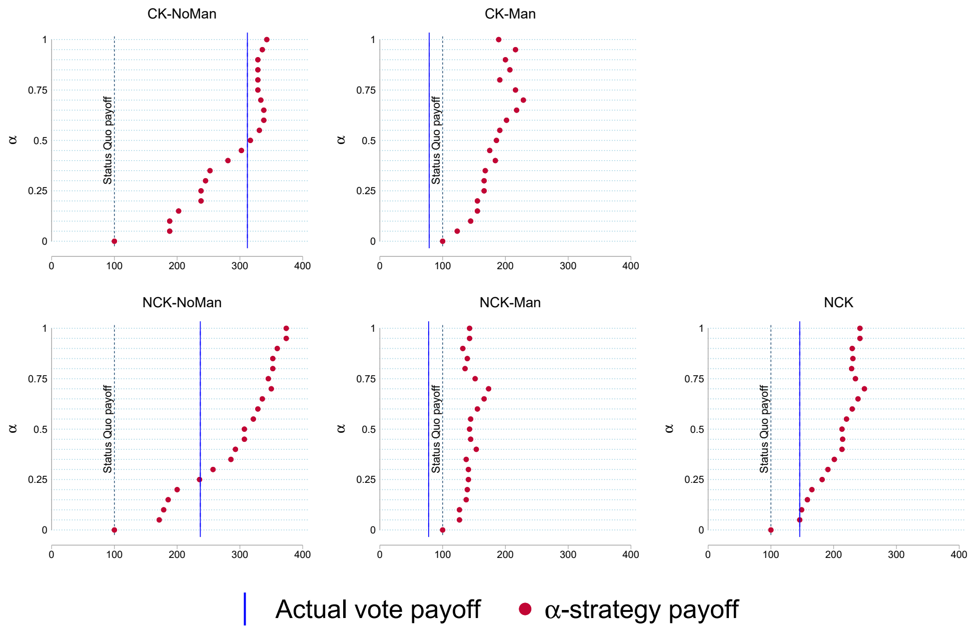

Note: dots mark the expected payoff from following the market if $P_2/P_1 \leq \alpha$ and voting for the status quo otherwise. The horizontal line marks the average payoff corresponding to the managers' votes.

Figure 4: Expected payoff for decisive votes.

For each market and value of the parameter α from zero to one in steps of 0.05, we calculate the price ratio P_2/P_1. If the price ratio is larger than α, the α-strategy dictates implementing the status quo for a payoff of 100. If the price ratio is lower (indicating high price differentiation), the payoff is 400 if the true security is the highest priced, and -400 otherwise.

Figure 4 plots the mean payoffs obtained by this procedure across markets, by treatment (in different panels) and the parameter α (indicated by the horizontal lines). In the *NCK* treatments, it is not clear ex-ante whether the managers can distinguish whether there are manipulators in the market. The figure, therefore, includes separate panels by manipulation treatments as well as a panel showing the combined results. The solid vertical line marks the mean actual payoff obtained in each treatment by determining the payoff of each manager based solely on her own actual vote (i.e., the payoff a manager would receive if her vote were always pivotal). Finally, the dashed vertical lines mark the status quo payoff of 100 ECU.

to S_1 on P_2/P_1 and the treatment, with robust standard errors clustered on sessions. The coefficients for P_2/P_1 are highly significant ($p < .001$) for all four treatments, and take values between -0.359 and -0.433.

4.3.1. Non-manipulation markets

In the *CK-NoMan* treatment, we see that even a little trust in the market can lead to substantial gains. The highly conservative strategy of following the market only if P_1 is at least 20 times larger than P_2 (i.e., $\alpha = 0.05$) leads to a mean payoff of 188 ECU, almost twice the status quo payoff. Moderate to high trust in the market ($\alpha \geq 0.5$) yields mean payoffs of more than 300 ECU, maximized at full trust in the market ($\alpha = 1$), with a mean payoff of 343 ECU.[32] Managers indeed trust the market, following the market in 89.3% of cases, ignoring the market in 2.4% of cases, and opposing the market in only 7.7% of cases, for a mean actual-vote payoff of 312 ECU.

The comparison to the *NCK-NoMan* treatment is illuminating. We see that—despite the lessened information aggregation in prices—prices are highly informative, with α-strategy payoffs as high as obtained in the *CK-NoMan* treatment. Blindly following the market (i.e., $\alpha = 1$) yields a high payoff of 374 ECU, not much less than the 400 ECU managers could obtain if they knew the True state for certain! Nevertheless, the actual voting behavior reveals low trust in the market. As noted above, the share of status quo votes (ignoring the market) increases from 2.4% in the *CK-NoMan* treatment to 13.1% in the *NCK-NoMan* treatment ($p = .045$). The share of managers who oppose the market also increases, from 8.3% in *CK-NoMan* to 11.3% in *NCK-NoMan*, although the difference is not statistically significant ($p = .532$). Consequently, the payoff based on actual votes is only 237 ECU, only 63% of the payoff obtainable by trusting the market fully, and a payoff comparable to the 235 ECU obtained from following the market only if P_1 security is at least *four* times larger than P_2 ($\alpha = 0.25$). We therefore conclude that the suboptimal voting observed in the previous section is not due to volatility in prices, but to mistrust in the market due to uncertainty regarding price manipulation. We first state the result regarding the ability to learn from markets without manipulators:

Result 6. *When there are no manipulators in the market, voting according to the security with the highest closing price extracts around* 80–90% *of the possible gains with respect to the status quo, regardless of whether the non-existence of manipulators is common knowledge.*

Next, we state the result regarding (mis)trust in the market:

[32] The payoff curve flattens above $\alpha = 0.5$ because the price ratio P_2/P_1 mostly falls below 0.5 in the *CK-NoMan* treatment markets.

Result 7. *When there is common knowledge that there are no manipulators in the market, managers trust the market, voting according to the observed transaction prices, and extracting most of the potential gains. Lack of common knowledge has a dramatic effect on trust, with many votes cast for the status quo, leading to suboptimal policies and considerable loss of efficiency.*

4.3.2. Manipulation markets

The effect of manipulators on voting behavior we saw in the previous section is evident in the middle column of Figure 4. The figure reveals that there are two separate effects in play. First, as can be expected based on the analysis of market prices, the amount of information in the market is substantially diminished with manipulators. Perhaps surprisingly, we see that there is nonetheless still much to gain from trusting the market. In the *CK-Man* treatment, the payoffs for high enough trust ($0.5 \leq \alpha \leq 1$) are in the range of 180 to 230 ECU. The situation is considerably worse when the other traders are not explicitly informed of the existence of manipulators in the market. The corresponding payoffs in the *NCK-Man* treatment are in the lower range of 130 to 175 ECU, though still substantially above the status quo payoff of 100 ECU.

The second effect is observed in the actual-vote payoffs, which are below the status quo payoff in both the *CK-Man* and *NCK-Man* treatments. That is, managers not only forgo the potential gains from *trusting* the market—which could be explained by risk aversion—but are even doing *worse* than they would by always *ignoring* the market. This implies that knowledge or suspicion of manipulation leads managers to oppose the market, even though high price differentiation typically indicates that the market was successful in reflecting the true state of the world. For example, when $P_2/P_1 \leq 0.25$ (i.e., P_1 is at least four times larger than P_2), S_1 is the True security in 14 of 15 (93.3%) markets in *CK-Man* and in 11 of 13 (84.6%) markets in *NCK-Man*.

Indeed, whereas 89.3% of votes in the *CK-NoMan* treatment and 75.6% of votes in *NCK-NoMan* treatments go to the policy corresponding to the S_1 security, these shares drop to 58.9% and 58.0% in the *CK-Man* and *NCK-Man* treatments, respectively ($p < .001$ for the separate and combined comparisons). The share of status quo votes significantly increases from 2.4% in the *CK-NoMan* treatment to 10.7% in the *CK-Man* treatment ($p = .004$) and from 13.1% in the *NCK-NoMan* treatment to 20.5% in *NCK-Man* treatment ($p < .07$). The share of managers opposing the market also increases significantly, from 8.3% in the *CK-NoMan* treatment to 30.4% in the *CK-Man* treatment ($p < .001$) and from 11.3% in the *NCK-NoMan* treatment to 21.4% in the *NCK-Man* treatment ($p = .008$).[33]

The considerable difference in voting strategies between the *NCK-NoMan* and *NCK-Man* treatments shows that the market activity provides enough information for managers to figure out (to a large extent) whether there are manipulators in the market, and to respond to the price ratio accordingly. Nonetheless, we can ask how a simple heuristic that conditions only on the price ratio and ignores all other market information fares. The bottom right panel in Figure 4 shows that such a heuristic can yield substantial gains, with full trust in the market yielding a payoff of 242 ECU. Actual behavior, however, reveals very low trust in the market, potentially yielding a payoff comparable to that obtained with $\alpha = 0.05$.

Result 8. *Manipulators affect managers' decisions via two channels. First, they manipulate market activity sufficiently to obscure the information reflected in market prices, though not sufficiently to eliminate the advantage in following the market completely. Second, managers who know or suspect manipulation tend to ignore or even vote against the market, and are, therefore, unsuccessful in utilizing the information conveyed by the market prices.*

4.3.3. Market heterogeneity

The results stated above refer to the average market behavior. While the overall performance is crucial for policy guidance, it is also important to understand whether the conclusions hold for individual markets as well as on average. Figure 5 plots the expected earnings from following the market (α=1 in Figure 4) against the expected earnings based on votes (solid vertical line in Figure 4) for each session, with and without manipulation. The dotted lines mark the status-quo payoff of 100. Observations below the diagonal indicate that managers' votes imply a lower average payoff than could be obtained by following the market.

We see that Results 6 and 7, regarding the non-manipulation markets, broadly hold for the individual sessions. Following the market guarantees, on average, more than the status quo payment, and is optimal in most markets. With common knowledge, managers largely extract the possible earnings from following the markets, even surpassing this benchmark in the market where prices were

[33]Note that this tendency to oppose the market rules out risk aversion as an explanation for the loss of potential gains.

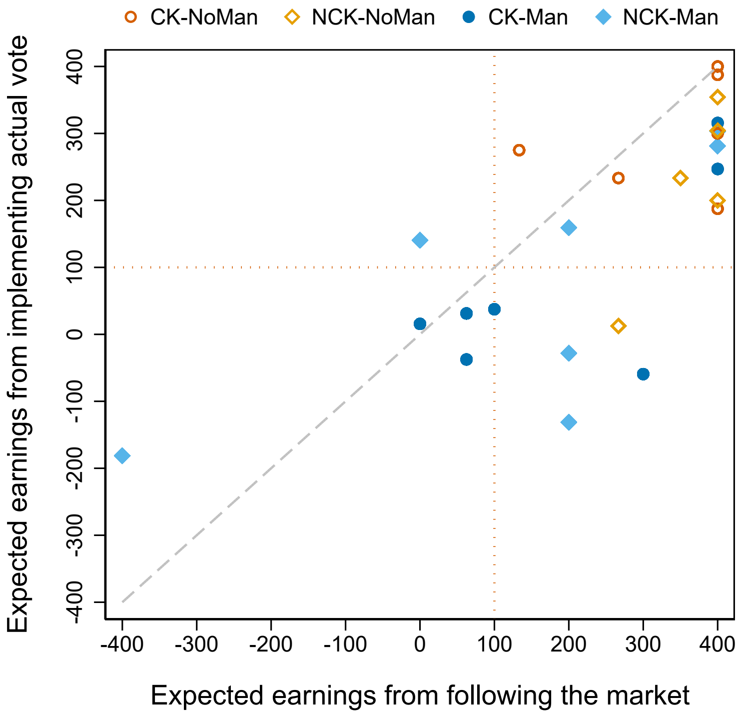

Note: Each point corresponds to one session. Points under the diagonal indicate that managers' actual votes correspond to a payoff lower than could be obtained by voting in line with the highest priced security. The dotted lines mark the status-quo payoff.

Figure 5: Expected earnings by sessions.

least informative. Without common knowledge, in contrast, earnings based on votes are generally further below the diagonal, indicating the erosion of trust under lack of common knowledge.

Manipulation markets similarly reveal a lack of trust in most markets. There are two noteworthy observations, however. First, in as many as five out of the fourteen markets, implementing the status quo is better than following the market. Second, managers are most likely to follow the market in the four sessions in which the prices were always informative (between 75% and 87.5% votes to follow the market, compared to 34.37% to 68.75% in the other ten markets). While the number of sessions provides insufficient power for formal tests, this suggests that managers can identify efficient markets. Taken together, these observations highlight the importance of market-specific attributes and investment in identifying successful manipulation.

5. Conclusion

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

Prediction markets proved successful in predicting real-world events (Wolfers and Zitzewitz, 2004). Accordingly, we constructed experimental markets where prices fully reveal the aggregate information in the market. Our design provides a new understanding of the potential effect of manipulators and the role of uncertainty on managerial decision making. Our results highlight the importance of trust in the market. When the market designer cannot guarantee that the market is free of manipulators, trade volatility increases. Market prices, however, still provide ample information for managers to reach close to optimal decisions. Nonetheless, managers lose trust in the market, leading to substantial loss of potential gains.

This result is reflected in the manipulation markets. While manipulators are able to manipulate the trading activity considerably, the market prices still reflect sufficient information to improve managerial decision making, in most markets and on average. Managers who are aware of the manipulation, however, mistrust the market and fail to utilize the information conveyed in the prices. Managers who are not explicitly informed of the manipulation nonetheless exhibit a similar distrust, often ignoring or voting against the market, indicating that managers are able to identify the existence of manipulation. This results in similar suboptimal policy decisions in the face of manipulation with and without common knowledge of the manipulation.

The experimental test of manipulation and trust in markets necessarily involves several design choices that may affect the results. Our markets are designed to capture core elements that are common to many information markets. Yet, the aggregation properties of the markets do not broadly generalize to other trading mechanisms or information structures, which may impact managerial trust. Managers' behavior may also vary depending on other aspects of the decision making environment, such as the voting mechanism or the ability to share information outside the market.

Notwithstanding these caveats, our findings point at general patterns and future research directions that bear important implications for the design of prediction markets and for decision making based on natural observations in financial markets. We find that existence or even suspicion of manipulation erodes trust in the market and managers' willingness to take risks based on information gathered from the market. Thus, the design of prediction markets should consider not only the information aggregation properties of the market—a topic that has received ample attention in the literature and in practice—but also ways to breed trust in predictions based on the market. Furthermore, while more trust improves managerial decision making overall, market prices are not informative in all markets. It would therefore also be beneficial to assist managers in identifying which markets merit more trust than others.

Regulating participation as a precaution to prevent manipulation can forward this cause not only by reducing actual manipulation, but also—if conveyed properly to decision makers—by increasing trust. Even if participation is not regulated, increasing familiarity and experience with markets can also improve trust and facilitate efficient managerial decision making (Maciejovsky and Budescu, 2020).

In this paper, we present an experimental paradigm to study manipulation in prediction markets. Our results affirm the usefulness of this paradigm as a testbed for studying managerial reliance on prediction markets. This paradigm can be extended in several ways to further study and develop tools to improve managerial decision making by promoting trust in the market. Future research will test the efficacy of various design tools. For example, studying managerial decision making under alternative institutions compared to the institutions implemented in the experiment (the continuous double auction in the trading stage or the plurality rule in the voting stage); providing managers with explicit information regarding the past performance of similar markets to enhance trust; implementing conflict of interest statements from traders to increase reliability and accountability in the market; or introducing an automated market maker such as the logarithmic market scoring rule (Hanson, 2003), which potentially restricts the possible scope of manipulation by placing constraints on price movements. Furthermore, the heterogeneous effects of manipulation across markets raise the importance of not only developing tools to identify manipulation (Blume, Luckner, and Weinhardt, 2010; Kloker and Kranz, 2017), but also of communicating to managers the estimated trustworthiness of specific markets. Finally, while various studies compared prediction markets to other methods of information aggregation (Healy et al., 2010), the focus is typically on the different institutions' ability to predict future events. In view of our results, more attention should be given to the ability of different institutions to promote trust in their predictions (Maciejovsky and Budescu, 2020).

References

Arrow, Kenneth J., Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz (2008). The promise of prediction markets. *Science* 320.(5878), pp. 877–878.

Blume, Michael, Stefan Luckner, and Christof Weinhardt (2010). Fraud detection in play-money prediction markets. *Information Systems and E-Business Management* 8.(4), pp. 395–413.

Bond, Philip and Itay Goldstein (2015). Government intervention and information aggregation by prices. *Journal of Finance* 70, pp. 2777–2812.

Chakraborty, Mithun and Sanmay Das (2016). Trading on a rigged game: outcome manipulation in prediction markets. In: *Ijcai*, pp. 158–164.

Chen, Kay-Yut and Charles R. Plott (2002). Information aggregation mechanisms: concept, design and implementation for a sales forecasting problem. Social Science Working Paper 1131. California Institute of Technology.

Chen, Yiling, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen (2010). Gaming prediction markets: equilibrium strategies with a market maker. *Algorithmica* 58.(4), pp. 930–969.

Chen, Yiling, Xi Alice Gao, Rick Goldstein, and Ian A. Kash (2015). Market manipulation with outside incentives. *Autonomous Agents and Multi-Agent Systems* 29.(2), pp. 230–265.

Chen, Yiling, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen (2007). Bluffing and strategic reticence in prediction markets. In: *Proceedings of the 3rd international conference on internet network economics*. Ed. by X. Deng and F. C. Graham. Springer, pp. 70–81.

Choo, Lawrence, Todd R. Kaplan, and Ro'i Zultan (2019). Information aggregation in Arrow-Debreu markets: an experiment. *Experimental Economics* 22, pp. 625–652.

Corgnet, Brice, Cary Deck, Mark DeSantis, and David Porter (2018). Information (non) aggregation in markets with costly signal acquisition. *Journal of Economic Behavior & Organization*.

Corgnet, Brice, Mark DeSantis, and David Porter (2015). Revisiting information aggregation in asset markets: reflective learning & market efficiency. *mimeo*.

Cowgill, Bo and Eric Zitzewitz (2015). Corporate prediction markets: evidence from Google, Ford, and Firm X. *The Review of Economic Studies* 82.(4), pp. 1309–1341.

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.

Dianat, Ahrash and Christoph Siemroth (2020). Improving decisions with market information: an experiment on corporate prediction markets. *Experimental Economics*, pp. 1–34.

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.

Fama, Eugene F. (1970). Efficient capital markets: a review of theory and empirical work. *The Journal of Finance* 25.(2), pp. 383–417.

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

Gandal, Neil, J. T. Hamrick, Tyler Moore, and Tali Oberman (2018). Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics* 95, pp. 86–96.

Gillen, Benjamin J., Charles R. Plott, and Matthew Shum (2017). A pari-mutuel-like mechanism for information aggregation: a field test inside Intel. *Journal of Political Economy* 125.(4), pp. 1075–1099.

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 (2003). Combinatorial information market design. *Information Systems Frontiers* 5.(1), pp. 107–119.

— (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, 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.

Hayek, Friedrich A. (1945). The use of knowledge in society. *American Economic Review* 35.(4), pp. 519–530.

Healy, Paul J, Sera Linardi, J Richard Lowery, and John O Ledyard (2010). Prediction markets: alternative mechanisms for complex environments with few traders. *Management Science* 56.(11), pp. 1977–1996.

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.

Jian, Lian and Rahul Sami (2012). Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution. *Management Science* 58.(1), pp. 123–140.

Kloker, Simon and Tobias T. Kranz (2017). Manipulation in prediction markets – chasing the fraudsters. In: *Proceedings of the 25th european conference of information systems.* 5th–10th June. Guimarães.

Lieli, Robert P. and Augusto Nieto-Barthaburu (2020). On the possibility of informative equilibria in futures markets with feedback. *Journal of the European Economic Association* 18.(3), pp. 1521–1552.

Maciejovsky, Boris and David V. Budescu (2013). Markets as a structural solution to knowledge-sharing dilemmas. *Organizational Behavior and Human Decision Processes* 120.(2), pp. 154–167.

— (2020). Too much trust in group decisions: uncovering hidden profiles by groups and markets. *Organization Science* 31.(6), pp. 1497–1514.

Muth, John F. (1961). Rational expectations and the theory of price movements. *Econometrica* 29.(3), pp. 315–335.

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

Palan, Stefan, Jürgen Huber, and Larissa Senninger (2019). Aggregation mechanisms for crowd predictions. *Experimental Economics*, pp. 1–27.

Plott, Charles R. (2000). Markets as information gathering tools. *Southern Economic Journal* 67.(1), pp. 1–15.

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.

Plott, Charles R. and Shyam Sunder (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.

Rosner, Bernard, Robert J. Glynn, and Mei-Ling Ting Lee (2003). Incorporation of clustering effects for the Wilcoxon rank sum test: a large-sample approach. *Biometrics* 59.(4), pp. 1089–1098.

— (2006). The Wilcoxon signed rank test for paired comparisons of clustered data. *Biometrics* 62.(1), pp. 185–192.

Veiga, Helena and Marc Vorsatz (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.