

# Bluffing and Strategic Reticence in Prediction Markets\*

Yiling Chen<sup>1</sup>, Daniel M. Reeves<sup>1</sup>, David M. Pennock<sup>1</sup>,  
Robin D. Hanson<sup>2</sup>, Lance Fortnow<sup>3</sup>, and Rica Gonen<sup>1</sup>

<sup>1</sup> Yahoo! Research

<sup>2</sup> George Mason University

<sup>3</sup> University of Chicago

**Abstract.** We study the equilibrium behavior of informed traders interacting with two types of automated market makers: market scoring rules (MSR) and dynamic parimutuel markets (DPM). Although both MSR and DPM subsidize trade to encourage information aggregation, and MSR is myopically incentive compatible, neither mechanism is incentive compatible in general. That is, there exist circumstances when traders can benefit by either hiding information (reticence) or lying about information (bluffing). We examine what information structures lead to straightforward play by traders, meaning that traders reveal all of their information truthfully as soon as they are able. Specifically, we analyze the behavior of risk-neutral traders with incomplete information playing in a finite-period dynamic game. We employ two different information structures for the logarithmic market scoring rule (LMSR): conditionally independent signals and conditionally dependent signals. When signals of traders are independent conditional on the state of the world, truthful betting is a Perfect Bayesian Equilibrium (PBE) for LMSR. However, when signals are conditionally dependent, there exist joint probability distributions on signals such that at a PBE in LMSR traders have an incentive to bet against their own information—strategically misleading other traders in order to later profit by correcting their errors. In DPM, we show that when traders anticipate sufficiently better-informed traders entering the market in the future, they have incentive to partially withhold their information by moving the market probability only partway toward their beliefs, or in some cases not participating in the market at all.

## 1 Introduction

The strongest form of the *efficient markets hypothesis* [1] posits that information is incorporated into prices fully and immediately, as soon as it becomes

---

\* An early version of this paper appeared at the Second Workshop on Prediction Markets. This version is much improved thanks to the insightful comments by Stanko Dimitrov, Paul J. Healy, Mohammad Mahdian, Rahul Sami, and the anonymous reviewers.

available to anyone. A *prediction market* is a financial market specifically designed to take advantage of this property. For example, to forecast whether a product will launch on time, a company might ask employees to trade a security that pays \$1 if and only if the product launches by the planned date. Everyone from managers to developers to administrative assistants with different forms and amounts of information can bet on the outcome. The resulting price constitutes their collective probability estimate that the launch will occur on time. Empirically, prediction markets like this outperform experts, group consensus, and polls across a variety of settings [2–10].

Yet the double-sided auction at the heart of nearly every prediction market is *not* incentive compatible. Information holders do not necessarily have incentive to fully reveal all their information right away, as soon as they obtain it. The extreme case of this is captured by the so-called *no trade theorems* [11]: When rational, risk-neutral agents with common priors interact in an unsubsidized (zero-sum) market, *the agents will not trade at all*, even if they have vastly different information and posterior beliefs. The informal reason is that any offer by one trader is a signal to a potential trading partner that results in belief revision discouraging trade.

The classic *market microstructure* model of a financial market posits two types of traders: rational traders and noise traders [12]. The existence of noise traders turns the game among rational traders into a positive-sum game, thereby resolving the no-trade paradox. However, even in this setting, the mechanism is not incentive compatible. For example, monopolist information holders will not fully reveal their information right away: instead, they will leak their information into the market gradually over time and in doing so will obtain a greater profit [13].

Instead of assuming or subsidizing noise traders, a prediction market designer might choose to directly subsidize the market by employing an *automated market maker* that expects to lose some money on average. Hanson’s *market scoring rule market maker* (MSR) is one example [14, 15]. MSR requires a patron to subsidize the market, but guarantees that the patron cannot lose more than a fixed amount set in advance, regardless of how many shares are exchanged or what outcome eventually occurs. The greater the subsidy, the greater the effective liquidity of the market. Since traders face a positive-sum game, even rational risk-neutral agents have incentive to participate. In fact, even a single trader can be induced to reveal information, something impossible in a standard double auction with no market maker. Hanson proves that *myopic* risk-neutral traders have incentive to reveal all their information, however forward-looking traders may not.

Pennock’s dynamic parimutuel market (DPM) [16, 17] is another subsidized market game that functions much like a market maker. Players compete for shares of the total money wagered by all players, where the payoff of each share varies depending on the final state of the system. Whereas in a standard prediction market for a binary outcome the payoff of every winning share is exactly \$1, the payoff in DPM is *at least* \$1, but could be more.

Though subsidized market makers improve incentives for information revelation, the mechanisms are still not incentive compatible. Much of the allure of prediction markets is the promise to gather information from a distributed group quickly and accurately. However, if traders have demonstrable incentives to either hide or falsify information, the accuracy of the resulting forecast may be in question.

In this paper, we examine the strategic behavior of (non-myopic) risk-neutral agents participating in prediction markets using two-outcome MSR and DPM mechanisms. We model the market as a dynamic game and solve for equilibrium trading strategies. We employ two different information structures for LMSR with incomplete information: conditionally independent signals and conditionally dependent signals. The equilibrium concept that we use is the Perfect Bayesian Equilibrium (PBE) [18]. We prove that with conditionally independent signals, a PBE of LMSR with finite players and finite periods consists of all players truthfully revealing their private information at their first chance to bet. With conditionally dependent information, we show that in LMSR there exist joint probability distributions on signals such that traders have an incentive to bluff, or bet against their own information, strategically misleading other traders in order to later correct the price. DPM is shown, via a two-player, two-stage game, to face another problem: traders may have incentives to completely withhold their private information or only partially reveal their information when they anticipate sufficiently better-informed agents trading after them. Due to lack of space, we omit or abridge some proofs of lemmas and theorems in this paper; full proofs can be obtained as an Appendix by request.

*Related Work* Theoretical work on price manipulation in financial markets [19, 13, 20] explains the logic of manipulation and indicates that double auctions are not incentive compatible. There are some experimental and empirical studies on price manipulation in prediction markets using double auction mechanisms; the results of which are mixed, some giving evidence for the success of price manipulation [21] and some showing the robustness of prediction markets to price manipulation [22–25]. The paper by Dimitrov and Sami [26], completed independently and first published simultaneously with an early version of this paper, is the most directly related work that we are aware of. Dimitro and Sami, with the aid of a projection game, study non-myopic strategies in LMSR with two players. By assuming signals of players are unconditionally independent and the LMSR market has infinite periods, they show that truthful betting is not an equilibrium strategy in general. Our study of LMSR with incomplete information in Sections 3 and 4 complements their work. Dimitro and Sami examine infinite periods of play, while we consider finite periods and finite players. On the one hand, the conditionally independent signals case that we examine directly implies that signals are unconditionally dependent unless they are not informative. On the other hand, the conditional dependence of signals assumption overlaps with Dimitro and Sami’s unconditional independence of signals.

## 2 Background

Consider a discrete random variable  $X$  that has  $n$  mutually exclusive and exhaustive outcomes. Subsidizing a market to predict the likelihood of each outcome, two classes of mechanisms, MSR and DPM, are known to guarantee that the market maker's loss is bounded.

### 2.1 Marketing Scoring Rules

Hanson [14, 15] shows how a proper scoring rule can be converted into a market maker mechanism, called market scoring rules (MSR). The market maker uses a proper scoring rule,  $S = \{s_1(\mathbf{r}), \dots, s_n(\mathbf{r})\}$ , where  $\mathbf{r} = \langle r_1, \dots, r_n \rangle$  is a reported probability estimate for the random variable  $X$ . Conceptually, every trader in the market may change the current probability estimate to a new estimate of its choice at any time as long as it agrees to pay the scoring rule payment associated with the current probability estimate and receive the scoring rule payment associated with the new estimate. If outcome  $i$  is realized, a trader that changes the probability estimate from  $\mathbf{r}^{\text{old}}$  to  $\mathbf{r}^{\text{new}}$  pays  $s_i(\mathbf{r}^{\text{old}})$  and receives  $s_i(\mathbf{r}^{\text{new}})$ .

Since a proper scoring rule is incentive compatible for a risk-neutral agents, if a trader can only change the probability estimate once, this modified proper scoring rule still incentivizes the trader to reveal its true probability estimate. However, when traders can participate multiple times, they might have incentives to manipulate information and mislead other traders.

Because traders change the probability estimate in sequence, MSR can be thought of as a sequential shared version of the scoring rule. The market maker pays the last trader and receives payment from the first trader. For a logarithmic market scoring rule market maker (LMSR) with the scoring function  $s_i(\mathbf{r}) = b \log(r_i)$  and  $b > 0$ , the maximum amount the market maker can lose is  $b \log n$ .

An MSR market can be equivalently implemented as a market maker offering  $n$  securities, each corresponding to one outcome and paying \$1 if the outcome is realized [14, 27]. Hence, changing the market probability of outcome  $i$  to some value  $r_i$  is the same as buying the security for outcome  $i$  until the market price of the security reaches  $r_i$ . Our analysis in this paper is facilitated by directly dealing with probabilities.

### 2.2 Dynamic Parimutuel Market

A dynamic parimutuel market (DPM) [16, 17] is a dynamic-cost variant of a parimutuel market. There are  $n$  securities offered in the market, each corresponding to an outcome of  $X$ . As in a parimutuel market, traders who wager on the true outcome split the total pool of money at the end of the market. However, the price of a single share varies dynamically according to a price function, thus allowing traders to sell their shares prior to the determination of the outcome for profits or losses.

From a trader's perspective, DPM acts as a market maker. A particularly natural way for the market maker to set security prices is to equate the ratio of prices of any two securities by the ratio of number of shares outstanding for the two securities. Let  $\mathbf{q} = \langle q_1, \dots, q_n \rangle$  be the vector of shares outstanding for all securities. Then the total money wagered in the market is

$$C(\mathbf{q}) = \kappa \sqrt{\sum_{j=1}^n q_j^2}, \quad (1)$$

while the instantaneous price is

$$p_i(\mathbf{q}) = \frac{\kappa q_i}{\sqrt{\sum_{j=1}^n q_j^2}} \quad \forall i, \quad (2)$$

where  $\kappa$  is a free parameter. When a trader buys or sells one or more securities, it changes the vector of outstanding shares from  $\mathbf{q}^{\text{old}}$  to  $\mathbf{q}^{\text{new}}$  and pays the market maker the amount  $C(\mathbf{q}^{\text{new}}) - C(\mathbf{q}^{\text{old}})$ , which equals the integral of the price functions from  $\mathbf{q}^{\text{old}}$  to  $\mathbf{q}^{\text{new}}$ . If outcome  $i$  occurs and the quantity vector at the end of the market is  $\mathbf{q}^f$ , the payoff for each share of the winning security is

$$o_i = \frac{C(\mathbf{q}^f)}{q_i^f} = \frac{\kappa \sqrt{\sum_j (q_j^f)^2}}{q_i^f}. \quad (3)$$

Unlike LMSR where the market probability of an outcome is directly listed, the market probability of outcome  $i$  in DPM with the above described cost, price, and payoff functions is given by  $\pi_i = \frac{p_i(\mathbf{q})}{C(\mathbf{q})/q_i}$  or, in terms of the shares directly,

$$\pi_i(\mathbf{q}) = \frac{q_i^2}{\sum_{j=1}^n q_j^2}. \quad (4)$$

For traders whose probabilities are the same as the market probabilities, they can not expect to profit from buying or selling securities if the DPM market liquidates in the current state.

A trader wagering on the correct outcome is guaranteed non-negative profit in DPM, because  $p_i$  is always less than or equal to  $\kappa$  and  $o_i$  is always greater than or equal to  $\kappa$ . Setting  $\kappa = 1$  yields a natural version where prices are less than or equal to 1 and payoffs are greater than or equal to 1. Because the price functions are not well-defined when  $\mathbf{q} = \mathbf{0}$ , the market maker needs to initialize the market with a non-zero quantity vector  $\mathbf{q}^0$  (which may be arbitrarily small). Hence, the market maker's loss is at most  $C(\mathbf{q}^0)$  whichever outcome is realized.

Compared with a parimutuel market, where traders are never worse off for waiting until the last minute to put their money in, the advantage of DPM is that it provides some incentive for informed traders to reveal their information earlier, because the price of a security increases (decreases) when more people buy (sell) the security. But it is not clear whether traders are better off by always and completely revealing their information as soon as they can.

### 2.3 Terminology

*Truthful betting* (TB) for a player in MSR and DPM is the strategy of immediately changing market probabilities to the player's probabilities. In other words, it is the strategy of always buying immediately when the price is too low and selling when the price is too high according to the player's information. *Bluffing* is the strategy of betting contrary to one's information in order to deceive future traders, with the intent of capitalizing on their resultant misinformed trading. *Strategic reticence* means withholding one's information; that is, delaying or abstaining from trading, or moving the market probabilities only partway toward one's actual beliefs. This paper investigates scenarios where traders with incomplete information have an incentive to deviate from truthful betting.<sup>4</sup>

## 3 LMSR with Conditionally Independent Signals

In this part, we start with simple 2-player 3-stage games and move toward the general finite-player finite-stage games to gradually capture the strategic behavior in LMSR when players have conditionally independent signals.

### 3.1 General Settings

$\Omega = \{Y, N\}$  is the state space of the world. The true state,  $\omega \in \Omega$ , is picked by nature according to a prior  $\mathbf{p}^0 = \langle p_Y^0, p_N^0 \rangle = \langle \Pr(\omega = Y), \Pr(\omega = N) \rangle$ . The prior is common knowledge to all players. A market, aiming at predicting the true state  $\omega$ , uses a LMSR market maker with initial probability estimate  $\mathbf{r}^0 = \langle r_Y^0, r_N^0 \rangle$ .

Players are risk neutral. Each player gets a private signal,  $c_i \in \mathbf{C}_i$ , about the state of the world at the beginning of the market.  $\mathbf{C}_i$  is the signal space of player  $i$  with  $|\mathbf{C}_i| = n_i$ . Players' signals are independent conditional on the state of the world. In other words, player  $i$ 's signal  $c_i$  is independently drawn by nature according to conditional probability distributions,

$$\Pr(c_i = \mathbf{C}_i\{1\} | Y), \Pr(c_i = \mathbf{C}_i\{2\} | Y), \dots, \Pr(c_i = \mathbf{C}_i\{n_i\} | Y) \quad (5)$$

if the true state is  $Y$ , and analogously if the true state is  $N$ .  $\mathbf{C}_i\{1\}$  to  $\mathbf{C}_i\{n_i\}$  are elements of  $\mathbf{C}_i$ . The signal distributions are common knowledge to all players. Based on their private signals, players update their beliefs. Then players trade in one or more rounds of LMSR.

### 3.2 Who Wants to Play First?

We first consider a simple 2-player sequence selection game. Suppose that Alice and Bob are the only players in the market. Alice independently gets a signal

---

<sup>4</sup> With complete information, traders should reveal all information right away in both MSR and DPM, because the market degenerates to a race to capitalize on the shared information first.

$c_A \in \mathbf{C}_A$ . Similarly, Bob independently gets a signal  $c_B \in \mathbf{C}_B$ . Let  $|\mathbf{C}_A| = n_A$  and  $|\mathbf{C}_B| = n_B$ .

In the first stage, Alice chooses who—herself or Bob—plays first. The selected player then changes the market probabilities as they see fit in the second stage. In the third stage, the other player gets the chance to change the market probabilities. Then, the market closes and the true state is revealed.

**Lemma 1.** *In a LMSR market, if stage  $t$  is player  $i$ 's last chance to play and  $\mu_i$  is player  $i$ 's belief over actions of previous players, player  $i$ 's best response at stage  $t$  is to play truthfully by changing the market probabilities to  $\mathbf{r}^t = \langle \Pr(Y|c_i, \mathbf{r}^{t-1}, \mu_i), \Pr(N|c_i, \mathbf{r}^{t-1}, \mu_i) \rangle$ , where  $\mathbf{r}^{t-1}$  is the market probability vector before player  $i$ 's action.*

*Proof.* When a player has its last chance to play in LMSR, it is the same as the player interacting with a logarithmic scoring rule. Because the logarithmic scoring rule is strictly proper, player  $i$ 's expected utility is maximized by truthfully reporting its posterior probability estimate given the information it has.  $\square$

**Lemma 2.** *When players have conditionally independent signals, if player  $i$  knows player  $j$ 's posterior probabilities  $\langle \Pr(Y|c_j), \Pr(N|c_j) \rangle$ , player  $i$  can infer the posterior probabilities conditionally on both signals. More specifically,*

$$\Pr(\omega|c_i, c_j) = \frac{\Pr(c_i|\omega) \Pr(\omega|c_j)}{\Pr(c_i|Y) \Pr(Y|c_j) + \Pr(c_i|N) \Pr(N|c_j)},$$

where  $\omega \in \{Y, N\}$ .

Lemma 2 is proved using Bayes rule. According to it, with conditionally independent signals, a player can make use of another player's information when knowing its posteriors, even if not knowing its signal distribution.

Let  $\mathbf{r}$  be the posteriors of player  $j$  that player  $i$  observes. For simplicity, let  $\mathbf{C}_j\{\mathbf{r}\}$  be a fictitious signal that satisfies  $\langle \Pr(Y|\mathbf{C}_j\{\mathbf{r}\}), \Pr(N|\mathbf{C}_j\{\mathbf{r}\}) \rangle = \mathbf{r}$ .  $\mathbf{C}_j\{\mathbf{r}\}$  does not necessarily belong to player  $j$ 's signal space  $\mathbf{C}_j$ . When  $\mathbf{r}$  is the true posteriors of player  $j$ ,  $\langle \Pr(Y|c_i, \mathbf{C}_j\{\mathbf{r}\}), \Pr(N|c_i, \mathbf{C}_j\{\mathbf{r}\}) \rangle$  is the same as  $\langle \Pr(Y|c_i, c_j), \Pr(N|c_i, c_j) \rangle$ . The following theorem gives a PBE of the sequence selection game.

**Theorem 1.** *When Alice and Bob have conditionally independent signals in LMSR, a PBE of the sequence selection game is a strategy-belief pair with strategies of  $(\sigma_A, \sigma_B)$  and belief  $\mu_B$ , where*

- Alice's strategy  $\sigma_A$  is (select herself to be the first player in the first stage, change the market probability to  $\langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$  in the second stage);
- Bob's strategy  $\sigma_B$  is (take current market prices  $\mathbf{r}$  as Alice's posteriors and change the market probability to  $\langle \Pr(Y|\mathbf{C}_A\{\mathbf{r}\}, c_B), \Pr(N|\mathbf{C}_A\{\mathbf{r}\}, c_B) \rangle$  when it's his turn to play);
- Bob's belief  $\mu_B$  is that  $\Pr(\text{in the second stage Alice changes market probabilities to } \langle \Pr(Y|c_A), \Pr(N|c_A) \rangle) = 1$ .

*Sketch of Proof:* Let  $EU_A^I$  be Alice's expected utility conditional on her signal when she selects herself as the first player and  $EU_A^{II}$  be Alice's expected utility conditional on her signal when she selects Bob as the first player. The proof reduces  $EU_A^I - EU_A^{II}$  to the Kullback-Leibler divergence (also called relative entropy or information divergence) [28] of two distributions, which is always non-negative.

### 3.3 The Alice-Bob-Alice Game

We now consider a 3-stage Alice-Bob-Alice game, where Alice plays in the first and third stages and Bob plays in the second stage. Alice may change the market probabilities however she wants in the first stage. Observing Alice's action, Bob may change the probabilities in the second stage. Alice can take another action in the third stage. Then, the market closes and the true state is revealed. We study the PBE of the game when Alice and Bob have conditionally independent signals.

Let  $\mathbf{r}^1 = \langle r_Y^1, r_N^1 \rangle$  be the market probabilities that Alice changes to in the first stage. Lemma 3 characterizes the equilibrium strategy of Alice in the third stage. Theorem 2 describes a PBE of the Alice-Bob-Alice game.

**Lemma 3.** *In a 3-stage Alice-Bob-Alice game in LMSR with conditionally independent signals, at a PBE Alice changes the market probabilities to  $\mathbf{r}^3 = \langle r_Y^3, r_N^3 \rangle = \langle \Pr(Y|\mathbf{C}_A\{k\}, \mathbf{C}_B\{l\}), \Pr(N|\mathbf{C}_A\{k\}, \mathbf{C}_B\{l\}) \rangle$  in the third stage, when Alice has signal  $\mathbf{C}_A\{k\}$  and Bob has signal  $\mathbf{C}_B\{l\}$ .*

**Theorem 2.** *When Alice and Bob have conditionally independent signals in LMSR, a PBE of the 3-stage Alice-Bob-Alice game is a strategy-belief pair with strategies  $(\sigma_A, \sigma_B)$  and beliefs  $(\mu_A, \mu_B)$  where*

- Alice's strategy  $\sigma_A$  is (change market probabilities to  $\mathbf{r}^1 = \langle \Pr(Y|c_A), \Pr(N|c_A) \rangle$  in the first stage, do nothing in the third stage);
- Bob's strategy  $\sigma_B$  is (take  $\mathbf{r}^1$  as Alice's posteriors and change market probabilities to  $\mathbf{r}^2 = \langle \Pr(Y|\mathbf{C}_A\{\mathbf{r}^1\}, c_B), \Pr(N|\mathbf{C}_A\{\mathbf{r}^1\}, c_B) \rangle$  in the second stage);
- Bob's belief of Alice's action in the first stage,  $\mu_B$ , is  $(\Pr(\text{Alice changes market probabilities to } \mathbf{r}^1 = \langle \Pr(Y|c_A), \Pr(N|c_A) \rangle \text{ in the first stage}) = 1)$ ;
- Alice's belief of Bob's action in the second stage,  $\mu_A$ , is  $(\Pr(\text{Bob changes market probabilities to } \mathbf{r}^2 = \langle \Pr(Y|\mathbf{C}_A\{\mathbf{r}^1\}, c_B), \Pr(N|\mathbf{C}_A\{\mathbf{r}^1\}, c_B) \rangle \text{ in the second stage}) = 1)$ ;

Theorem 2 states that at a PBE of the Alice-Bob-Alice game, Alice truthfully reports her posterior probabilities in the first stage, Bob believes that Alice is truthful and reports his posterior probabilities based on both Alice's report and his private signal in the second stage, and Alice believes that Bob is truthful and does nothing in the third stage because all information has been revealed in the second stage. It's clear that Bob never wants to deviate from being truthful by Lemma 1. To prove that Alice does not want to deviate from being truthful either, we show that deviating is equivalent to selecting herself as the second



player in a sequence selection game, while being truthful is equivalent to selecting herself as the first player in the sequence selection game. Alice is worse off by deviating.

### 3.4 Finite-Player Finite-Stage Game

We extend our results for the Alice-Bob-Alice game to games with a finite number of players and finite stages in LMSR. Each player can change the market probabilities multiple times and all changes happen in sequence.

**Theorem 3.** *In the finite-player, finite-stage game with LMSR, if players have conditionally independent signals, a PBE of the game is a strategy-belief pair where each player reports their posterior probabilities in their first stage of play and all players believe that other players are truthful.*

*Proof.* Given that every player believes that all players before it act truthfully, we prove the theorem recursively. If it's player  $i$ 's last chance to play, it will truthfully report its posterior probabilities by Lemma 1. If it's player  $i$ 's second to last chance to play, there are other players standing in between its second to last chance to play and its last chance to play. We can combine the signals of those players standing in between as one signal and treat those players as one composite player. Because signals are conditionally independent, the signal of the composite player is conditionally independent of the signal of player  $i$ . The game becomes an Alice-Bob-Alice game for player  $i$  and at the unique PBE player  $i$  reports truthfully at its second to last chance to play according to Theorem 2. Inferring recursively, any player should report truthfully at its first chance to play.  $\square$

## 4 LMSR with Conditionally Dependent Signals

We now introduce a simple model of conditionally dependent signals and show that bluffing can be an equilibrium. In our model, Alice and Bob each see an independent coin flip and then participate in an LMSR prediction market with outcomes corresponding to whether or not both coins came up heads. Thus  $\omega \in \{\text{HH}, (\text{HT}|\text{TH}|\text{TT})\}$ . We again consider an Alice-Bob-Alice game structure.

**Theorem 4.** *In the Alice-Bob-Alice LMSR coin-flipping game, where the probability of heads is  $p$ , truthful betting (TB) is not a PBE. Now restrict Alice's first round strategies to either play TB or as if her coin is heads ( $\hat{H}$ ). A PBE in this game has Alice play TB with probability  $1 + \frac{p}{(1-(1-p)^{-1/p})(1-p)}$ , and otherwise play  $\hat{H}$ .*

*Proof.* TB cannot be an equilibrium because if Bob trusted Alice's move in the first round then her best response would be to pretend to have heads when she has tails. By doing so Bob would, when he has heads, move the probability of

HH to 1. Alice would then move the probability to 0 in the last round and collect an infinite payout.

To show that bluffing is a PBE in the restricted game, we show that Bob's best response makes Alice indifferent between her pure strategies. Bob's best response is, if he has heads, to set the probability of HH to the probability that Alice has heads given that she plays  $\hat{H}$ , or  $\Pr(\text{HH} \mid \hat{\text{HH}})$ . If Bob has tails he sets the probability of HH to zero. Assuming such a strategy for Bob, we can compute Alice's expected utility for playing TB and  $\hat{H}$ . It turns out that Alice's expected utility is the same whether she plays TB or  $\hat{H}$ . Thus in a PBE Alice should, with probability  $\frac{p}{(1-(1-p)^{-1/p})(1-p)}$ , pretend to have seen heads regardless of her actual information.  $\square$

Note that conditional dependence of signals is not a sufficient condition for bluffing in LMSR. Taking an extreme example, suppose that Alice and Bob again predict whether or not two coins both come up heads. Alice observes the result of one coin flip, but Bob with probability 1/2 observes the *same* coin flip as Alice and otherwise observes nothing. Then Alice will want to play truthfully and completely reveal her information in the first stage.

## 5 Withholding Information in DPM

Suppose Alice has the opportunity to trade in a two-outcome DPM with initial shares  $\mathbf{q}^0 = \langle 1, 1 \rangle$  for outcomes  $\{Y, N\}$ . According to equation (2), the initial market prices for the two outcomes are  $\langle p_Y^0, p_N^0 \rangle = \langle \kappa/\sqrt{2}, \kappa/\sqrt{2} \rangle$ . The initial market probabilities, according to equation (4), are  $\langle \pi_Y^0, \pi_N^0 \rangle = \langle 1/2, 1/2 \rangle$ .

Let  $p$  be Alice's posterior probability of outcome  $Y$  given her private information. If there are no other participants and  $p > 1/2$  then Alice should buy shares in outcome  $Y$  until the market probability  $\pi_Y$  reaches  $p$ . Thus, Alice's best strategy is to change market probabilities to  $\langle p, 1-p \rangle$  when  $p > 1/2$ .

We now show that if Alice anticipates that a sufficiently better-informed player will bet *after* her, then she will not fully reveal her information.

**Theorem 5.** *Alice, believing that outcome  $Y$  will occur with probability  $p > 1/2$ , plays in a two-outcome DPM seeded with initial quantities  $\langle 1, 1 \rangle$ . If a perfectly-informed Oracle plays after her, Alice will move the market probability of outcome  $Y$  to  $\max(p^2, 1/2)$ .*

*Proof.* Alice's expected utility is:

$$\kappa \left( px \frac{\sqrt{(1+x+g)^2+1}}{1+x+g} - \left( \sqrt{(1+x)^2+1} - \sqrt{2} \right) \right). \quad (6)$$

where  $x$  and  $g$  are the quantities of shares of  $Y$  purchased by Alice and the Oracle, respectively. Without loss of generality, suppose the true outcome is  $Y$ . Since the Oracle knows the outcome with certainty, we take the limit of (6) as  $g$  approaches infinity, yielding:  $\kappa(px - \sqrt{(1+x)^2+1} + \sqrt{2})$ . We find

the maximum using the first-order condition. This yields a function of  $p$  giving the optimal number of shares for Alice to purchase,  $x^* = \max(0, \frac{p}{\sqrt{1-p^2}} - 1)$ , which is greater than zero only when  $p > 1/\sqrt{2} \approx 0.707$ . The new numbers of shares are  $\mathbf{q} = \langle x^* + 1, 1 \rangle$ , yielding the market probability of outcome  $Y$  equal to  $\max(p^2, 1/2)$ .  $\square$

By assuming that the second player is perfectly informed, we mimic the scenario where a prediction market closes after the true outcome is revealed.

## 6 Conclusion

We have investigated the strategic behavior of traders in the MSR and DPM prediction markets using dynamic games. Specifically, we examine different scenarios where traders at equilibrium bet truthfully, bluff, or strategically delay.

Two different information structures, conditional independence and conditional dependence of signals, are considered for LMSR with incomplete information. We show that traders with conditionally independent signals may be worse off by either delaying trading or bluffing in LMSR. Moreover, truthful betting is a PBE strategy for all traders in LMSR with finite traders and finite periods. On the other hand, when the signals of traders are conditionally dependent there may exist probability distributions on signals such that truthful betting is not an equilibrium strategy; traders have an incentive to strategically mislead other traders with the intent of correcting the errors made by others in a later period; such bluffing can be a PBE strategy. DPM with incomplete information is shown to face another problem: traders may have an incentive to completely or partially withhold their private information if they anticipate sufficiently better-informed traders in later periods.

## References

1. Fama, E.F.: Efficient capital market: A review of theory and empirical work. *Journal of Finance* **25** (1970) 383–417
2. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Anatomy of an experimental political stock market. *American Economic Review* **82**(5) (1992) 1142–1161
3. Forsythe, R., Rietz, T.A., Ross, T.W.: Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization* **39** (1999) 83–110
4. Oliven, K., Rietz, T.A.: Suckers are born, but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market. *Management Science* **50**(3) (2004) 336–351
5. Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A.: Results from a dozen years of election futures markets research. In Plott, C.A., Smith, V., eds.: *Handbook of Experimental Economic Results* (forthcoming). (2001)
6. Berg, J.E., Rietz, T.A.: Prediction markets as decision support systems. *Information Systems Frontier* **5** (2003) 79–93

7. Gandar, J.M., Dare, W.H., Brown, C.R., Zuber, R.A.: Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance* **LIII**(1) (1999) 385–401
8. Thaler, R.H., Ziemba, W.T.: Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives* **2**(2) (1988) 161–174
9. Debnath, S., Pennock, D.M., Giles, C.L., Lawrence, S.: Information incorporation in online in-game sports betting markets. In: *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC'03)*, San Diego, CA (2003)
10. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper No. 1131, California Institute of Technology, Division of the Humanities and Social Sciences (2002)
11. Milgrom, P., Stokey, N.L.: Information, trade and common knowledge. *Journal of Economic Theory* **26**(1) (1982) 17–27
12. Kyle, A.S.: Continuous auctions and insider trading. *Econometrica* **53**(6) (1985) 1315–1336
13. Chakraborty, A., Yilmaz, B.: Manipulation in market order models. *Journal of Financial Markets* **7**(2) (2004) 187–206
14. Hanson, R.D.: Combinatorial information market design. *Information Systems Frontiers* **5**(1) (2003) 107–119
15. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* **1**(1) (2007) 1–15
16. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC'04)*. (2004)
17. Mangold, B., Dooley, M., Dornfest, R., Flake, G.W., Hoffman, H., Kasturi, T., Pennock, D.M.: The tech buzz game. *IEEE Computer* **38**(7) (2005) 94–97
18. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomics Theory*. Oxford University Press, New York, NY (1995)
19. Allen, F., Gale, D.: Stock-price manipulation. *The Review of Financial Studies* **5**(5) (1992) 503–529
20. Kumar, P., Seppi, D.J.: Futures manipulation with cash settlement. *Journal of Finance* **47**(1992) 1485–1502
21. Hansen, J., Schmidt, C., Strobel, M.: Manipulation in political stock markets - preconditions and evidence. Technical Report (2001)
22. Camerer, C.F.: Can asset markets be manipulated? A field experiment with race-track betting. *Journal of Political Economy* **106**(1998) 457–482
23. Hanson, R.D., Oprea, R., Porter, D.: Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization* **60**(4) (2007) 449–459
24. Rhode, P.W., Strumpf, K.S.: Historical presidential betting markets. *Journal of Economic Perspectives* **18**(2) (2004) 127–142
25. Rhode, P.W., Strumpf, K.S.: Manipulating political stock markets: A field experiment and a century of observational data. Working Paper (2007)
26. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: *The Second Workshop on Prediction Markets*, San Diego, CA (2007)
27. Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. In: *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, Vancouver, BC Canada (2007) 49–56
28. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. John Wiley & Sons, Inc. (1991)