Gaming Dynamic Parimutuel Markets

Qianya Lin[1,⋆] and Yiling Chen[2]

[1] City University of Hong Kong, Hong Kong SAR
[2] Harvard University, Cambridge, MA, USA

Abstract. We study the strategic behavior of risk-neutral non-myopic agents in *Dynamic Parimutuel Markets* (DPM). In a DPM, agents buy or sell shares of contracts, whose future payoff in a particular state depends on aggregated trades of all agents. A forward-looking agent hence takes into consideration of possible future trades of other agents when making its trading decision. In this paper, we analyze non-myopic strategies in a two-outcome DPM under a simple model of incomplete information and examine whether an agent will truthfully reveal its information in the market. Specifically, we first characterize a single agent's optimal trading strategy given the payoff uncertainty. Then, we use a two-player game to examine whether an agent will truthfully reveal its information when it only participates in the market once. We prove that truthful betting is a Nash equilibrium of the two-stage game in our simple setting for uniform initial market probabilities. However, we show that there exists some initial market probabilities at which the first player has incentives to mislead the other agent in the two-stage game. Finally, we briefly discuss when an agent can participate more than once in the market whether it will truthfully reveal its information at its first play in a three-stage game. We find that in some occasions truthful betting is not a Nash equilibrium of the three-stage game even for uniform initial market probabilities.

1 Introduction

Prediction markets are used to aggregate dispersed information about uncertain events of interest and have provided accurate forecasts of event outcomes, often outperforming other forecasting methods, in many real-world domains [1–8]. To achieve its information aggregation goal, a prediction market for an uncertain event offers contracts whose future payoff is tied to the event outcome. For example, a contract that pays off $1 per share if there are more than 6,000 H1N1 flu cases confirmed in U.S. by August 30, 2009 and $0 otherwise can be traded to predict the likelihood of the specified activity level of H1N1 flu.

Most market mechanisms used by prediction markets, including *continuous double auctions* (CDA) and *market scoring rules* (MSR) [9, 10], trade contracts whose payoff in each state is fixed, as in the above example. Contracts in *dynamic parimutuel markets* (DPM) [11, 12], however, have variable payoff that depends on the aggregated trades of all market participants. The payoff uncertainty makes DPM a mechanism that admits more speculation and strategic play.

⋆ Part of this work was done while Qianya Lin was visiting Harvard University.

As the goal of prediction markets is to aggregate information, it is important to understand whether and how participants reveal their information in the market. In this paper, we study the strategic behavior of risk-neutrual non-myopic agents in a two-outcome DPM under a simple setting of incomplete information, with the intent to understand how forward-looking agents reveal their information in DPM and whether they will reveal their information truthfully. We first characterize a single agent's optimal trading strategy given payoff uncertainty. Then, we consider a two-player two-stage dynamic game where each player only participates once in DPM, to examine whether the first player has incentives to misreport its information to mislead the second player and obtain higher profit even if it can only play once. We prove that truthful betting is a Nash equilibrium of the two-stage game for uniform initial market probabilities. We show that there exists some initial market probabilities at which the first player has incentives to mislead the other agent in the two-stage game. Finally, we discuss a three-stage game in which an agent can participate more than once. we find that the truthful betting is not a Nash equilibrium of the three-stage game in some occasions even for uniform initial market probabilities.

Related Work. Chen et al. [13] provide a specific example of a two-player two-stage game in DPM where the second player is perfectly informed and show that the first player may sometimes choose not to trade. Our work does not assume perfectly informed agents and we characterize non-myopic strategies in more general settings. The example of Chen et al. is a special case of our results for the two-player two-stage game. Nikolova and Sami [14] use a projection game to study DPM. They show that a rational agent will never hold shares of both outcomes in a two-outcome DPM when short sales are not allowed. This is consistent with our characterization of a single agent's optimal strategy given payoff uncertainty. Bu et al. [15] study the strategies of a myopic agent who believes that the contract payoff in the future is the same as the payoff if the market closes right after the agent's trade in a DPM. Our work focuses on forward-looking agents who take into consideration of the payoff uncertainty when making their trading decisions.

Some theoretical attempts have been made to characterize non-myopic strategies in other markets, including logarithmic market scoring rule (LMSR) [13, 16, 17], financial markets (i.e. CDA) [18–20], and parimutuel markets [21]. In all these markets, agents may have incentives to misreport their information. Ostrovsky [22] provides a separability condition that contracts need to satisfy to guarantee market convergence to full information aggregation at a perfect Bayesian equilibrium in LMSR and CDA.

2 Dynamic Parimutuel Markets

A dynamic parimutuel market (DPM) [11, 12] is a dynamic-cost variant of a parimutuel market. Suppose an uncertain event of interest has n mutually exclusive outcomes. Let Ω denote the outcome space. A DPM offers n contracts, each corresponding to an outcome. 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, hence incentivizing traders to reveal their information earlier.

DPM operates as a market maker. Let q_ω be the total number of shares of contract ω that have been purchased by all traders. We use $\boldsymbol{q}$ to denote the vector of outstanding shares for all contracts. The DPM market maker keeps a cost function, $C(\boldsymbol{q}) = \sqrt{\sum_{\omega \in \Omega} q_\omega^2}$, that captures total money wagered in the market, and an instantaneous price function for contract ω, $p_\omega = \frac{q_\omega}{\sqrt{\sum_{\psi \in \Omega} q_\psi^2}}$. A trader who buys contracts and changes the outstanding shares from $\boldsymbol{q}$ to $\tilde{\boldsymbol{q}}$ pays the market maker $C(\tilde{\boldsymbol{q}}) - C(\boldsymbol{q})$. The market probability on outcome ω is $\pi_\omega = \frac{q_\omega^2}{\sum_{\psi \in \Omega} q_\psi^2}$. In DPM, market price of a contract does not represent the market probability of the corresponding state. Instead, $\pi_\omega = p_\omega^2$.

If outcome ω is realized, each share of contract ω gets an equal share of the total market capitalization. Its payoff is $o_\omega = \frac{\sqrt{\sum_{\psi \in \Omega} (q_\psi^f)^2}}{q_\omega^f}$, where q_ω^f is the outstanding shares of contract ω at the end of the market. All other contracts have zero payoff. As the value of $\boldsymbol{q^f}$ is not known before the market closes, o_ω is not fixed while the market is open. The relation of the final market price, final market probability, and the contract payoff when outcome ω is realized, is $o_\omega = \frac{1}{p_\omega^f} = \frac{1}{\sqrt{\pi_\omega^f}}$, where p_ω^f and π_ω^f denote the last market price and market probability before the market closes.

As a market maker mechanism, DPM offers infinite liquidity. Because the price function is not defined when $\boldsymbol{q} = 0$, the market maker subsidizes the market by starting the market with some positive shares. The subsidy turns DPM into a positive-sum game and can circumvent the *no-trade theorem* [23] for zero-sum games. Tech Buzz Game [12] used DPM as its market mechanism and market probabilities in the game have been shown to offer informative forecasts for the underlying events [24].

3 Our Setting

We consider a simple incomplete information setting for a DPM in this paper. There is a single event whose outcome space contains two discrete mutually exclusive states $\Omega = \{Y, N\}$. The eventual event outcome is picked by Nature with prior probability $\mathrm{P}(Y) = \mathrm{P}(N) = \frac{1}{2}$. The DPM offers two contracts, each corresponding to one outcome. There are two players in the market. Each player i receives a piece of private signal $c_i \in \{y_i, n_i\}$. The signal is independently drawn by Nature conditional on the true state. In other words, signals are conditionally independent, $\mathrm{P}(c_i, c_j|\omega) = \mathrm{P}(c_i|\omega)\mathrm{P}(c_j|\omega)$. The prior probabilities and the signal distributions are common knowledge to all players.

We further assume that player's signals are symmetric such that $\mathrm{P}(y_i|Y) = \mathrm{P}(n_i|N)$ for all i. With this, we define the *signal quality* of player i as $\theta_i = \mathrm{P}(y_i|Y) = \mathrm{P}(n_i|N)$. The signal quality θ_i captures the likelihood for agent i to receive a "correct" signal. Without loss of generality, we assume $\theta_i \in (\frac{1}{2}, 1]$.

The above conditions, together with conditional independence of signals, imply that for two players i and j we have $\mathrm{P}(Y|c_i, y_j) > \mathrm{P}(Y|c_i, n_j)$, $\mathrm{P}(N|c_i, n_j) > \mathrm{P}(N|c_i, y_j)$, $\mathrm{P}(y_i|y_j) > \mathrm{P}(n_i|y_j)$, and $\mathrm{P}(n_i|n_j) > \mathrm{P}(y_i|n_j)$ for $c_i \in \{y_i, n_i\}$.

Short sell is not allowed in the market. Agents are risk neutral and participate in the market sequentially. We also assume that they do not possess any shares at the beginning of the market and have unlimited wealth.

4 Optimal Trading Strategy of A Single Agent

We first consider a single agent's optimal trading strategy given the payoff uncertainty in DPM. We assume that the agent only trades once in DPM. Let $\mathrm{P}(\omega, s)$ be the agent's subjective probability that the event outcome will be ω and the set of information available to the last trader is s. The market probability at the end of the market will reflect all available information. Hence the final price for contract ω when s is available is $\pi^f_\omega(s) = \mathrm{P}(\omega|s)$.

The agent compares the current price of a contract, p_ω, with its expected future payoff. Note that the future payoff of contract ω in state ω only relates to the final market probability π^f_ω and does not relate to the process of reaching it. The expected future payoff of contract ω is $\varphi_\omega = \sum_s \frac{\mathrm{P}(\omega,s)}{\sqrt{\pi^f_\omega(s)}} = \sum_s \frac{\mathrm{P}(\omega,s)}{\sqrt{\mathrm{P}(\omega|s)}} = \sum_s \mathrm{P}(s)\sqrt{\mathrm{P}(\omega|s)}$. We have the following lemma.

Lemma 4.1. $\sum_\omega \varphi^2_\omega \leq 1$.

Suppose the agent purchases $\boldsymbol{\Delta q}$ and changes the outstanding shares from $\boldsymbol{q}$ to $\tilde{\boldsymbol{q}} = \boldsymbol{q} + \boldsymbol{\Delta q}$. The market prices before and after the trade are $\boldsymbol{p}$ and $\tilde{\boldsymbol{p}}$ respectively. Theorem 4.2 characterizes the agent's optimal purchases when it attempts to maximize its expected profit.

Theorem 4.2. *In a two-outcome DPM, if a risk-neutral agent maximizes its expected profit by purchasing $\boldsymbol{\Delta q} \geq 0$, the following conditions must satisfy:*

1. *For any contract ω, if $p_\omega < \varphi_\omega$, then $\Delta q_\omega > 0$ and $\tilde{p}_\omega = \varphi_\omega$.*
2. *For any contract ω, if $p_\omega > \varphi_\omega$, then $\tilde{p}_\omega \geq \varphi_\omega$ and when the inequality is strict, $\Delta q_\omega = 0$.*
3. *For any contract ω, if $p_\omega > \varphi_\omega$, $\tilde{p}_\omega = \varphi_\omega$, and $\Delta q_\omega > 0$, there exists an equivalent $\boldsymbol{\Delta q'} \geq 0$ with $\Delta q'_\omega = 0$ that satisfies conditions 1 and 2 and have the same expected profit as $\boldsymbol{\Delta q}$.*
4. *If $p_Y > \varphi_Y$ and $p_N > \varphi_N$, $\boldsymbol{\Delta q} = \mathbf{0}$.*

Theorem 4.2 means that, in a two-outcome DPM, when $\sum_\omega \varphi^2_\omega < 1$, the optimal strategy for an agent is to buy shares of the contract whose current price is lower than its expected payoff and drive its price up to its expected payoff. When $\sum_\omega \varphi^2_\omega = 1$, it's possible to achieve the desired market prices by purchasing both contracts, but this is equivalent, in terms of expected profit, to the strategy that only purchases the contract whose current price is lower than its expected payoff. Thus, the optimal strategy of an agent is to buy shares for the contract whose current price is too low. We now give the optimal shares that an agent would purchase and its optimal expected profit in the following theorem.

Theorem 4.3. *In a two-outcome market, when* $\frac{q_\omega}{\sqrt{q_\omega^2+q_{\bar\omega}^2}} < \varphi_\omega \le 1$*, a trader with expected payoff* φ_ω *for contract* ω *will purchase* $\Delta q_\omega^* = \frac{\varphi_\omega}{\sqrt{1-\varphi_\omega^2}} q_{\bar\omega} - q_\omega$ *to maximize his expected profit, where* q_ω *is the current outstanding shares for outcome* ω *in the market, and* $q_{\bar\omega}$ *is the outstanding shares for the other outcome. His optimal expected profit is* $U(\Delta q_\omega^*) = \sqrt{q_\omega^2 + q_{\bar\omega}^2} - q_\omega \varphi_\omega - q_{\bar\omega}\sqrt{1-\varphi_\omega^2}$*. When* $\frac{q_\omega}{\sqrt{q_\omega^2+q_{\bar\omega}^2}} > \varphi_\omega$ *and* $\frac{q_{\bar\omega}}{\sqrt{q_\omega^2+q_{\bar\omega}^2}} > \varphi_{\bar\omega}$*, the trader does not purchase any contract.*

Because of the payoff uncertainty, an agent who only trades once in DPM will not change the market price to its posterior probability as in CDA or MSR, but will change the market price to $(\varphi_Y, \sqrt{1-\varphi_Y^2})$ if purchasing contract Y and to $(\sqrt{1-\varphi_N^2}, \varphi_N)$ if purchasing contract N. The corresponding market probabilities are $(\varphi_Y^2, 1-\varphi_Y^2)$ and $(1-\varphi_N^2, \varphi_N^2)$ respectively. It is possible that the agent's optimal strategy is to not trade in the market. This happens when the current price $\frac{q_\omega}{\sqrt{q_\omega^2+q_{\bar\omega}^2}}$ is greater than φ_ω for all ω.

5 Two-Player Games

When analyzing a single agent's optimal strategy in the previous section, we do not consider the possibility that the agent's behavior may affect the set of information available to later traders. In DPM, as agents can infer information of other agents from their trading decisions, an agent who plays earlier in the market may mislead those who play later and affects the information sets of the later traders. In this section, we use two-player games to study this issue.

5.1 Two-Player Two-Stage Game

We first consider a two-player two-stage game, the Alice-Bob game, to examine whether an agent will try to affect the expected contract payoff by misreporting its own information. In the Alice-Bob game, Alice and Bob are the only players in the market. Each of them can trade only once. Alice plays first, followed by Bob. We are interested in whether there exists an equilibrium at which Alice fully reveals her information in the first stage if she trades and Bob infers Alice's information and acts based on both pieces of information in the second stage. In particular, at the equilibrium, when having signal $c_A \in \{y_A, n_A\}$, Alice believes that $\varphi_\omega(c_A) = \sum_{c_B} \mathrm{P}(c_B|c_A)\sqrt{\mathrm{P}(\omega|c_A, c_B)}$ and plays her optimal strategy according to Theorem 4.3. Bob with signal c_B believes that $\varphi_\omega(c_A, c_B) = \sqrt{\mathrm{P}(\omega|c_A, c_B)}$ if Alice trades in the first stage and $\varphi_\omega(c_B) = \sqrt{\mathrm{P}(\omega|c_B, \text{Alice doesn't trade})}$ if Alice doesn't trade, and plays his optimal strategy according to Theorem 4.3. We call such an equilibrium a *truthful betting* equilibrium. Note that the truthful betting equilibrium does not guarantee full information aggregation at the end of the game, because if Alice does not trade her information is not fully revealed. Since we are interested in the strategic behavior of agents, we also assume that $\varphi_\omega(y_A) \neq \varphi_\omega(n_A)$ to rule out degenerated cases.

In the rest of this section, we show that truthful betting equilibrium exists when the initial market probability is uniform, but does not exist with some other initial market probabilities.

Truthful Betting Equilibrium We assume that the market starts with uniform initial market probability, i.e. $\pi_Y = \pi_N = \frac{1}{2}$. This means that the initial market prices are $p_Y = p_N = \frac{1}{\sqrt{2}}$. As signals are symmetric, we have the following lemma.

Lemma 5.1. *Suppose all information is revealed after Bob's play. Alice's expected payoff for contract Y when she has signal y_A equals her expected payoff for contract N when she has signal n_A. That is, $\varphi_Y(y_A) = \varphi_N(n_A)$, where $\varphi_Y(y_A) = \sum_{c_B} P(c_B|y_A)\sqrt{P(Y|c_B, y_A)}$ and $\varphi_N(n_A) = \sum_{c_B} P(c_B|n_A)\sqrt{P(N|c_B, n_A)}$.*

We use φ to denote both $\varphi_Y(y_A)$ and $\varphi_N(n_A)$. Theorem 5.1 characterizes the truthful betting equilibrium for the game with uniform initial market probability.

Theorem 5.2. *In a two-outcome DPM with uniform initial market probability, truthful betting is a Bayesian Nash equilibrium for the Alice-Bob game. At the equilibrium, Alice does not trade if $\varphi \leq \frac{1}{\sqrt{2}}$. If $\varphi > \frac{1}{\sqrt{2}}$, Alice purchases contract Y and changes the price for Y to φ if she has y_A, and purchases contract N and changes the price for N to φ if she has n_A. If Alice trades, Bob infers her signal and changes the market probability to the posterior probability conditional on both signals. If Alice does not trade, Bob changes the market probability to the posterior probability conditional on his own signal.*

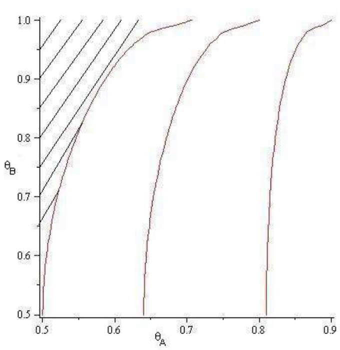

Fig. 1. Signal Qualities and Alice's Expected Payoff

Figure 1 plots the iso-value lines of φ as a function of θ_A and θ_B. The leftmost curve is $\varphi(\theta_A, \theta_B) = \frac{1}{\sqrt{2}}$. The value of φ increases as the curve moves toward the right. As the initial market price is $\frac{1}{\sqrt{2}}$ for both outcomes, the curve $\varphi(\theta_A, \theta_B) = \frac{1}{\sqrt{2}}$ gives the boundary that at the equilibrium Alice trades in the first stage. The shaded area gives the range of signal qualities that Alice is better off not trading at the equilibrium. When $\theta_B = 1$, that is Bob is perfectly informed, Alice won't trade if her signal quality θ_A is less than $\frac{1}{\sqrt{2}}$. This is consistent with the example given by Chen et al. [13].

Non-existence of the Truthful Betting Equilibrium In Alice-Bob game, truthful betting is not a Nash equilibrium for arbitrary initial market probability.

Theorem 5.3. *In a two-outcome DPM, there exists some initial market probabilities where truthful betting is not a Nash equilibrium for the Alice-Bob game.*

The intuition is that if the initial market price for one contract is very low that Alice will purchase the the contract no matter which signal she gets, Alice may pretend to have a different signal by purchasing less when she should buy more if being truthful. If Bob is mislead, this increases the expected payoff per share of the contract and hence can increase Alice's expected total profit even if she purchases less. This is very different from other market mechanisms. In both CDA and MSR, if a player only plays once in the market, disregard of whether there are other players behind it, the player will always play truthfully.

5.2 Two-Player Three-Stage Game

In the Alice-Bob game, Alice may not play truthfully in order to mislead Bob and achieve a higher expected payoff per share, but she does not directly make profits from Bob's uninformed trades. Now we consider a three-stage game, the Alice-Bob-Alice game, where Alice can play a second time after Bob's play. Truthful betting equilibrium in this game means that both players play their truthful betting equilibrium strategies of the Alice-Bob game in the first two stages and Alice does nothing in the third stage. Clearly, if Alice has incentives to deviate from truthful betting in the Alice-Bob game, she will also deviate in the Alice-Bob-Alice game, because playing a second time allows Alice to gain more profit by capitalizing on Bob's uninformed trades. Even for settings where truthful betting is a Bayesian Nash equilibrium for the Alice-Bob game, a truthful betting equilibrium may not exist for the Alice-Bob-Alice game. For example, with uniform initial market probability, if $\theta_A = 0.6$, $\theta_B = 0.8$ and Alice has y_A, Alice is better off pretending to have n_A given Bob believes that she plays truthfully. In contrast, in LMSR when agents have conditionally independent signals, truthful betting is the unique perfect Bayesian equilibrium [13, 17].

6 Conclusion

Using a simple setting of incomplete information, we show that DPM admits more gaming than several other prediction market mechanisms due to its payoff uncertainty. We show that even when a player only participates once in the market, e.g. in an Alice-Bob game, it still has incentives to bluff and pretend to have a different signal. The bluffing behavior exists more generally when traders participate the market more than once.

References

1. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Forecasting elections: A market alternative to polls. In Palfrey, T.R., ed.: Contemporary Laboratory Experiments in Political Economy. University of Michigan Press (1991) 69–111
2. 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. (2001)

3. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper (2002)
4. Cowgill, B., Wolfers, J., EricZitzewitz: Using prediction markets to track information flows: Evidence from Google. Working paper (2008)
5. Thaler, R.H., Ziemba, W.T.: Anomalies: Parimutuel betting markets: Racetracks and lotteries. Journal of Economic Perspectives **2**(2) (1988) 161–174
6. Pennock, D.M., Lawrence, S., Giles, C.L., Nielsen, F.A.: The real power of artificial markets. Science **291** (February 2002) 987–988
7. Debnath, S., Pennock, D.M., Giles, C.L., Lawrence, S.: Information incorporation in online in-game sports betting markets. In: The 4th ACM Conference on Electronic Commerce (EC). (2003)
8. Chen, Y., Chu, C.H., Mullen, T., Pennock, D.M.: Information markets vs. opinion pools: An empirical comparison. In: The 6th ACM Conference on Electronic Commerce (EC). (2005)
9. Hanson, R.D.: Combinatorial information market design. Information Systems Frontiers **5**(1) (2003) 107–119
10. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. Journal of Prediction Markets **1**(1) (2007) 1–15
11. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: The 5th ACM Conference on Electronic Commerce (EC). (2004)
12. 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
13. Chen, Y., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R.: Bluffing and strategic reticence in prediction markets. In: Lecture Notes in Computer Science, Workshop on Internet and Network Economics (WINE). (2007)
14. Nikolova, E., Sami, R.: A strategic model for information markets. In: The Seventh Annual ACM Conference on Electronic Commerce (EC). (2007)
15. Bu, T.M., Deng, X., Lin, Q., Qi, Q.: Strategies in dynamic pari-mutuel markets. In: Lecture Notes in Computer Science, Workshop on Internet and Network Economics (WINE). (2008)
16. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: The Nineth ACM Conference on Electronic Commerce (EC). (2008)
17. Chen, Y., Dimitrov, S., Sami, R., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R.: Gaming prediction markets: Equilibrium strategies with a market maker. Algorithmica (2009)
18. Allen, F., Gale, D.: Stock-price manipulation. The Review of Financial Studies (5) (1992) 503–529
19. Chakraborty, A., Yilmaz, B.: Manipulation in market order models. Journal of Financial Markets **7**(2) (2004) 187–206
20. Kumar, P., Seppi, D.J.: Futures manipulation with cash settlement. Journal of Finance (47) (1992) 1485–1502
21. Koessler, F., Noussair, C., Ziegelmeyer, A.: Parimutuel betting under asymmetric information. In: Journal of Mathematical Economics. (July 2008) 733–744
22. Ostrovsky, M.: Information aggregation in dynamic markets with strategic traders. Working paper (2009)
23. Milgrom, P., Stokey, N.L.: Information, trade and common knowledge. Journal of Economic Theory **26**(1) (1982) 17–27
24. Chen, Y., Pennock, D.M., Kasturi, T.: An empirical study of dynamic pari-mutuel markets: Evidence from the Tech Buzz Game. In: Lecture Notes in Computer Science, Workshop on Web Mining and Web Usage Analysis (WebKDD). (2008)

Appendix

A Proofs

A.1 Proof of Lemma 4.1

$\sum_{\omega} \varphi_{\omega}^2 = \sum_{\omega} \left(\sum_s \mathrm{P}(s) \sqrt{\mathrm{P}(\omega|s)} \right)^2 \leq \sum_{\omega} \sum_s \mathrm{P}(s)\mathrm{P}(\omega|s) = 1$. The inequality is due to Jensen's inequality and the fact that the function $f(x) = x^2$ is convex.

A.2 Proof of Theorem 4.2

According to Lemma 4.1, $\sum_{\omega} \varphi_{\omega}^2 \leq 1$. Because $\sum_{\omega} p_{\omega}^2 = 1$, if there exists some state such that $p_{\omega} < \varphi_{\omega}$, there must exist some other state where the inequality is reversed.

We first prove that after the agent's purchase, for any outcome ω, it must be $\tilde{p}_{\omega} \geq \varphi_{\omega}$. Suppose the contrary, $\tilde{p}_{\omega} < \varphi_{\omega}$. The agent can always purchase more shares for contract ω to increase his expected profit because the expected payoff exceeds the cost of purchase. This contradicts that purchasing $\boldsymbol{\Delta q}$ maximizes the expected profit of the agent.

Because the price of contract ω increases with q_{ω} and decreases with q_{ψ} for $\psi \neq \omega$. If $p_{\omega} < \varphi_{\omega}$ and $\tilde{p}_{\omega} \geq \varphi_{\omega}$, it must be that $\Delta q_{\omega} > 0$. Moreover, when $\Delta q_{\omega} > 0$, it can not be that $\tilde{p}_{\omega} > \varphi_{\omega}$ because the trader can buy less and increase his expected profit since the cost of purchasing is higher than the expected payoff. Thus, if $p_{\omega} < \varphi_{\omega}$, we must have $\Delta q_{\omega} > 0$ and $\tilde{p}_{\omega} = \varphi_{\omega}$. This proves condition 1.

When $p_{\omega} > \varphi_{\omega}$ and $\tilde{p}_{\omega} > \varphi_{\omega}$, Δq_{ω} can not be positive, because if it is, by decreasing it the agent can get higher expected profit. This proves condition 2.

To prove condition 3, without loss of generality, assume that $p_Y > \varphi_Y$, $\tilde{p}_Y = \varphi_Y$, and $\Delta q_Y > 0$. Because $p_Y > \varphi_Y$, $p_N < \varphi_N$. According to condition 1, we have $\tilde{p}_N = \varphi_N$. Thus, $\tilde{p}_Y^2 + \tilde{p}_N^2 = \varphi_Y^2 + \varphi_N^2$=1. Thus, starting with $\boldsymbol{\Delta q}$, we can simultaneously reduce purchases of both contracts while keeping the price $\tilde{\boldsymbol{p}}$ unchanged, until the holding for contract Y drops to 0. This doesn't change the expected profit of the trader since the sell price equals the expected payoff.

When the price is higher than the expected payoff for both contracts, it is clearly optimal for the agent to not buy any contract as buying will result in loss in expectation. This gives condition 4.

A.3 Proof of Theorem 4.3

As $\varphi_{\omega} > \frac{q_{\omega}}{(q_{\omega})^2 + (q_{\bar{\omega}})^2}$, according to Theorem 4.2, the trader will purchase Δq_{ω}^* shares of contract ω and change the market price of ω to φ_{ω}, i.e. $\frac{q_{\omega} + \Delta q_{\omega}^*}{\sqrt{(q_{\omega} + \Delta q_{\omega}^*)^2 + q_{\bar{\omega}}^2}} = \varphi_{\omega}$. Solving this equation, we have $\Delta q_{\omega}^* = \frac{\varphi_{\omega}}{\sqrt{1 - \varphi_{\omega}^2}} q_{\bar{\omega}} - q_{\omega}$. The trader's expected utility is $U(\Delta q_{\omega}^*) = \varphi_{\omega} \Delta q_{\omega}^* - \sqrt{(q_{\omega} + \Delta q_{\omega}^*)^2 + q_{\bar{\omega}}^2} + \sqrt{q_{\omega}^2 + q_{\bar{\omega}}^2}$ for purchasing Δq_{ω}^* shares of contract ω. Plugging in the expression of Δq_{ω}^*, we get $U(\Delta q_{\omega}^*) = \sqrt{q_{\omega}^2 + q_{\bar{\omega}}^2} - q_{\omega} \varphi_{\omega} - q_{\bar{\omega}} \sqrt{1 - \varphi_{\omega}^2}$.

A.4 Proof of Lemma 5.1

When Alice has y_A,

$$\begin{aligned}
\varphi_Y(y_A) &= \sum_{c_B} \mathrm{P}(c_B, Y|y_A)\frac{1}{\sqrt{\pi_Y^f}} = \sum_{c_B} \mathrm{P}(c_B|y_A)\mathrm{P}(Y|c_B, y_A)\frac{1}{\sqrt{\mathrm{P}(Y|c_B, y_A)}} \\
&= \sum_{c_B} \mathrm{P}(c_B|y_A)\sqrt{\mathrm{P}(Y|c_B, y_A)} \\
&= \sqrt{\mathrm{P}(y_B|Y)\mathrm{P}(y_A|Y) + \mathrm{P}(y_B|N)\mathrm{P}(y_A|N)}\sqrt{\mathrm{P}(y_B|Y)\mathrm{P}(y_A|Y)} \\
&+ \sqrt{\mathrm{P}(n_B|Y)\mathrm{P}(y_A|Y) + \mathrm{P}(n_B|N)\mathrm{P}(y_A|N)}\sqrt{\mathrm{P}(n_B|Y)\mathrm{P}(y_A|Y)} \\
&= \sqrt{\theta_A\theta_B + (1-\theta_A)(1-\theta_B)}\sqrt{\theta_A\theta_B} \\
&+ \sqrt{\theta_A(1-\theta_B) + (1-\theta_A)\theta_B}\sqrt{\theta_A(1-\theta_B)}.
\end{aligned}$$

When Alice has n_A,

$$\begin{aligned}
\varphi_N(n_A) &= \sum_{c_B} \mathrm{P}(c_B, N|n_A)\frac{1}{\sqrt{\pi_N^f}} = \sum_{c_B} \mathrm{P}(c_B|n_A)\mathrm{P}(N|c_B, n_A)\frac{1}{\sqrt{\mathrm{P}(N|c_B, n_A)}} \\
&= \sum_{c_B} \mathrm{P}(c_B|n_A)\sqrt{\mathrm{P}(N|c_B, n_A)} \\
&= \sqrt{\mathrm{P}(y_B|N)\mathrm{P}(n_A|N) + \mathrm{P}(y_B|Y)\mathrm{P}(n_A|Y)}\sqrt{\mathrm{P}(y_B|N)\mathrm{P}(n_A|N)} \\
&+ \sqrt{\mathrm{P}(n_B|N)\mathrm{P}(n_A|N) + \mathrm{P}(n_B|Y)\mathrm{P}(n_A|Y)}\sqrt{\mathrm{P}(n_B|N)\mathrm{P}(n_A|N)} \\
&= \sqrt{\theta_A(1-\theta_B) + (1-\theta_A)\theta_B}\sqrt{\theta_A(1-\theta_B)} \\
&+ \sqrt{\theta_A\theta_B + (1-\theta_A)(1-\theta_B)}\sqrt{\theta_A\theta_B}.
\end{aligned}$$

In conclusion, we have $\varphi_Y(y_A)$=$\varphi_N(n_A)$.

A.5 Proof of Theorem 5.2

If Alice gets y_A, her expected payoff of contracts Y and N are $\varphi_Y(y_A) = \sum_{c_B} \mathrm{P}(c_B|y_A)\sqrt{\mathrm{P}(Y|c_B, y_A)}$ and $\varphi_N(y_A) = \sum_{c_B} \mathrm{P}(c_B|y_A)\sqrt{\mathrm{P}(N|c_B, y_A)}$ respectively.

$$\begin{aligned}
\varphi_Y(y_A) - \varphi_N(y_A) &= \sqrt{\mathrm{P}(y_B|y_A)}(\sqrt{\mathrm{P}(y_B|Y)\mathrm{P}(y_A|Y)} - \sqrt{\mathrm{P}(y_B|N)\mathrm{P}(y_A|N)}) \\
&+ \sqrt{\mathrm{P}(n_B|y_A)}(\sqrt{\mathrm{P}(n_B|Y)\mathrm{P}(y_A|Y)} - \sqrt{\mathrm{P}(n_B|N)\mathrm{P}(y_A|N)}) \\
&= \sqrt{\mathrm{P}(y_B|y_A)}(\sqrt{\theta_A\theta_B} - \sqrt{(1-\theta_A)(1-\theta_B)}) \\
&+ \sqrt{\mathrm{P}(n_B|y_A)}(\sqrt{\theta_A(1-\theta_B)} - \sqrt{(1-\theta_A)\theta_B}).
\end{aligned}$$

As signals are symmetric, $\mathrm{P}(y_B|y_A) > \mathrm{P}(n_B|y_A)$ and $\theta_A\theta_B > \theta_A(1-\theta_B)$, $(1-\theta_A)(1-\theta_B) < (1-\theta_A)\theta_B$. We have $\varphi_Y(y_A) > \varphi_N(y_A)$.

If Alice gets n_A, her expected payoff of contracts Y and N are $\varphi_Y(n_A) = \sum_{c_B} \mathrm{P}(c_B|n_A)\sqrt{\mathrm{P}(Y|c_B, n_A)}$ and $\varphi_N(n_A) = \sum_{c_B} \mathrm{P}(c_B|n_A)\sqrt{\mathrm{P}(N|c_B, n_A)}$ respectively.

$$\begin{aligned}\varphi_N(n_A) - \varphi_Y(n_A) &= \sqrt{\mathrm{P}(y_B|n_A)}(\sqrt{\mathrm{P}(y_B|N)\mathrm{P}(n_A|N)} - \sqrt{\mathrm{P}(y_B|Y)\mathrm{P}(n_A|Y)}) \\ &+ \sqrt{\mathrm{P}(n_B|n_A)}(\sqrt{\mathrm{P}(n_B|N)\mathrm{P}(n_A|N)} - \sqrt{\mathrm{P}(n_B|Y)\mathrm{P}(n_A|Y)}) \\ &= \sqrt{\mathrm{P}(y_B|n_A)}(\sqrt{(1-\theta_B)\theta_A} - \sqrt{(1-\theta_A)\theta_B}) \\ &+ \sqrt{\mathrm{P}(n_B|n_A)}(\sqrt{\theta_A\theta_B} - \sqrt{(1-\theta_A)(1-\theta_B)}).\end{aligned}$$

Similarly, $\varphi_N(n_A) > \varphi_Y(n_A)$.

According to Lemma 5.1, $\varphi_Y(y_A) = \varphi_N(n_A) = \varphi$. If the market starts with uniform initial probability, $p_Y = p_N = \frac{1}{\sqrt{2}}$. If $\varphi > \frac{1}{\sqrt{2}}$, we know that $\varphi_Y(y_A) > p_Y$, $\varphi_N(y_A) < p_N$, $\varphi_N(n_A) > p_N$, and $\varphi_Y(n_A) < p_Y$ because of Lemma 4.1. Alice will purchase Y if she has y_A and purchase N if she has n_A and change the market price for the corresponding contract to φ if she plays truthfully, according to Theorems 4.2 and 4.3. Without loss of generality, assume that Alice gets y_A.

We show that Alice won't pretend to have n_A when she get y_A if $\varphi > \frac{1}{\sqrt{2}}$. If she attempts to mislead Bob by acting as if she has n_A, she will buy on N until the market price of N reaches $\varphi_N(n_A)$, her actual expected payoff by playing this bluffing strategy is:

$$\varphi_N^B = \sum_{c_B} \mathrm{P}(N, c_B|y_A)\frac{1}{\sqrt{\mathrm{P}(N|n_A, c_B)}} < \sum_{c_B} \mathrm{P}(N, c_B|y_A)\frac{1}{\sqrt{\mathrm{P}(N|y_A, c_B)}} = \varphi_N(y_A).$$

The inequality comes from $\mathrm{P}(N|n_A, c_B) > \mathrm{P}(N|y_A, c_B)$. Comparing the expected final payoff, $\varphi_N^B < \varphi_N(y_A) < p_N$. This suggests that if Alice bluffs her expected final payoff per share is lower than her purchasing price. Alice is worse off bluffing.

Next, we show that Alice does not want to deviate to change the market price to any other values if $\varphi > \frac{1}{\sqrt{2}}$. At the equilibrium, Bob's belief is that Alice gets y_A if she buys contract Y and gets n_A if she buys contract N. If Alice buys contract N and changes its price to something different from $\varphi_N(n_A)$, her expected payoff is still φ_N^B which is lower than her purchasing price. If Alice buys contract Y, her expected payoff is $\varphi_Y(y_A)$ and truthful betting is her optimal strategy.

Finally, we prove that when $\varphi \leq \frac{1}{\sqrt{2}}$, Alice does not trade not matter what signal she has and Bob will change the final market probability to the posterior probability given his own signal. If Alice does not trade, Bob believes that she has y_A or n_A each with probability 0.5. Thus, Bob's best response is to change the market probability to his posterior probability conditional only on his signal. Alice does not want to deviate to trade. This is because, if she has y_A and purchases Y, Bob will then believe that she has y_A, and the expected payoff of Y for Alice becomes $\varphi_Y(y_A)$ which is lower than Alice's purchasing price. If she has y_A and purchases N, Bob will then believe that she has n_A, and the expected payoff of N for Alice becomes $\varphi_N(n_A)$ which is also less than Alice's purchasing price.

A.6 Proof of Theorem 5.3

For non-uniform initial market probability, we prove by giving an example where truthful betting is not the best response for Alice. When the initial market price for Y is low such that $\varphi_Y(y_A) > \varphi_Y(n_A) > p_Y$, Alice will buy contract Y no matter which signal she gets if she plays truthfully, although the amount that she will buy when having signal n_A is less than the amount that she will buy when having signal y_A. Assume $p_Y = \frac{1}{3}$, $\theta_A = \frac{3}{5}$, and $\theta_B = \frac{2}{5}$, which suggest $\varphi_Y(y_A) > \varphi_Y(n_A) > \frac{1}{3}$. The difference of Alice's expected profit between when she plays truthfully with a signal y_A, denoted U^T, and when she pretends to have a signal n_A and Bob believes it, denoted U^B, is $U^T - U^B = q_Y((\varphi_Y^B - \varphi_Y(y_A)) + \sqrt{8}(\frac{1-\varphi_Y(n_A)\varphi_Y^B}{\sqrt{1-(\varphi_Y(n_A))^2}} - \sqrt{1-(\varphi_Y(y_A))^2})) = -1.56q_Y < 0$. φ_Y^B is the expected payoff of Y when Alice pretends to have n_A. The difference is negative, showing that Alice would want to deviate from truthful betting.