AN AUTOMATED INFORMATION-BASED MARKET MAKER IN PREDICTION MARKETS

By

Aseem Brahma

A Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

MASTER OF SCIENCE

Major Subject: COMPUTER SCIENCE

Approved:

____
Sanmay Das, Thesis Adviser

Rensselaer Polytechnic Institute
Troy, New York

May 2010
(For Graduation May 2010)

© Copyright 2010
by
Aseem Brahma
All Rights Reserved

CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGMENT

I would like to thank Sanmay for being an excellent advisor, always providing direction to the research at every step. A big thanks also to Prof. Malik for many valuable discussions, and also for allowing for the deployment of the trading platform in his classrooms thereby providing us with a sizable pool of beta testers. I'd also like to thank Naveen, Mithun and Meenal for their companionship and making the office a fun place to work, and my family for their love and support.

This thesis describes joint work with Sanmay Das and Malik Magdon-Ismail.

ABSTRACT

This thesis describes and evaluates an information based market making algorithm, the Zero-profit Information based market maker (ZPI MM) within a prediction market setting. The market maker operates in a market with asymmetric information and attempts to learn the true value of the market by updating its belief based on the trades it receives. It provides an extension of the existing Zero-profit market maker to include features desirable in a practical implementation. We present a method of dealing with quantities, relaxing the unit quantity constraint from the theoretical model. The ZPI market maker is also adaptive to shocks to the true value of the market. We present results from experiments conducted in simulation to compare its performance with a logarithmic market scoring rule based market maker (LMSR), in terms of convergence properties as well as profits and liquidity. A significant contribution is a comparison of market makers in a set of live trading experiments. The novel experimental design presented allows for their comparison simultaneously in symmetric markets. These human subject experiments help validate several properties of the ZPI market maker. In particular it provides tight convergence of mean beliefs and in general earns a profit, making it a suitable choice in a real-world prediction market, where subsidizing the market to elicit information would prove costly.

CHAPTER 1
Introduction

Prediction markets (or information, decision markets) are markets designed for the primary purpose of forecasting the outcome of future events by aggregating information dispersed among traders participating in the market. This ability of prediction markets is based on the efficient markets hypothesis [10], which posits that the price of a security reflects all the available information, and it can hence be used as a predictor for uncertain events. In real stock markets, futures and options for example, aggregate information on the future value of stocks, commodities and exchange rates. Prediction markets are thus often referred to as event futures. The market representing the event creates a financial security where the final payoff is linked to the outcome. The final trading price of the market should ideally reflect the aggregate of the initial information of all traders. Prediction markets therefore are a valuable tool to provide informed traders incentives to truthfully reveal their knowledge, and help aggregate opinions.

The first practical application of prediction markets were the Iowa Electronic Markets set up in 1988 as small scale real-money markets to predict election outcomes. Berg et al [2] in summarizing their results over a period of 13 years, showed how the markets dominated the election polls used in conjunction to forecast the outcome. Prediction markets have thus been shown to equal or better traditional methods of information aggregation. With increasing interest in this area in the last few years they have been used by researchers, policy and decision makers, and corporations in the private sector to aggregate internal information and opinions. Within the public domain, they have been used in the Foresight Exchange for social, political and scientific events, and the Hollywood Stock Exchange to predict box office performance. Within the private sector, they have been used by Google [6] to determine the number of people using different online applications, by HP [24] to forecast printer sales, and others to predict future product launch dates and sales forecasts. Indeed several companies such as InklingMarkets and Crowdcast offer prediction market platforms to client organizations in order for them to make better informed decisions.

Wolfers and Zitzewitz [28, 30] discuss the growth in the use of prediction markets in recent years and provide a summary of the types of contracts used in predicting such quantitative statistics as the probability of an event or its mean or median expectation. An advantage of these markets over traditional mechanisms is that they continually aggregate information, keeping decision makers informed of the participants' fluctations in opinions over the duration of the market. As more information is made available and traders become better informed or more adept at using the system, the market forecasts also improve over time.

In addition to providing continuous information aggregation, prediction markets have been shown to outperform traditional methods of information aggregation on several occasions. For example, in the Iowa Electronic Markets, the markets were closer to the eventual outcomes 74% of the time, as compared to the 964 polls over the five presidential elections [3]; in the Hewlett-Packard markets, the predictions bested those of traditional printer sales forecasts 75% of the time.

However, prediction markets are not without their limitations. While empirical evidence suggests that prediction markets outperform traditional alternatives, Goel et al [15] find that simple statistical forecasting techniques result in comparable accuracy on several markets from the domain of sports and box-office movie revenues. Other limiting factors are biases among traders, which affect how they perceive information and interact with the market. Such biases may include the longshot bias (overestimating the probability of low-probability events occurring) or its opposite, under-pricing extreme outcomes, as found in the prediction markets at Google.

The operators of prediction markets often utilize cash and non-cash incentives such as prizes to elicit participation. Traditional belief about such incentives was that real-money markets, i.e. markets where traders would risk their own money, would produce better forecasts. Empirically however, Servan-Schreiber et al [25] found play-money markets to yield predictions approximately equaling those from real-money markets. Wolfers and Zitzewitz [27] describe five issues that must be addressed by designers of prediction markets. First among these is the problem of attracting uninformed traders. Liquidity is a key issue within prediction markets as it greatly influences the volume of trade and hence the resulting accuracy of forecasts.

Liquidity and the thin markets problem

Liquidity within a market refers to the ease with which an asset can be bought or sold without affecting the price adversely, and thus constitutes a measure of the level of participation in the market. Typically a market is organized as a continuous double auction, in which arriving traders enter the prices at which they are willing to buy or sell shares. Orders are matched where possible, i.e. a buy order is executed against the lowest sell order, and a sell order against the highest buy order, conditional on the trader specified price limits not being violated. Unfulfilled orders are placed on bid and ask queues, collectively called the order book. Such orders are called *limit orders*, and by their very nature their execution is dependent on the level of activity. (The bid-ask spread in this model is the difference between the lowest sell order price and the highest buy order price). This "chicken and egg" scenario where the lack of participation is caused by fewer traders and limited trading, constitutes the thin markets problem. This problem is exacerbated in combinatorial markets where a fixed number of traders are faced with an exponential (in the number of outcomes) number of markets. In addition to submitted orders not being fulfilled, traders might also face an active incentive not to post a trade in an illiquid market as it would reveal their private information with little chance of a benefit (*no-trade theorems* [20]). Common measures of liquidity include the quoted bid-ask spread, the volume of trading, and the depth of the trade order book. A thin market generally stagnates without reaching equilibrium.

1.1 Market making

An alternative to the mechanism of trade execution previously described, involves the use of a market maker. A market maker is an agent who is always willing to trade. Since it will always fulfill the other side of the trade, arriving traders do not have to endure a waiting period before another party becomes willing to accept their trade at their quoted price. The market maker posts bid and ask prices, and incoming orders which are executed at these prices against the market maker are called *market orders.* Market makers may incur a loss by taking on this intermediary role, but this loss may be seen as a subsidy towards increasing the liquidity in the market.

While the role of a market maker in major stock exchanges is assigned to major financial institutions, prediction markets in particular warrant the use of automated market maker functions which can guarantee to intermediate in a potentially exponential number of markets.

1.2 Inventory based market making

Scoring rules have been used in decision theory to obtain information by eliciting probabilities from individuals about the outcome of an event. They provide participants with an incentive to make an accurate prediction, and have been used for weather and economic forecasting. Hanson [16] describes *market scoring rules* as a combination of the advantages of information markets and traditional scoring rules, which can be used in the form of an automated market maker.

Under a scoring rule $\vec{s}(\vec{r})$, an arriving trader reports a probability distribution $\vec{r}$ over a complete set I of disjoint events i, and receives a cash payment $c_i = s_i(\vec{r})$ if i turns out to be the actual outcome. *Proper* scoring rules are those which maximize the predictor's expected reward for reporting her true personal belief. Under *market* scoring rules however, the arriving trader receives a payment according to the amount by which she improves on the prediction, i.e. if r^{t-1} was the last report made, the arriving trader receives a reward of $c_i = s_i(r^t) - s_i(r^{t-1})$ for her report of r^t.

Market scoring rules thus operate in a sequential manner and are inventory-based. The cost to the market maker depends not on the number of trading agents or the frequency of trades, but only on the informativeness of the last trade with respect to the initial probability estimate. The total cost therefore is dependent only on the initial and final reports made.

Logarithmic Market Scoring Rule

For each market scoring rule, an automated market maker can be constructed which implements the underlying proper scoring rule. Of these the logarithmic scoring rule is advantageous when dealing with a large outcome space because it preserves conditional independence relations (i.e. a bet on an event A given event B does not change the probability $p(B)$ of B or the conditional probabilities $p(C|AB)$, $p(C|\bar{A}B)$, $p(C|\bar{B})$ of some event C). Due to these properties making it especially suited for combinatorial markets, logarithmic market scoring rule (*LMSR*) based market makers have been widely adopted as the *de facto* standard in several real-world prediction market implementations [17].

The LMSR market maker is based on the logarithmic scoring rule,

$$s_i = a_i + b\log(r_i)$$

As previously described, the market maker can be thought of as a sequential shared scoring rule. It can be equivalently be formulated into the more conventional setting where the market maker posts the prices of the security traded and updates those prices after every trade. Each trade is made with the market maker as the buyer or seller, which is equivalent to the probability report sequence model. Chen and Pennock [4] formulate the cost function which records the total amount spent by traders as a function of the total quantities of shares outstanding. We briefly review this process below [23].

The market maker maintains a probability distribution over the outcome space Ω which consists of a security corresponding to each possible outcome. Let q_j be the quantity of security j held by all traders combined and $\vec{q}$ the vector of all quantities held. An arriving trader who wishes to purchase Q shares of security j must pay an amount $C(q_1, \ldots, q_j + Q, \ldots, q_{|\Omega|}) - C(\vec{q})$. In general, when trading any bundle of assets which changes the vector of all quantities from $\vec{q}_{old}$ to $\vec{q}_{new}$, the cost to the the trader is given by $C(\vec{q}_{new}) - C(\vec{q}_{old})$. When outcome i happens, the market maker pays $1 per share to the holders of the winning security i.

The cost function under this model is given by

$$C(\vec{q}) = b \log \sum_j e^{q_j/b}$$

where the parameter b represents the loss parameter of the market. Additionally the price per share for purchasing an infinitesimal quantity of shares for a security q_j is given by

$$\partial C/\partial q_j = e^{q_j/b} / \sum_k e^{q_k/b}$$

This price is posted as the spot price in the security. The market maker begins with the the quantity vector $\vec{q}$ set to 0, yielding an initial price of 0.5 in each security. A negative value of the quantity traded Q is used to encode sell orders, and a negative total cost represents an earning for the trader.

Bounded Loss

As mentioned before, proper scoring rules maximize a trader's reward when she reveals her true probability estimate. The LMSR-based market maker is therefore myopically incentive compatible. In addition, the loss of the market maker is bounded by the difference of the maximum payment possible to the last trader and that to the first trader, i.e. $C(q^0) = b \log |\Omega|$.

The parameter b represents the effective liquidity or depth of the market, i.e. for a larger value of b, more shares can be purchased near the current price without driving it up significantly. Therefore, selecting the value of b involves a trade-off between higher liquidity and lower worst-case loss. A smaller value of b causes the market maker to be more adaptive but also increases the bid-ask spread thereby decreasing liquidity. In order to evaluate the effect on spreads, one can consider the spread for a typical trade size Q, i.e. the difference between the average price paid in order to purchase Q shares and the average price received when selling Q shares. In a setting with a single security, the cost function is given by $b \ln(1 + e^{\vec{q_t}/b})$ and the maximum loss reduces to $b \ln 2$. If the current inventory level is given by q_t, then the spread $\delta(Q)$ is given by,

$$\delta(Q) = \frac{b}{Q} \ln \left(\frac{\cosh q_t/b + \cosh Q/b}{2\cosh^2 q_t/2b} \right).$$

In a setting where the posterior distribution of beliefs among traders does not converge over time and has a non-zero variance, the spot price of the LMSR based market maker will not converge. Convergence will occur in a rational expectations equilibrium (REE). However this relies on the strong assumption that traders are hyper-rational. These fluctuations about the equilibrium spot price make it hard to extract a quantitative probability estimate of the true price of the security. A smaller value of b therefore leads to low liquidity causing prices to fluctuate after each trade. Also the magnitude of the change in prices for a fixed bet remains the same.

At the beginning of the trading period, the market maker has no knowledge of traders' bets or of factors such as the number of traders which might affect the amount of liquidity required. Since the parameter b is set to a fixed value before the market maker has this knowledge, the choice of b remains an open problem. Othman *et al* [22] describe an adaptation of the basic LMSR market maker which adjusts the elasticity of prices with the level of activity. Abramowicz [1] describes another shortcoming of the LMSR market maker - the profit earned by a trader for improving the prediction varies greatly across the probability spectrum - and describes an implementation of the quadratic market scoring rule instead, to provide uniform liquidity.

1.3 Information based market making

Since the LMSR based market maker, and indeed others based on market scoring rules are purely inventory based, they are extremely uninformed. An alternative to this is an information based market maker which maintains at all times a distribution over trader valuations. The zero-profit market maker presented by Das and Magdon-Ismail [7, 8, 9] is an example this, which operates under an extension of the Glosten-Milgrom model of price-setting under asymmetric information [13]. Under this model, the market maker intermediating between trades faces an adverse selection problem because those traders willing to trade at the posted bid and ask prices, do so because they know something the market maker does not. The market maker must therefore offset its losses from these informed traders with gains from the uninformed (liquidity traders).

The key differences between the model presented in the above mentioned literature and the operation of real stock markets are the existence of limit orders and the restriction that an arriving trader chooses to place a trade against the bid or ask price only for a unit quantity. The focus is maintained only on pure market orders without considering limit orders because even though such orders can impact transaction price dynamics, incorporating them into the model would require traders to optimally choose the type of limit order (specifying an optimal reservation price and quantity for example). This in turn would require a detailed description of individual behavior when traders interact with a combination of market and limit orders, a subject of potential future research.

Additionally Das [7] treats each arriving market order as placed for a unit quantity because it allows an analysis of the model without restricting the form of traders' information. We later describe how to relax this restriction in a practical prediction market implementation operating with the zero-profit market maker.

We briefly describe the market model under which the zero-profit market maker operates, leaving a more detailed description of its parameters and their dynamics for the subsequent chapter. At time t the market maker has a prior probability density $p_t(v)$ over the value of the security. We assume $p_0(v)$ to be correct, so the realized value $V \sim p_t(v)$. At each time step a single trader arrives with a signal s drawn from a distribution with expectation V. The variance of this distribution represents the uncertainty in traders' signals. The information asymmetry between the market maker's knowledge of the true value V and the traders' knowledge, can be measured by the information disadvantage, the ratio of the variance in the market maker's prior beliefs and the variance in the trader's uncertainty.

For given bid price B and ask price A, the optimal decision of the trader is to: buy if $s > A$, and sell if $s < B$. It is assumed that the market maker does not face any transaction costs. In a competitive setting where two competing market makers face the same trading population and have the same information, the second will rationally undercut the prices of the first, forcing them to a zero expected profit equilibrium. This condition is satisfied by the market maker solving the following two fixed point equations.

$$\text{ask} = E_{p_t(v)}[v|s > \text{ask}],$$

$$\text{bid} = E_{p_t(v)}[v|s < \text{bid}].$$

After setting prices, the market maker can observe the actions of the trader. This action then gives the market maker information regarding the trader's signal s and therefore information regarding the realization V. It can then update its prior beliefs $p_t(v)$ to $p_{t+1}(v)$, and await the next signal. Hence the new bid and ask prices reflect the revised expectations of V if the market maker buys or sells.

Glosten and Milgrom suggest that the market maker should set the bid and ask prices respectively to $E[V|\text{Sell}]$ and $E[V|\text{Buy}]$. If there exist such bid and ask prices satisfying the zero profit conditions, then ask $\geq E[V] \geq$ bid [13]. Das [7] described how this can be extended to a practical setting through maintaining the probability distribution $p_t(v)$ described in the market model previously.

Das and Magdon-Ismail [9] subsequently present approximate algorithms for performing the updates to the market maker's probability distribution under such conditions as zero-profit (ZP), optimal (maximum discounted profit), and myopically optimal (maximizing the expected profit from the next trade). They formulate the market maker's decision problem as a reinforcement learning problem where the bid and ask prices serve as actions which both elicit information and generate reward. The decisions of trading agents to buy or sell serve as the observations allowing the market maker to update its beliefs. In the algorithm thus presented, trader signals are drawn from a Gaussian distribution and the initial belief of the market maker is also Gaussian. Starting from a state of high information disadvantage, these algorithms result in near or maximum expected profit when the initial belief is nearly correct. They therefore present clear advantages over the convergence and profit characteristics of the LMSR based market maker.

A disadvantage of the algorithms described is that they are slow to adapt to a shock in the underlying true value of the market, after the initial convergence. The convergence is exponential in the magnitude of the shock. To this effect we present in the subsequent chapter an adaptive mechanism that can be added to the existing algorithms, which takes a window of trade history into account in order to adapt to an unlikely sequence of trades.

One can thus combine the quick convergence properties of the existing information based market making algorithms with the adaptive nature of Hanson's LMSR based market makers.

CHAPTER 2
Evaluating the information based MM

2.1 Market model

The market is organized as a pure dealer market where the market maker is the sole price setter. Also the market maker does not provide any brokerage services in terms of matching orders as in an order book structure. Traders do not have the option of specifying a reservation price for the orders they wish to execute, instead all trader orders are treated as pure market orders, to be transacted at the prices set by the market maker. At all times the market maker publishes a mean price, which is the price to execute an infinitesimal quantity of the security being traded.

An arriving trader sees the published price and based on her valuation of the security, can query the market maker for the actual cost per share to execute a trade of a specified quantity. Based on this offered price, the trader may choose to execute or cancel the trade.

We begin by describing the Zero-profit (ZP) and Optimal (OPT) market makers, both described in detail by Das and Magdon-Ismail [9]. These information based market-makers maintain, at all times, a belief about the mean of the distribution of trader types (valuations). We enforce normality of this belief by extracting the mean and variance of the true posterior, and then using those as the mean and variance of the Gaussian belief. Suppose the market-maker's belief about V is normal with mean μ_t and standard deviation σ_t at time t, and the market-maker quotes μ_t (its current estimate of the true mean of the stock) as the infinitesimal price. This is also the price which an arriving trader observes as the spot price of the market.

We assume that the arriving trader receives a noisy signal w_t of the true value V of the security, which is drawn from $N(V, \sigma_\epsilon)$. The trader's decision is then based upon the quoted price and her valuation. If $w_t > \mu_t$ the trader's decision is to buy (as her valuation indicates that stock is undervalued) and to sell when $w_t < \mu_t$. The market maker assumes in general that the demand d may be a quantity randomly chosen or linear in the difference of the market maker and the trader's estimations of the true value. Thus for a trader's valuation of w_t and the market maker's quoted price of μ, the trader may choose to place an order for $\gamma.(w_t - \mu)$ units of the instrument.

An important parameter in the update procedure is the information disadvantage of the market maker, represented as the ratio of the uncertainties in belief, of the market maker and the incoming traders.

$$\rho_t = \frac{\sigma_t}{\sigma_\epsilon}$$

The signal x_t received by the market maker conveys the direction of the incoming trade, with values -1, 0, +1 for sell, cancel, and buy orders respectively. We define variables z^+ and z^- such that the arriving trader's realized signal is known to lie between the two values. For example, if she chooses to sell, her signal w_t must be less than the market maker's bid price b_t. Hence the bounds z^+ and z^- will take on the values b_t and $-\infty$. They therefore form the bounds for the Gaussian integral for updating the market-maker's belief. If a_t and b_t are the ask and bid prices, the z^+ and z^- values respectively equal $+\infty$, a_t, b_t and a_t, b_t, $-\infty$ when the signal x_t = +1, 0, -1.

Also, the bid and ask prices presented to the trader are given by, $b_t = \mu_t - \delta_t$ and $a_t = \mu_t + \delta_t$. The value δ_t thus represents the magnitude of the update to the mean belief, and is defined as,

$$\delta_t = q\sigma_\epsilon\sqrt{1 + \rho_t^2}$$

The value q used in calculating the price is computed by iterating over a fixed point equation for either the myopic or optimal market maker in order to solve the Bellman equation for the value function, as described by Das and Magdon-Ismail [9]. Considering unit quantities, the zero-profit market maker's quoted price gives zero expected profit, given that the trade is executed. The optimal MM on the other hand, aims to maximize the expected discounted return. Figure 2.1 contrasts the profit characteristics of the market makers. (Unless otherwise specified, the loss parameter for the LMSR market maker is assumed to be set at a value of 120.0, to

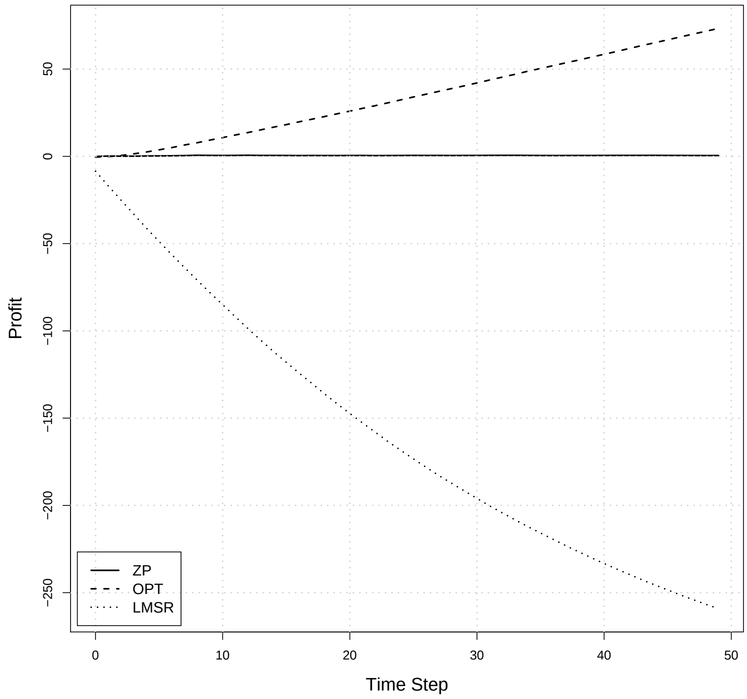

Figure 2.1: Total profit over time, with unit quantities. The Optimal (OPT) and Zero-profit (ZP) market makers begin with the variance in uncertainty set at $\sigma_0 = 12$, and have their estimate of noise in trader signal as $\sigma_\epsilon = 5$. The loss parameter b for the LMSR market maker is set at 120 to maintain similar average spreads. The unit quantities constraint follows the theoretical model presented in [13] and extended in [8]. The Zero-profit market maker makes zero profit in expectation.

produce comparable average spreads as the ZP market maker). The profit values are calculated at each trade as the difference between the transaction price per share and the true value of the security.

The spread values shown in Figure 2.2 indicate the difference in cost to the trader, between the sale and purchase of a single share in the market. As can be seen in the evolution of spreads over time, both the zero-profit and optimal market makers begin with a period of high spreads, which is then subsequently lowered as they receive more information and thereby become more certain about their beliefs.

Lastly, the update to the MM's belief consists of individual updates to its mean and variance, as below,

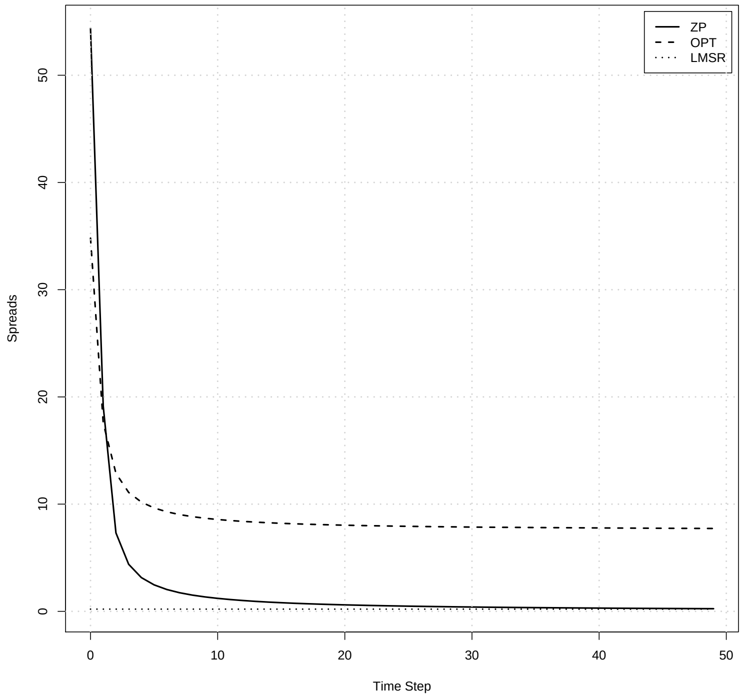

Figure 2.2: Spread per share for unit quantities. The spread presented is measured as the difference of the prices when buying and selling a single quantity. The information based market makers begin in a regime of high spreads. Once they true value of the market is learned, the spreads are lowered.

$$\mu_{t+1} = \mu_t + \sigma_t.\frac{B}{A},$$

$$\sigma_{t+1}^2 = \sigma_t^2\left(1 - \frac{AC + B^2}{A^2}\right)$$

where A, B, and C are expressions involving Gaussian integrals used in the derivation of the state updates [9].

Once the trader has been quoted a price for the quantity demanded, she can choose to accept the offer or cancel. If she chooses to cancel, the market maker performs a single price update that accounts for the added information received, i.e. the trader's valuation must lie between the original mean belief μ and the price quoted. The levels of inventory maintained by each market maker over time, is

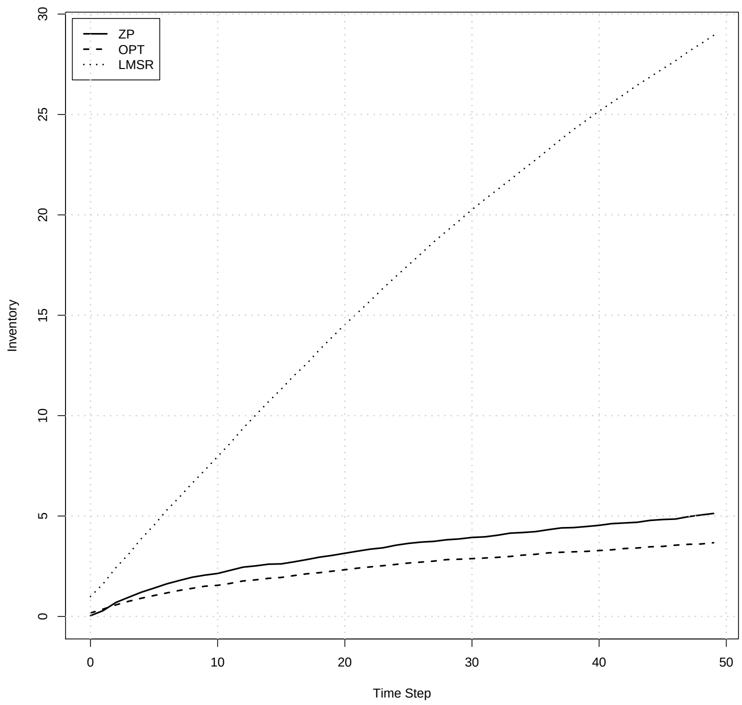

Figure 2.3: Inventory level under a unit quantity model

shown in Figure 2.3.

2.2 Quoting prices for a quantity of shares

While the market making model presented in [9] viewed the market as a series of arriving traders, each transacting a single share of the market, in a real-world setting it is essential for the market to be able to quote a price for an arbitrary number of shares to be bought or sold. To this effect, we assume that an arriving trader places an order for a quantity k of shares. An arriving trader sees the mean price in the market and places an inquiry into the price of buying or selling k shares of the event (if the arriving trader chooses not to do anything, the state of the world does not change). The market-maker's first task is to determine an appropriate price for the k shares. This is done by treating an order for k shares as a sequence of k/α unit orders, and calculating the volume weighted average price. Essentially this implements the heuristic of treating an order for a larger quantity of shares as multiple smaller (fixed-size) independent orders. Specifically, the order for k shares is treated as a sequence of orders of quantities $\alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_n$, where

$$\alpha_1, \alpha_2, \ldots, \alpha_{n-1} = \alpha,$$

$$\sum \alpha_i = k$$

$$n = \lceil \frac{k}{\alpha} \rceil$$

The parameter α governs the conversion of the order placed by the trader to the number of unit orders seen by the market maker. It therefore decides the number of updates to the market maker's belief of the distribution of valuations. When a market maker receives an order for k shares, the market maker quotes a single price that offers α shares at the single update price, α shares at the twice updated price and so on, till the demand is satisfied. Hence, for k shares, the process performs $\lceil \frac{k}{\alpha} \rceil$ updates. The update function follows the Bayesian update for the market maker's belief state, parameterized by the mean and variance, (μ_t, σ_t), as described in [9].

The overall update and price calculation procedure is shown in Algorithm 1. The mini-trades involved are not real but are fictitious trades used to elicit the series of updates. The final price quoted is therefore the sum of the costs of all the mini-trades divided by the quantity k ordered.

As is evident from the algorithm described, a smaller value of the parameter α leads to a larger number of fictitious updates used in calculating the quoted price. The larger number of updates leads to wider spreads, as shown in Figure 2.4.

2.3 Performance characteristics

2.3.1 Convergence of beliefs

The information-based market makers are quick to converge from the initial time-step, and, the decreased variance in beliefs restricts the magnitude of further mean price changes, thus causing less volatility in the prices in comparison to the LMSR market maker. This characteristic is illustrated in Figure 2.5, and further emphasized in Figure 2.6.

```
Input: Signal and Quantity for incoming trade
Output: Quoted price
i = 0;
qtyRemaining = 0;
while qtyRemaining > 0 do
    q = QValue(ρ);
    δ = q . σ_ϵ √(1 + ρ_t^2);
    if signal == 1 then
        z+ = ∞;
        z− = μ + δ;
    end
    if signal == -1 then
        z+ = μ - δ;
        z− = −∞;
    end
    blockCost = μ + signal . δ;
    blockQty = min(qtyRemaining, α);
    totalCost = totalCost + blockCost . blockQty;
    qtyRemaining = qtyRemaining - blockQty;
    (μ_{t+1}, σ_{t+1}) = updateBelief(μ_t, σ_t, z+, z−);
end
costPerShare = totalCost / totalQuantity;
```

Algorithm 1: Calculating the price per share for k shares, where the k shares are broken up into blocks of shares

To simulate the ability of arriving traders to place trades involving arbitrary quantities, we assume that the quantity to be traded is drawn from an exponential distribution. Thus, the probability distribution function for a quantity x, $(x \geq 0)$ is given by,

$$f(x; \lambda) = \lambda e^{-\lambda x}$$

Here the parameter λ equals the inverse of the mean quantity used in the market simulations. (Based on live trading experiments, discussed later, the mean quantity is typically set to 40, unless otherwise specified). Figure 2.7(a) and Figure 2.7(b) show the convergence of the mean belief from the market maker's initial state, until convergence, for different final true values. Thus in all cases observed, irrespective of the difference between the market's true value (the market maker's final belief)

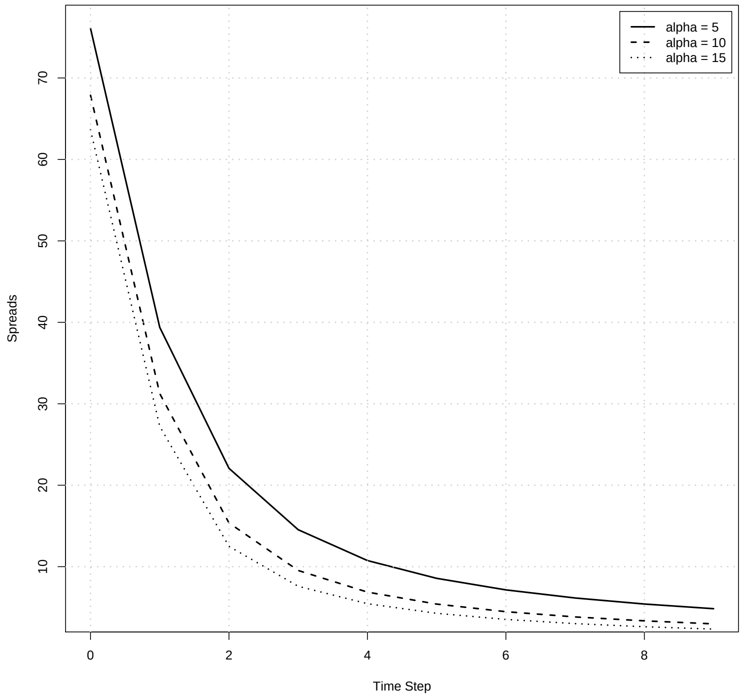

Figure 2.4: Spread values measured as the difference per share of bid and ask values for 40 shares, under different values of the stock block size α, for a zero-profit market maker. Larger values of α correspond to fewer updates for the same quantity of shares, thereby reducing the spread

and the market maker's initial mean belief (initialized to 50), the information-based market makers are quick to converge to the true value and exhibit low volatility of the mean price.

2.3.2 Spreads and variance in belief

Under the unit quantity model, the bid and ask prices to the trader are given by, $b_t = \mu_t - \delta_t$ and $a_t = \mu_t + \delta_t$, where the value δ_t represents the magnitude of the update to the mean belief. Thus the overall spread there would be given by 2δ. Under the extended model involving quantities, the spread is typically calculated for a fixed quantity of x shares, as,

$$\text{spread}(x) = \frac{\text{ask}(x) - \text{bid}(x)}{x}$$

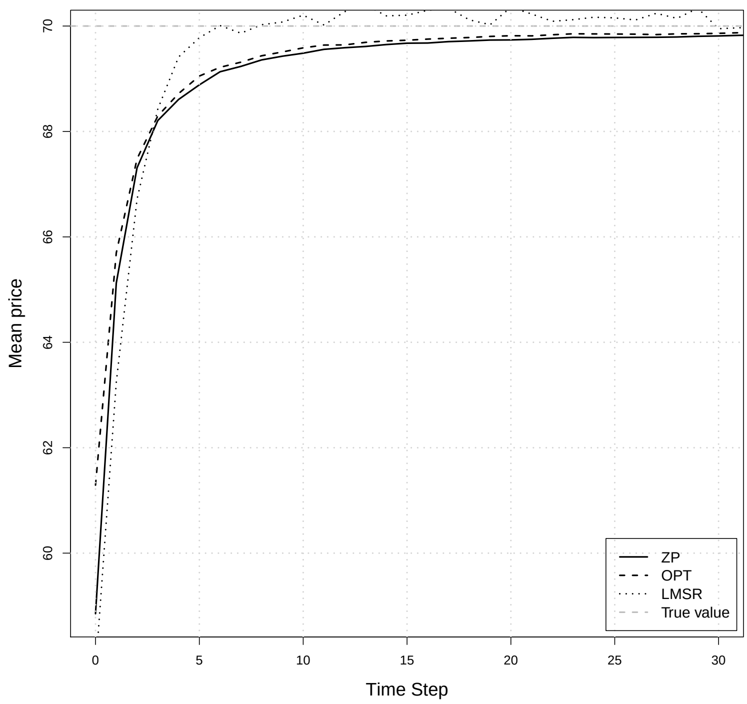

Figure 2.5: Convergence of mean prices under the different market makers, with a true market value of 70. Both information based market mmakers are quick to learn the true value, with the Optimal market maker converging faster

where ask(x) and bid(x) represent the quoted price for the purchase or sale respectively, of x shares. Figure 2.8 shows the spreads for different quantities at an early and a later stage in the convergence of mean beliefs.

The difference in the magnitude of spreads from Figure 2.8(a) to 2.8(b) is in accordance with the two-regime behavior established in [8] and reproduced in [9], i.e. when learning a new price the market maker goes through a period of high spreads which corresponds to the price-discovery regime. This is then followed by an equilibrium regime of lower spreads when the market maker learns the true valuation of the market.

Another evident feature of the value of the spread as it increases with the quantity demanded, is the impact of the parameter α which defines the number of shares that are treated as a single unit in terms of price updates. As the size of α

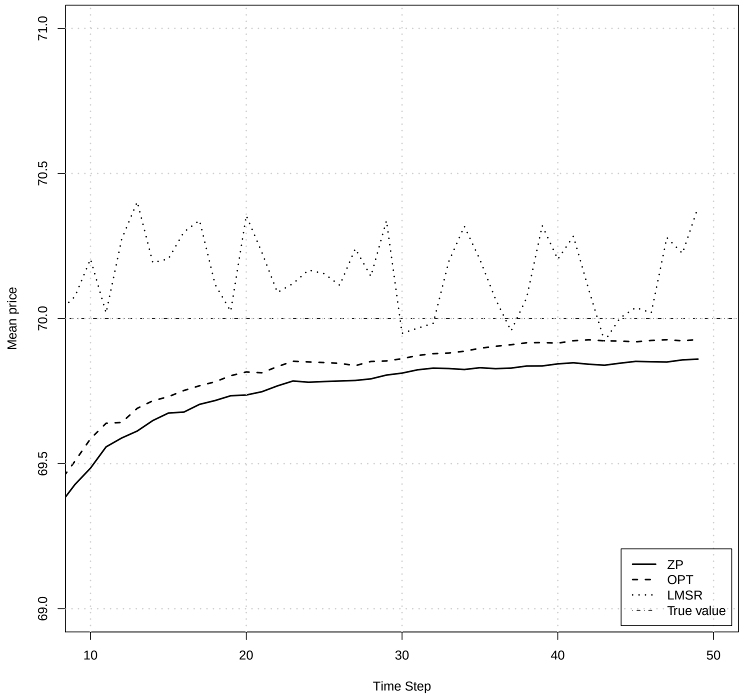

Figure 2.6: Volatility of mean prices after convergence under the different market makers, with a true market value of 70. The uncertainty in the beliefs of the Optimal and Zero-profit market makers reduces with time, giving a tight convergence. The LMSR market maker on the other hand, experiences significant fluctuations.

used in this case is 10, for any quantity x of shares $x < 10$, the bid and ask prices revert to simply $b_t = \mu_t - \delta_t$ and $a_t = \mu_t + \delta_t$, and hence the spread for each such quantity is the same, ie. 2δ.

Also, the remainder of the curve for quantities $x > \alpha$, follows the same shape, repeated every α shares. To understand the shape of the curve, recall that the computation of prices for a quantity of shares involves treating the trade order of size x as a series of $\lceil x/\alpha \rceil$ orders of size at most α. The price is computed for each such mini-order, and the final cost is the volume weighted average of the costs of the mini-orders. The returned cost per share is then the total cost averaged over the original quantity ordered x. Thus, the cost per share for buying, or equivalently selling, and therefore the difference in the bid and ask costs (i.e. the spread), can

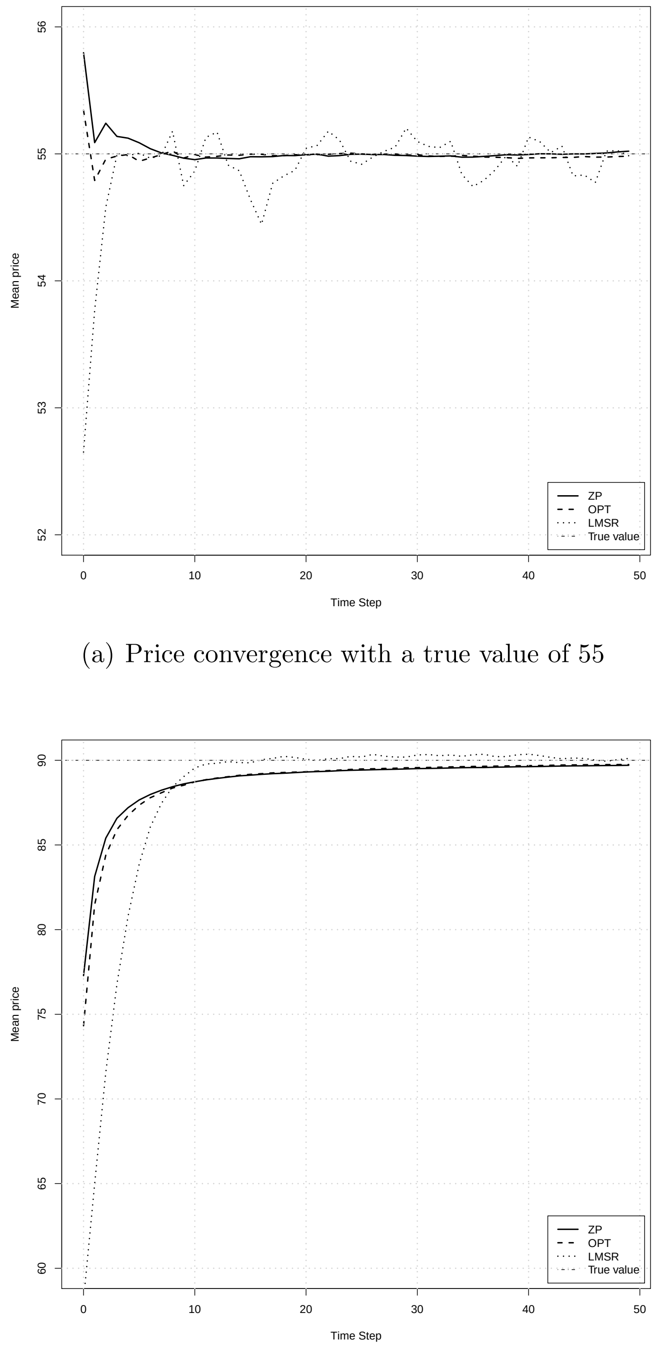

(a) Price convergence with a true value of 55

(b) Price convergence with a true value of 90

Figure 2.7: Convergence of mean belief under different underlying true values. The difference in the time to convergence between the information based and the LMSR market makers increases with increasing magnitude of the difference between the true value and the initial belief.

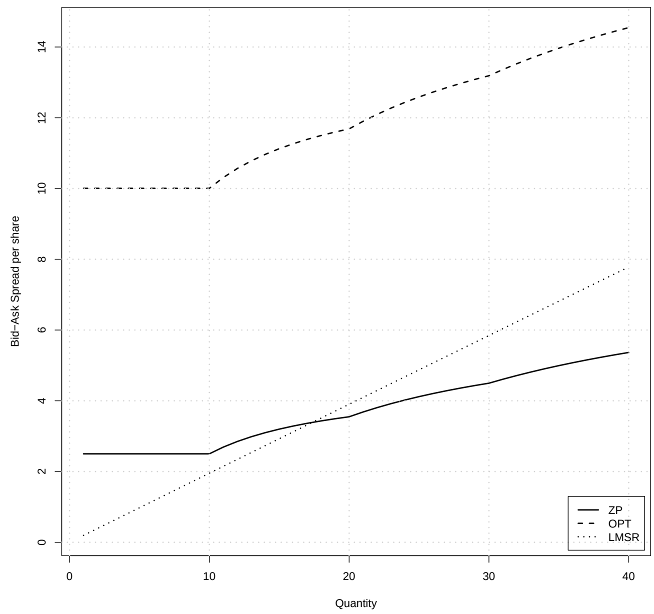

(a) Comparison at time 1

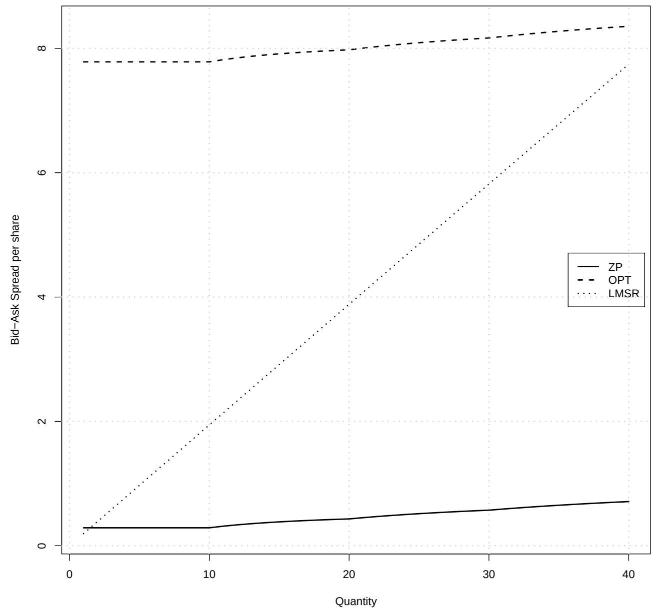

(b) Comparison at time 40

Figure 2.8: Spreads per quantity of shares before and after convergence. The LMSR spreads undergo little change as the market maker learns the true value of the market. The Optimal market maker begins with lower spreads and suffers initial losses in order to learn the true value quicker. As the Optimal and Zero-profit market makers converge, the uncertainty of beliefs is reduced leading to lower spreads.

be represented as

$$\text{spread}(x) = \frac{2\delta_1\alpha + 2\delta_2\alpha + \ldots + 2\delta_{n+1}(x - n\alpha)}{x}$$

where δ_i represents the cost on the buy (or sell) side for the mini-order i in the sequence, and $\lceil x/\alpha \rceil = n+1$. When considering possible quantities x which involve the same block of shares, they will undergo the same number of updates, and we can abbreviate the above spread to

$$\text{spread}(x) = \frac{a\alpha + b(x - n\alpha)}{x} = \frac{a\alpha}{x} + b(1 - n\frac{\alpha}{x})$$

where a, b are positive constants representing the difference of buy, sell costs over all but the last update iteration, and the last update, respectively. To take a simple example, let $\alpha = 10$, and the list of possible quantities range from 10 through 20, so that $n = 1$. Figure 2.9 plots the curve of costs according to the above abbreviated spread equation for example values of $a = 5$, $b = 6$.

As the information based market maker receives more information from traders in the form of confirmed buy or sell orders or cancel orders, it reduces its variance of beliefs in accordance with its learning of the new valuation and hence increased certainty about the true value. The influence of the variance on the spreads in the market can be gauged by the change in the spread δ within each mini-order, which is given by,

$$\delta_t = q\sigma_\epsilon\sqrt{1 + (\frac{\sigma_t}{\sigma_\epsilon})^2}$$

Therefore, as the variance σ_t decreases over time, the spreads decrease too. Figure 2.10 shows the decrease in the variance and spreads over time.

2.3.3 Combining Optimal and Zero-profit behavior

The Optimal information based market maker takes early losses by offering lower spreads than the Zero profit market maker, in order to learn the market price quicker and make significantly higher profits later. The Zero profit market maker operates in an environment of perfect competition where the bid and ask prices are

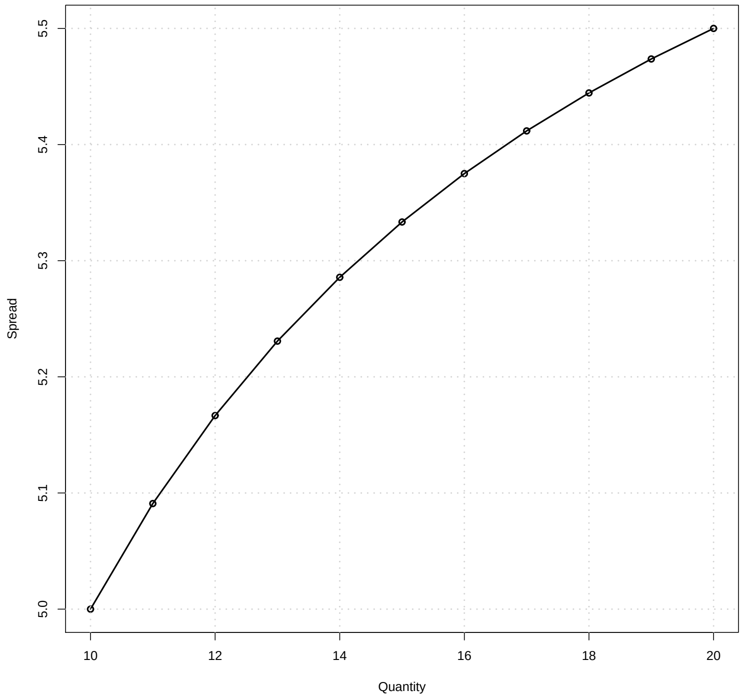

Figure 2.9: The increase in spreads for a quanitity of shares involving the same number of stock blocks as defined by the parameter α, and therefore involving the same number of fictitious updates to the market maker's mean belief and variance

pushed towards yielding a zero profit in expectation.

One can combine the properties of both models by switching from the Optimal to the Zero-profit market maker once the variance in the belief of the true value is reduced closer to its asymptotic limit. The resulting behavior thus exhibits low initial spreads in accordance with the optimal behavior while learning the true price quicker, and subsequently lowers its spreads to the level of the Zero-profit market maker, below those of the original Optimal behavior. Figure 2.11 shows the influence on spreads and profits, of the switching behavior described above.

2.3.4 Profit

In computing the profit at the end of the period of trading, we assume that the stock can be liquidated at the true value of the underlying security. The profit

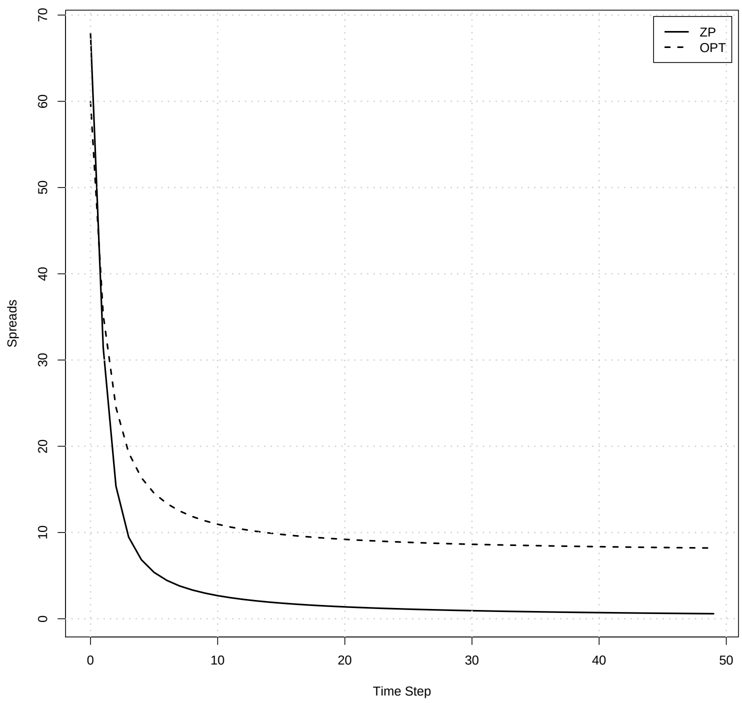

(a) Spreads over time

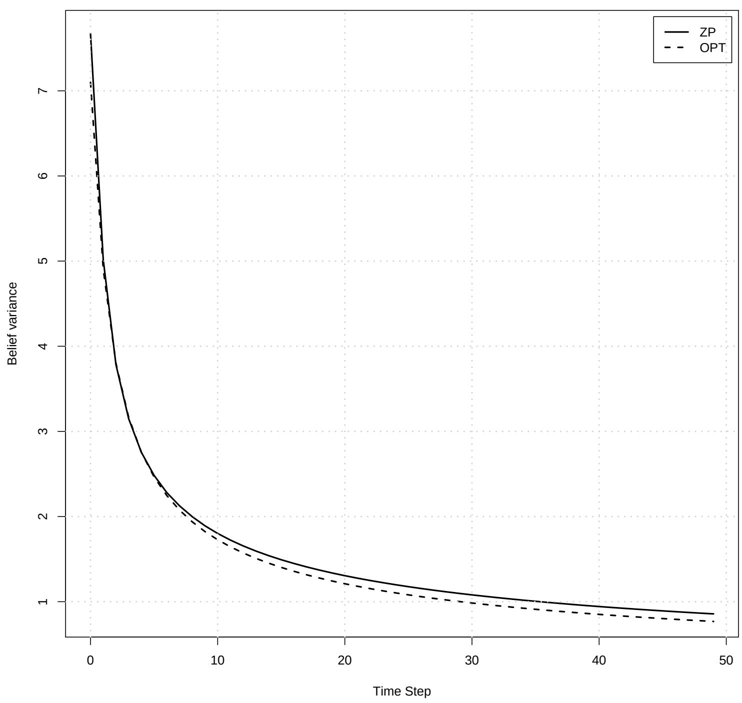

(b) Variance in beliefs over time

Figure 2.10: Decrease in spread per share (averaged over 40 shares) and belief variance iterated over simulations with random true means

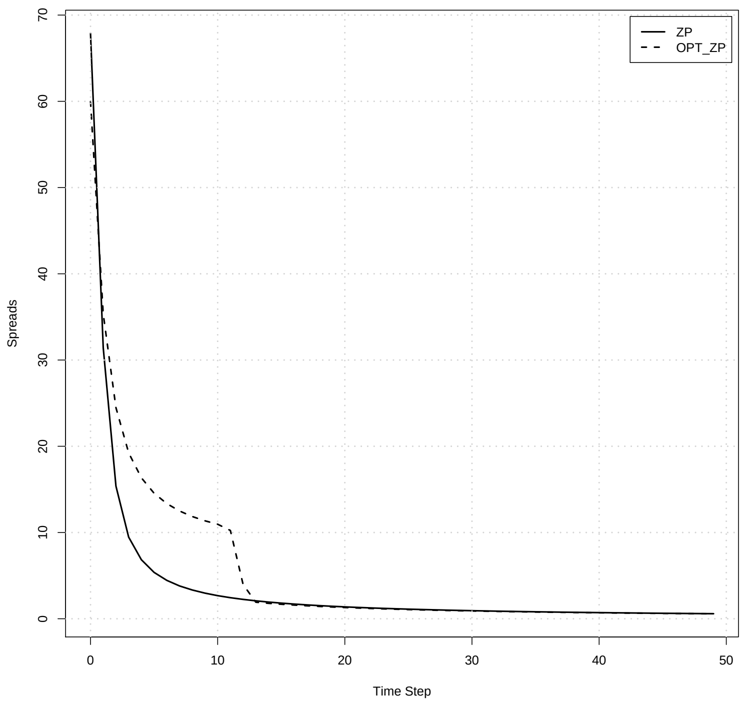

(a) Spreads over time

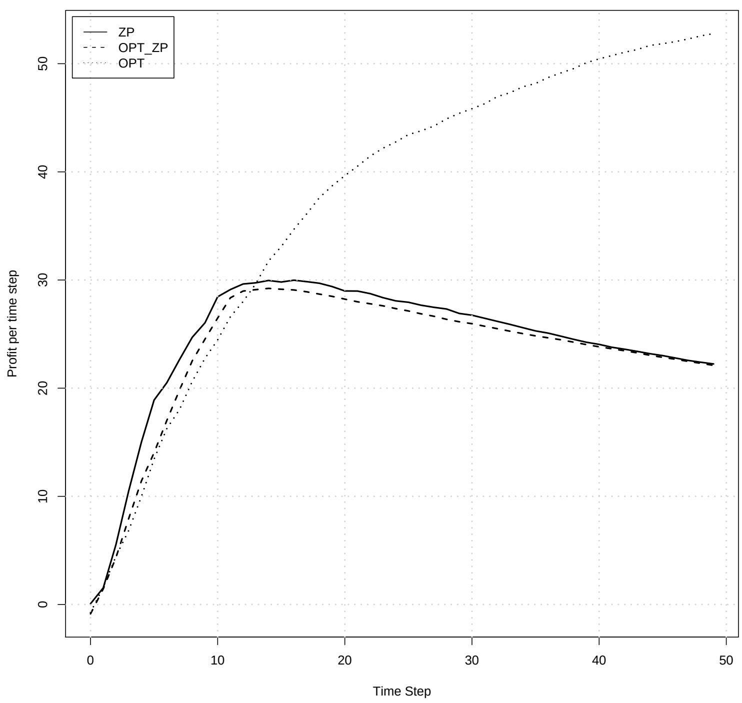

(b) Profit per time step over time

Figure 2.11: Switching from Optimal to Zero-profit market making

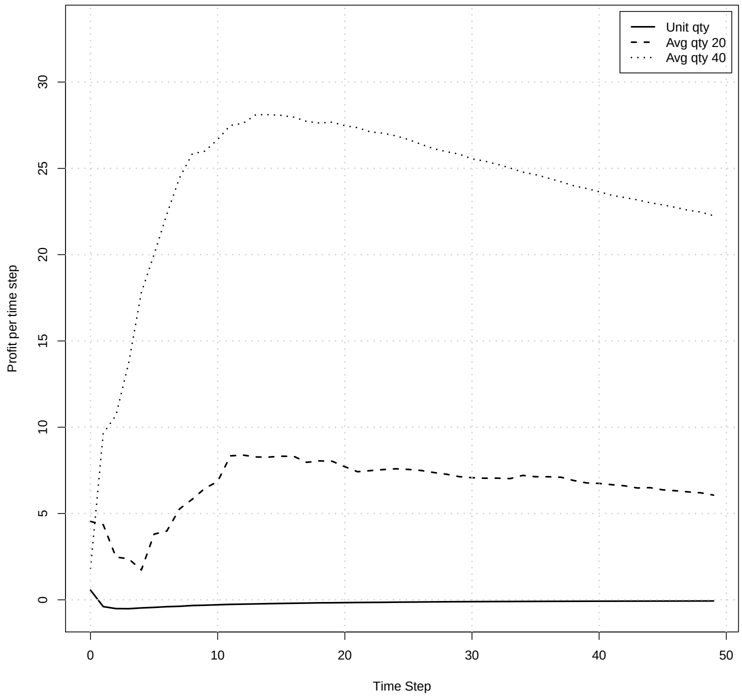

Figure 2.12: The change in profits per time step of the Zero-profit market maker, with increasing size of average trade quantities. Under unit quantities the market maker makes zero mean profit over time. When the unit quantity constraint is relaxed, and a block of shares of size $\alpha = 10$ is used for each update, the market maker makes a profit over the trading period.

π_t for a trade of quantity q can then be computed as,

$$\pi_t = q(c_t - V)x_t$$

where c_t is the cost per share quoted for the quantity q, and $x_t = +1$, 0, -1 is the signal corresponding to a confirmed buy, cancel, or sell trade.

For single unit demands, the Zero Profit market maker has an expected profit per time step of zero. This behavior is slightly distorted when quantities are not placed under such a restriction, as seen in Figure 2.12.

2.4 Robustness of the informational model: Mis-estimating trader signals

During each trading period, at each time step, we assumed that traders received a signal $w_t \sim N(V, \sigma_\epsilon)$ where V is the true value of the market. The information based market maker in addition to the belief parameters μ_t and σ_t, also maintains its own estimate of the noise in the signals that the arriving traders receive. As the market maker can only observe the buy, sell or cancel actions of each trader, it is able to revise only its estimate of the belief, while the estimate of the traders' valuations remains constant. This limitation does not significantly affect the number of updates required for the Zero-profit and Optimal market makers to converge to the true mean, even when this parameter is incorrectly estimated by the market maker.

Figure 2.13 shows the impact on the convergence and profit properties of the Zero-profit market maker, of mis-estimating the noise variance in the signals which arriving trader's receive, and therefore mis-estimating its information disadvantage.

2.5 Linear demand

In previous simulation experiments, we had assumed that the quantity that each trader picked to trade, was drawn randomly from an exponential distribution. This is consistent with literature which treats the arrival rate of incoming trades as a Poisson process, for example [11]. A section of economics literature is devoted to the pricing of assets, looks at the relationship between the wealth of investors and their utility functions, towards constructing an overall demand function for risky assets. Verrecchia [26] for example, in their description of a market for information acquisition assume that all traders have a utility for wealth implied by constant absolute risk tolerance (the inverse of risk aversion). Thus, the utility for wealth is given by the negative exponential utility function, $U(w) = -e^{-w/r}$ where the constant r is the trader's level of risk tolerance. The exponential function implies a linear demand for the risky asset which does not depend on wealth. In addition, Wolfers and Zitzewitz [29] indicate that traders with log utility have demand which is linear in their expected returns, for a binary prediction market security, and

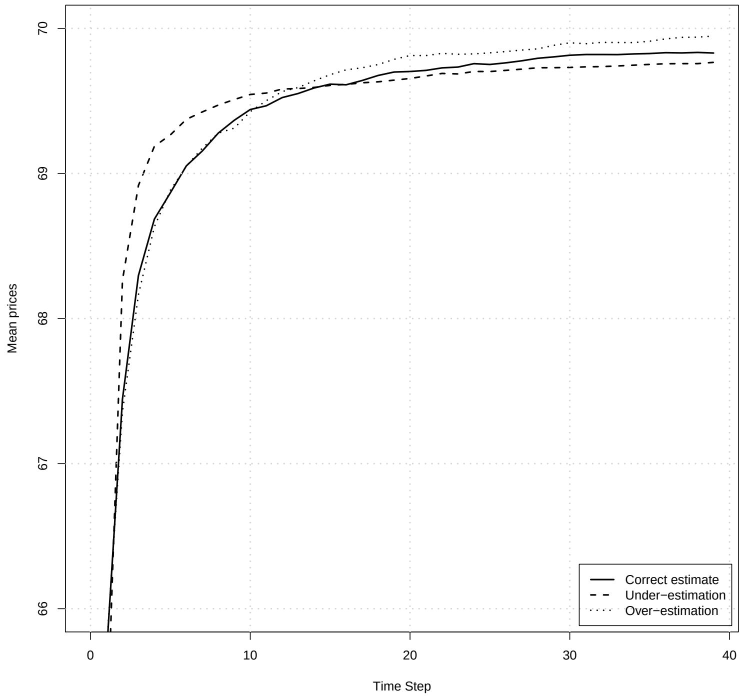

(a) Impact on convergence of mean beliefs over time

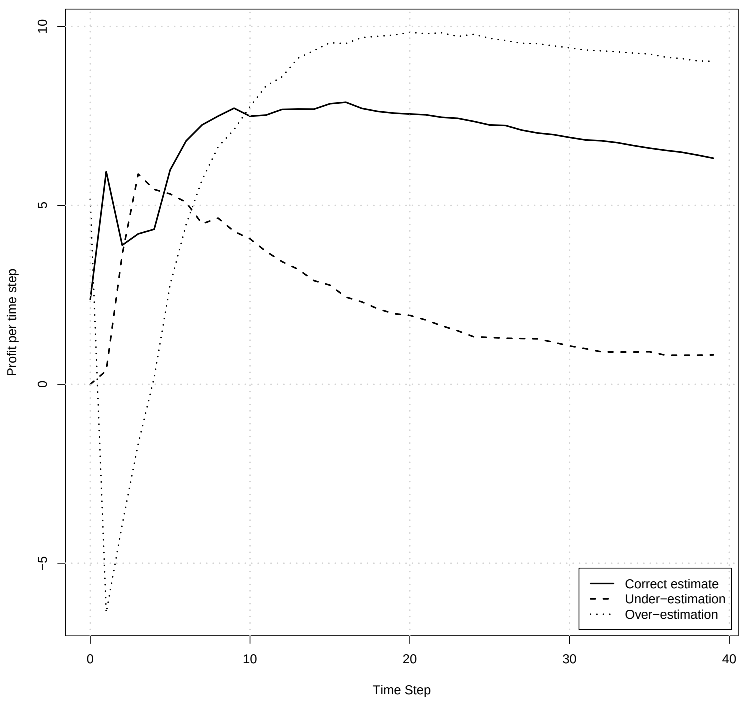

(b) Impact on profit per time step over time

Figure 2.13: Effect of mis-estimating the noise variance in traders' signals. Traders' signals are drawn from the distribution $N(V, \sigma_\epsilon)$ where $\sigma_\epsilon = 5$. With underestimation, the market maker's estimate of this variance is set at $\sigma'_\epsilon = 2.5$, and $\sigma'_\epsilon = 7.5$ when overestimating

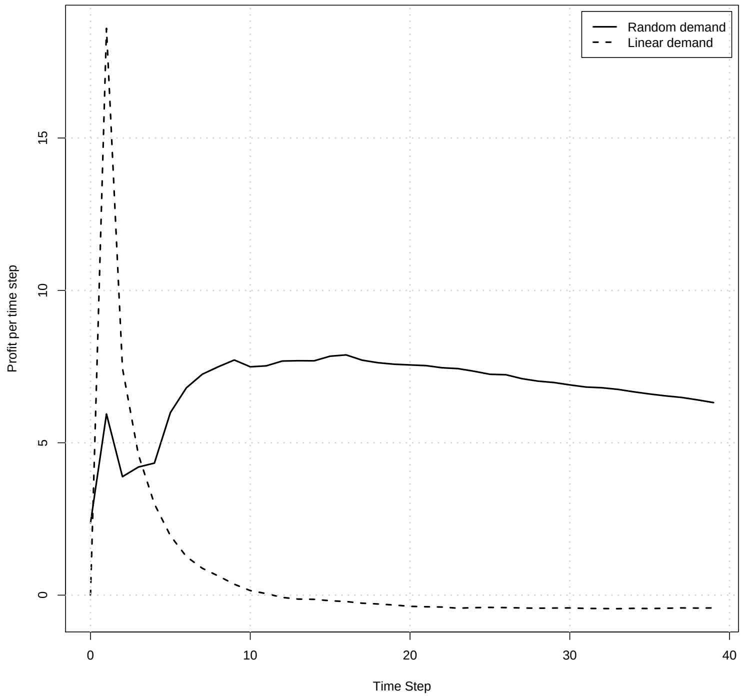

Figure 2.14: Profit per time step under a linear demand model, for the Zero-profit market maker. In the random demand case, the quantities are drawn from an exponential distribution with $\lambda = 0.05$ (i.e. mean quantity 20). In the linear demand case, the linear co-efficient is set to $\gamma = 1.0$

approximately linear for reasonable assumptions about risk aversion.

We therefore looked to test the properties of the market maker under a linear demand model which assumes that the quantities demanded by a trader are directly proportional to the difference in her valuation and the market maker's mean belief.This follows intuitively from the notion that the more undervalued (overvalued) the trader believes the stock is, the more she will be willing to buy (sell). Therefore, if a trader arriving at time t observes the mean price of the market to be μ_t and receives a signal w_t as her valuation of the stock, the quantity demanded k is given by $k = \lfloor \gamma |w_t - \mu_t| \rfloor$ where the parameter γ is the linear demand coefficient characterizing the nature of traders' demands.

As 2.14 illustrates, when the market maker has learned a better estimate of the true value of the market, the difference in its published mean price and the traders' signals reduces, thereby reducing the size of the arriving orders in accordance with the linear demand model. This translates into reduced profits as compared to the random exponential demand model.

2.5.1 α-γ scaling of profits

Depending on the estimated γ parameter that characterizes the trader demand, the definition of a block of stocks α (which thereby governs the number of updates for a given quantity of shares) may be varied to alter the profit made by the market maker.

Therefore, we aim to demonstrate that α is simply a parameter which can be adjusted depending on the perceived demand function of the trading population, to control profits. In order to analyze the relationship between these two parameters, we simulate trading periods under fixed values of $\gamma = 1.0$ and $\gamma = 2.0$ for varying values of α. To keep the simulation tractable, we compute the profit at only the first time step. (Since we only consider the first time step, it is essential to ensure that there are a substantial number of data points relating to confirmed buy and sell orders, as cancel orders yield no profit. To this effect, the simulations were run over 200 million iterations)

The relationship between these parameters and the expected profit π exhibits a scaling relation as illustrated in Figure 2.15. 2.15(b) displays the data for $\gamma = 1.0$ rescaled, such that each original point at (α_i, π_i) is now displayed at $(2\alpha_i, 2\pi_i)$. As the new points now lie along the $\gamma = 2.0$ curve, this empirically demonstrates the scaling relationship between the two parameters.

$$2\pi_{t=0}(\alpha, \gamma) = \pi_{t=0}(2\alpha, 2\gamma)$$

In a simulated experiment for the first time step, the true value V is drawn from a normal distribution $N(50, \sigma)$, and trader valuations w are drawn from $N(V, \sigma_{epsilon})$. One can compute the expected profit given V, assuming the buy and sell side to be symmetric, as below.

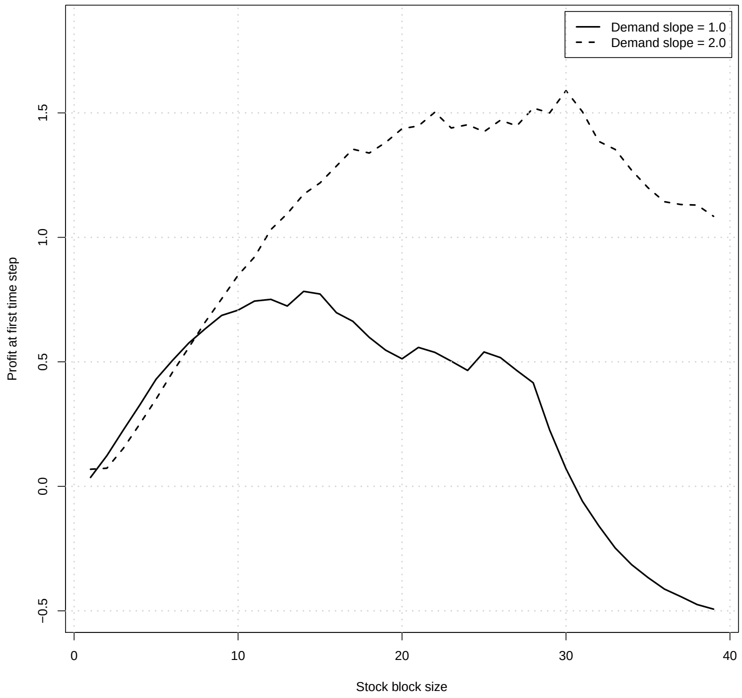

(a) Profit at first time step, for two different linear demand coefficients

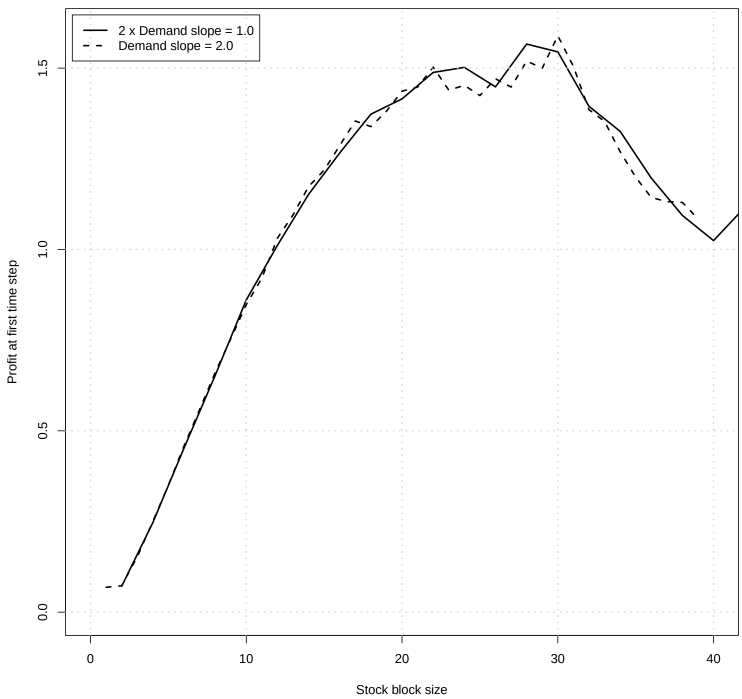

(b) Scaling of profits with linear demand

Figure 2.15: Relationship between the linear demand and the stock block size, for $\gamma = 1.0$ and $\gamma = 2.0$. Profits are computed at the first time step. In (b), the points under $\gamma = 1.0$ are re-plotted from (α_i, π_i) to (2 α_i, 2 π_i). The re-plotted points fall along the curve for $\gamma = 2.0$, following the scaling property discussed

$$\begin{aligned} E[\pi|V] &= \int_0^{100} dw N(V,\sigma)\pi(V,w) \\ &= \int_0^{\mu} dw N(V,\sigma_\epsilon)\pi(w|V) + \int_\mu^{100} dw N(V,\sigma_\epsilon)\pi(w|V) \end{aligned}$$

As before, we consider only the profits for trades occurring at time 0. The expression $\pi(w|V)$ denotes the profit given the true value, and the trader's signal (the buy/sell decision is based on these two parameters as well). Assuming the sell decision side to be symmetric at time 0 when $\mu_0 = 50$, we can fold profits under both cases to twice the profits under the trader's decision to buy. The quantity to be traded under the linear demand model, is given by $d = \lceil \gamma |w - \mu| \rceil$. Once a price per share C is quoted to the trader, the trader will choose to cancel if $C > w$ (resulting in no profit), and choose to confirm execution of the buy trade if $C < w$, resulting in a profit of $d \cdot (C - V)$.

Since we are considering only trades occurring at time 0, the updates to the mean belief can be pre-computed. For example, if a single update to μ_0 yields a_0, and the update to a_0 yields a_1, the price per share for two blocks of shares is given by $(a_0 + a_1)/2$. In general, if j updates to the mean belief starting at $\mu_0 = 50$, yield the value a_j, the price per share for n blocks of shares is given by $f(n) = \frac{\sum_{j=1}^{j=n} a_j}{n}$.

Therefore, the price per share for a number of such blocks can be pre-computed. The parameter α denotes the size of a block of shares. A demanded quantity d can then be broken down as $i = \lfloor d/\alpha \rfloor$ complete blocks of shares and $(d - i\alpha)$ remaining shares. Hence the overall cost per share, using the pre-computed prices, is

$$C = \frac{[(i+1)f(i+1) + if(i)] \cdot (d - i\alpha) + f(i) \cdot i\alpha}{d}$$

As noted earlier, a non-zero profit occurs only if $C < w$. Substituting the value of the cost per share, the resulting condition can be expressed as

$$\frac{d}{\alpha} < i(i+1) \cdot \frac{f(i+1) - f(i)}{(i+1)f(i+1) - (i + w/f(i))f(i)}$$

This condition imposes certain limits within which the signal w must lie in order for the trade to execute and the market maket to earn a profit. The value of d/α can be expressed as a ratio of γ/α. Therefore all the terms of the above condition are dependent only on the ratio of γ and α but independent of their actual values, and keeping the ratio the same implies the same limits on the signal w.

In addition, one can substitute the value for the cost per share into the expression for the profit, as shown below.

$$\begin{aligned}
\pi(V,w) &= d \cdot (C - V) \\
&= \alpha(\frac{d}{\alpha}C - \frac{d}{\alpha}V) \\
&= \alpha(\frac{[(i+1)f(i+1) + if(i)] \cdot (d - i\alpha) + f(i) \cdot i\alpha}{\alpha} - \frac{d}{\alpha}V) \\
&= \alpha([(i+1)f(i+1) + if(i)] \cdot (\frac{d}{\alpha} - i) + f(i) \cdot i - \frac{d}{\alpha}V)
\end{aligned}$$

As before terms within the inner expression involve only the factor d/α, dependent on the ratio γ/α. The overall profit however changes proportional to α. Therefore, doubling the value of α, keeping the ratio of γ/α, also doubles the profits at the first time step, as seen empirically.

2.6 Adapting to jumps

With each arriving trade the market maker receives new information which is used to update its mean belief and narrow down the variance in the belief. Thus, after a number of trades have been processed the variance and therefore the spreads are significantly reduced. While this increases liquidity and encourages further trading towards the true stock valuation, it also poses a restriction on changes to the true value. As the magnitude of each mean belief update is directly proportional to the standard deviation of the market-maker's belief, large jumps in the true underlying value coupled with a small belief variance, lead to very small update values.

After a jump occurs, the sequence of trades will be inconsistent with the market maker's belief of the old valuation along with its low belief variance (high

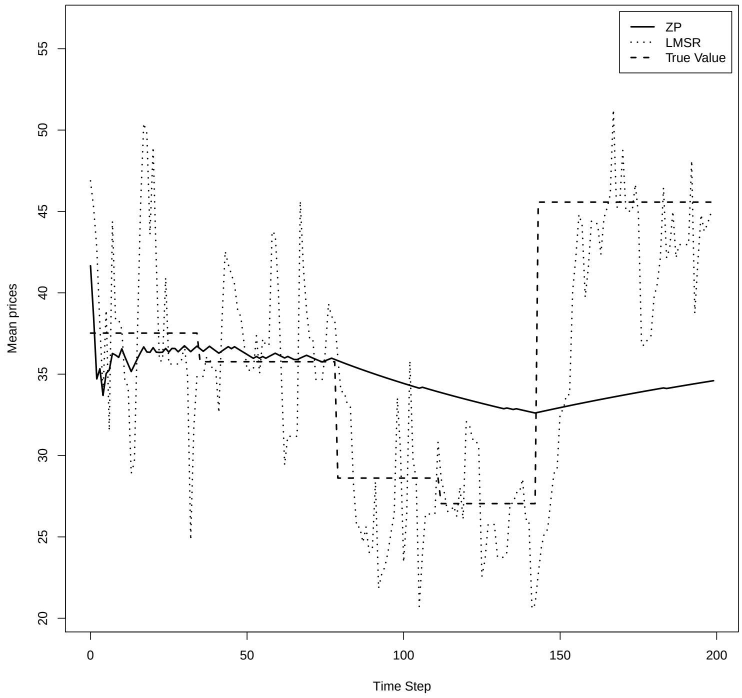

Figure 2.16: Mean prices with jumps in the true value. The jumps are uniform in nature, i.e. the new true value is chosen uniformly at random between 0 and 100. The probability of a jump at each time step is 0.01.

confidence). Therefore, an intuitive solution to incorporate changes in the underlying stock value involves increasing the variance when a stream of incoming trades appear to lead the market-maker in a different direction than what the mean-belief had previously converged to. The market-maker keeps track of a fixed window of previously occurring trades (including canceled trades), along with their corresponding update integral limits, i.e. the z^+ and z^- values which were inferred from those trades. At a particular time step, the probability of a sequence of trades represented by a window of size s, is computed as shown below.

$$L(\mu, \sigma) = \int_{-\infty}^{\infty} N\big(v, \mu, \sigma\big). \prod_{i=1}^{s} \bigg(\Phi\big(z_i^+, v, \sigma_\epsilon\big) - \Phi\big(z_i^-, v, \sigma_\epsilon\big) \bigg) \, dv$$

$L(\mu, \sigma)$ represents the likelihood of the market maker observing the previous $|W|$ trades under its current belief state (μ, σ). One can then compare this probability against a fixed threshold. If the probability is small, the market maker is in and inconsistent belief state and its variance (level of uncertainty) must be increased. However, this solution is problematic because the choice of the threshold and the choice of the increase in uncertainty, are highly sensitive to the choice of the window size. Instead we make a relative comparison between the probability computed at the current state and that computed at twice the level of uncertainty. Hence we can define a consistency index as

$$C(\text{history}) = L(\mu_t, 2\sigma_t) - L(\mu_t, \sigma_t)$$

If $L(\mu_t, 2\sigma_t) > L(\mu_t, \sigma_t)$ i.e. if $C > 0$, we increase the the variance of the market maker's belief, $\sigma_{t+1} = 2\sigma_t$. In general σ_t can be increased by any multiplier greater than 1. Since the consistency index is defined as a relative measure, the choice of the multiplier is arbitrary. During our simulations and trading experiments, we used a multiplier of 2, but we expect the results to be robust for other choices, as compared to a fixed probability threshold.

In simulation experiments, we consider two different types of jumps in the true value of the market. The first type is a Gaussian shock, a more realistic scenario, where the magnitude of the jump is drawn from a Gaussian distribution with mean 0 and variance σ_j^2. In the second type of market shock, the new value is drawn uniformly at random between 0 and 100. This represents a very problematic case for the market maker. At any point of time the probability of a jump occurring was set to $p_{jump} = 0.01$. The following table summarizes results from 1000 simulations under the above two cases.

The root mean square deviation of the infinitesimal price from the true value is used as a measure of price discovery. The maximum loss represents the worst loss suffered by either market maker in any single simulation. The loss parameter for the LMSR market maker was set to 125 in order to maintain similar average spreads with the ZPI market maker, and the maximum loss observed is close to the theoretical bound of 8664.34. The ZPI market maker performs better in general, in terms of both price discovery as well as earning a profit. However as the probability of a jump goes up, and if the shocks to the market are uniform ones, the loss suffered

	Gaussian Shocks		Uniform Shocks	
	ZPIMM	LMSR	ZPIMM	LMSR
Profit	2081.35	−2457.30	603.40	−1897.98
Max Loss	9479.82	8662.32	50183.77	8384.42
Spread	1.42	1.35	1.79	1.40
RMSD	2.92	5.38	8.78	10.79

Table 2.1: Performance of ZPIMM and the LMSR MM in simulated trading

increase. Therefore it might not be the best choice in highly unstable markets. Table 2.1 summarizes our results for both types of jumps.

Figures 2.17 and 2.18 compare results under uniform shocks to the market, for windows 5 and 10. It is evident from the changes in spreads over time, that the window size W becomes the dominant factor in measures like the average spread, so that the particular value of the estimate of variance in trader signals σ_ϵ becomes unimportant under the adaptive mechanism. The probability value thus computed is compared against a threshold probability, and the market-maker's variance is increased by a constant factor accordingly.

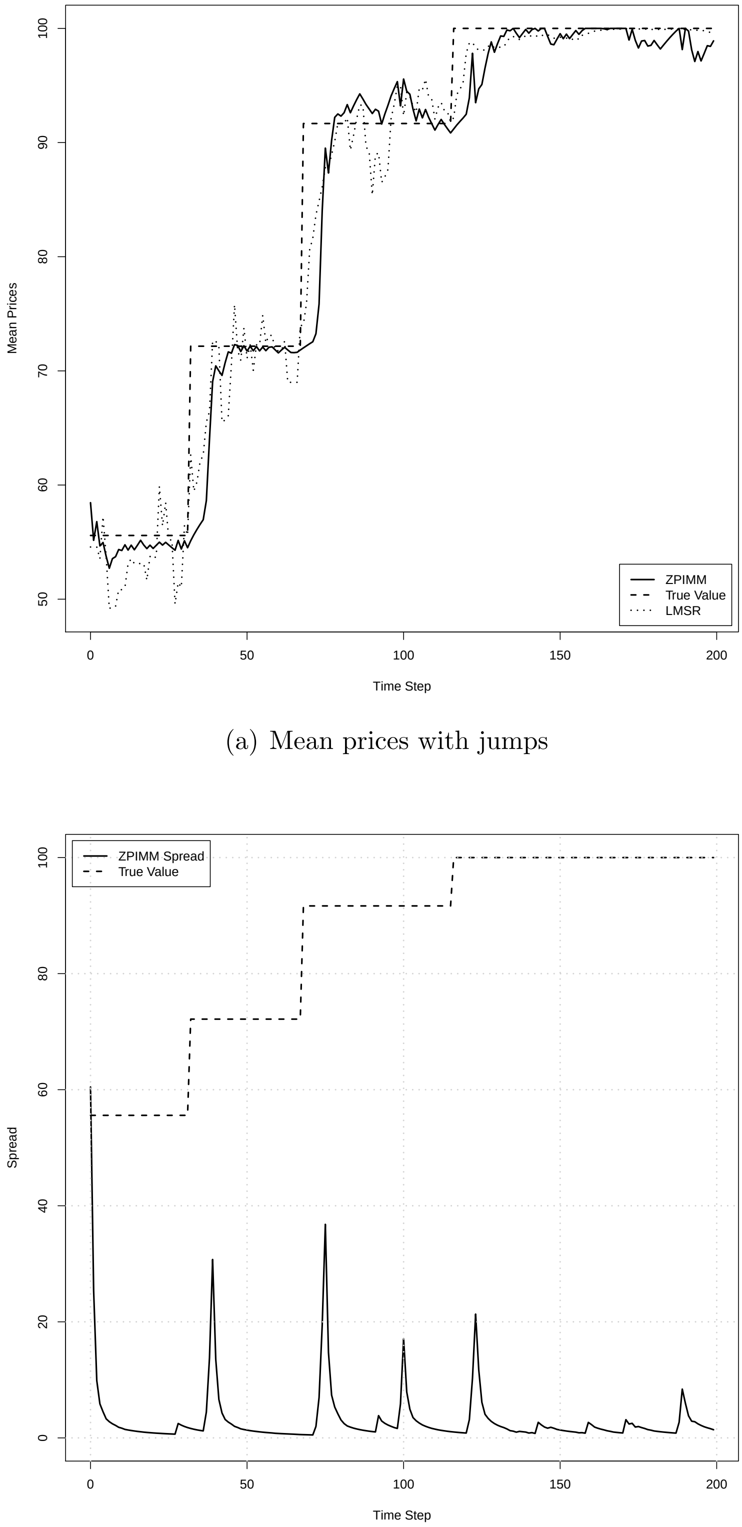

(a) Mean prices with jumps

(b) Change in spreads under adaptive ZPIMM

Figure 2.17: Adaptive behavior of ZPIMM with window 5. The probability of a jump at each time step is 0.05. Quantities are drawn randomly from an exponential distribution with mean quantity 20

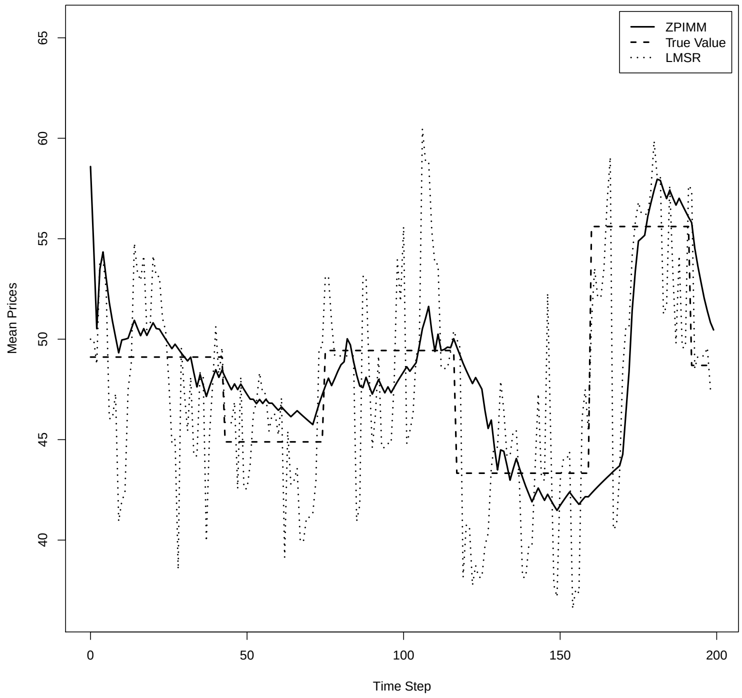

(a) Mean prices with jumps

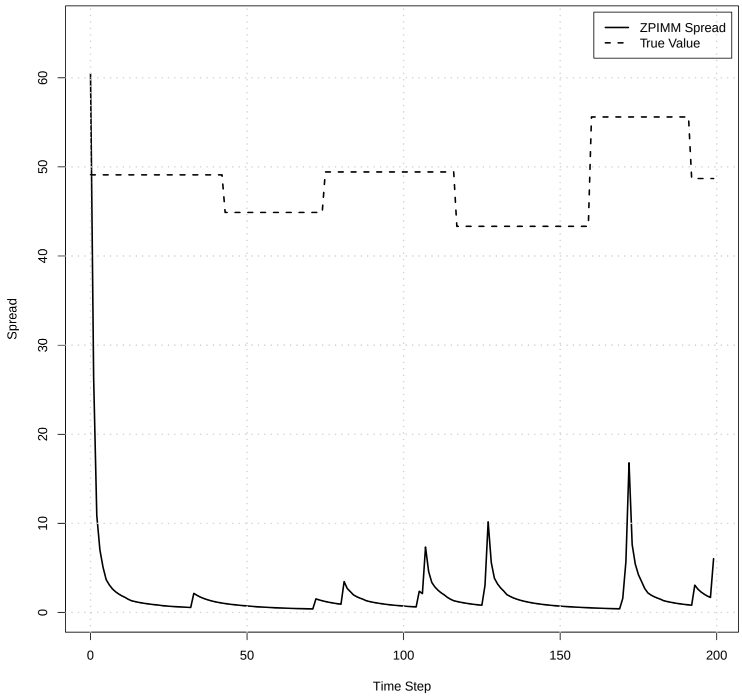

(b) Change in spreads under adaptive ZPIMM

Figure 2.18: Adaptive behavior of ZPIMM with window 10. The probability of a jump at each time step is 0.05. Quantities are drawn randomly from an exponential distribution with mean quantity 20

CHAPTER 3
Live Experiments

Prediction markets act as a platform for aggregating information, and in many real-world scenarios, they have been deployed and found to outperform traditional mechanisms such as surveys and polls. The experiments conducted in simulation serve to analyze the behavior of the market makers in a controlled environment, where traders are completely clear on their valuation (i.e. their signals about the price of the security, are clearly defined), and we assume that their corresponding action is rational. In a real-world deployment however, both these constraints must be relaxed when designing a market maker which can efficiently and effectively aggregate the information of the trading crown to arrive at a consensus.

In a university setting, there are traditionally many methods of feedback used - classroom feedback for lecturers, student body polls on decisions made by the administration, etc. It follows therefore that prediction markets would be an invaluable tool to take advantage of the individual information and opinions each student holds about such issues - a pertinent example is the prediction market constructed to predict the opening of two new buildings within the Computer Science department at Carnegie Mellon [21].

In an initial effort towards achieving such a goal, and also to be able to compare the behavior of real traders under different market makers, we designed a web-based trading interface, initially made available to students in selected courses, and mostly during timed experiments.

3.1 Features

The trading interface presented to users resembles an online broker (Figure 3.1), displaying the price-volume history and the current mean price of the market. Traders can also view the number of shares they currently hold across all markets. This is relevant because students participating in trading experiments would typically have holdings in each of the markets being run under the different market

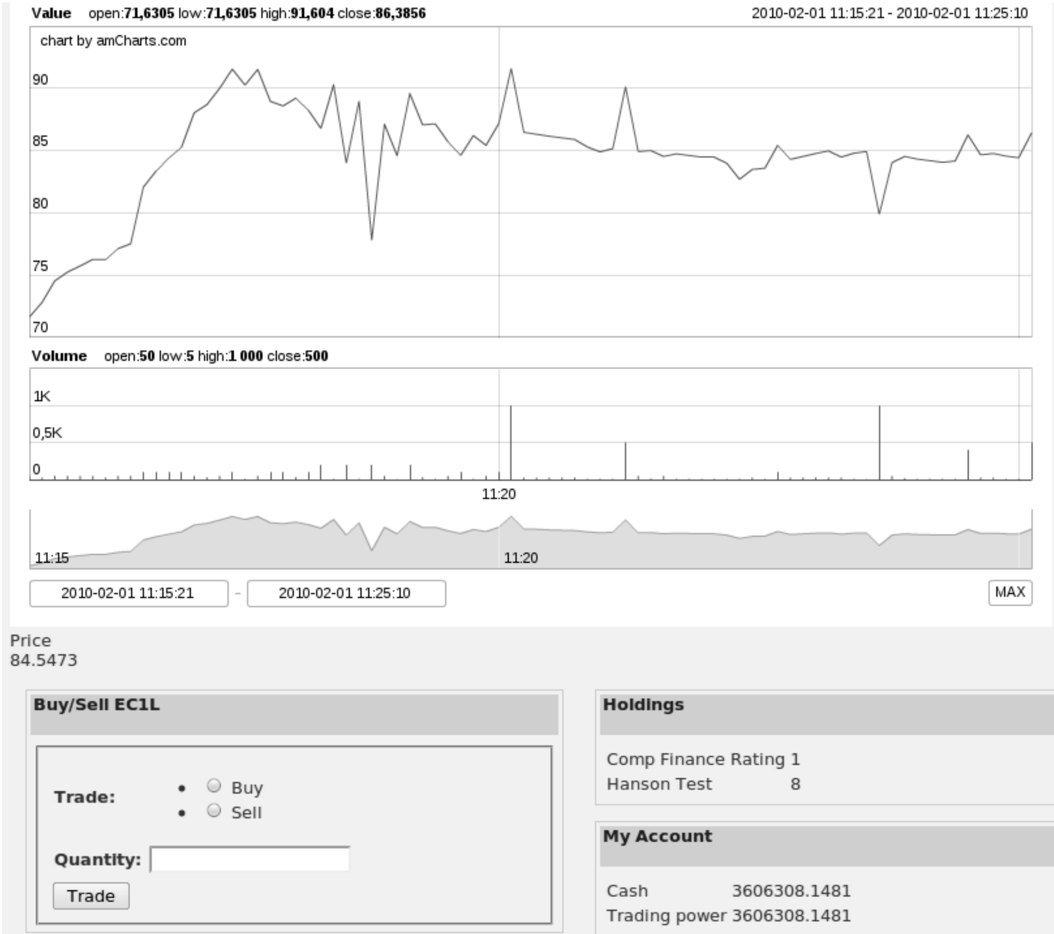

Figure 3.1: Trading interface

makers. Also, they may begin the experiment with an initial endowment, which is typically in the form of cash, but which may also include holdings in one or more markets.

Each trading account is assigned a margin value. A margin is defined as the amount of collateral that the trader must deposit to cover the counter-party's credit risk. In effect, this allows a trader to use the existing holdings as collateral in order to increase her position and buy more shares in the security. Therefore the 'trading power' available to the user is given no longer by the cash amount or the current market values of the holdings alone, but is expressed as,

$$tradingpower = \frac{cash + |L| - |S|}{m} - (|L| + |S|)$$

where L, S are the current market values of the long and short positions respectively, and m is the margin (ratio between 0 and 1). For any trade to execute, the magnitude of the total cost of the trade, i.e. the 'cost per share' times the quantity demanded, must lie within the trading power.

Using the trading interface, users are allowed to view the spot price and place only buy/sell market orders for a desired quantity. Once an order is placed, they receive a notification with the price per share which the market maker quotes, which they can choose to accept or reject (a reject leads to a 'cancel' signal sent to the market maker). If the trader chooses to accept the quoted price, the trader may receive either a confirmation of trade execution, a notice that the trade could not be executed due to trading power constraints, or, notification that the price quote is no longer valid due to other incoming trades which have changed the state of the market maker. In the event that the market maker's state has changed, the trader must then issue a fresh order for a new trade.

Deployment of the trading system

The trading interface was initially made available to students in the Computational Finance course (Fall 2009) for educational purposes and to conduct trading games through the duration of the course. There were two main markets set up, a ratings market to reflect weekly lecture ratings as percentages, and a homework market to predict the mean grade on class homeworks. Across the group of 23 trading accounts registered within the system, the ratings market received 3263 trades, while the homework market received 583 trades.

In addition to trading each week, students would receive a small percentage of interest daily, based on their cash holdings. In order to maintain the appearance of activity in periods of low trading (and thereby encourage traders to participate) a random trader was placed in the ratings market, which would place a random trade (random in both the direction and quantity of the trade) at periodic intervals a few hours apart.

Extras

The ratings market used in the Computational Finance course also had a section with polls for students to vote on each lecture, assigning it a score on a scale of 1 to 5. Based on the observed poll average, the trading interface allowed for the administration of cash dividends to users, which represented the payoff for the week's trading.

While students had individual opinions of each lecture, and expressed this individual information through their poll votes, the market was used as a means to aggregate this individual information. The objective therefore was to observe the evolution of prices each week from the previous week's rating to the collective knowledge of the latest lecture's rating.

Although unused in the experiments so far, the web-interface was designed to incorporate user comments, and thoughts in the form of user blogs and wiki entries. This is a useful feature of real prediction markets because it encourages interaction among participants and discussion about trading strategies. It could therefore serve as a medium of information dissemination where informed traders may choose to explain why they traded the way they did (a feature found for example at inklingmarkets.com) and uninformed traders may become better informed.

3.2 Experiment Design

In order to compare the Information Based market maker with the LMSR based market maker, we ran two markets sequentially till the mean prices converged. Given the limited nature of human subject experiments in terms of the number of participants, it would be cumbersome (and wasteful) to extract only the behavior of a single market maker on a single experiment, from each group of participants. In addition, different (human) participants are very likely to exhibit different behaviours and tendencies in reacting to their individual signals (or valuation) and placing trades accordingly. Even with repeating the experiment within the same group, subjects are likely to form biases, generalizing based on the outcome of the previous experiment. Therefore, it is necessary to devise a mechanism to test the trading behavior and the reaction of the market makers in a setting where partici-

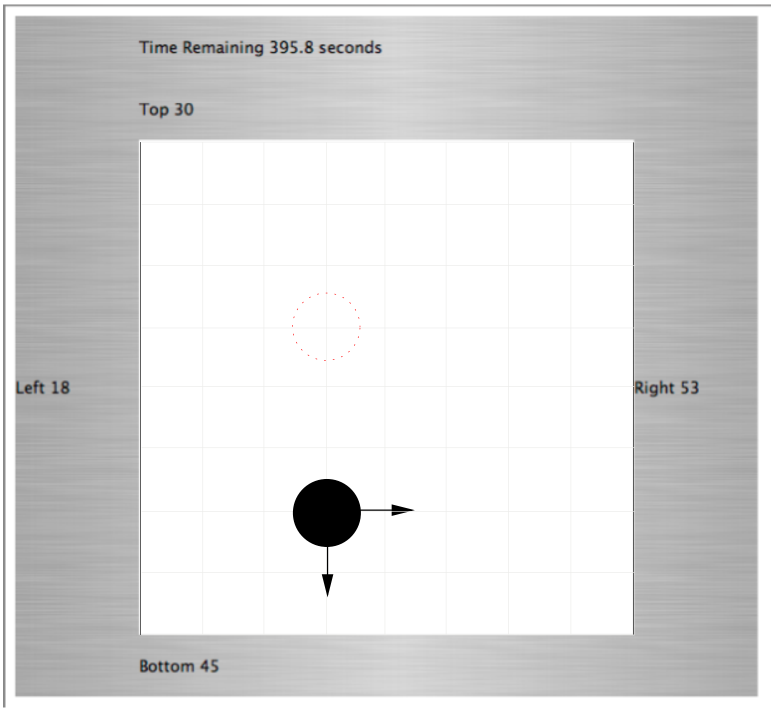

Figure 3.2: Random walk along 2-dimensions

pants can trade against both markets simultaneously. Also it is necessary that the two markets be constructed in a symmetric manner and that the traders be unaware of the market makers underlying them.

The design used for our experiments consisted of a ball moving in a two dimensional random walk, across a grid, as shown in 3.2. The walk consists of two 1-dimensional random walks, along each of the axes, horizontal or left-right (LR), and vertical or top-bottom (TB). The initial position of the ball, on the grid, is given by (x_0, y_0). At each time step, the ball has a probability p_{LR} of moving right along the horizontal axis, and a probability p_{TB} of moving down along the vertical axis. The position (x, y) of the ball is restricted to the boundary of the grid of size S, i.e. $[-S, S]$ along each axis. If $|x| \geq S$, the x co-ordinate is reset to x_0, and similarly if $|y| \geq S$, the y co-ordinate is reset to y_0. (In each case the other co-ordinate remains unchanged). In each of these cases (where the position off the ball goes beyond the grid and has to be reset), there are counters maintained which keep track of the number of times that the position (x, y) has exceeded the top, bottom, left, or right boundary of the grid.

3.2.1 Trader signals

Participating traders view a realization of a random walk over a fixed period of time. Along with the walk itself, they are shown the values of the counters which maintain the number of times the ball has hit the left, right, top, or bottom edge of the grid, and the time remaining in the experiment. The trader can estimate the probability that the ball hits the right (respectively bottom) edge before it hits the left (respectively top) edge, V_{LR} (respectively V_{TB}), from these counter values. For example, from the partial realization shown in 3.2, one can estimate $V_{LR} \approx \frac{53}{71} = 0.75$ and $V_{TB} \approx \frac{45}{75} = 0.60$. As is evident from these estimated values, even though the same random process operates along each axis, the traders receive their individual noisy signals of the true value. This noisy signal improves with time, and as $t \to \infty$, traders have perfect information. Thus, the evolution of the process is similar to that observed in a real prediction market, where information becomes available to traders gradually, and they become more certain about the true value over time.

3.2.2 Solving for the random walk probability

The true probabilities V_{LR} (resp. V_{TB}) of the ball hitting the right (resp. bottom) edge before it hits the left (resp. top) edge can be computed analytically in terms of p_{LR}, p_{TB}, x_0, and y_0.

The movement of the ball on the grid is governed by a random walk in each dimension. Let us consider the horizontal dimension (the analysis remains the same for the vertical dimension).

The ball moves a single step right with a probability p, and a single step left with a probability $1-p$. At the beginning of the experiment, and each time the ball falls off either side of the grid, the position of the ball is set to x_0. The size of the grid is $[-S, S]$. The true value under this model, is given by the probability of the number of times the ball falls off the right edge before it falls off the left edge. This is similar to the Gambler's Ruin problem described by Feller [12]. We trace out its derivation below.

Let this probability of the ball falling off the right edge before the left edge, given that the ball starts at a position of n, be given by P_n. Then P_n can be described as follows,

$$P_n = \begin{cases} 0 & \text{if } x = -S \\ 1 & \text{if } x = +S \\ pP_{n+1} + (1-p)P_{n-1} & \text{otherwise} \end{cases}$$

Re-ordering the terms of the recurrence relation and substituting $P_n = x^n$ in order to solve the difference equations, yields the following.

$$\begin{aligned} pP_{n+1} - P_n + (1-p)P_{n-1} &= 0 \\ px^{n+1} - x^n + (1-p)x^{n-1} &= 0 \\ px^2 - x + (1-p) &= 0 \end{aligned}$$

Solving the quadratic equation gives the roots $\frac{1-p}{p}$, 1. The probability P_n is therefore given by a linear combination of the roots.

$$\begin{aligned} P_n &= A\left(\frac{1-p}{p}\right)^n + B(1)^n \\ &= A\left(\frac{1-p}{p}\right)^n + B \end{aligned}$$

Using the boundary conditions, $P_{-S} = 0$ and $P_{+S} = 1$, we get two equations in the two co-efficients A and B. Solving,

$$\begin{aligned} A &= \left(\left(\frac{p}{1-p}\right)^S - \left(\frac{p}{1-p}\right)^{-S}\right)^{-1} \\ B &= -A\left(\frac{p}{1-p}\right)^S \end{aligned}$$

Substituting these values back into the equation for P_n, and evaluating for $n = x_0$, we get

$$P_{x_0} = \frac{\lambda^{S-x_0} - \lambda^{2S}}{1 - \lambda^{2S}}$$

where $\lambda = \frac{p}{1-p}$.

Therefore the true values for each of the two markets, Top-Bottom and Left-Right, are as follows.

$$V_{LR} = \frac{\lambda_{LR}^{S-x_0} - \lambda_{LR}^{2S}}{1 - \lambda_{LR}^{2S}}$$

$$V_{TB} = \frac{\lambda_{TB}^{S-y_0} - \lambda_{TB}^{2S}}{1 - \lambda_{TB}^{2S}}$$

$$\lambda = \frac{p}{1-p}$$

3.2.3 Experiments

For each of the experiments conducted, there were two markets set up, *Top Bottom* and *Left Right*, i.e. one along each dimension, with one of the two operating under the LMSR market maker and the other under the Information Based market maker with fixed values of α the stock block size (set to 10), and the window size of trade history used in the adaptive model. The random walk was run for a fixed period of time (usually around 10 minutes), with each time step (when the ball moves) occurring at intervals of 200 milliseconds. The parameters governing the random walk were set as $p_{LR} = p_{TB} = p$, and $x_0 = y_0 = z$, in order to ensure that both markets were completely symmetric in nature. In each experiment, the subjects traded within the same room. In all but one experiment (common information), the traders would view their own individual realization of the random walk on their computer screens as a random walk. As subjects were among students enrolled in the 'Computational Finance' and 'E-Commerce, Social Networks and Collective Wisdom' courses, they were all familiar with basic trading concepts (including for example, selling short). The final liquidation values of the markets were in general the true analytical probabilities computed for that experiment based on the parameters p (probability), z (starting location). In one of the experiments, the liquidation value was based on the final realized values of the random walk, so that traders' information would eventually be completely correct. Experiments can also include shocks to the true value of the market. This can be done by altering either one of the random walk parameters, to reflect different types of shocks. These shocks become apparent to the traders either through uncharacteristic changes to the counter values (of the number of times the ball hits the edges) if for example, p is changed, or through visual cues if z is changed. Due to the nature of our experimental design though (the position of the ball shifting every 200ms), visual cues are less noticeable, and traders have to rely on the change in the counter values to infer whether a shock has occurred.

3.3 Experimental Results

We ran six experiments under different settings, and differing values of the parameters governing the random walk as well as the adaptive capacity of the information based market maker. Experiments 1 and 2 were conducted with participants from the Computational Finance course. In both of these, traders were incentivized with a \$15 value gift-certificate for the trader with the largest return, \$10 for the second largest, and three \$5 value certificates for the three next best returns. The remainder of the data was collected from an educational deployment amongst participants from the E-commerce, Social Networks and Collective Wisdom course. (They received no explicit incentive other than extra credit points towards the course grade!) (With respect to the duration of each experiment, note that each time step in the random walk occurred at intervals of 200 milliseconds, or at a rate of 300 time steps a minute)

Experiment Setup

The LMSR based market maker was configured with the loss parameter b set to 125. The information-based market maker was configured to begin with belief $\mu_t = 50$, $\sigma_t = 12$, and estimation of the trader noise given by $\sigma_\epsilon = 5$. The window of trade history for the adaptive mechanism was set to 5 for the first two experiments and to 10 for the later three.

	p	S	z	V
Equilibrium	0.600	4	-1	0.7322
CommonInformationShock	0.600	2	-1	0.5846
jump to	0.600	2	+1	0.1231
LimitedInformation	0.764	4	-3	0.6912
Equilibrium(2)	0.533	4	-3	0.1897
Equilibrium(3)	0.866	4	-3	0.8453
IndividualInformationShock	0.826	4	-3	0.7890
jump to	0.516	4	-2	0.7890

Table 3.1: Configuration parameters for the live experiments

In the first two experiments traders started with $100,000$ units of currency, and were allowed to take both long and short positions in each market. In the three subsequent experiments, traders began with $10,000$ units of currency, and received 200 shares within each of the Top-Bottom and Left-Right markets. In addition they were restricted from taking short positions in any market.

Table 3.1 summarizes the individual parameters of the random walk for each experiment. (For experiments other than the first two, the random walk parameters were chosen as follows: a coin was flipped in order to determine if there would be a jump or not. If a jump occurs, it does at a randomly generated time between 3 and 7 minutes. In each experiment the traders had no knowledge of whether the market would sustain a shock in the underlying parameters.)

Experiment 1: Equilibrium

The purpose of the first experiment was to analyze the evolution of the mean prices of the market makers without any shocks to the underlying parameters. The experiment was run for a duration of 10 minutes, and each trader viewed their own realization of the random walk and therefore their own information and price dynamics based on the parameters governing the random walk which other traders were seeing. Figure 3.3 shows the convergence of the spot price (i.e. the mean price for an infinitesimal number of shares of the stock) over the course of the experiment, and the change in the average spreads.

Although there are initial fluctuations in the mean belief of the ZPI market maker, it converges quite well to the true mean price, as does the LMSR market maker. (The sharp spike in the mean price of the ZPIMM was due to a rogue, irrational trader choosing to buy at a price of 100)

Experiment 2: Market Shock under common information

Traders in this experiment traded under common information as all traders were viewing the same realization of the random walk. The market makers underlying the Top-Bottom and Left-Right markets were switched from the previous experiment, i.e. the LMSR market was now operating the Left-Right market, while the ZPIMM was the market maker for the Top-Bottom market. This was done in order to ensure that there was no bias along the horizontal axis versus the vertical, based on the price dynamics of the previous experiment. (This interchanging of market makers in consecutive experiments was also done for the remainder of the human subject experiments described). The payoffs for this experiment were changed from the previous experiment in that the ratio used was now the number of times the ball fell off the top edge before the bottom, and the final payoff of the market was based not on the analytical solution but on the final ratio in the realization of the random walk. The random walk was set to run for a period of 10 minutes, where halfway (i.e. at the 5 minute mark) the market received a shock with a change in the reset position of the ball. Since the shock occurs halfway through the experiment and random walk spends 50% of its time in each state, the expected value of the market (i.e. the fraction of time the ball hits the right edge of the grid before the left edge), is the average, $V = 0.3538$. The markets were liquidated at the actual realized ratios for the Top-Bottom (0.3370) and Left-Right (0.3756) markets. By doing so, the information available to the traders gradually becomes completely correct, so the market maker is then trading with perfectly informed traders.

As can be seen in the price convergence in 3.4(a) the ZPIMM adapts well to the market shock and is quick to adjust its mean belief. While the LMSR mean price also moves towards the new true value, it displays characteristic fluctuations throughout the duration of the experiment. As seen in the experiments in simulation, the variance in the information based market maker's belief, and hence its spreads,

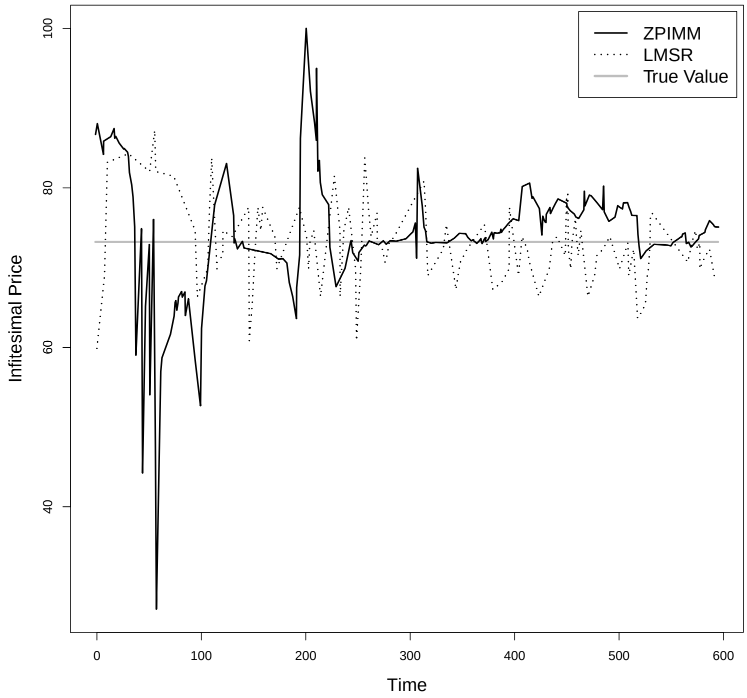

(a) Price convergence over time

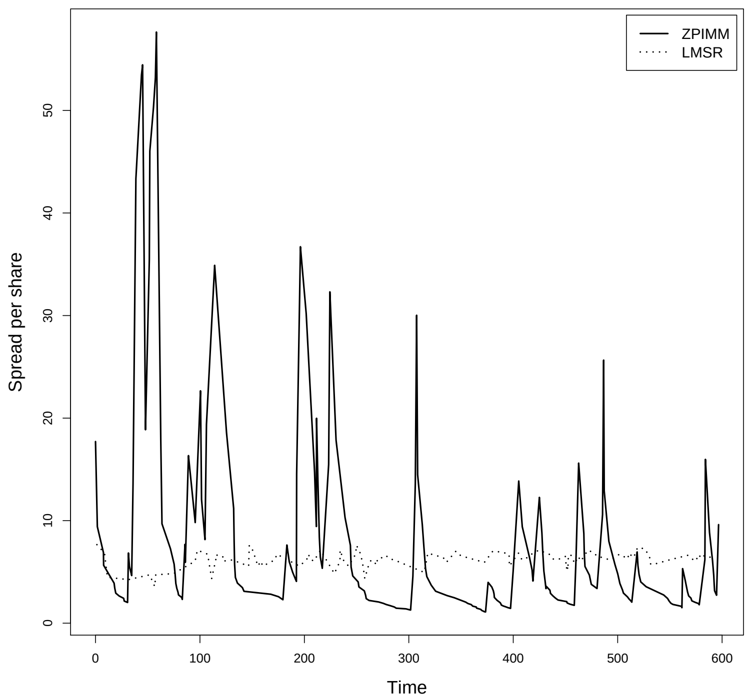

(b) Change in average spreads over time

Figure 3.3: Experiment 1: Equilibrium

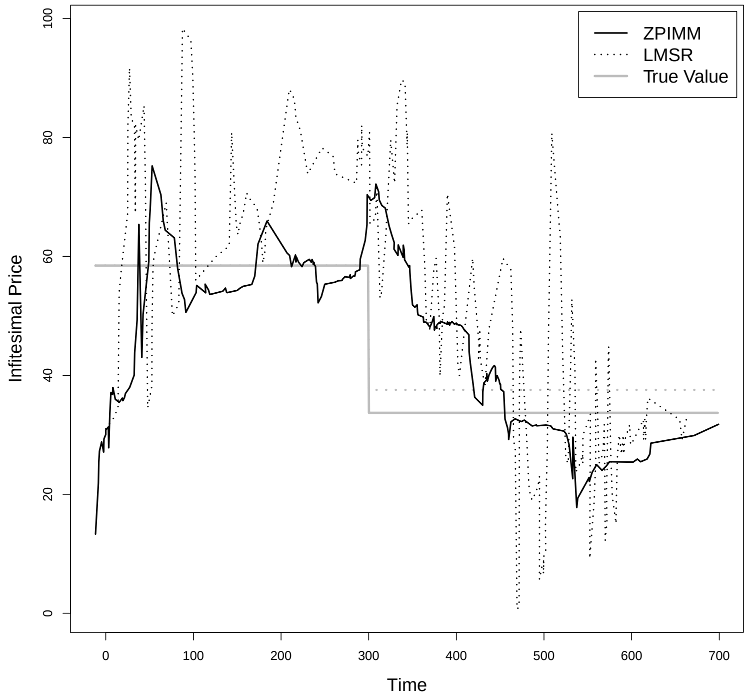

(a) Price convergence over time

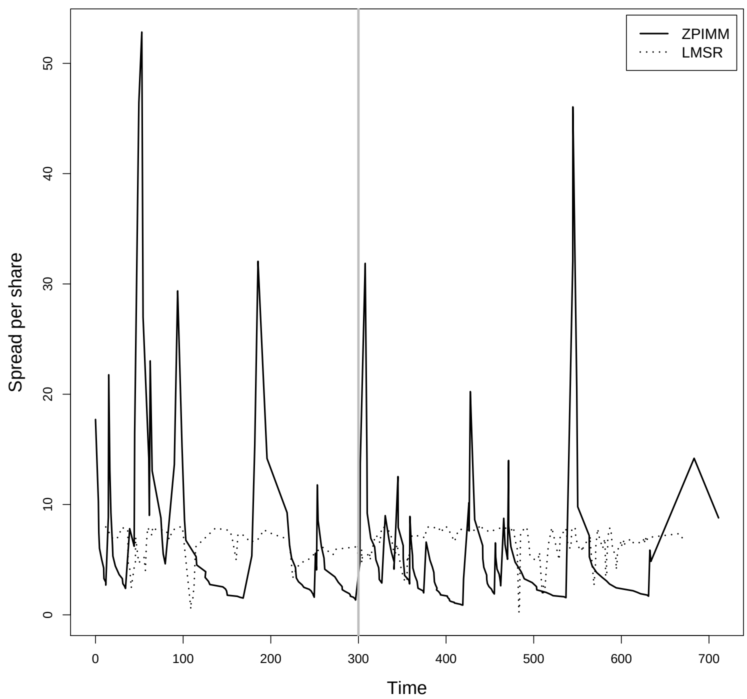

(b) Change in average spreads over time

Figure 3.4: Experiment 2: Market Shock under common information

decrease over time. This behavior is modified with the addition of the adaptive mechanism which allows for changes to the belief variance in order to better react to market shocks. The spikes in the spreads of the ZPIMM 3.4(b) therefore, indicate points in time where the adaptive mechanism increased the belief variance and with it the spreads, of the market maker.

Experiment 3: Equilibrium under limited information

While in the previous two experiments, traders could view an individual or common realization of the random walk for the entire duration of trading (10 minutes), in this experiment the random walk ran only for a period of 2 minutes, while trading continued for the entire 10 minutes. With only 2 minutes of the random walk, or 600 steps taken by the ball on the grid, each trader had a smaller and restricted set of information by themselves. The objective of the experiment therefore was to observe the evolution of the mean prices under conditions of individually limited, but collectively sufficient (with 17 participating traders), information required for the mean belief to converge to the true value.

Experiment 4: Equilibrium (2)

This experiment was similar to the previous one in its setup, and also had no shock to the market value during the trading process. However as Figure 3.6(a) shows, both market makers performed poorly in terms of price discovery. This demonstrates an example where inconsistent trades caused the market maker to be misled during the first half of the trading period. While both market makers moved towards the true value during the latter half, the prices had not converged before the close of trading. Due to this the ZPI market maker ended up with a greater loss as compared to the LMSR market maker, a contrast from the rest of the trading experiments.

Experiment 5: Equilibrium (3)

The window of trade history for the first two experiments was set to a size of 5. A smaller window corresponds to a quicker reaction by the ZPIMM by doubling the belief variance of the market maker. Even a small set of improbable trades

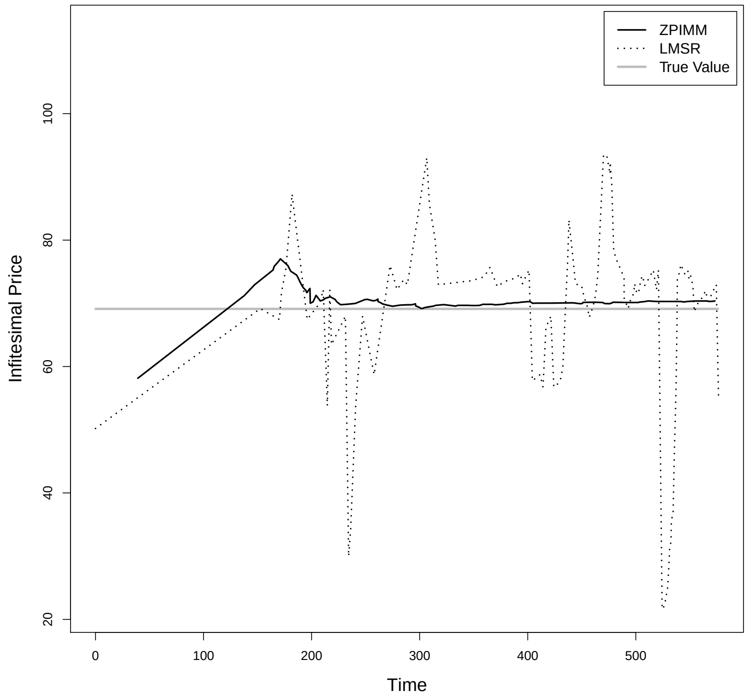

(a) Price convergence over time

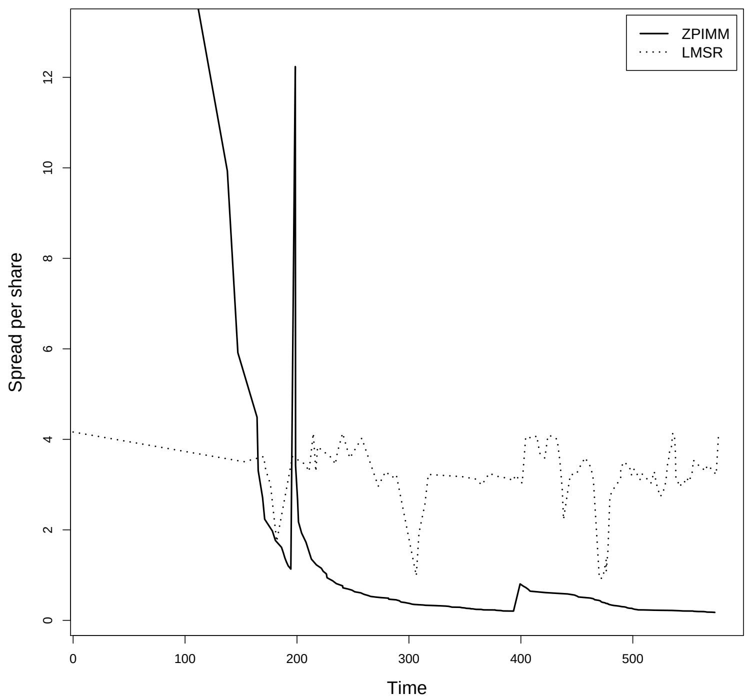

(b) Change in average spreads over time

Figure 3.5: Experiment 3: Equilibrium under limited information

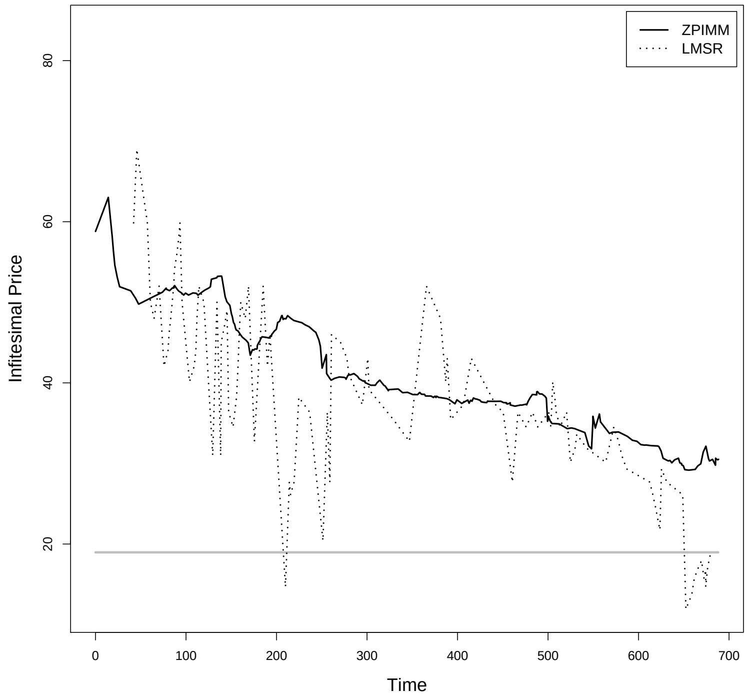

(a) Price convergence over time

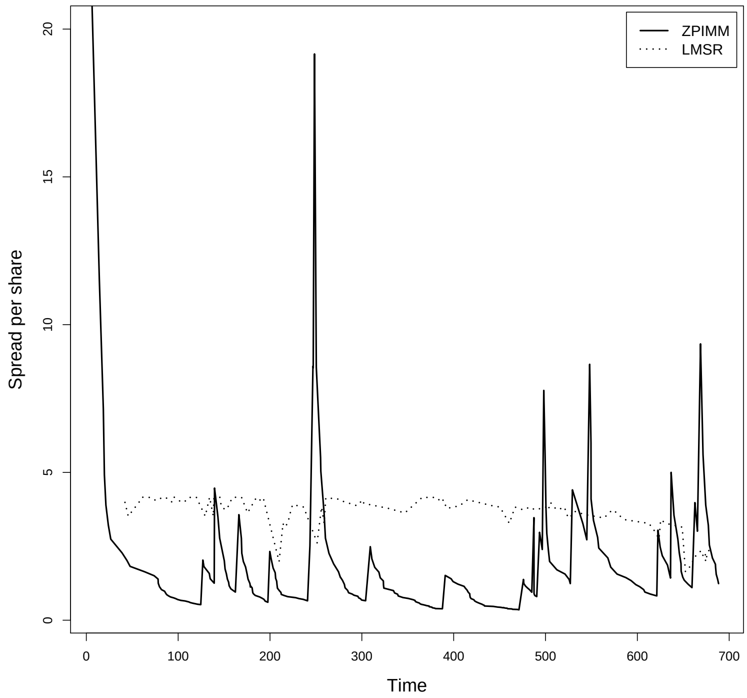

(b) Change in average spreads over time

Figure 3.6: Experiment 4: Equilibrium(2)

(with respect to the market maker's current belief) therefore is capable of causing a sharp rise in the spreads. This results in frequent fluctuations in the mean belief. Increasing the window to 10 reduces the oscillatory behavior of the mean and spikes in the spreads.

As seen in Figure 3.6 the price curve under the ZPIMM exhibits little fluctuation, with fewer changes in the spreads over time, in comparison to the equilibrium in Experiment 1.

Experiment 6: Market Shock under individual information

As with the previous two experiments, the size of the window of trade history was set to 10. After a duration of 5 minutes and 27 seconds, the market underwent a shock with the probability p and reset position on the grid z, both undergoing a change. The market shock therefore appeared visually in terms of the reset position each time the ball hit either boundary of the grid, and also in terms of the ratio of the number of times the ball fell off the right (resp. bottom) edge over the left (resp. top) edge. The final payoff was to be based on the final true value of the market (i.e. $V = 0.300$) and not the actual realized values. Since the period before the jump covered about half of the 10-minute trading time, the counts of the number of times the ball fell off the edges of the grid reflected a ratio (of reaching the right edge over the left edge for example) of about 0.600. Most traders therefore mistakenly assumed that the payoff of the market would be based on this value of $V' = 0.600$. Figure 3.7 indicates both values (the true mean and the trader belief) along with the mean prices and spread curves. As with the previous market shock experiment, the ZPIMM is quick to adapt to the market shock and exhibits little oscillation in mean belief thereafter. The bid-ask spreads under ZPIMM too exhibit few shocks in the curve once the market maker has learned the true value and reached a regime of low spreads.

Results

Table 3.2 provide a summary of the comparison of the LMSR and ZPIMM market makers in terms of the number of *confirmed* buy or sell trades executed in each experiment. In each experiment both markets received a comparable number of

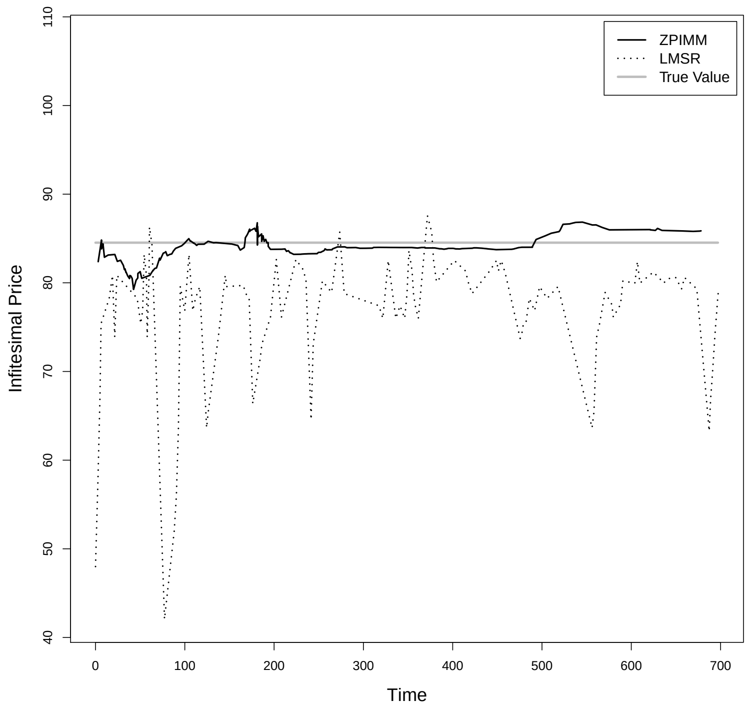

(a) Price convergence over time

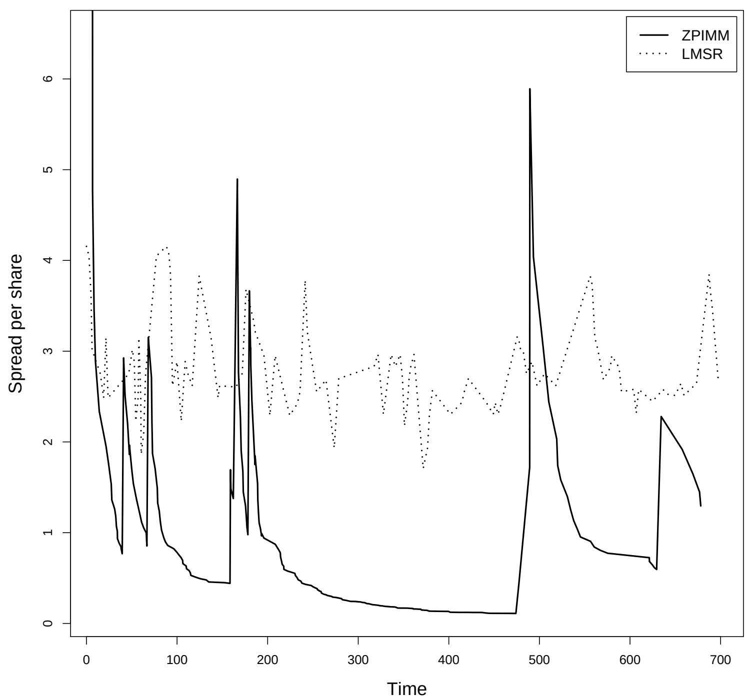

(b) Change in average spreads over time

Figure 3.7: Experiment 5: Equilibrium(3)

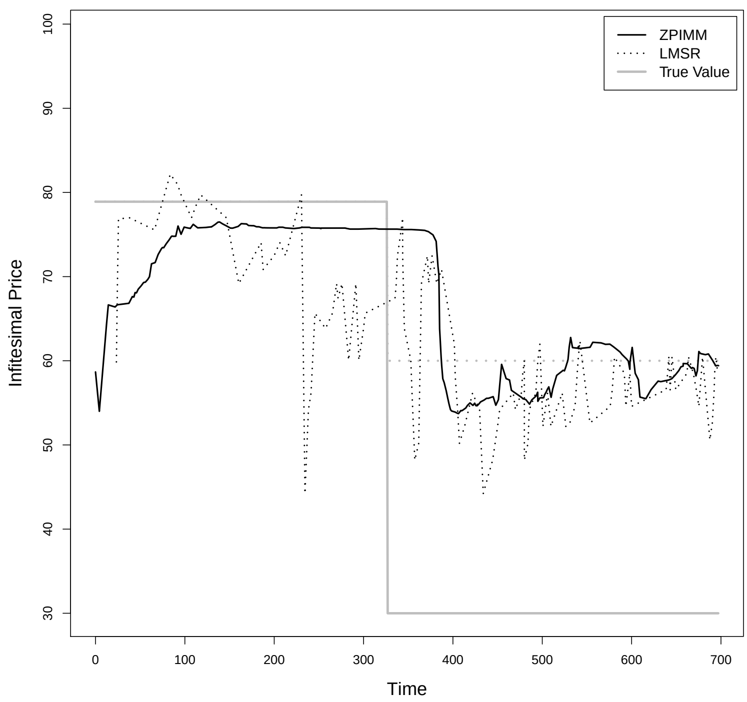

(a) Price convergence over time

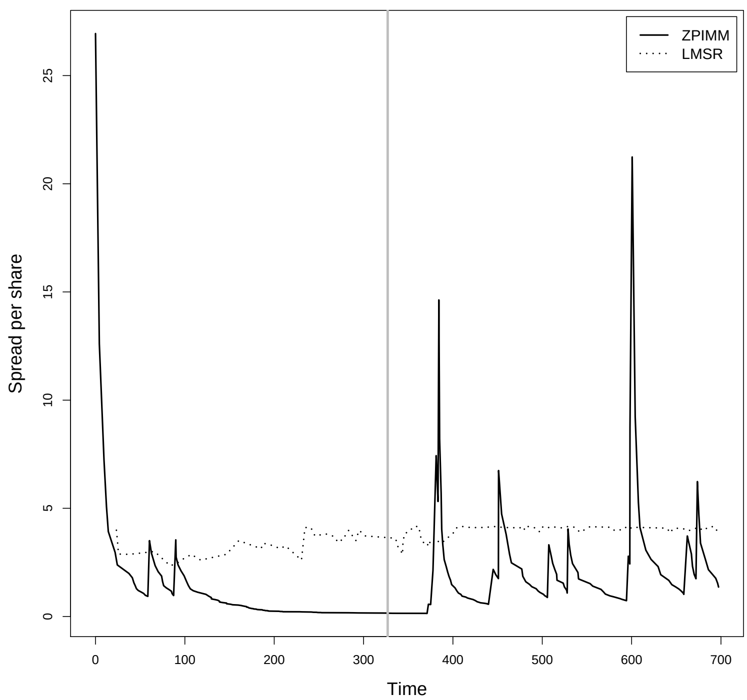

(b) Change in average spreads over time

Figure 3.8: Experiment 6: Market Shock under individual information

	# Buy/Sell		# Traders
	LMSR	ZPIMM	
Equilibrium	119	146	11
CommonInformationShock	213	131	9
LimitedInformation	116	63	17
Equilibrium(2)	109	113	17
Equilibrium(3)	118	85	17
IndividualInformationShock	113	121	17

Table 3.2: Number of participating traders and executed trades

	Mean		Median	
	LMSR	ZPIMM	LMSR	ZPIMM
Equilibrium	26.26	47.23	20	25
CommonInformationShock	46.24	57.79	20	40
LimitedInformation	32.34	60.59	10	30
Equilibrium(2)	30.09	36.72	20	20
Equilibrium(3)	27.75	42.81	20	10
IndividualInformationShock	21.97	38.88	10	20

Table 3.3: Mean and median quantities among confirmed (executed) trades

confirmed orders. Also traders had no knowledge of the underlying market makers, which in turn were alternated between the Top-Bottom and Left-Right markets. Hence there was no apparent bias in the competing markets in terms of the trading activity.

Table 3.3 compares the nature of the quantities within the confirmed trades of each market. The mean quantity size traded over all markets over all experiments is 39.06. The spread calculations which follow therefore include the difference of the bid and ask price per share for a quantity of 40 shares.

From the experiments conducted, the profit earned by the ZPIMM vastly exceeds that of the LMSR market maker in all but one. The ZPIMM profit in the first equilibrium experiment is a little exaggerated and therefore misleading due to a single rogue trader, as mentioned before. The LMSR based market maker suffers a substantial loss in each experiment. In the fourth experiment however, the ZPI market maker performs worse in terms of profit. This occurred possibly due to the

	Profit		Spread	
	LMSR	ZPIMM	LMSR	ZPIMM
Equilibrium	-1350.12	47231.77	3.12	4.04
CommonInformationShock	-1510.89	8972.50	3.06	3.21
LimitedInformation	-1602.14	4083.95	1.61	0.49
Equilibrium(2)	-2619.07	-10588.86	1.81	0.95
Equilibrium(3)	-3168.55	9134.58	1.42	0.51
IndividualInformationShock	-92.29	20226.44	1.89	0.89

Table 3.4: Performance of ZPIMM and the LMSR MM in live trading experiments

	RMSD		RMSDeq	
	LMSR	ZPIMM	LMSR	ZPIMM
Equilibrium	4.92	7.66	3.73	3.01
CommonInformationShock	20.67	15.98	16.51	6.76
LimitedInformation	14.43	2.15	14.56	0.93
Equilibrium(2)	21.67	23.13	14.82	17.05
Equilibrium(3)	11.18	1.5	8.15	1
IndividualInformationShock	8.87	6.47	6.88	6.6

Table 3.5: Convergence characteristics of ZPIMM and the LMSR MM in live trading

market maker being misled a number of times during the initial part of the trading period.

As in the simulation experiments, one can compare the accuracy of the price discovery process by means of the root mean squared deviation between the market maker's infinitesimal price and the true value of the market. The quantity $RMSD_{eq}$ is the same metric evaluated once the market maker converges - in computing these values we consider the time period which begins halfway between the last change to the true value of the market and the end of trading, till the end of the trading period. The LMSR market maker suffers from characteristic fluctuations in the spot price even after it has attained equilibrium. The ZPI market maker provides a tighter belief once it has converged.

The data gathered from experiments conducted and from an educational deployment of the trading platform, help validate several properties of the Zero-profit

Information based market maker. It provided an environment where the traders were completely unconstrained and both market makers could be compared simultaneously and in a symmetric manner. We see that in general the ZPI market maker provides a tight belief over the true value learned and also earns a profit in the process of doing so.

CHAPTER 4
Conclusion

4.1 Summary

The major contribution of this thesis is the extension of an information based market maker, towards applying it within a practical prediction market setting. The Zero-profit market maker was presented by Das [8] and built upon the information theoretic model of Glosten and Milgrom [13]. We extend the market maker to incorporate pricing of arbitrary quantities of shares, and compare its characteristics with the LMSR market maker, the *de facto* standard in current commercial prediction market platforms under simulated trading experiments.

The Zero-profit Information based market maker (ZPI MM) exhibits faster convergence towards the true value. While the original zero-profit market maker makes zero profit in expectation, the ZPI MM in general earns a profit when the unit quantity assumption is relaxed. The LMSR based market maker on the other hand suffers a loss during the course of each experiment. The loss incurred is often considered as a cost towards providing liquidity in the market. However this cost also makes it unsuitable for use in real-money markets, especially where a set of possible outcomes can require the creation of an exponential number of markets. Although the LMSR loss parameter does provided a bound on the worst case loss, this parameter must often be set in advance before the market maker has any knowledge of the level of trading or liquidity. Parameters to the ZPI market maker include the initial level of uncertainty and an estimate of the noise variance in traders' signals. We show that the characteristics of the ZPI MM are robust to mis-estimation of this noise variance. In addition, when the demand function of traders is known to be linear in the difference of their valuation and the spot price, profits can be controlled based on the selected stock block size. An added advantage of the information based market maker is that it provides a belief distribution over the inferred true value, rather than a single spot price. This allows the market creator to gauge the degree of convergence based on the market maker's uncertainty at any point of time, and make a decision to close the market accordingly or allow trading till the desired level of convergence is attained.

We also describe an adaptive mechanism which utilizes a window of trade history in order to respond to an unlikely sequence of trades. When the sequence is more likely under an increased level of uncertainty, the ZPI market maker increases the variance of its belief. An increase in the size of the window leads to slower adaptivity but tighter convergence during periods of certainty.

Lastly, we describe a set of human subject experiments which provide real world validation of several characteristics observed in simulation. The experimental design allows for the simultaneous comparison of both the ZPI and LMSR market makers. Under cases of limited, common and individual knowledge, and under conditions of shocks to the market, the ZPI MM obtains a greater profit, and provides tighter convergence to the true value.

4.2 Discussion

While the ZPI market maker in general exhibits better liquidity and price discovery characteristics as compared to the LMSR market maker, the latter provides a theoretical bound on the loss that it can occur over the course of the market. Experiment 4 also provided an instance where the ZPI MM was misled in the initial period of trading and was unable to learn the true value by the end of trading, suffering a greater loss than the LMSR market maker in the process. An open issue therefore is analyzing its performance in markets of extreme volatility.

In general this falls under the issue of trader manipulation. The subjects in the live experiments were completely unconstrained. Traders were allowed to query the market maker repeatedly and for an arbitrary quantity of shares, before they decided to execute an order. Indeed, some of the trades which misled the ZPI market maker might have been in an attempt to exploit its behavior, although the existence of any such advantage was not obvious to us.

The deployment of the trading platform among human subjects also raised the issue of offering the right payment schemes which offered the best incentive to trade, and which might also improve the accuracy of predictions. Gneezy and Rustichini

[14] find however, that offering monetary compensation does not always produce an improvement. Intrinsic motivational factors other than profit maximization are found to powerfully influence the behavior of participants in experiments [18]. Christiansen [5] also finds there can be sufficient incentives to encourage participation in prediction markets without requiring the market operator to subsidize traders with the opportunity cost of the time they contribute to the experiment. In addition they show empirically that even markets with very small number of participants maintain their ability to accurately estimate probabilities.

Addressing the specific issue of incentive schemes in prediction markets, Luckner and Weinhardt [19] compare the accuracy of prediction markets under fixed incentive (a fixed amount paid to all participants), rank order (payments to the top three for example), and performance compatible (payments to all traders, linearly linked to performance) schemes, and find the rank order payments to outperform the others. This is also the the scheme we followed when compensating traders in Experiments 1 and 2.

In addition to the issue of picking the right incentives to motivate traders to participate and reveal their knowledge, the underlying model of how traders react to the information presented to them also remains unexplored. In the experimental design for the live experiments, each trader observes an individual or collective signal at each time step (in our experiments, every $200ms$). For each trader, analyzing the frequency of trades as well as the quantities traded can help refine the underlying model of trader behavior.

Lastly, our assumed model of a market allowed only for pure market orders which were fulfilled by the market maker. A more sophisticated model which incorporates limit orders and potentially the existence of competing market makers, and their influence on liquidity and accuracy of predictions, still remains an issue largely unaddressed.

REFERENCES

[1] Michael Abramowicz. The hidden beauty of the quadratic market scoring rule: A uniform liquidity market maker, with variations. *Journal of Prediction Markets*, 1(2):111–125, 2007.

[2] J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research. *Handbook of Experimental Economic Results*, pages 486–515, 2001.

[3] Joyce E. Berg, Forrest D. Nelson, and Thomas A. Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300, 2008.

[4] Y. Chen and D.M. Pennock. A utility framework for bounded-loss market makers. In *Uncertainty in Artificial Intelligence (UAI)*, pages 49–56, 2007.

[5] Jed D. Christiansen. Prediction markets: Practical experiments in small markets and behaviours observed. *Journal of Prediction Markets*, 1(1):17–41, February 2007.

[6] Bo Cowgill, Justin Wolfers, and Eric Zitzewitz. Using prediction markets to track information flows: Evidence from google. In *Auctions, Market Mechanisms and Their Applications (AMMA)*, page 3, 2009.

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

[8] Sanmay Das. The effects of market-making on price dynamics. In *Autonomous Agents and Multiagent Systems (AAMAS)*, May 2008.

[9] Sanmay Das and Malik Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. In *Neural Information Processing Systems (NIPS)*, pages 361–368, 2008.

[10] Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417, May 1970.

[11] Doyne J. Farmer, Paolo Patelli, and Ilija I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences (PNAS)*, 102(11):2254–2259, 2005.

[12] William Feller. *An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition*. Wiley, January 1968.

[13] L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.

[14] Uri Gneezy and Aldo Rustichini. Pay enough or don't pay at all. *The Quarterly Journal of Economics*, 115(3):791–810, 2000.

[15] S. Goel, D.M. Reeves, D.J. Watts, and D.M. Pennock. Prediction without markets. In *11th ACM Conference on Electronic Commerce (EC)*, 2010.

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

[17] Robin Hanson. On market maker functions. *Journal of Prediction Markets*, 3(1):61–63, April 2009.

[18] George Loewenstein. Experimental economics from the vantage-point of behavioural economics. *The Economic Journal*, 109(453):F25–F34, 1999.

[19] S. Luckner and C. Weinhardt. How to Pay Traders in Information Markets? Results from a Field Experiment. *Journal of Prediction Markets*, 1(2):147–156, 2007.

[20] P. Milgrom and N. Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

[21] A. Othman and T. Sandholm. Automated market making in the large: The gates hillman prediction market. In *11th ACM Conference on Electronic Commerce (EC)*, 2010.

[22] A. Othman, T. Sandholm, D.M. Pennock, and D.M. Reeves. A practical liquidity-sensitive automated market maker. In *11th ACM Conference on Electronic Commerce (EC)*, 2010.

[23] D. Pennock and R. Sami. Computational aspects of prediction markets. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[24] Charles R. Plott and Kay-Yut Chen. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Papers 1131, California Institute of Technology, Division of the Humanities and Social Sciences, March 2002.

[25] E. Servan-Schreiber, J. Wolfers, D.M. Pennock, and B. Galebach. Prediction markets: does money matter? *Electronic Markets*, 14(3):243–251, 2004.

[26] Robert E Verrecchia. Information acquisition in a noisy rational expectations economy. *Econometrica*, 50(6):1415–30, November 1982.

[27] J. Wolfers and E. Zitzewitz. Five open questions about prediction markets. Stanford GSB Working Paper, 2004.

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

[29] Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. CEPR Discussion Papers 5676, C.E.P.R. Discussion Papers, May 2006.

[30] Justin Wolfers and Eric Zitzewitz. Prediction markets in theory and practice. IZA Discussion Papers 1991, Institute for the Study of Labor (IZA), March 2006.