An Empirical Study of Dynamic Pari-mutuel Markets: Evidence from the Tech Buzz Game

Yiling Chen[1], David M. Pennock[2], and Tejaswi Kasturi[3]

[1] Harvard University, Cambridge, MA
[2] Yahoo! Research, New York, NY
[3] Yahoo! Research, Burbank, CA

Abstract. A dynamic pari-mutuel market (DPM) is a hybrid between a continuous double auction (CDA) and a pari-mutuel market. Like a CDA, a DPM incentivizes traders to reveal their information early. Like a pari-mutuel market, a DPM has infinite liquidity, allowing trades at any time. In this paper, we examine empirical questions related to DPMs: Do prices in DPMs predict events of interests? How do traders behave in DPMs? Leveraging a data set from the Yahoo!/O'Reilly Tech Buzz Game, a live system using the DPM, we show that prices offer informative forecasts of future event trends. At the agent level, we find that on average human traders outperform robot traders who randomly place trades. Those human traders who are seemingly more rational, buying when the implied market probability is low and selling when it is high, obtain higher profit on average than those who appear less rational. We examine other aspects of the game, including incentives to manipulate the market.

1 Introduction

Prediction markets, markets that trade securities whose future payout is tied with outcomes of some event of interest, are becoming a promising method for aggregating relevant information. Prediction markets, operated as continuous double auctions (CDAs), pari-mutuel markets, or automated market makers [1, 2], yield accurate forecasts in lab tests, field studies, and practice [3–9]. Both CDAs and pari-mutel markets have disadvantages. CDAs may suffer from illiquidity when there are few traders, because every buy order must be matched with a sell order. Pari-mutuel markets have infinite liquidity but do not incentivize traders to reveal their information early. Traders in a pari-mutuel market is never worse off by waiting until the last minute to place their bets.

Dynamic pari-mutuel markets (DPM), proposed by Pennock [10], is a hybrid between CDAs and pari-mutuel markets, and can also be thought as a form of automated market makers. DPMs in theory provide both infinite liquidity, as in pari-mutuel markets, and incentives for traders to reveal their information early, as in CDAs. However, we are not aware any empirical study of DPMs that evaluate how well they work in practice. In this paper, we empirically examine whether prices in DPMs have predictability for the event of interests and how traders behave in DPMs.

We obtain our data from the Yahoo!/O'Reilly Tech Buzz Game, a fantasy prediction market that uses DPM as its trading mechanism. Applying two methods to analyze two phases of the game, we provide evidence on DPM's predictability. Examining individual behavior in the market, we show that human traders obtain more profit than naive robot traders. Rational traders, as defined by a myopically optimal trading strategy, achieve higher profit than traders who are less rational. We also investigate traders' attempts to manipulate the market.

2 Dynamic Pari-mutuel Markets

In a dynamic pari-mutuel market [10] for an event with n exclusive and exhaustive outcomes, n securities are offered, each corresponding to an outcome. Traders wager on the outcomes by buying corresponding securities. After the true outcome is revealed, traders who wager on the true outcome split the total money pool at the end of the market as in a pari-mutuel market. Traders who wager on other outcomes lose their wagers. However, the price of a single share changes dynamically according to a price function such that it costs less to buy a share of a security earlier than later if many traders are buying the same security.

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 at any time of the market. Let $\boldsymbol{q} = \langle q_1, \ldots, q_n \rangle$ be the vector of shares outstanding for all securities. Then the total money wagered in the market is captured by the share-ratio cost function:

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

while the instantaneous price for security i is given by the share-ratio price function:

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

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

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

Prices in a DPM however do not directly correspond to probabilities of outcomes. The market probability of outcome i is given by

$$\pi_i(\boldsymbol{q}) = \frac{q_i^2}{\sum_{j=1}^{n} q_j^2} = \frac{p_i^2}{\sum_{j=1}^{n} p_j^2}. \tag{4}$$

This is the implied probability assuming that the market closes at the current state.

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

3 Tech Buzz Game

Yahoo! Research and O'Reilly Media jointly launched the Tech Buzz Game (http://buzz.research.yahoo.com), a fantasy prediction market for forecasting high-tech events and trends, on March 15, 2005. Tech Buzz Game uses DPM as its trading mechanism. In the first week, participants of the game invested \$88,480,577 fantasy dollars in the market [11].

Tech Buzz Game consists of multiple sub-markets, each of which includes several stocks for rival technologies. For example, the Video Game Consoles market has three stocks representing Xbox, Nintendo, and Playstation. Players have access to the current "search buzz" around each technology, measured by the number of users searching information on it at Yahoo! Search. A stock's *buzz score* is the number of buzz searches over the past seven days, as a percentage of all searches in the same market. Thus, if the searches for Nintendo make up 50% of all Yahoo! searches in the Video Game Consoles market, Nintendo's buzz score is 50. The objective of the game is to predict future buzz scores via buying or selling corresponding stocks.

A stock can bring revenue to its shareholders in two ways, dividends and cash-out settlements. Dividends are paid every Friday. Cash-out events liquidate shares for fantasy dollars at long-term intervals. Both paid dividends and final cash-out settlements are in proportions to buzz scores at the time of payment. During a dividend payment event, a certain amount of total dividend is allocated to a market. Within a market, the money is first divided among stocks according to the buzz scores. For example, if the total dividend for the Video Game Consoles market is \$10,000, and buzz scores for Xbox, Nintendo, and Playstation are 30, 50, and 20 respectively, then \$3,000 will be allocated to Xbox, \$5,000 to Nintendo, and \$2,000 to Playstation. The money allocated to each stock is then distributed to shareholders as dividends, with each share receives an equal portion. During a cash-out event, all money invested in every market is distributed to shareholders. Similar to dividend payments, within a market the money is first allocated to stocks according to buzz scores. Then, every share gets an equal portion of the money allocated to the stock.

Tech Buzz game has two phases that have slightly different rules. Phase One was the contest period, from March 15, 2005 to July 29, 2005. In this period, Tech Buzz Game had 47 markets. Each market had 2 to 12 stocks for competing technologies, resulting in a total of 305 stocks across all markets. Every Friday, \$10,000 fantasy dollars were paid as dividend to each market. Cash-out events happened at the end of the market on July 29, 2005. Each participant of the contest received an initial \$10,000 in fantasy dollars in his/her portfolio, with which trades could be made. A participant may borrow fantasy dollars if his/her portfolio net value drops below \$1,000. At the end of the contest, the top 3 players ranked by the largest change in portfolio net value won prizes. The contest period initially used the DPM *money-ratio* price function, which defines the ratio of any two stock prices in the same market as always equal to the ratio of money invested in the stocks [10]. However, this price function is not arbitrage-free. The arbitrage opportunities were exploited by market participants in the second week of the contest, causing prices of all stocks in some markets quickly drop toward zero. The contest was paused and reopened on April 1, 2005 with the *share-ratio* price function defined in (2), which is arbitrage-free.

The Tech Buzz Game entered Phase Two, the non-contest period, on August 22, 2005. It started with 44 markets with a total of 224 stocks[4]. The DPM share-ratio price function is used. Each market receives a dividend of 3% of its capitalization every Friday. So far, there have not been cash-out events in Phase Two except for a couple of markets that have been closed.

Note that the Tech Buzz Game is not designed in a way that securities in a market correspond to mutually exclusive outcomes as introduced in the previous section. With mutually exclusive outcomes, one security will get all the money wagered in the market in a cash-out event. In Tech Buzz Game, all stocks share the total money in the same market according to buzz scores. Thus, the share-ratio cost and price functions in Tech Buzz Game are the same as defined in (1) and (2). The payout of stock i during a cash-out event becomes

$$o_i(\boldsymbol{q}^f, \boldsymbol{b}^f) = \frac{C(\boldsymbol{q}^f) b_i^f}{100 q_i^f} = \frac{\kappa b_i^f \sqrt{\sum_j (q_j^f)^2}}{100 q_i^f}, \tag{5}$$

when the share-ratio price function is used. b_i^f is the buzz score for stock i at the time of cash-out. Similarly, dividend of stock i is

$$d_i(\boldsymbol{q}) = \frac{D b_i}{100 q_i}, \tag{6}$$

where D is the total dividend for the market and b_i is the current buzz score. Similar as (4), market prices in the game implicitly define what the market thinks

[4] We excluded an experimental market for hurricanes.

buzz scores should be, denoted as *implied buzz*

$$\hat{b}_i(\boldsymbol{q}) = \frac{100q_i^2}{\sum_{j=1}^{n} q_j^2} = \frac{100p_i^2}{\sum_{j=1}^{n} p_j^2}. \tag{7}$$

Implied buzz is the buzz score at which a trader is indifferent of buying or selling the stock if the market liquidates at the current states, i.e. $o_i(\boldsymbol{q}, \hat{\boldsymbol{b}}) = p_i(\boldsymbol{q})$.

4 Data Sets

We obtained transaction data for all 47 markets and 305 stocks in Tech Buzz Game Phase One for the period from April 1, 2005 to July 29, 2005, spanning 17 weeks. In this period, the share-ratio price function was used. We also acquired the corresponding daily buzz scores for the same time period. For Tech Buzz Game Phase Two, we obtained both transaction data and daily buzz scores for the period from August 22, 2005 to January 27, 2006, 22 weeks in total. The data have a record for every transaction that includes anonymous trader id, date and time of the transaction, the stock traded, number of shares bought or sold, cost of the transaction, and price of the stock before the transaction. With these data sets, we can calculate the prices of stocks, portfolio value of participants, dividends, and cash-out settlements at any time in the game.

5 Predictability of Tech Buzz Game

The first question we are interested in investigating is whether prices in the Tech Buzz Game predict the future buzz scores, a proxy for popularity of the underlying technologies. We evaluate the prediction accuracy for both phases of the Tech Buzz Game separately.

5.1 Predictability Analysis of Phase One

Tech Buzz Game Phase One, the contest period, had a cash-out event at the end of the contest on July 29, 2005. We first calculate the implied buzzes for all stocks right before the close of the market according to (7). If the market offers perfect predictions, the implied buzz of a stock should equal its buzz score in the cash-out event, since it is the only source of future payment. We get the absolute difference between the implied buzz and actual buzz scores for every stock, the distribution of which is shown in Table 1. 162 out of 305 stocks have an absolute error that is less than 10, indicating that implied buzz in many markets offer some insights into the actual buzz score.

In Phase One of the game, each market received $10,000 dividends every Friday. Both dividends and cash-out settlement are cash flows to a shareholder. Thus, we use a *discounted future cash flows* approach to estimate the value of a stock at some time in the market. We get the ratio of the price to the value (calculated using discounted future cash flows) of stocks at the end of week 1,

Table 1. Distribution of absolute errors in Tech Buzz Game Phase One on July 29, 2005, right before the market close.

$\lvert\hat{b}_i - b_i^f\rvert$	Number of Stocks
[0, 10)	162
[10, 20)	53
[20, 30)	32
[30, 40)	21
[40, 50)	16
[50, 100]	21

and examine how prices change in later weeks for stocks with different price-value ratios. The purpose of this examination is to check whether prices follow the underlying value of stocks in the game. There are 80 stocks whose average buzz scores over the contest period are less than 1, making their underlying value very small. We exclude these stocks in the analysis. For the remaining 159 stocks, we follow the procedure described below:

1. Calculate the value of each stock at the end of the first week of the game (on April 7, 2005),

$$v_i = \frac{o_i(\boldsymbol{q}^f, \boldsymbol{b}^f)}{(1+r)^{n-1}} + \sum_{j=2}^{n} \frac{d_i^j(\boldsymbol{q}^j)}{(1+r)^{j-1}},$$

   where r is the discount rate, n is the total number of weeks of the contest, which is 17 in our case, $d_i^j(\boldsymbol{q}^j)$ is the dividend of stock i at the end of week j. The discount rate r in theory should be the prevailing interest rate in the system. We tried some reasonable range of r, from 0 to 5%. It does not significantly affect our analysis. The results reported in this section are based on $r = 4\%$.
2. Calculate the ratio of price over value, p_i/v_i, for each stock at the end of the first week.
3. Allocate stocks to six buckets according to their values of p_i/v_i. The six buckets have stocks whose price-value ratios fall into $[0, 0.52]$, $(0.52, 0.75]$, $(0.75, 1.1]$, $(1.1, 1.72]$, $(1.72, 3.7]$, and $(3.7, +\infty)$ respectively. We choose the value ranges such that the six buckets are roughly balanced.
4. For stocks in each bucket, calculate their average weekly closing prices for week 1 to 17. Denote P_i^j as the average closing price of bucket i at week j.
5. Calculate the relative average price changes for all buckets and all weeks,

$$\Delta P_i^j = \frac{100(P_i^j - P_i^1)}{P_i^1}.$$

Figure 1 plots the relative average price changes, ΔP_i^j, for the six buckets. We can see that average price increases for buckets with a low price-value ratio and

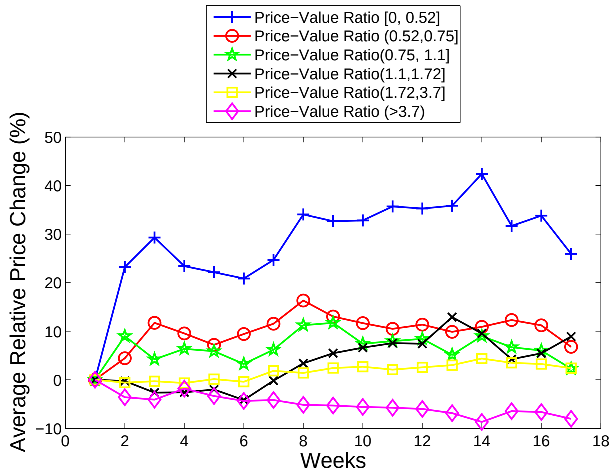

Fig. 1. Relative average price changes for stocks with different price-value ratios in Tech Buzz Game Phase One

decreases for buckets with a high price-value ratio, indicating that prices in the Tech Buzz Game gradually incorporate information and reflect the underlying values of stocks.

5.2 Predictability Analysis of Phase Two

Phase Two of the Tech Buzz Game does not have any cash-out event, making it hard to evaluate the underlying value of stocks. We take another approach to examine the predictability of the game. At any time in the game, the current status of the market provides implied buzzes as captured by (7). If the market has predictability, high implied buzz of a stock means that the market predicts the future buzz score of the stock is high, even if its current buzz score is low. Thus, for stocks that have a high ratio of implied buzz over actual buzz, their actual buzz should increase in the future if the market prediction is correct. We exclude 49 stocks whose average buzz scores are below 1 from our analysis, because very low buzz scores can cause the ratio of implied buzz over actual buzz extremely high. Because all markets mandatorily started with equal implied buzz for all stocks in the same market, we further exclude data from the first 4 weeks, which should give traders enough time to change market probabilities to reflect their information. For the remaining 175 stocks and 18 weeks, we carry out the analysis as follow:

1. Calculate the implied buzz for every stock at the market close of every Friday according to (7).
2. Let $\hat{b}_i^j$ and b_i^j represent the implied buzz and actual buzz scores of stock i for the Friday of week j respectively. Calculate the implied buzz-actual buzz ratio, $\hat{b}_i^j / b_i^j$.

3. For every week j from 1 to 18, allocate stocks into six buckets according to their values of $\hat{b}_i^j / b_i^j$. The six buckets have stocks whose implied buzz-actual buzz ratios at week j fall into $[0, 0.5]$, $(0.5, 0.7]$, $(0.7, 1.1]$, $(1.1 - 1.8]$, $(1.8, 4.1]$ and $(4.1, +\infty)$ respectively. Since we allocate stocks to buckets every week, buckets may not have the same stocks from week to week.
4. For every week $j = 1...18$ and for every bucket $k = 1...6$, for stocks allocated to bucket k at week j calculate the average buzz scores at week $j + d$ for $0 \leq d \leq (18 - j)$. Denote $B_k^j(d)$ as the average buzz scores at week $j + d$ for those stocks that are allocated to bucket k at week j.
5. For every week $j = 1...17$ and for every bucket $k = 1...6$, calculate $\Delta_k^j(d) = 100(B_k^j(d) - B_k^j(d-1))/B_k^j(d-1)$ for all $1 \leq d \leq (18 - j)$. $\Delta_k^j(d)$ represents the one-week percentage change of actual buzz scores for those stocks that are allocated to bucket k at week j, measured d weeks ahead of the bucket formation.
6. For every bucket $k = 1...6$ and every $d = 1...17$, calculate the average one-week percentage change of actual buzz scores measured d weeks ahead of the bucket formation, $\Delta_k(d) = \frac{\sum_{j=1}^{18-d} \Delta_k^j(d)}{18-d}$. Note that the summation contains less terms as d increase because we only have 18 weeks data.
7. Calculate the cumulative percentage change of average buzz scores d weeks ahead of the bucket formation for every bucket $k = 1...6$ and for every $d = 1...17$, $S_k(d) = \sum_{i=1}^{d} \Delta_k(d)$. Set $S_k(0) = 0$.

The cumulative change of average buzz scores, $S_k(d)$, represents the average rate of change of actual buzz scores for bucket k from the time the bucket is formed to d weeks ahead. It captures the general trend of buzz score changes for the bucket.

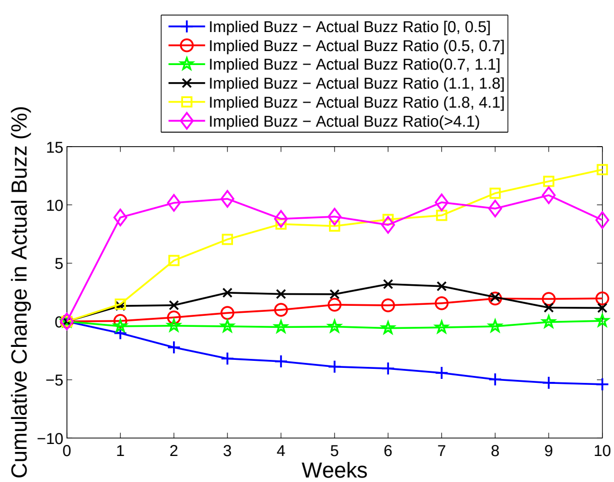

Fig. 2. Cumulative average buzz score changes for stocks with different implied buzz - actual buzz ratios in Tech Buzz Game Phase Two

Figure 2 plots $S_k(d)$ for $d = 0...10$ and for all six buckets[5]. It is clear that for buckets with a low implied buzz-actual buzz ratio, cumulative change of average buzz scores decreases over time. For buckets with a high implied buzz-actual buzz ratio, it increases over time. This shows that implied buzzes to some degree predict future buzz scores.

We also applied the discounted future cash flow approach to the Phase Two data set. Since in Phase Two, each market gets 3% of its capitalization as dividends each week, we use 3% as the discount rate. Because there is no cash-out event, we make the simplified assumption that future weekly dividend equals the average of the 22 dividend payments that we observed for any stock and stocks have infinite life span. Other aspects of the analysis is the same as those for Phase One. Figure 3 plots our results, showing that prices tend to follow underlying values of stocks in Phase Two of the game.

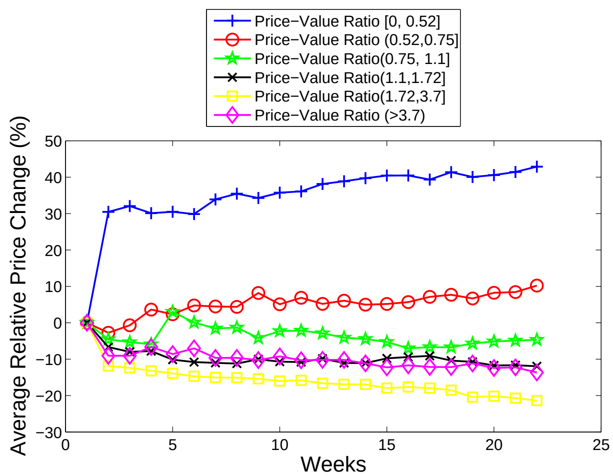

Fig. 3. Relative average price changes for stocks with different pric-value rations in Tech Buzz Game Phase Two

6 Trader Behavior in Tech Buzz Game

Dynamic pari-mutuel markets differ from other popular market mechanisms in many aspects. Prices change continuously in DPMs while in CDAs prices are typically fixed for some shares of transactions. Moreover, prices in DPMs do not directly correspond to probabilities. Instead, implied probabilities can be calculated based on them. These make it less obvious to identify profitable transactions in DPMs. Thus, we are interested in investigating how traders behave

[5] We only provide the figure for the first 10 weeks because $\Delta_k(d)$ averages over $j = 1$ to $j = 18 - d$ and hence as d increases it averages over an increasingly smaller range.

in DPMs. We use the data from Tech Buzz Game Phase One to peek into the behavior of individual traders.

6.1 Human Traders vs. Robot Traders

Tech Buzz Game Phase One had 100 robot traders played in the market. These robot traders randomly buy or sell some randomly picked stock mainly for the purpose of increasing activities in the market. However, robot traders can serve as a very rudimentary benchmark to examine human traders. The very first thing we want to verify is that human traders in Tech Buzz Game are not completely confused. They should at least outperform naive robot traders in terms of making profit.

We computed the net profit for every trader for the period of April 1, 2005 to July 29, 2005 including all dividends and the cash-out settlements. We calculated the average profit per agent for human traders and robot traders separately. The results are listed in Table 2. Human traders make money on average, while robot traders lose money on average, which is not surprising. It seems that human traders were able to exploit the mistakes made by robot traders.

Table 2. Human vs. robot traders in Tech Buzz Game Phase One

	Human	Robots
Number of Traders	4,819	100
Average number of transactions per agent	22	8,782
Average profit per agent	$1,853	-$11,615

6.2 Gaming the Market

Tech Buzz Game is a play-money market. In the contest phase, only the top 3 players ranked by largest portfolio net values can get prizes. Since every account is seeded with $10,000 fantasy dollars, the game gives participants incentives to open multiple accounts and transfer money among them. Despite that the game rule explicitly specifies that traders are not allowed to open multiple accounts, we still find that there were a small number of traders who attempted to game the system.

Suppose a trader opens two accounts, A and B. He can transfer money from account B to account A by conducting a sequence of transactions in a market: Account A buys q_A shares of stock i; account B buys q_B shares of stock i; account A sells q_A shares of stock i; and account B sells q_B shares of stock i. After the 4 transactions, both accounts do not hold any stock i, but account A now has more fantasy dollars than before. This is because that account B's purchase of stock i drives the price of stock i up and when account A sells its holdings of stock i it sells at a higher price than his purchasing price.

Thus, when a trader buy and then sell the same number of shares of the same stock within a short period of time. He may be transferring money between accounts. We count the number of such behaviors for each user. The threshold for determining suspicious behavior is set as 5 minutes. We categorize a trader as dishonest when he behaves in this manner for more than 5 times. With this criteria, 175 traders are identified as dishonest. Table 3 shows that these dishonest traders have very high net profit compared with other traders. We note that some of the money-supplying accounts that commit the suspicious behavior less than five times may be left in the Other Traders category, contributing to its negative average profit.

Table 3. Honest vs. dishonest human traders in Tech Buzz Game Phase One

	Number	Average Profit
Dishonest traders	175	$68,567
Other traders	4,644	-$661

6.3 Rational Behavior

There is no dominant strategy when playing in DPMs. However, a myopically rational trader may want to buy when the implied buzz is lower than his expected buzz score and sell when the implied buzz is higher than his expected buzz score. We don't know a trader's expected buzz score, but if we make the simplified assumption that a trader's expected buzz score is affected by current and past buzz scores and buzz score changes in general are random walks, then a myopically rational trader often will buy when the implied buzz is lower than the current buzz and sell vice versa.

We examine whether trader's behavior in the Tech Buzz Game is consistent with what is described above and whether traders who follow such trading heuristics more often do indeed make more profit than those who do not. To answer these, for each transaction, we label it as rational if it is a purchase and the implied buzz is lower than the buzz score or if it is a sale and the implied buzz is higher than the buzz score at the time of the transaction. For each trader, we calculate the percentage of transactions that are labeled as rational. If the percentage is greater than or equal to 2/3 for a trader, he is considered as a rational trader.

Table 4 summarizes our results. Only 867 out of 4644 honest human traders are considered rational in the market. On average these rational traders make more profit in the market than traders who are less rational. How traders make their trading decisions in DPMs is still unclear. However, it seems that this myopically optimal trading strategy gives its users an edge in the market.

Table 4. Rational vs irrational traders in Tech Buzz Game Phase One

	Number Of Traders	Average Profit
Rational traders	867	$7,864
Irrational traders	3777	-$2,617

7 Conclusion

With data from the Yahoo!/O'Reilly Tech Buzz Game, a live system using the dynamic pari-mutuel market (DPM) mechanism, we empirically examine how well DPM works in practice. Specifically, we investigate whether prices in DPMs predict events of interest and how traders behave in DPMs.

We use two methods to evaluate the predictability of DPMs. The first method examines whether prices of stocks follow their underlying values. The second method explores whether buzz scores increase (decrease) over time in the future when the market implied buzz scores are higher (lower) than the current actual buzz scores. Both methods give positive evidence that DPMs can provide valuable information for event trends.

Looking at the agent level, we find that human traders obtain more profits than robot traders who randomly place trades in DPMs. Those human traders who follow a myopically optimal trading strategy achieve higher profits than those who don't. Moreover, we look into traders' attempts to game the market. We are able to identify traders who create multiple accounts and transfer money among them. These traders achieve much higher average profit than other traders in the game. The manipulation behavior of traders however is largely inherent to play-money markets where prizes are awarded only to top players.

References

1. Hanson, R.D.: Combinatorial information market design. Information Systems Frontiers **5**(1) (2003) 107–119
2. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. Journal of Prediction Markets **1**(1) (2007) 1–15
3. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Anatomy of an experimental political stock market. American Economic Review **82**(5) (1992) 1142–1161
4. 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
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. 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

7. Thaler, R.H., Ziemba, W.T.: Anomalies: Parimutuel betting markets: Racetracks and lotteries. Journal of Economic Perspectives **2**(2) (1988) 161–174
8. 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)
9. 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)
10. 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) 170–179
11. 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