

Price Manipulation in Prediction Markets: Analysis and Mitigation

Eric H. Huang
Computer Science Department
Stanford University
ehuang@stanford.edu

Yoav Shoham
Computer Science Department
Stanford University
shoham@stanford.edu

ABSTRACT

We consider the possible existence of a manipulator in a prediction market, whose incentive is to maximally increase the predicted probability of an event, and for whom profit or loss in the market is immaterial. We characterize the equilibria in a single-round market scoring rule (MSR), showing that the manipulator will play a strategy that mixes between pretending to have received one of the top signals. We propose a modification to the MSR in the form of trade limits, a maximum amount by which the price of the security can change at a given round. We show analytically that without a manipulator, this process converges to the true posterior, and computationally that in a market with a manipulator, the limits help reduce the distortion by the manipulator when traders do not know about the manipulator's existence. Specifically, we show through simulations that with high probability the honest traders will fully reveal their signals before the manipulator does, and that the price at this point of full revelation by the honest traders can be a significantly better approximation of the true posterior than the ultimate price reached, suggesting a rule by which the market should be stopped at that point.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

Keywords

Economics; Game Theory; Economically-Motivated Agents

1. INTRODUCTION

In many scenarios the opinions of a group of informed experts on the outcome of a future event can be aggregated to form an accurate prediction. One way to elicit and aggregate these opinions is via prediction markets, which are trading mechanisms that provide incentives for participants to share their information and in the process aggregate that information. Prediction markets have been shown to be very successful in several real-world implementations (Iowa Electronic Market [1], Google [5], HP [10], etc.).

Most work in prediction markets to date has assumed that the traders are motivated purely by their profit or loss in

the market. However, there are often externalities that introduce additional incentives. In particular, traders may have an external reason for preferring the aggregated prediction to turn out a certain way, and may bias their trading accordingly. In extreme cases they may care *only* about the aggregate prediction, with the profit or loss in the market being immaterial. In this paper we will assume such profit-indifferent agents, and furthermore assume that their specific goal is to maximize the aggregate prediction. As a concrete motivation, imagine a prediction market for citation counts of scientific papers (which is in fact our original motivation for this work), with a typical security paying off \$1 if paper x has more than y citations to it within z years. The author of paper x (perhaps an assistant professor facing a tenure decision) would like to push the price of the security as close to \$1 as possible, and surely would be willing to pay more than \$1 for this.

To model this formally, we start with a popular class of prediction market mechanisms, the market scoring rules (MSR) [7], which are sequential versions of scoring rules. A scoring rule is a function of a set of event outcomes and their probability; for example, (“the paper will get more than 100 citations in two years”, 0.8). In MSRs, the market starts with some prior belief regarding the event, and the traders enter trades sequentially, each time moving the probability of the event and agreeing to pay a price that is the difference between the scoring rules applied to the event with the two probabilities. The prediction is the final probability arrived at.

MSR mechanisms are well understood and (with so-called “proper” scoring rules) are also well behaved. But when a manipulator is present, things change. One way to think about it is that the honest traders receive private signals from a joint distribution of outcomes and signals, while the manipulator pretends to have some signal, hoping to mislead the honest traders in order to drive up price.

We consider a restricted yet subtle setting with two traders: an honest trader and a manipulator, where the honest trader has a prior belief on the manipulator's existence, the value of which may correspond to describing three different real-world situations, i.e. the manipulator's existence being 1) public where the manipulator's existence is known to the honest trader, 2) secret where the honest trader does not believe that there is a manipulator, or 3) private where the honest trader is unsure of the manipulator's existence. We analyze equilibria under these three different belief structures of the honest trader, and show that 1) with public information, the manipulator's trade is simply ignored by the honest

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

trader, 2) with secret information, the manipulator will pretend to be an honest trader with the highest possible signal, and 3) with private information, the manipulator will mix between pretending to have one of the highest signals.

We then propose a modification to the market scoring rule in order to reduce the magnitude of price distortion in the worst setting where the manipulator’s existence is *secret* information, which is when the price can be most distorted. Specifically, we impose *trade limits*; that is, a maximum amount by which the market probability can change at a given round.

We first show that in a normal market without a manipulator, and without suspicion of there being a manipulator, the market with large enough trade limits converges to the true posterior. But of course inventing a slower MSR is hardly our goal, and this is just a sanity check that we have not broken something fundamental in the logic of the mechanism.

The trade limits pay off when there is indeed a manipulator. The intuition is that while the traditional market scoring rule allows every trader to express his entire information all at once (including the manipulator with his fake information), the trade limits delay the expressing of information over several rounds. When a trader makes a trade that does not use up all of the limit, the trader has fully revealed his signal. The key observation is that the manipulator is more likely to be the last trader to fully reveal his (fake, extremely high) signal than the truthful traders, in which case the rest of the trading can be ignored (or indeed stopped).

We show computationally that indeed with high probability the manipulator is the last one to fully reveal his signal, in which case the early-stopping rule is obviously beneficial, and sometimes lead to predictions that are substantially closer to the true posterior of the honest agents than the ultimate prediction if trading is allowed to continue to the end. Furthermore, we show that in the relatively rare cases in which the manipulator is not the last one to fully reveal his signal, the resulting loss from the early-stopping rule is low (and indeed frequently zero).

The rest of the paper is organized as follows. In section 2, we go over related work. In section 3, we review the traditional model and mechanism for prediction markets. In section 4, we introduce and analyze the new setting where a manipulator may be present. In section 5, we introduce our modified mechanism with trade limits and discuss the market behavior with and without a manipulator. Moreover, we show via simulations how trade limits can reduce manipulability. Finally, we conclude in section 6.

2. RELATED WORK

Many prior works have studied the effect of outside incentives on prediction markets. Chen et al. [4] consider a two-player market scoring rule where a manipulator with outside incentives in the form of a monotone function of the final market probability may exist. They employ a fixed setting where the manipulator trades first, followed by the truthful trader, who has a prior on the manipulator’s existence, and reasons about the observed move. They consider cases where the manipulator’s incentive is common knowledge or not, and characterize the conditions for which a separating equilibrium may or may not exist, as well as the associated information loss. Our work is different in that the manipu-

lator in our model does not care about profits in the market, and our model includes more than two signals. In addition, we propose a new mechanism to reduce manipulability in the market under some conditions.

Boutilier [2] models a scenario where a decision maker tries to elicit forecasts from self-interested experts, who are incentivized to misreport to steer decisions in their favor. He proposes a compensation rule that, when combined with the expert’s utility, induces a net scoring rule that behaves like a traditional proper scoring rule. This, however, relies on knowledge of the expert’s utility, as well as the ability to compensate the expert enough. We are interested in the setting where the manipulator’s profit or loss in the market is immaterial, thus cannot be compensated with extra incentives from the market maker.

Hanson and Oprea [8] show an interesting result for Kyle-style markets. Namely, the manipulator’s mean target price has no effect on prices, but increasing the variance of the manipulator’s target price actually increases average price accuracy. However, they assume that the mean and variance of the distribution from which the target price is drawn is common knowledge. Also, the results apply to *average* price accuracy, so it is not that surprising that traders can adjust prices knowing the mean target price, canceling out its effect on average. Finally, the results do not apply to MSR’s.

Other types of manipulation in prediction markets include influencing the outcome of the event associated with the market [11, 9] as well as misleading other traders in order to profit at later trades in the same market [3, 6].

3. REVIEW: TRADITIONAL MODEL AND MECHANISM

We will review the traditional model of market scoring rules for two players, but note that MSRs are general mechanisms for any number of traders.

3.1 Information Structure

Consider a Boolean random variable X . We would like to aggregate information from two agents to predict its realization $x \in \{0, 1\}$. Each participant i in the market receives a private signal s_i , which is the realization of a discrete random variable S_i . We assume that the joint prior probability distribution of X, S_1, S_2 is common knowledge. To simplify notation, we will sometimes denote $P(X = 1)$ as $f_{0,0}$, $P(X = 1|S_1 = s_1)$ as $f_{s_1,0}$, $P(X = 1|S_2 = s_2)$ as f_{0,s_2} , and $P(X = 1|S_1 = s_1, S_2 = s_2)$ as f_{s_1,s_2} (notations borrowed from Chen et al.[4]). Let Σ_i denote the set of possible values S_i can take. Each signal in each player’s signal space is *unique*, i.e. $P(X = 1|S_i = u, S_{-i} \in \Sigma'_{-i}) \neq P(X = 1|S_i = v, S_{-i} \in \Sigma'_{-i})$ for any two different signals $u, v \in \Sigma_i$, and any subsets Σ'_{-i} of Σ_{-i} , where $-i$ denotes the player that is not i . Assume also that the signals have a *strict ordering*, that is, $f_{s_1^{(1)},s_2} > \dots > f_{0,s_2} > \dots > f_{s_1^{(m)},s_2}$ for all m signals in Σ_1 and any $s_2 \in \Sigma_2$; and similarly for player 2. We say that a signal u is *higher* than another signal v if the posterior given u is higher than that given v .

3.2 Traditional Mechanism: Market Scoring Rule

A market scoring rule is the sequential version of a scoring rule. To predict the outcome of a binary event X , the market starts with a probability estimate p_0 of the outcome $x =$

1. Traders then interact with the market one by one, each time changing the current market probability estimate p to some new probability q , and agreeing to pay the difference $s(x, q) - s(x, p)$, where s is the scoring rule. There are many proper scoring rules, of which the logarithmic market scoring rule is one of the more popular: $s(x, p) = b \log(p)$ if $x = 1$, and $s(x, p) = b \log(1 - p)$ if $x = 0$, where b is a parameter.

3.3 Participants and Strategies

We assume participants to be rational and risk-neutral utility maximizers whose utility function is derived from the profit or loss in the market, as determined by the scoring rule. Market scoring rules are attractive because it is myopically incentive compatible when a proper scoring rule is used. This means that a trader will always move the market probability to the posterior given his received private signal and other traders' previous moves in the market. In a single-round MSR, this truthful reporting is in fact the optimal strategy. However, as shown by Chen et al. [3] and Gao et al. [6], when traders can participate for multiple rounds in a MSR, equilibrium strategies depend on the specific information structure.

In the analysis that follows, we will consider a MSR with a single round of trading. Because the joint prior probability distribution of the outcome variable and signals is common knowledge, when a trader moves the market probability, he essentially reveals his private signal fully to the other trader. Thus, let p_0 denote the initial market probability, the market probability will move as follows: $p_0 \rightarrow f_{s_1,0} \rightarrow f_{s_1,s_2}$, at which point information is said to be fully aggregated.

4. A MANIPULATOR IN THE TRADITIONAL MARKET

In this section, we will first introduce a modified setting with a manipulator, and then analyze how a manipulator could affect the market price in the traditional market scoring rule.

4.1 Participants and Payoffs

In the new setting, the market may consist of two types of traders: market-oriented traders and a manipulator. Market-oriented traders receive a private signal and derive payoffs from the profit or loss in the market and are essentially the risk-neutral utility maximizers as described in the traditional model above. On the other hand, the manipulator does not receive any signal and his payoff is derived *purely* from an external source and is increasing in the final market probability. We assume that the honest trader possesses a prior belief p_M on the event that there exists a manipulator, while the manipulator knows that he is the only manipulator and that the other trader is of the market-oriented type, as well as p_M .

4.2 Strategies

The strategies are trivial in the case when the market-oriented trader moves first. He will simply move the price to the posterior given his signal, then the manipulator will move the price as high as possible. We will analyze the case where the first trader is a manipulator, followed by the market-oriented trader.

The manipulator needs to decide which report to make that can maximize his expected payoff, so the manipulator's

strategy is a probability distribution over the interval $[0, 1]$. Let σ denote the manipulator's strategy, and $\sigma(r)$ denote the manipulator's probability of reporting $r \in [0, 1]$.

In contrast to the market without manipulators, trader 2 can no longer simply trust the signal as expressed by trader 1's report. Rather, trader 2 needs to form beliefs over trader 1's type (if he is a manipulator, or his private signal if he is not). Let $\mu : \Sigma_2 \times [0, 1] \times (\Sigma_1 \cup M) \rightarrow [0, 1]$ denote trader 2's belief, and let $\mu_{s_2, r_1}(s_1)$ be trader 2's belief on the probability of the event that trader 1 is not a manipulator and has received signal s_1 , having observed trader 1's report r_1 , and received signal s_2 . Similarly, let $\mu_{s_2, r_1}(M)$ be trader 2's belief on the probability of the event that trader 1 is a manipulator, having observed trader 1's report r_1 , and received signal s_2 . These beliefs, together with the observed report made by trader 1 and trader 2's own private signal, determine trader 2's optimal move: he will move the market probability to $r_2 = \sum_{s_1 \in S_1} \mu_{s_2, r_1}(s_1) f_{s_1, s_2} + \mu_{s_2, r_1}(M) f_{0, s_2}$.

At equilibrium, trader 2's belief must be derived from trader 1's strategy, which can be done by mechanically applying Bayes' rule. We can divide trader 1's possible reports into two sets: 1) the reports that correspond to some posterior of the outcome given some signal, i.e. one of the possible reports a market-oriented trader could make, and 2) arbitrary reports not in the first set. Let $s_{r_1} \in \Sigma_1$ be the signal such that $r_1 = f_{s_{r_1}, 0}$. Then the above two sets correspond to 1) s_{r_1} exists, and 2) s_{r_1} does not exist. Trader 2's beliefs are as follows: if s_{r_1} exists,

$$\mu_{s_2, r_1}(s) = \begin{cases} \frac{(1-p_M)P(s_{r_1}|s_2)}{p_M\sigma(r_1)+(1-p_M)P(s_{r_1}|s_2)} & \text{if } s = s_{r_1} \\ 0 & \text{if } s \neq s_{r_1} \end{cases}$$

$$\mu_{s_2, r_1}(M) = \frac{p_M\sigma(r_1)}{p_M\sigma(r_1) + (1-p_M)P(s_{r_1}|s_2)}$$

And if s_{r_1} does not exist, i.e. r_1 is an arbitrary report,

$$\mu_{s_2, r_1}(s) = 0 \quad \forall s \in \Sigma_1$$

$$\mu_{s_2, r_1}(M) = 1$$

Having described trader 2's beliefs given trader 1's strategy, we now turn to look at properties of the manipulator's equilibrium strategies. To gain some intuition on what the manipulator would do, let us first consider an obvious strategy of always reporting the posterior given the highest signal, which one might expect to be an equilibrium strategy; however, this is not always true.

PROPOSITION 4.1. *The obvious strategy of always pretending to have the highest signal,*

$$\sigma(r) = \begin{cases} 1 & \text{if } r = f_{s_{max}, 0} \\ 0 & \text{otherwise} \end{cases}$$

where s_{max} denotes the highest signal in Σ_1 , is not always an equilibrium strategy.

PROOF. We will prove by a counterexample. Suppose $\Sigma_1 = \{s_a, s_b, s_c\}$ where $s_a < s_b < s_c$ and $\Sigma_2 = \{s_2\}$. The manipulator's payoff for reporting the highest signal, $f_{s_c, 0}$, is

$$\sum_{s_1 \in \Sigma_1} \mu_{s_2, f_{s_c, 0}}(s_1) f_{s_1, s_2} + \mu_{s_2, f_{s_c, 0}}(M) f_{0, s_2}$$

$$= \frac{(1-p_M)P(s_c|s_2)}{p_M + (1-p_M)P(s_c|s_2)} f_{s_c, s_2} + \frac{p_M}{p_M + (1-p_M)P(s_c|s_2)} f_{0, s_2}$$

Suppose $p_M = 1 - \epsilon$, $P(s_c|s_2) = \epsilon^2$, $f_{0,s_2} = 0.5$, $f_{s_b,s_2} = 1 - \epsilon$, $f_{s_c,s_2} = 1$. We see that when ϵ is small, the manipulator's payoff for reporting $f_{s_c,0}$ is close to $f_{0,s_2} = 0.5$; however, reporting $f_{s_b,0}$ would give him a payoff of $f_{s_b,s_2} = 1 - \epsilon$ because trader 2 would believe that he is not a manipulator, so he would deviate to report $f_{s_b,0}$. \square

Given this insight, we now turn to characterize the manipulator's equilibrium strategy. The following lemma establishes which reports cannot be made by the manipulator at equilibrium.

LEMMA 4.2. *Any posterior $r \in \{f_{s_1,0}|f_{s_1,0} < f_{0,0}, s_1 \in \Sigma_1\}$ cannot be in the support of an equilibrium strategy. And any arbitrary report $r \notin \{f_{s,0}|s \in \Sigma_1\}$ cannot be in the support of an equilibrium strategy.*

PROOF. Suppose for contradiction that $r \in \{f_{s_1,0}|f_{s_1,0} < f_{0,0}, s_1 \in \Sigma_1\}$ is in the support of some equilibrium strategy σ . The expected payoff of reporting r is

$$\begin{aligned} & \mathbb{E}_{s_2} \left[\frac{p_M \sigma(r)}{p_M \sigma(r) + (1 - p_M) P(s_1|s_2)} f_{0,s_2} \right. \\ & \left. + \frac{(1 - p_M) P(s_1|s_2)}{p_M \sigma(r) + (1 - p_M) P(s_1|s_2)} f_{s_1,s_2} \right] < f_{0,0} \end{aligned}$$

But there exists at least one higher posterior $f_{s',0} > f_{0,0}$, which cannot be in the support of σ because its expected payoff is always greater than $f_{0,0}$, and higher than the manipulator's expected payoff if σ includes reporting r , making him better off to deviate.

Similarly for the second claim, reporting any arbitrary report gives away the fact that trader 1 is a manipulator, so the manipulator's expected payoff of reporting the arbitrary report is $f_{0,0}$, and there must be a higher posterior that the manipulator can deviate to. \square

Let $\Sigma_{high} = \{s_1|f_{s_1,0} > f_{0,0}, s_1 \in \Sigma_1\}$. We have now shown that only the posteriors corresponding to signals in Σ_{high} could be in the support of an equilibrium strategy. Now, we prove several lemmas that will be useful in our main theorem. The following lemma establishes how the expected payoff of a particular report is affected by the manipulator's probability of making that report.

LEMMA 4.3. *Let $V_\sigma(r)$ be the manipulator's expected payoff for reporting some r in the support of a strategy σ and $r = f_{s_r,0}$ for some $s_r \in \Sigma_1$. $V_\sigma(r)$ is increasing in $\sigma(r)$ if $f_{s_r,0} < f_{0,0}$, and is decreasing in $\sigma(r)$ if $f_{s_r,0} > f_{0,0}$.*

PROOF.

$$\begin{aligned} V_\sigma(r) = & \mathbb{E}_{s_2} \left[\frac{p_M \sigma(r)}{p_M \sigma(r) + (1 - p_M) P(s_r|s_2)} f_{0,s_2} \right. \\ & \left. + \frac{(1 - p_M) P(s_r|s_2)}{p_M \sigma(r) + (1 - p_M) P(s_r|s_2)} f_{s_r,s_2} \right] \end{aligned}$$

If $f_{s_r,0} < f_{0,0}$, then $f_{s_r,s_2} < f_{0,s_2}$ for any s_2 by the strict ordering assumption, so the expression in the expectation is increasing in $\sigma(r)$ for every s_2 , and so is $V_\sigma(r)$. The proof of the second claim is analogous. \square

Let $R = [0, 1]^n$ be the support of some strategy where n is the size of the support, let σ^R denote a strategy under which the manipulator's expected payoff of reporting every $r \in R$ is the same, and let $\pi(\sigma^R)$ denote the manipulator's expected payoff when he plays the strategy σ^R .

LEMMA 4.4. *If σ^R exists for some $R = \{f_{s,0}|s \in \Sigma_{high}\}$, $\pi(\sigma^R) < f_{s,0}$ for all $f_{s,0} \in R$.*

PROOF. Consider any $f_{s,0} \in R$. The payoff of reporting $f_{s,0}$ under σ^R must be less than $f_{s,0}$, since $V_\sigma(f_{s,0}) = f_{s,0}$ when $\sigma(f_{s,0}) = 0$ and is decreasing in $\sigma(f_{s,0})$ according to Lemma 4.3. \square

LEMMA 4.5. *Let $R = \{f_{s,0}|s \in \Sigma_{high}\}$. If there exists $s' \in \Sigma_{high}$ such that $f_{s',0} > \pi(\sigma^R)$ and $f_{s',0} \notin R$. Then, $\pi(\sigma^R) < \pi(\sigma^{R \cup f_{s',0}}) < f_{s',0}$. On the other hand, if there exists $s' \in \Sigma_{high}$ such that $f_{s',0} < \pi(\sigma^R)$ and $f_{s',0} \notin R$, $\pi(\sigma^{R \cup f_{s',0}})$ does not exist.*

PROOF. Suppose for now that the size of R is one. $\sigma^{R \cup f_{s',0}}$, where $f_{s',0} > \pi(\sigma^R)$, always exists: we start with σ^R and add $f_{s',0}$ to the support by shifting the probability mass from $\sigma(R)$ to $\sigma(f_{s',0})$, which increases the payoff of reporting R and decreases the payoff of reporting $f_{s',0}$ according to Lemma 4.3. We do this until the payoff of reporting R equals the payoff of reporting $f_{s',0}$. So we have $\pi(\sigma^R) < \pi(\sigma^{R \cup f_{s',0}}) < f_{s',0}$.

If the size of R is greater than one, we can still add $f_{s',0}$ to the support and equate all payoffs by shifting probability mass from reporting something in R to reporting $f_{s',0}$ while preserving equal payoffs of reporting something in R .

For the second part of the lemma, since all posteriors in R correspond to signals in Σ_{high} , decreasing $\sigma(r)$ for any $r \in R$, increases the payoff of reporting r , while increasing $\sigma(f_{s',0})$ decreases the payoff of reporting $f_{s',0}$. Therefore, there is no way to equate the payoff of reporting something in R and reporting $f_{s',0}$. \square

We can now characterize the manipulator's equilibrium strategy in the following theorem.

THEOREM 4.6. *At any equilibrium, the manipulator mixes among $R^* = \{f_{s',0}|s' \geq s_t\}$ for some $s_t \in \Sigma_{high}$, i.e. the posteriors given any of the signals at least as high as s_t .*

PROOF. First note that the equilibrium strategy with the support R^* always exists and we can compute it by starting with the support $R = \{f_{s_{max},0}\}$. If this is not an equilibrium strategy, $f_{s_{max},0}$ is higher than $\pi(\sigma^R)$, where $s_{max}-1$ is the second highest signal. From Lemma 4.5, we know that $\pi(\sigma^{R \cup f_{s_{max}-1,0}}) > \pi(\sigma^R)$. We can repeatedly add the posterior corresponding to the next highest signal to R until $\pi(\sigma^R)$ is greater than $f_{s_{R-1},0}$, where s_{R-1} is the highest signal whose posterior is not in R . By construction, $\pi(\sigma^{R^*}) > f_{s',0}$ for all $s' \notin R^*$, so adding the posterior given any one of the signals not in R^* cannot be an equilibrium strategy according to Lemma 4.5.

Now consider the set $R = \{f_{s',0}|s' \in \Sigma'_1\}$ where $\Sigma'_1 \subset \{s|s \geq s_t \text{ and } s \in \Sigma_1\}$ and s_t is the lowest signal in Σ'_1 . Let s_{t+1} be the signal corresponding to the lowest report in R such that $\exists s_{nt}$ where $f_{s_{nt},0} \notin R$ and $f_{s_t,0} < f_{s_{nt},0} < f_{s_{t+1},0}$. By Lemma 4.4, we know that $\pi(\sigma^R) < f_{s',0}$ for any $f_{s',0} \in R$. However, since $f_{s_t,0} < f_{s_{nt},0}$, $\pi(\sigma^R) < f_{s_{nt},0}$, so σ^R cannot be an equilibrium strategy. We have now shown that at equilibrium, the manipulator has to mix between the posteriors of all $s \geq s_t$. \square

4.3 Prior Probability on Manipulator’s Existence

Here we note the richness of our model and the three real-world scenarios to which different values of p_M may correspond. First, when $p_M = 1$, every trader knows that there exists a manipulator as well as his identity. We call this the *public* information case, which could model a situation where a particular trader’s interest in the prediction is commonly known, e.g. a pharmaceutical company participating in a prediction market for virality of a new flu. The honest trader will simply ignore the manipulator’s trade at an equilibrium and will move the price to the posterior given his private signal. The manipulator cannot affect the final probability in this case, unless he is the last trader to act, in which case he can move the market probability as high as possible.

Second, when $p_M \in (0, 1)$, only the manipulator knows that he himself is the manipulator, while others are not entirely sure that a manipulator exists. We call this the *private* information case. An example of this might be a situation where a junior faculty under review for tenureship participates in a prediction market for citations on his own paper, but others only know that some trader may be under tenure review with probability p_M .

Finally, the *secret* information case corresponds to when $p_M = 0$: only the manipulator himself knows that he is the manipulator, while others do not believe that there is a manipulator. Since the market-oriented trader has no knowledge of the existence of a manipulator, he will trade as if there is no manipulator. As described before, at his turn, a market-oriented trader will move the market probability to the posterior given his private signal. We assume that trades that are inconsistent with the commonly-known prior distribution would be ignored by the market-oriented trader. To maximize the final market probability, the manipulator will trade as if he has the highest signal in his signal space, reporting $r_1 = f_{s_{max}, 0}$. The market-oriented trader believes the signal as expressed by the trade, and will report $r_2 = f_{s_{max}, s_2}$, which is the maximal r_2 .

5. MARKET SCORING RULE WITH TRADE LIMITS

We have described how a manipulator can distort the market probability under three different cases corresponding to different values of the prior probability on the manipulator’s existence. We will now consider the secret information case which is the worst scenario where price can be most distorted by the manipulator. In order to mitigate this kind of manipulation when the honest trader does not know about the manipulator’s existence, we propose to modify the market scoring rule by imposing a limit $L > 0$ by which a trader can maximally move the market probability at a time. Formally, given the current market probability p , the trader can only change it to a new probability $q \in [p - L, p + L]$. In the rest of the paper, we modify the market to have multiple repeated rounds of trading with a fixed trading sequence. We assume that the market-oriented trader is myopic and will maximize the payoff of his trade in the current round, and that the manipulator knows that the market-oriented trader is myopic.

5.1 Strategies and Market Probability without a Manipulator

Let the current market probability be p , and trader i ’s belief of the posterior be b_i , and q be the probability to which he chooses to change p . Since trader i is myopic, he would want to choose q so that his expected payoff from this round is maximized with respect to his belief. The following lemma follows:

LEMMA 5.1. *At his turn, a myopic trader will maximize his expected payoff by moving the market probability maximally towards his belief.*

Definition A trader *fully reveals* his signal when he moves the market price from p to $q \in (p - L, p + L)$ (i.e. does not use up the full limit), as the other traders have become knowledgeable of his exact private signal.

We present the following theorem on the convergence of the market probability when no manipulator is present.

THEOREM 5.2. *In a MSR with a large enough trade limit, myopic traders will fully reveal their signal one by one as they make trades. Eventually, the market probability will converge to the posterior given all traders’ private signals, at which point no more trades would occur.*

PROOF. The market’s state can be described by two variables. The first is the current market probability, and the second is the set of possible signals of each trader. We let this set be described by a random variable S . Initially, the market price is p_0 and $S = \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_N$. Let S_{-i} denote the set of possible signals of all traders except for trader i ’s.

Let p denote the current market price and D the current set of possible signals. If trader i changes the price from p to $q \in (p - L, p + L)$, the other traders observe this move and learn that $P(X = 1 | S_i = s_i, S_{-i} \in D_{-i}) = q$. Since every signal is the signal space is unique by assumption, all traders can infer trader i ’s private signal from this move.

If trader i changes p to $p + L$, it implies that $P(X = 1 | S_i = s_i, S_{-i} \in D_{-i}) \geq p + L$, and signals to other traders that S_i is at least some signal s'_i where s'_i is the lowest signal such that $P(X = 1 | S_i = s'_i, S_{-i} \in D_{-i}) \geq p + L$, updating the set of possible signals to $D' = D \cap (S_i \geq s'_i)$. Note that it is possible that $D' = D$, and that no information is learned from the trade, which is the reason of potential non-convergence to the true posterior.

Let us first consider the condition of non-convergence. If the market probability does not converge to the true posterior, the market must enter a cycle where trader i faces the same set of possible signals D with the same market probability p over and over. D has to remain the same because it cannot increase, and if the market probability does not converge to the true posterior, it means that D does not decrease, so it must remain the same in this cycle. p also has to remain the same because if D is not updated, every trader would make the same trade (either increase or decrease by L), and p will always move in the same direction, eventually reaching a value where some trader’s trade can update D . Therefore, every trader has to face the same market probability every time, if p does not converge to the true posterior.

Now, if we set the trade limit to $L_c = \max_{i, u, v, D_{-i}} |P(X = 1 | S_i = u, S_{-i} = D_{-i}) - P(X = 1 | S_i = v, S_{-i} = D_{-i})|$,

every trader will always be able to update the set of possible signals with a trade. Thus, the set of possible signals will eventually collapse to a point where everyone’s signal is common knowledge. All traders will agree on the posterior $P(X = 1|S_1 = s_1, \dots, S_N = s_N)$ and they will change the market probability to exactly that value. Note that L_c as defined above assumes that the worst case is reached in the process of trading. The market probability could converge for smaller trade limits as well.¹ \square

We have shown that in a normal market without a manipulator, the market probability in a market scoring rule with trade limits will still converge to the posterior given all private signals, so there is no loss in information, but only in the speed of the process.

5.2 Strategies and Market Probability when a Manipulator is Present

We now turn to the scenario that inspired the trade limits in the first place, the setting with a manipulator. We show that the cost of delayed information aggregation can be justified when a manipulator is present, as we can obtain a more accurate estimate. First, note that with a manipulator, the best estimate we can get is the posterior given all of the truthful participants’ signals. We call this the true posterior.

Recall again that in the setting with secret information, the market-oriented trader has no knowledge of the manipulator, and will trade as if the other trader is market-oriented, so Lemma 5.1 still holds for the market-oriented trader. Knowing that other trader is market-oriented and his strategy, the manipulator’s optimal strategy is to act like a market-oriented trader with the highest signal.

5.3 Benefits of Trade Limits

We know now that as traders trade in the market, they fully reveal their signals one by one, i.e., their private signals would become public, and eventually the market probability will converge to $P(X = 1|S_1 = s_1, \dots, S_i = s_i^{max}, \dots, S_N = s_N)$ if trader i is a manipulator. We hypothesize that the manipulator is likely to be the last trader to fully reveal his signal because his signal is extreme, and it should take more rounds to collapse the set of the possible signals he has. When this is the case, it is obvious that the price right after the second-to-last trader revealed his signal is closer to the true posterior given the $N - 1$ traders than all later prices. We call this price right after all but one trader fully revealed their signals the *early-stopping price*.

However, the manipulator is not always the last trader to fully reveal his signal, which depends on the specific probability distribution. This raises several interesting questions: When the manipulator is the last one to fully reveal his signal, how much better is the early-stopping price than the fully converged price? How often is the manipulator *not* the last one to fully reveal his signal? And when this occurs, by how much is the early-stopping price worse than the fully converged price? How does the number of rounds it takes to reach the early-stopping price compared to the number of rounds to reach convergence? How do these quantities relate to the size of trade limit we impose? How do they

¹In the case with continuous signal spaces, the market probability always converges regardless of the size of the trade limit.

change with different distributions? We will explore these questions computationally.

5.3.1 Experimental Setup

One way to generate the joint probability distributions over outcome and signals is to first generate the joint probability distribution of signals, then generate the posterior of outcomes given signals. For the joint probability distribution of signals, we assume a symmetric multivariate normal distribution. We let there be 100 signals in the signal space, and map the signals uniformly onto $[-1, 1]$, we then generate the probabilities of the signals according to a truncated normal distribution. We generate distributions with several different values of the mean ($\{-0.8, 0, 0.8\}$) and variance ($\{1, 0.1, 0.01\}$). Finally, we generate the posterior given signals, $P(X = 1|S_1 = s_1, \dots, S_N = s_N)$ with a function $f(S_1, \dots, S_N)$, which is increasing in S_1, \dots, S_N . For simplicity, we use a simple linear function, $f(S_1, \dots, S_N) = S_1 + \dots + S_N$. Then, map the values onto $[0, 1]$. To explore the effect of different sizes of trade limits, for each distribution generated, we run simulations for $L \in \{0.01, 0.02, \dots, 0.5\}$. We run simulations for different number of traders (from two to four).

5.3.2 Results

In this section, we discuss the interesting characteristics of the mechanism observed in the simulations. We found that the results are qualitatively similar across the different number of traders, and we hypothesize them to hold true for more traders. Below are the main takeaways:

- The accuracy of early-stopping price are non-monotonic in trade limit size; as we decrease trade limit size, the accuracy typically increases.
- Trade limits work best if the signal distributions are centered far away from the highest signal.
- The number of rounds to convergence increases exponentially as we decrease trade limit size; however, we do not need a very small trade limit to benefit from the early-stopping rule.
- The probability that the manipulator is not the last trader to fully reveal signal is low. When this does occur, the differences between the early-stopping prices and the fully converged prices are negligible.

We will discuss each of them in detail below, and we present data for simulations of markets with two traders. In the figures presented in this section, the values are expectations over different signal realizations.

Trade limit vs. accuracy of early-stopping price.

For small trade limits, the difference between the true posterior and early-stopping price can be as small as 10% of that the difference between the true posterior and the converged price, as shown in Figure 1. For all trade limits, the early-stopping price is closer to the true posterior than the converged price on average.

Note that with large enough trade limits, it is effectively the same as without trade limits, and the early-stopping price is the price after the first trader moves. As we decrease the trade limit, we typically see that the accuracy of the early-stopping price increases. This agrees with our

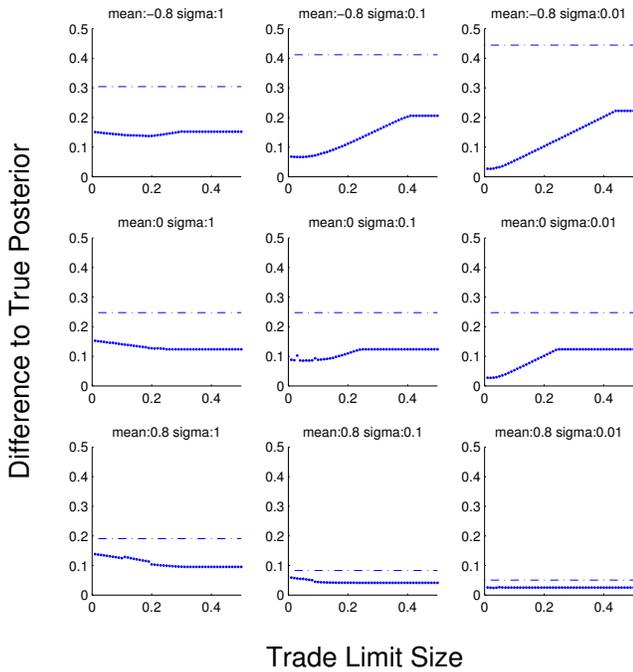


Figure 1: Trade limit size vs. accuracy of the early-stopping prices and the fully converged prices across different means and variances of the signal distribution. The dashed line is the fully converged price, and the dots are the early-stopping prices.

intuition that smaller trade limits limit how much the manipulator can distort the price before the truthful traders fully revealed their signals. However, as the figure also shows, decreasing trade limits does not always increase accuracy, but sometimes decreases accuracy. We can see how this occurs with the following example. Consider a market where the manipulator trades first. He first moves the price from p_0 to $p_0 + L \leq P(X = 1 | S_1 \geq s'_1, S_2 \in \Sigma_2)$ for some $s'_1 \in \Sigma_1$. If L is too small, it could be the case that $P(X = 1 | S_1 \geq s'_1, S_2 = s_2^{min}) \geq p_0 + 2L$, so when trader 2 makes his move, he has not revealed any information. Thus, the manipulator can reveal more information before the other traders fully revealed their signals, effectively decreasing the accuracy of the early-stopping price.

Figure 1 also shows that the accuracy of prices vary across signal distributions with different means and variances. When the signals are distributed around moderate or low signals (zero or negative mean), and with the optimal trade limit size, we see a trend that the smaller the variance of the signal distribution is, the closer the early-stopping prices get to the true posterior, relative to the fully converged price. This makes sense because small variances imply that signals are less likely to be extremely high, so truthful traders will be able to fully express their signals in fewer rounds.

Across signal distributions with different means, we see that the benefit of trade limits diminishes as we increase the mean. This is because with increasing means, even truthful traders begin to look more like the manipulator as their signals are likely to be high as well. Thus, it makes sense that trade limits would not help distinguishing the manipulator as much, because the difference between the manipulator

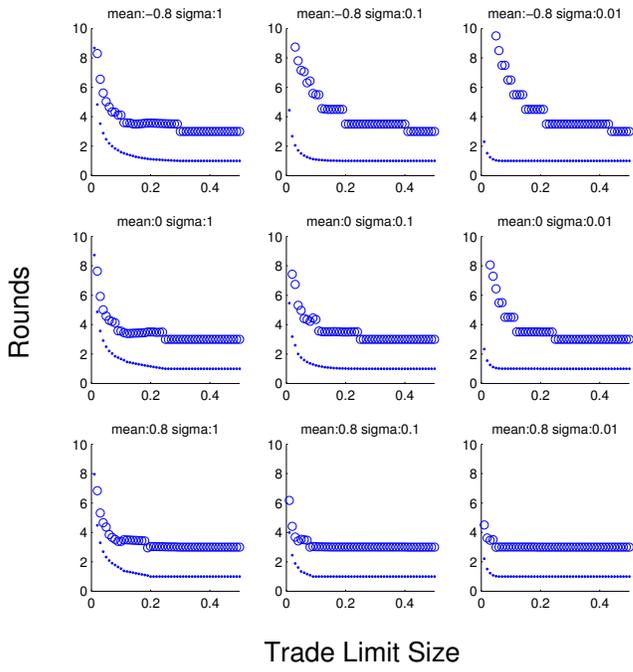


Figure 2: Trade limit size vs. the number of rounds it takes to reach early-stopping prices and fully converged prices across different means and variances of the signal distribution. The circles are the number of rounds to fully convergence, and the dots are the number of rounds to the early-stopping prices.

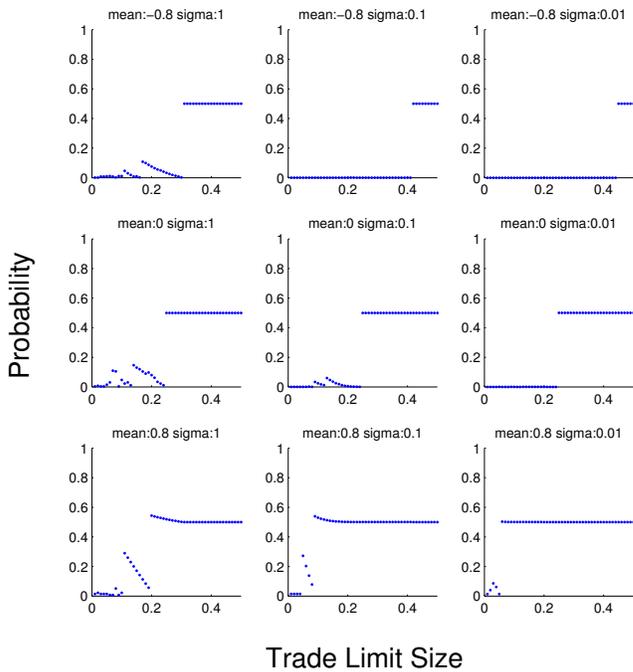


Figure 3: Trade limit size vs. the probability that the manipulator is not the last trade to fully reveal his signal across different means and variances of the signal distribution.

and truthful traders is inherently smaller. These observations suggest that trade limits work best when the signal distribution are centered far away the manipulator’s extreme signal, so that the manipulator’s signal is close to an outlier in the signal distribution.

Finally, we also looked at the case when there is no manipulator and found that the expected loss in accuracy is negligible, which is mostly due to the normal distributions that are used in our experiments: when signals are closer to the mean, there is not much loss in accuracy.

Trade limits vs. number of rounds.

The results in figure 2 agree with the intuition that the number of rounds to early-stopping price and convergence both increase as trade limit decreases, which happens at an exponential rate. However, note that the optimal trade limit sizes that give the most accurate early-stopping price often correspond to fewer than 10 rounds to convergence. Also, it can be seen that the number of rounds to convergence and early-stopping price depends mostly on the trade limit size, and is fairly robust to changes in the signal distribution.

Probability that the manipulator is not the last trader to fully reveal signal.

As figure 3 shows, with small enough trade limit size, the probability that the manipulator is the last trader to fully reveal his signal is very high ($> 80\%$ mostly), with the exception of signal distributions with high means. Note that without trade limits, 50% is the chance that the manipulator is the last trader as there are two traders in our simulations. We see that with high means, the trade limit size required to distinguish between truthful traders and the manipulator decreases, because the signals of truthful traders are closer to that of the manipulator as mentioned previously.

However, even though the manipulator is not always the last trader to fully reveal his signal, we found that in the cases where the manipulator is not the last trader to reveal signal, the early-stopping prices are nearly identical to the fully converged price. This could happen if posterior given the truthful trader’s signal lies precisely at the trade limit boundary, in which case he has not fully revealed his signal until convergence, so the early-stopping price in this case is the converged price.

6. CONCLUSION

We analyzed strategic behaviors in a MSR where a profit-indifferent manipulator exists and the honest trader has a prior belief on his existence. We introduced a modification to the market scoring rule in the form of trade limits, in order to reduce manipulability of prediction markets in the worst situation where the honest trader does not know about the manipulator’s existence. In the case where there is no manipulator, we showed that the market price converges to the true posterior and information gets fully aggregated although delayed. We showed via simulations that trade limits indeed help reduce manipulability by delaying the manipulator’s revelation of his signal.

There are many potential areas for future work. An obvious direction is to extend the equilibrium analysis to more than two agents, and multi-round MSR. Another direction is a rigorous theoretical analysis of the mechanism with trade limits. It would be interesting to characterize the condi-

tion under which the manipulator is the last trader to fully reveal his signal. It is also useful to provide a bound on the difference between the early-stopping price and the true posterior, as well as the number of rounds to convergence. Finally, one can work to reduce manipulability in mechanisms beyond market scoring rules.

7. ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation under Grant No. IIS0963478. The first author is supported by a Stanford Graduate Fellowship.

8. REFERENCES

- [1] J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research, 2001.
- [2] C. Boutilier. Eliciting forecasts from self-interested experts: scoring rules for decision makers. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS ’12, pages 737–744, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
- [3] Y. Chen, S. Dimitrov, R. Sami, D. M. Reeves, D. M. Pennock, R. D. Hanson, L. Fortnow, and R. Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010.
- [4] Y. Chen, X. A. Gao, R. Goldstein, and I. A. Kash. Market Manipulation with Outside Incentives. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI’11)*, 2011.
- [5] B. Cowgill, J. Wolfers, W. U. Penn, E. Zitzewitz, U. Malmendier, K. M. Murphy, M. Ostrovsky, P. Oyer, P. Pathak, and T. Rosenblat. Using prediction markets to track information flows: Evidence from google, 2008.
- [6] X. A. Gao, J. Zhang, and Y. Chen. What you jointly know determines how you act: strategic interactions in prediction markets. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, EC ’13, pages 489–506, New York, NY, USA, 2013. ACM.
- [7] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1:2007, 2002.
- [8] R. Hanson and R. Oprea. A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314, 2009.
- [9] M. Ottaviani and P. N. Sørensen. Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association*, 5(2-3):554–563, 04-05 2007.
- [10] C. R. Plott and K.-Y. Chen. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Papers 1131, California Institute of Technology, Division of the Humanities and Social Sciences, Mar. 2002.
- [11] P. Shi, V. Conitzer, and M. Guo. Prediction mechanisms that do not incentivize undesirable actions. In *In WINE*, 2009.