An Empirical Game-Theoretic Analysis of Price Discovery in Prediction Markets

Elaine Wah
Computer Science & Engineering
University of Michigan
ewah@umich.edu

Sébastien Lahaie
Microsoft Research
New York, NY
slahaie@microsoft.com

David M. Pennock
Microsoft Research
New York, NY
dpennock@microsoft.com

Abstract

In this paper, we employ simulation-based methods to study the role of a market maker in improving price discovery in a prediction market. In our model, traders receive a lagged signal of a ground truth, which is based on real price data from prediction markets on NBA games in the 2014–2015 season. We employ empirical game-theoretic analysis to identify equilibria under different settings of market maker liquidity and spread. We study two settings: one in which traders only enter the market once, and one in which traders have the option to reenter to trade later. We evaluate welfare and the profits accrued by traders, and we characterize the conditions under which the market maker promotes price discovery in both settings.

1 Introduction

Prediction markets offer a marketplace in which traders can buy or sell securities whose payoffs are based on the realizations of unknown future events [Wolfers and Zitzewitz, 2004]. Whereas traditional financial markets exist to facilitate investment, speculation, and hedging, the express purpose of a prediction market is to forecast future events. Prices in a prediction market have direct interpretations as event probabilities, and therefore *price discovery* in this context becomes synonymous with information aggregation.

The most common prediction market structure in the literature relies on a centralized market maker (MM) to provide *liquidity*, or the prevalence of opportunities to trade at current market prices. Central among these designs is the Logarithmic Market Scoring Rule (LMSR) devised by Hanson [2007], which uses a cost function to assign charges to trades and prices to securities. The cost function is parametrized by a liquidity parameter to control the sensitivity of prices to trade volume, and can be adapted to implement a spread between buy and sell prices. These parameters have an important effect on prediction performance and MM loss but can be difficult to set in practice, especially given the challenge of anticipating trader behavior [Brahma *et al.*, 2012; Othman and Sandholm, 2010; Othman *et al.*, 2013].

In this paper, we combine agent-based simulation, equilibrium computation, and real-world data to characterize the conditions that promote price discovery in a prediction market. We introduce the use of *empirical game-theoretic analysis* (EGTA) in the context of prediction markets to uncover trader strategies in equilibrium, which allows us to quantify the impact of different market parametrizations and environments. EGTA has been used to study multi-agent systems in financial markets [Wah and Wellman, 2015], and related approaches have examined continuous double auctions [Phelps *et al.*, 2004; Walsh *et al.*, 2002]. Our study is the first to employ EGTA to study prediction markets, and offers a proof-of-concept of the technique in this context.

In our model, there is a single security traded in a market with an LMSR market maker. Our prediction market is populated by informed and less informed traders that each receive a lagged signal of a ground truth, as well as noise traders who create profit opportunities. The ground truth we use is based on prediction market data from Betfair on NBA games in the 2014–2015 season. Strategic traders formulate a belief on the value of the security by forming a convex combination of the MM's quoted price and their own information. We focus on trader behavior in equilibrium, where market participants are strategically responding to each other. We compare two different modes of market entry: one where traders enter only once, and one where they are allowed multiple entries.

Our methodology is as follows. We first verify that our ground-truth data is well-calibrated—that is, for any probability level, the proportion of price series at that level matches the proportion of realized outcomes. This confirms that the ground truth is a consistent signal of payout. To obtain payoff data for EGTA equilibrium computations, we simulate our prediction market model using a discrete-event simulation system that samples the ground truth with replacement. Once we identify equilibria, we evaluate market performance in terms of MM loss and price discovery on a hold-out sample of the ground truth, after confirming that in- and out-of-sample metrics such as player profits closely agree.

As expected, our results show that price discovery is sensitive to the liquidity and spread settings, as noted in practice. Our analysis also reveals that the mode of trader entry (single or multiple) impacts not only sensitivity, but also the overall relationships between liquidity, spread, price discovery, and MM loss. Informed traders prefer to trade as late as possible to fully exploit their information advantage—which is more valuable near the end of the trading horizon—but do not have this option in the single-entry setting. Price discovery therefore benefits from increased liquidity only in the multi-entry setting; otherwise, low liquidity improves price discovery by rapidly incorporating the information from trades.

2 Related Work

Much prior work examines different prediction market making mechanisms [Abernethy *et al.*, 2011; Hanson, 2003], including extensions to ordinary prediction markets such as incorporating limit orders [Chakraborty *et al.*, 2015; Heidari *et al.*, 2015], employing pari-mutuel techniques for payoffs [Pennock, 2004] and predicting combinations of events [Dudík *et al.*, 2012; 2013].

Price discovery is of primary interest in prediction markets, but the true probability of an event is unknown, rendering evaluation of real-world information aggregation over time near impossible. To address this, several studies employ agent-based modeling (ABM) and simulation to study price discovery, most using a random process as a basis for the ground truth probability [Brahma *et al.*, 2012; Chakraborty *et al.*, 2015; Slamka *et al.*, 2013]. Jumadinova and Dasgupta [2011] use prediction market data from the Iowa Electronic Market to run simulations characterizing the role of information on trader behavior. Slamka et al. [2013] compare price discovery under four automated MMs, including LMSR and the dynamic pari-mutuel market [Pennock, 2004].

Other prior work focuses on deriving equilibria for heterogeneously informed traders, but many of these studies [Hanson and Oprea, 2009] have been restricted to simple models extending the classic model by Kyle [1985]. Pennock and Sami [2007] describe how a market employing LMSR converges to the rational expectations equilibrium price. Chen et al. [2007] study incentives to bluff strategically when trader information is conditionally dependent on the ground truth. Dimitrov and Sami [2008] construct a model with partially informed traders observing independent signals, showing that the myopically optimal strategy profile is not a weak perfect Bayesian equilibrium for an LMSR MM. Ostrovsky [2012] demonstrates that information about separable securities in certain markets with partially informed strategic traders is always aggregated in equilibrium.

To our knowledge, prior work does not explicitly compare equilibrium outcomes under different entry schemes. The theoretical literature typically relies on simple finite-period models [Dimitrov and Sami, 2008; Gao *et al.*, 2013]; previous ABM studies (which do not compare equilibrium outcomes) generally implement single-entry schemes, with agents arriving randomly [Chakraborty *et al.*, 2015], sequentially [Slamka *et al.*, 2013], or on every time step [Jumadinova and Dasgupta, 2011].

3 Market Maker

In our model, there is a single security traded in a prediction market that operates as a continuous trading mechanism. The security pays off \$1 if event E occurs and \$0 if it does not. There are two possible outcomes: one where event E occurs and one where it does not (i.e., event $\overline{E}$). There is a single market maker in our prediction market who acts as a counterparty for all transactions. The MM uses the standard *logarithmic market scoring rule* (LMSR) to charge for trades and to price securities [Hanson, 2007]. The MM has two available parameters: spread δ and liquidity parameter b. Both spread and the amount of liquidity are set a priori and are fixed for the duration of trading. Smaller b reflects lower liquidity, as prices will increase *faster* as traders buy shares from the market maker. Larger b reflects higher liquidity, as prices will increase *slower* as traders buy shares from the MM.

Cost Function We implement our model under a no-selling scheme in order to prevent traders from arbitraging the MM by selling a security paying out a guaranteed \$1 for more than \$1 [Othman *et al.*, 2013]. The two-element vector $\mathbf{q}_t = (q_{E,t}, q_{\overline{E},t})$ represents the total number of shares traders have bet at time t on event E and on $\overline{E}$, respectively. From the market maker's perspective, $q_{E,t}$ is the number of shares the MM has sold, whereas $q_{\overline{E},t}$ is the number of shares of the security the MM has bought. All elements of $\mathbf{q}_t$ are positive.

In the single-security setting where the LMSR market maker can set a nonnegative spread δ, the cost function is

$$C(\mathbf{q}_t) = (1+\delta) \cdot b \cdot \ln\left(e^{q_{E,t}/b} + e^{q_{\overline{E},t}/b}\right).$$

The cost charged to a trader to buy $x > 0$ shares is

$$\rho(x) = C\left(q_{E,t-1} + x, q_{\overline{E},t-1}\right) - C\left(q_{E,t-1}, q_{\overline{E},t-1}\right).$$

Cost is negative for sellers, as the trader will be compensated for its shares. The cost of selling x shares is

$$\rho(x) = C\left(q_{E,t-1}, q_{\overline{E},t-1} + x\right) - C\left(q_{E,t-1}, q_{\overline{E},t-1}\right).$$

Due to the spread, a trader cannot buy some quantity and sell it back for zero cost: it pays $(1+\delta)$ per share to buy from and sell back to the MM, receiving a guaranteed payout of \$1.

Quoted Prices At time t, BID_t corresponds to the price at which the MM offers to buy the security on event E, and ASK_t is the quoted price at which the MM offers to sell the security. The quote on complementary event $\overline{E}$ is comprised of $\overline{BID}_t$ and $\overline{ASK}_t$, with $BID_t < ASK_t$ and $\overline{BID}_t < \overline{ASK}_t$. The MM sets a spread $\delta \geq 0$, which is the difference between the BID and ASK. The midpoint of the spread can be interpreted as the market probability of event E. The current instantaneous price at which a trader arriving to the market can buy a share (i.e., bet on event E) is the cost of buying an infinitesimal amount from the MM:

$$ASK_t = p(\mathbf{q}_t) = \frac{(1+\delta)\exp(q_{E,t}/b)}{\exp(q_{E,t}/b) + \exp(q_{\overline{E},t}/b)}.$$

Similarly, $\overline{ASK}$ is the price at which an entering trader can bet on event $\overline{E}$ (or equivalently, on event E *not* occurring):

$$\overline{ASK}_t = \frac{(1+\delta)\exp(q_{\overline{E},t}/b)}{\exp(q_{E,t}/b) + \exp(q_{\overline{E},t}/b)}.$$

The BID price is defined as $BID = ASK - \delta$, and $\overline{BID}$ is defined similarly, with $\overline{BID} = \overline{ASK} - \delta$. The ASK prices for the two outcomes (E and $\overline{E}$) are given by vector $\mathbf{p}_{ASK} = \left[ASK_t, \overline{ASK}_t\right]$, with $ASK_t + \overline{ASK}_t = 1 + \delta$. The BID prices are given by $\mathbf{p}_{BID} = \left[BID_t, \overline{BID}_t\right]$, with $BID_t + \overline{BID}_t = 1 - \delta$. Because these are complementary events, offering to buy (sell) from incoming traders on event E at price ASK_t (BID_t) is equivalent to offering to sell (buy) on event $\overline{E}$ at $\overline{BID}_t$ ($\overline{ASK}_t$). Therefore, the following conditions hold for all times t: $\overline{ASK}_t = 1 - BID_t$ and $ASK_t = 1 - \overline{BID}_t$, and the no-arbitrage condition holds.

The MM liquidates its inventory at time T, which denotes the end of the trading horizon. Its payoff is the revenue from trading, minus the value of liquidation:

$$C(\mathbf{q}_T) - C(\mathbf{q}_0) - \begin{cases} q_{E,T} & \text{if the event occurs, or} \\ q_{\overline{E},T} & \text{if the event does not occur.} \end{cases}$$

Here $\mathbf{q}_0 = (0, 0)$ is the initial quantity, and $C(\mathbf{q}_T) - C(\mathbf{q}_0)$ is the total amount paid to the MM.

4 Traders

We include strategic traders, whose payoffs are used in equilibrium computation, and noise traders in our model. Trader profit is the cash flow from trading.

4.1 Strategic Traders

Each strategic trader has an individual valuation for the security that is based on its private belief, which is a lagged signal of the ground truth. We include two types of traders in our model: *informed* traders receive information that is more recent, on average, than *less informed* traders.

Ground Truth The ground truth v_t is the underlying true value of the security. It represents the probability at time t that the underlying event E will ultimately occur at time T. The probability of event E not happening is given by v'_t, and the true values of the two complementary events always sum to 1: $v'_t = 1 - v_t$. At the end of the trading horizon, one of the two states of the world will be realized, so $v_T \in \{0, 1\}$.

Private Values Each trader j has an individual private belief $w_{j,t}$ regarding the probability of event E, with $w_{j,t} \in (0, 1)$. An agent's belief $w_{j,t}$ is a lagged signal of the true value v_t at time t, that is, $w_{j,t} = v_{t-\Delta_t}$ with lag drawn from an exponential distribution, or $\Delta_t \sim \text{Exp}(\lambda)$. Less informed traders have a lower lag rate (and higher mean lag time) than informed traders. The lag rate of informed and less informed traders is λ_{INF} and λ_{LESS}, respectively, with $\lambda_{\text{INF}} > \lambda_{\text{LESS}}$.

Utility Function The traders in our model are risk neutral. The utility π_j of a trader j who pays (receives) ρ to buy (sell) q_j shares is the surplus when liquidating at value v_T:

$$\pi_j = \begin{cases} v_T q_j - \rho(q_j) & \text{for buy transactions, or} \\ \rho(q_j) - v_T q_j & \text{for sell transactions.} \end{cases}$$

At time T, a trader who bought (sold) the security will receive \$1 (\$0) per share if E occurs and \$0 (\$1) if it does not.

We impose a budget constraint c on the traders: the total cost of any buy order and the total redemption value (or liability) of any sell order must not exceed c. More precisely, the following conditions must hold for quantity $q_j > 0$:

$$c \geq \begin{cases} C\left(q_{E,t} + q_j, q_{\overline{E},t}\right) - C\left(q_{E,t}, q_{\overline{E},t}\right) & \text{if buying} \\ C\left(q_{E,t}, q_{\overline{E},t} + q_j\right) - C\left(q_{E,t}, q_{\overline{E},t}\right) & \text{if selling.} \end{cases}$$

The maximum quantity for a buyer is therefore

$$q_j^* = b \cdot \ln\left(\exp\left(\frac{c + C(\mathbf{q}_t)}{b \cdot (1+\delta)}\right) - \exp\left(\frac{q_{\overline{E},t}}{b}\right)\right) - q_{E,t}.$$

The quantity for a seller is analogous.

Single Entry vs. Multiple Entry There are two modes of market entry: single and multiple. Regardless of entry mode, each agent only participates in one trade. In the single-entry mode, traders enter the prediction market once, with time of entry $t_e \sim \mathcal{U}[0, T-1]$, where T is the length of the trading horizon in time steps. In the multi-entry mode, traders arrive to the prediction market over time, with time of arrival (and subsequent reentries) determined by an exponential distribution with rate λ_r. If a potential trade may not be profitable (i.e., $BID_t < w_j < ASK_t$), the agent does not submit an order and waits until its next reentry to reevaluate market conditions. It continues to reenter until it has successfully traded.

Strategies Traders in our model play a parametrization that we call the *weighted belief update* strategy, similar to that described by Jumadinova and Dasgupta [2011]. Trader j computes its belief as a convex combination of the current market price (the midpoint of the BID-ASK spread) and its individual private signal of the ground truth:

$$x_{j,t} = (1-\theta)\left(\frac{ASK_t - BID_t}{2}\right) + \theta\, w_{j,t}.$$

Traders optimize the weighting parameter θ (which is selected a priori and fixed for the entire trading duration) to maximize their payoffs. If $x_{j,t} \geq ASK_t$, the trader submits a buy order; if $x_{j,t} \leq BID_t$, the trader submits a sell order. If $BID_t < x_{j,t} < ASK_t$, then the trader does not submit an order (but may have the opportunity to reenter later to trade, depending on the type of market entry). To move the ASK to its belief $x_{j,t}$, the optimal quantity $q_j^\dagger > 0$ for a buyer j is:

$$q_j^\dagger = b \cdot \ln\left(\frac{x_{j,t}}{1 + \delta - x_{j,t}}\right) + q_{\overline{E},t} - q_{E,t}.$$

The optimal quantity for a seller is analogous. The trader either exhausts its budget or the price exceeds its private belief, so the actual quantity submitted is $q = \min\{q_j^*, q_j^\dagger\}$.

4.2 Noise Traders

Noise traders determine a priori whether or not to buy or sell, each with equal probability. They arrive into the market only once, with entry times uniformly distributed across the trading horizon. They buy or sell as much as permitted by their budget, so a buyer j computes quantity q_j^* as above.

5 Empirical Game-Theoretic Analysis

Empirical game-theoretic analysis (EGTA) is a methodology for performing strategy selection by comparing the payoffs of different combinations of trader-strategy assignments. EGTA allows one to compare trader behavior in equilibrium under different market conditions. We give a high-level overview of EGTA here and refer to Wellman [2006] for complete details.

EGTA is an iterative process that involves discretizing the strategy space and analyzing the empirical game model induced by payoff data from simulations. Players are divided into *roles* and strategies are symmetric within roles. We generate equilibrium candidates by analyzing *complete subgames*, which are sets of strategies (one per role) for which we have collected data for all profile combinations. This gives us *role-symmetric Nash equilibria* (RSNE) of the subgames, which are equilibrium candidates for the full game. A subgame RSNE is confirmed as a full-game equilibrium if there are no beneficial deviations in the full strategy set. If deviations are found, they are added to the candidate's support, creating a new subgame. The process repeats until quiescence.

There are limitations that must be taken into consideration given this methodology. As game size grows exponentially with the number of players, we rely on *deviation-preserving reduction* (DPR) to construct a reduced-game approximation of the full game. DPR preserves payoffs from single-player deviations and has been shown to produce good approximations in other problems [Wiedenbeck and Wellman, 2012], but equilibrium estimates from DPR are not guaranteed. In addition, as we are unable to evaluate all profiles given the size of the game, we cannot guarantee that we have found all equilibria (even in the reduced game), although our process seeks to evaluate all promising equilibrium candidates.

There are two roles in our prediction market game, representing the informed and less informed traders, respectively. For our simulations we used 42 players, with 21 in each role. These specific numbers were chosen because they conveniently reduce via DPR to a $(3, 3)$-reduced game. We simulate our prediction market model using a discrete-event simulation system [Wah and Wellman, 2013] and we manage our experiments via the EGTAOnline infrastructure [Cassell and Wellman, 2013]. We mitigate sampling error by collecting payoff data over many simulation runs: a minimum of 10,000 samples per strategy profile evaluated, with 20,000 for the majority of profiles and at least 19,587 samples per profile on average. We use the RSNE computed to characterize the conditions under which the MM promotes price discovery.

5.1 Market Environment

We utilize real-world prediction market data from Betfair, an Internet betting exchange based in the U.K., as a basis for the ground truth, or the underlying true value of the security. We include data on NBA games from the 2014–2015 season, including the post-season. The price series for a given game is constructed by querying Betfair prices approximately every 2 to 3 minutes for the duration of the game. Excluding the price series with incomplete data, we have price information from 1126 NBA games. Each basketball game is associated with two price series: one reflecting the price of the security on the event the home team wins, and one for the security on the event the home team loses. Each simulation run entails drawing a sample ground truth (with replacement) from a pool of ground-truth price series.

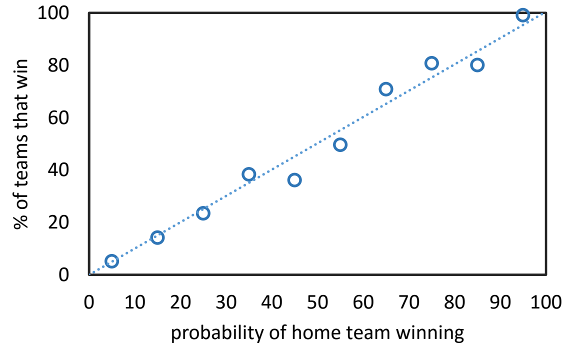

Figure 1: Calibration of the Betfair NBA game price series.

We assess calibration of our data by binning forecasts to the nearest 10% mark according to the price at each time step (Figure 1). For each bin, we determine the percentage of time series for which the team in question wins. Since each bin represents a forecast, the data is perfectly calibrated when the average fraction of teams that win for every bin is equal to the average forecast probability of winning. We observe that the Betfair NBA data is well calibrated over time, which is typical of real betting markets. This confirms that traders in our model have a reason to take into account their signals when trading, since the ground truth correctly reflects the probability of a payoff at different time steps.

There are 14 noise traders in our prediction market, each with a budget of \$0.10, which permits a maximum change in price of \$0.05 when the market midquote price is \$0.50 and the liquidity parameter is 1. Strategic traders (informed and less informed) have a budget of \$0.50. There are 21 informed traders and 21 less informed traders. In the multi-entry setting, strategic traders arrive to the market according an exponential distribution with rate $\lambda_r = 0.125$. We set the trading horizon T to 80 time steps to incorporate price data from the full duration of the game. The informed-trader lag rate is $\lambda_{\text{INF}} = 0.5$; the less-informed lag rate is $\lambda_{\text{LESS}} = 0.1$.

5.2 Experiments

We characterize performance of our prediction market primarily via out-of-sample price discovery, which we measure using the root-mean-square deviation (RMSD) between the midquote price from the MM and the ground truth at every time step [Brahma *et al.*, 2012; Chakraborty *et al.*, 2015]. Lower RMSD indicates better price discovery. We also examine how price discovery changes relative to MM loss, total welfare, and trader profit. We use four liquidity settings $b \in \{1, 2, 5, 10\}$ and four spread settings $\delta \in \{0, 0.01, 0.05, 0.1\}$ for the MM, giving a total of 16 games for each market entry mode, or 32 games total.

We separate the Betfair data into an in-sample dataset of 845 games, or approximately 75% of the full dataset, and an out-of-sample dataset of 281 games. We use data from the in-sample set for the empirical game-theoretic analysis. We analyze price discovery and welfare for the out-of-sample

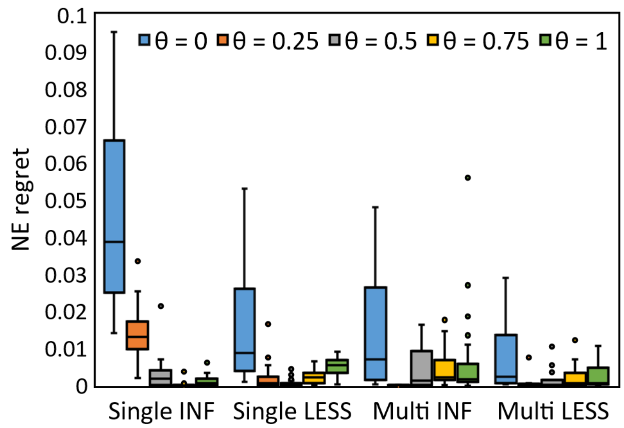

Figure 2: Regret of strategies by role, aggregated across all RSNE found, in the two market entry settings. The five boxes in each group shows from left to right the regret for θ increasing from 0 to 1. Plots show medians (solid horizontal line), first and third quartiles (box outline), 90th percentile values (whiskers), and outliers (dots).

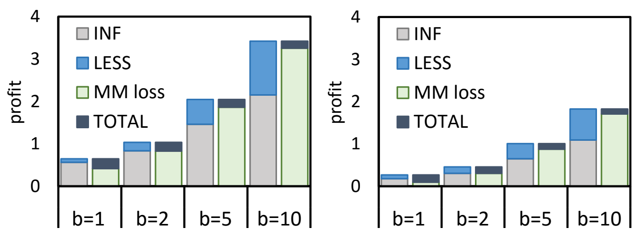

(a) Single entry (b) Multiple entry

Figure 3: Trader profits at spread $\delta = 0.01$ for varying levels of liquidity. There is a pair of stacked bars for each liquidity setting: the stacked bar on the left is comprised of the profits of informed traders (bottom) and less informed traders (top); the stacked bar on the right is comprised of MM loss (bottom) and total welfare (top). Welfare is equivalent to noise trader losses. The bars are always of equal height, reflecting the fact that total welfare equals trader profits net of MM loss. For games with multiple equilibria, we plot the profits from the maximum-welfare equilibrium.

dataset via 10,000 simulation runs over the mixture probabilities in the role-symmetric equilibria found. Mixed-strategy RSNE are approximated by profiles with trader population proportions corresponding to the strategy probabilities. Our out-of-sample profit results are very close to those achieved on the in-sample dataset. The mean absolute error (MAE) between the in- and out-of-sample informed trader profits is 0.068 and 0.038 for the single and multiple entry cases, respectively (with profit values ranging from 0.158 to 2.247); the less-informed MAE is 0.063 for single-entry and 0.060 for multiple entries (with values between 0.026 and 0.944). The MAE in RMSD is 0.002 for both types of market entry; single-entry RMSD ranges from 0.099 to 0.158 and multi-entry RMSD ranges from 0.123 to 0.184. As such, we focus our evaluation on out-of-sample profit and RMSD.

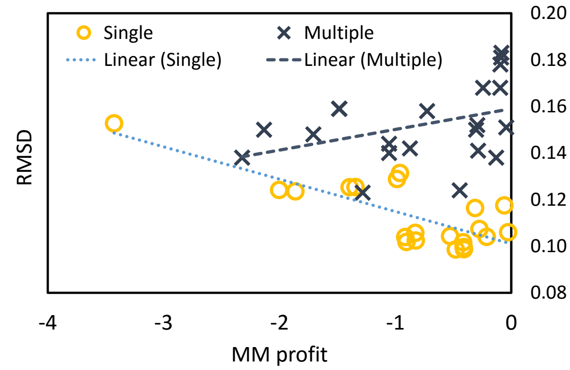

Figure 4: RMSD plotted against MM profit. The O's represent single-entry equilibria; the X's are multi-entry equilibria. The light blue dotted line is the trend line ($R^2 = 0.603$) for single-entry games; the dark blue dashed line is the trend line ($R^2 = 0.147$) for multi-entry games.

Equilibrium Analysis We find at least one and up to three RSNE[1] within each game. In cases with multiple RSNE, differences in trader and MM performance, as well as price discovery, are generally not significant. Figure 2 provides a visual summary of the *regret* for each strategy, which is defined as a player's loss in utility if deviating from a Nash equilibrium to the given strategy. We see that informed traders in the single-entry case generally minimize regret with $\theta = 0.75$, whereas those in the multi-entry case weigh their own information less heavily, with θ in equilibrium of 0.25. These values correspond to the most common pure-strategy RSNE. In both the single- and the multi-entry cases, the regret is higher for informed than less informed traders, indicating that the informed traders stand to lose more by deviating from equilibrium. The equilibrium strategies of less informed traders are similar across the two market entry scenarios.

The computed equilibria reveal substantially different behavior depending on the possibility of reentry. When traders only enter once, informed traders almost universally weigh their own lagged signal more than or equivalent to the less informed traders. This makes sense, as informed traders receive a more accurate signal of the ground truth than less informed traders. When traders can reenter, however, informed traders do not weigh their own information more heavily in equilibrium. This is because informed traders do not have a significant information advantage over the less informed traders early on: the ground truth price series tend to be most volatile near the end of the trading horizon, and in many cases, the price is fairly stable at the beginning. Consequently, informed traders are incentivized to trade as late as possible, and they avoid trading on their first entry into the market by weighing the market quote more heavily than their own information.

Trader Profit, MM Loss, and Welfare Figure 3 shows trader profits, MM loss, and overall welfare at spread 0.01; the qualitative trends here as liquidity increases are the same

[1]Full details of all RSNE found are available in an online appendix (http://hdl.handle.net/2027.42/117580).

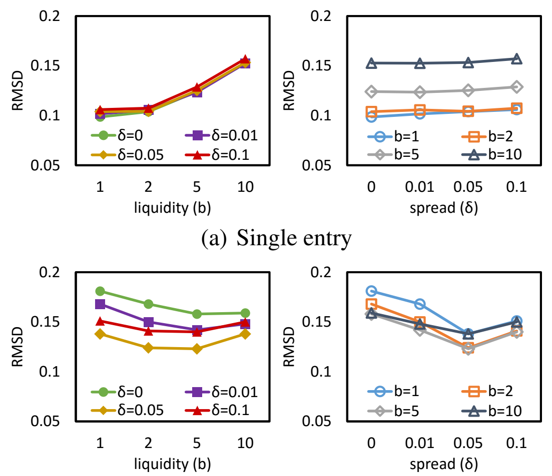

(a) Single entry

(b) Multiple entry

Figure 5: RMSD vs. liquidity and spread for the two market entry settings.

for the other spread settings. In all games, regardless of market entry type, informed traders have higher aggregate profit in equilibrium than less informed traders. Traders generate the highest profits when liquidity is high (i.e., for high b). Multi-entry traders obtain the highest profits when spread is 0.05; single-entry traders do the best when spread is 0. For both entry settings, the MM incurs the greatest losses at high levels of liquidity, and maximum welfare is obtained with the widest spread and the lowest liquidity.

Aggregate trader profit and welfare is similar across the two market entry modes when the MM spread is large. In the single-entry case, trader profits decline as spread increases, for fixed liquidity. This is consistent with what might be expected, as a wider spread results in less trade. In the multi-entry setting, however, aggregate trader profits improve as the spread widens. This is because the larger the spread, the more likely agents will be able to defer trading to a later entry into the market, and information closer to the end of the trading horizon is more valuable, as previously mentioned.

Price Discovery We find that market conditions and trader reentry play a significant role in price discovery. Figure 4 illustrates the relationship between price discovery and MM profit (which is negative as the MM incurs a loss in all RSNE found). MM loss is typically around 10% of the maximum possible subsidy ($b \log 2$), which increases with liquidity. In the multi-entry case, improvement in price discovery (i.e., lower RMSD) comes at additional cost to the MM, as traders are compensated for providing information to the market. The single-entry setting runs counter to this intuition, however: price discovery improves as MM profit increases.

In the multi-entry case, price discovery improves as liquidity increases. This follows intuition since higher liquidity means that prices do not change as quickly, so market prices are not as sensitive to noise trader activity. Optimal price discovery is achieved with spread $\delta = 0.05$ and with liquidity setting $b = 5$, as we see from Figure 5(b). In fact, $\delta = 0.05$ is the optimal spread at all liquidity levels. Price discovery is particularly sensitive to spread: RMSD deteriorates by 15% to 35% when moving from $\delta = 0.05$ to $\delta = 0$ (for all liquidity settings). In contrast, Figure 5(a) shows that spread has little impact on price discovery in the single-entry case. Nonetheless, RMSD remains sensitive to the liquidity setting, with an improvement of up to 6% when moving from $b = 2$ to $b = 5$.

In the single-entry setting, price discovery is improved when liquidity is low. RMSD worsens with liquidity, regardless of spread setting. To understand why, note that when traders only enter once, they are forced to trade on their information irrespective of entry time. As a result, price discovery is best when liquidity is low and prices change quickly to reflect the information incorporated by any new trades.

6 Conclusions

In this paper, we employed empirical game-theoretic analysis to compare price discovery, as well as trader and market maker performance, in equilibrium. Our model was comprised of informed and less informed traders, who submit orders to a prediction market mediated by an LMSR market maker. We considered two market entry schemes: one where traders arrive once, and one where traders can reenter.

Our results demonstrate that trader reentry and market conditions are pivotal in promoting price discovery in a prediction market. When traders are allowed multiple entries, price discovery improves (up to a point) as liquidity increases—but contrary to what might be expected, this relationship is inverted when traders can only enter once. Forecasting accuracy only benefits from higher liquidity when traders can be strategic about when they submit an order, as an information advantage near the end of the trading horizon is more valuable. The equilibria found also reflect this phenomenon: single-entry informed traders weigh their private information more, but are incentivized to postpone trading to a later entry in the multi-entry scenario.

The simulation-based methodology we use to elucidate the drivers behind price discovery offers a promising approach for evaluating other market design choices, such as a liquidity-adaptive MM [Othman *et al.*, 2013] or a market making mechanism that learns from historical activity [Brahma *et al.*, 2012; Das, 2005]. While our results depend on our specific modeling choices, in general we based our model on prior works in the literature. We explored a limited set of trader strategies in this study, and our equilibrium findings could be altered by the inclusion of additional strategies such as Zero-Intelligence [Gode and Sunder, 1993] or other types of traders and utility functions. Another open question regards the market maker parameters in equilibrium, and the impact of a strategic MM on price discovery.

Acknowledgments

We are grateful to Dean Foster and Michael Wellman for valuable discussions, David Rothschild for providing the BetFair data, and the anonymous reviewers for helpful feedback.

References

[Abernethy *et al.*, 2011] Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. An optimization-based framework for automated market-making. In *12th ACM Conference on Electronic Commerce*, pages 297–306, 2011.

[Brahma *et al.*, 2012] Aseem Brahma, Mithun Chakraborty, Sanmay Das, Allen Lavoie, and Malik Magdon-Ismail. A Bayesian market maker. In *13th ACM Conference on Electronic Commerce*, pages 215–232, 2012.

[Cassell and Wellman, 2013] Ben-Alexander Cassell and Michael P. Wellman. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent-Based Simulation XIII*, volume 7838 of *Lecture Notes in Artificial Intelligence*. Springer, 2013.

[Chakraborty *et al.*, 2015] Mithun Chakraborty, Sanmay Das, and Justin Peabody. Price evolution in a continuous double auction prediction market with a scoring-rule based market maker. In *29th AAAI Conference on Artificial Intelligence*, pages 835–841, 2015.

[Chen *et al.*, 2007] Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *Internet and Network Economics*, pages 70–81. Springer, 2007.

[Das, 2005] Sanmay Das. A learning market-maker in the Glosten–Milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.

[Dimitrov and Sami, 2008] Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *9th ACM Conference on Electronic Commerce*, pages 200–209, 2008.

[Dudík *et al.*, 2012] Miroslav Dudík, Sébastien Lahaie, and David M. Pennock. A tractable combinatorial market maker using constraint generation. In *13th ACM Conference on Electronic Commerce*, pages 459–476, 2012.

[Dudík *et al.*, 2013] Miroslav Dudík, Sébastien Lahaie, David M. Pennock, and David Rothschild. A combinatorial prediction market for the U.S. elections. In *14th ACM Conference on Electronic Commerce*, pages 341–358, 2013.

[Gao *et al.*, 2013] Xi Alice Gao, Jie Zhang, and Yiling Chen. What you jointly know determines how you act: Strategic interactions in prediction markets. In *14th ACM Conference on Electronic Commerce*, pages 489–506. ACM, 2013.

[Gode and Sunder, 1993] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[Hanson and Oprea, 2009] Robin Hanson and Ryan Oprea. A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314, 2009.

[Hanson, 2003] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

[Hanson, 2007] Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.

[Heidari *et al.*, 2015] Hoda Heidari, Sébastien Lahaie, David M. Pennock, and Jennifer Wortman Vaughan. Integrating market makers, limit orders, and continuous trade in prediction markets. In *16th ACM Conference on Economics and Computation*, pages 583–600, 2015.

[Jumadinova and Dasgupta, 2011] Janyl Jumadinova and Prithviraj Dasgupta. A multi-agent system for analyzing the effect of information on prediction markets. *International Journal of Intelligent Systems*, 26(5):383–409, 2011.

[Kyle, 1985] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

[Ostrovsky, 2012] Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 20(6):2595–2647, 2012.

[Othman and Sandholm, 2010] Abraham Othman and Tuomas Sandholm. Automated market-making in the large: The Gates Hillman prediction market. In *11th ACM Conference on Electronic Commerce*, pages 367–376, 2010.

[Othman *et al.*, 2013] Abraham Othman, David M. Pennock, Daniel M. Reeves, and Tuomas Sandholm. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation*, 1(3):377–386, 2013.

[Pennock and Sami, 2007] David M. Pennock and Rahul Sami. Computational aspects of prediction markets. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.

[Pennock, 2004] David M. Pennock. A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *5th ACM Conference on Electronic Commerce*, pages 170–179, 2004.

[Phelps *et al.*, 2004] Steve Phelps, Simon Parsons, and Peter McBurney. An evolutionary game-theoretic comparison of two double-auction market designs. In P. Faratin and J. A. Rodríguez-Aguilar, editors, *Agent-Mediated Electronic Commerce VI. Theories for and engineering of distributed mechanisms and systems*, pages 101–114. Springer, 2004.

[Slamka *et al.*, 2013] Christian Slamka, Bernd Skiera, and Martin Spann. Prediction market performance and market liquidity: A comparison of automated market makers. *IEEE Transactions on Engineering Management*, 60(1):169–185, 2013.

[Wah and Wellman, 2013] Elaine Wah and Michael P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *14th ACM Conference on Electronic Commerce*, pages 855–872, 2013.

[Wah and Wellman, 2015] Elaine Wah and Michael P. Wellman. Welfare effects of market making in continuous double auctions. In *14th International Conference on Autonomous Agents and Multiagent Systems*, pages 57–66, 2015.

[Walsh *et al.*, 2002] William E. Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, pages 109–118, 2002.

[Wellman, 2006] Michael P. Wellman. Methods for empirical game-theoretic analysis (extended abstract). In *21st National Conference on Artificial Intelligence*, pages 1552–1555, 2006.

[Wiedenbeck and Wellman, 2012] Bryce Wiedenbeck and Michael P. Wellman. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pages 931–938, 2012.

[Wolfers and Zitzewitz, 2004] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.