

A Multi-Agent System for Analyzing the Effect of Information on Prediction Markets

Janyl Jumadinova, Prithviraj Dasgupta

April 7, 2010

Abstract Prediction markets have been shown to be a useful tool for forecasting the outcome of future events by aggregating public opinion about the event's outcome. In this paper, we investigate an important aspect of prediction markets - the effect of different information related parameters on the behavior of the traders in the market. We have developed a multi-agent based system that incorporates different information-related aspects including the arrival rate of information, the reliability of information, the penetration or accessibility of information among the different trades and the perception or impact of information by the traders. We have also performed extensive simulations of our agent-based prediction market for analyzing the effect of information related parameters on the traders' behaviors expressed through their trading prices. Our results provide a better understanding of the effect of information on the operation of prediction markets and on the strategies used by the prediction market's traders. (150 words)

Keywords Agent based modeling and simulation, prediction market, effect of information.

1 Introduction

Over the past few years prediction markets have emerged as an effective mechanism to predict the possible outcome of popular events such as presidential elections, performance of movies at the box office, outcome of sporting events, etc. A prediction market consists of human traders that bet their money on the possible future outcome of real-world events. The amount of money betted by a trader on a particular event depends on the trader's current belief about the outcome of the event. The aggregate value of the monetary bets made by different traders on an event dynamically determines the price of future bets related to the event. Prediction markets are considered to be an efficient aggregation mechanism for public opinion on the event because the dynamic price fluctuations of the bets related to an event is claimed to be an indicator of the public opinion or belief about the outcome of the event. Prediction markets were initially introduced as social research tools for aggregating the opinions of a large number of people on the future outcome of imminent events. The following success of prediction markets as an effective aggregator of public opinion has led to their adoption in various domains ranging from academic research¹ to commercial betting markets for popular events such as sporting events and Hollywood movies^{2,3} and predicting the performance or sales of products by software companies⁴. Recently, several researchers have also attempted to understand and explain the behavior of prediction markets by building analytical models^{5,6,7,8,9}. In addition to these approaches, the effect of information on prediction markets is another crucial factor that affects the behavior of the traders in the market. Information about an event that the traders receive affects their belief values about the outcome of an event, influences the prices corresponding to the event and finally determines the outcome of the event. Therefore, it makes sense to analyze the behavior of the traders in response to different information-related parameters in a prediction market.

In this paper, we have developed an agent-based prediction market consisting of multiple events and multiple trading agents. The multi-agent implementation of a prediction market allows us to easily analyze and verify the trading agents' behavior while

varying different market and agent related internal parameters of the prediction market, as well as external parameters related to the information about events arriving at the market. To build our multi-agent system, we have used modeling parameters obtained from various sources such as existing analytical models of financial markets, empirical evidence and data from real prediction markets, and agent utility and belief theory. We have used our multi-agent prediction market to analyze the effect of different information-related aspects including the arrival rate of the information, the reliability of the information, the penetration or accessibility of the information among the different trades and the perception or impact of the information by the traders. We have provided extensive empirical results using data sets from the Iowa Electronic Marketplace to understand how different values of information-related parameters impact the belief values and utilities of trading agents. We have also compared the utility maximizing strategy used by the trading agents in our prediction market with another pricing strategy used in prediction markets called zero intelligence (ZI) and showed that agents are able to set prices more accurately and obtain higher utilities when they use the available market information in the form of past prices, consider their own risk attitudes, and consider the impact of information about an event while setting their prices. The rest of the paper is organized as follows. We begin by describing previous work on prediction markets in Section 2. We then describe the operation of a prediction market using software agents in Section 3. Section 4 discusses the performance of our agent-based prediction market on data obtained from the Iowa Electronic Marketplace¹ and shows the effect of various market-based and agent-based parameters on the prediction market and on the trading agents' behavior. In Section 5 we discuss and interpret our results, and, finally, in section 6 we conclude.

2 Related Work

Prediction markets were started in 1988 at the Iowa Electronic Marketplace¹ to investigate whether betting on the outcome of geo-political events (e.g. outcome of presidential elections, possible outcome of international political or military crises, etc.) using real

money could elicit more accurate information about the event's outcome than regular polls. Following the success of prediction markets in eliciting information about events' outcomes, several other prediction markets have been started that trade on events using either real or virtual money. For example, the Hollywood Stock Exchange³ is used to predict performance of movies at the box-office, Tradesports¹⁰ is mostly used to predict the outcome of sports related events, while BetFair² is used to predict the outcome of sports or finance-related events. Recently, several companies such as Google⁴ and Yahoo's Yootles¹¹ ran prediction markets to analyze different technology trends such as information processing practices in organizations and product management strategies. Wolfers and Zitzewitz give an excellent overview of prediction markets¹². Continuous double auctions (CDA) have been used most frequently as a model for analyzing prediction markets¹². In one of the earliest attempts at analyzing prediction markets, Manski⁸ had a negative result showing that the predictions of event outcomes made in prediction markets do not closely correspond to the actual probability beliefs of the traders in the market unless the probability is near either 0 or 1. However, Gjerstad⁷ extended Manski's model to show that predictions are typically very close to the mean belief of the traders in the market if the distribution of beliefs is smooth and risk aversion is modeled into the traders' behavior. Wolfers and Zitzewitz⁹ have shown that prices of binary prediction markets are closely related to the actual frequencies of events in the real world. Feigenbaum⁶ and Chen⁵ have modeled the prediction market as a Shapley-Shubik market game and theoretically analyzed the aggregation function and convergence in prediction markets. Chen *et. al.*¹³ characterized the uncertainty of market participants' private information by incorporating aggregate uncertainty in their market model. However, both these models consider traders that are risk-neutral, myopic and truthful. Hanson¹⁴ developed a market scoring rule that is used to reward traders for making and improving a prediction about the outcome of an event. He further showed how any proper scoring rule can serve as an automated market maker. Hanson¹⁵ used market scoring rules in combinatorial information markets, where a probability distribution over all combinations of values of many variables is estimated.

Dimitrov and Sami¹⁶ have also studied the effect of non-myopic revelations by trading agents in a prediction market and concluded that myopic strategies are almost never optimal in the market with the non-myopic traders and there is a need for discounting. Recently, researchers have started using software agent mediators in prediction markets. Das¹⁷ studied the effect of specialized agents called market-makers which behave as intermediaries to absorb price shocks in the market. Das empirically studied different market-making strategies and concludes that a heuristic strategy that adds a random value to zero-profit market-makers improves the profits in the markets. Othman¹⁸ used Manski's model of prediction markets⁸ and analyzed trading agents' belief functions. He concluded that zero-intelligence agents can perform comparably with agents having more sophisticated behavior.

Most of the previous work on prediction markets provides theoretical models for the behavior of the human traders in a prediction market and analyzes the effect of those behaviors on the market. Our objective in this paper is to obtain insights into the behavior of traders in a prediction market in response to different parameters related to the information arriving at the prediction market. To achieve this, we use software trading agents to model the human traders in a prediction market. The behavior pattern of these trading agents are modeled using human trading patterns observed in financial markets. We develop an analytical model of our agent-based prediction market including different market-related parameters that affect the behavior of the trading agents. To analyze the behavior of our agent-based prediction market, we have performed extensive empirical simulations to quantify the effect of information-related parameters on the prices predicted by our trading agents and on the utilities received by them. We have also compared the results from our agent-based market with those obtained in the Iowa Electronic Marketplace while using data sets corresponding to the same events in both markets. Finally, we have compared the utility maximizing strategy used by our agents with a zero-intelligence (ZI) strategy that has been used formerly for strategic pricing in prediction markets.

3 Prediction Market Preliminaries

A prediction market is run by a company such as BetFair² or an institution such as the University of Iowa¹ and consists of a set of N traders and a set of E future events whose outcome has not yet been determined. Each event $e \in E$ is created on day $d_{e,0}$, lasts for a duration or lifetime of τ_e days, and the event's outcome is determined at the end of its lifetime. The outcome of each event $e \in E$ is considered as a binary variable with the outcome being 1 if the event happens and the outcome being 0 if it does not. Each event e is associated with an information set $I_{e,d}$ on day d which accumulates over time as new information about the event emerges from sources external to the market. The expected probability that an event $e \in E$ will happen (that is, e will have an outcome = 1) at the end of its duration varies dynamically as new information about the event becomes available to the market and as the traders update their beliefs about the outcome of an event. A *security* is a contract related to a specific event in the prediction market between the traders and the company or institution that runs the prediction market. Securities can be purchased or sold by traders at any time during the lifetime of the security's event. A single event can have multiple securities associated with it. Traders can purchase or sell one or more of the securities for each event at a time. The traders that hold a security when the lifetime of its associated event ends receive a payment based on the outcome of the event.

The market clearing rules of a prediction market are similar to the continuous double auction (CDA) protocol. Each event e is associated with two prices in the market - a *bid* price $\pi_{e,d}^{buy}$ at which traders can purchase a security of the event and a *ask* price $\pi_{e,d}^{sell}$ at which traders can sell a security of the event during day d . Both of these prices are normalized to the range of $[0, 1]$ and traders calculate the current market price of an event e during day d , $\pi_{e,d}$, as the average of the current bid and ask prices of the event during the day. That is, $\pi_{e,d} = 0.5 \times (\pi_{e,d}^{buy} + \pi_{e,d}^{sell})$. Traders then interpret the current market price $\pi_{e,d}$ for the event e on day d as the expected probability of the event e happening (that is, its outcome being 1).

A trader $n \in N$ observes the information set $I_{e,d}$ associated with event e on day d and decides to place an order to buy the security associated with event e at a price $\pi_{n,e,d}^{buy}$ or to sell the security associated with event e at a price $\pi_{n,e,d}^{sell}$. A *market order* for trading a security contains the identifier and number of units of the security the trader wishes to buy or sell. Market orders are executed instantly in the market, but the price of $\pi_{n,e,d}^{buy}$ (for purchase orders) or $\pi_{n,e,d}^{sell}$ (for sell orders) is not guaranteed to trader n for buying or selling the security. Instead, if a trader wishes to purchase the security, it has to pay a price of $\pi_{e,d}^{buy}$, corresponding to the maximum price that has been offered by any trader for purchasing a security related to the event during the day d . $\pi_{e,d}^{buy}$ is called the *bid (buy) price* of the event and is given by $\pi_{e,d}^{buy} = \max_{n \in N} \pi_{n,e,d}^{buy}$. Similarly, when a trader wishes to sell the security, it has to pay a price of $\pi_{e,d}^{sell}$, where $\pi_{e,d}^{sell} = \min_{n \in N} \pi_{n,e,d}^{sell}$. $\pi_{e,d}^{sell}$ is called the *ask (sell) price* of the event and corresponds to the minimum price that a trader has offered for selling a security related to the event during day d . In contrast to market orders, traders could also place *limit orders* where the orders are entered into the order book. Each limit order contains the the identifier and number of units of the security the trader wishes to buy or sell, along with the price at which it wishes to buy or sell. The order book is arranged in ascending order of sell prices and descending order of buy prices of securities. A trade is made for those securities for which the buy prices are greater than the sell prices and the number of units is less than or equal to number of units offered for sell. Traders are allowed to either buy or sell (but not both) securities corresponding to an event during a day. The difference between the bid and ask prices of an event during a day is called the event's *spread* for that day. The company running the prediction market usually absorbs this spread. Without loss of generality, we only consider market orders in our agent-based prediction market.

A security expires when the event e associated with it happens at the end of the event's duration τ_e . At this point the outcome of the event has just been determined and all traders are notified of the event's outcome. The traders that currently hold the security then get paid \$1 if the event happens with an outcome of 1, or, they do not

Market	Parameters
N	Number of trading agents in the market
E	Number of events in the market
H	Number of past days in history
$\pi_{e,d}$	Current market price of the event e during day d
$\pi_{e,d}^{sell}$	Ask price of the event e on day d
$\pi_{e,d}^{buy}$	Bid price of the event e on day d
τ_e	Event e 's duration
System	Parameters
ρ	Rate of a new event generation
$\zeta_{i,e}$	Impact factor of new information i about an event e of the security for an event e
$Poisson(\lambda_e)$	Poisson distribution with a new information arrival rate λ_e for an event e
D	Number of days over which the prediction market runs
Agent	Parameters
$w_{n,d}$	Wealth level of trading agent n on day d
$u_{n,d}$	Trading agent n 's utility on day d
θ_n	Risk coefficient of agent n
$Beta(\alpha, \beta)$	Belief beta-distribution parameters
$b_{n,e,d}$	Belief of trading agent n about event e on day d
$\mu_{n,e,d}$	Agent n 's weight factor in its beliefs for event e on day d
$r_{n,e}$	Agent n 's information reliability for event e
$q_{n,e,d}$	Actual quantity of securities for an event e which agent n holds on day d
$Q_{n,e,d}$	Optimum quantity of securities for an event e which agent n needs to hold to maximize its utility during day d
$a_{n,e}$	Agent n 's information availability for event e

Table 1 Parameters used in our agent-based prediction market.

get paid anything and lose the money they had spent on buying the security if the event happens with an outcome of 0.

A prediction market lasts for D days, starting on day $d = d_0$ and ending on day $d = d_{D-1}$. On each day, a trader makes a decision to buy some securities related to ongoing events in the market, or to sell or hold some of the securities the trader has already purchased.

3.1 Agent-based Prediction Market

We have used a multi-agent prediction market to analyze its behavior in response to different market parameters and trader behavior. In our prediction market, each human

trader is modeled as a software agent called a trading agent that embodies the behavior of a human trader. The different parameters used in our prediction market to define the market characteristics and specify the trading agents behavior are shown in Table 1 and described below.

We assume that events are generated in our prediction market at a rate ρ , where ρ follows a uniform distribution. New information about an event arrives at discrete points in time during the lifetime of the event. Following a study of information arrival in stock markets¹⁹, we assume that the rate at which new event information arrives follows a Poisson distribution with a rate parameter of λ_e .

3.1.1 Trading Agents: Functionality and Parameters

One of the most important parameters in a prediction market is the perception of the probability of occurrence of an event by the human traders. This parameter affects the dynamic values of the prices of securities corresponding to the different events in the market as well as the quantities of different securities each trader buys or sells. To model the perception of event occurrence probabilities, each trading agent n in our prediction market uses a private belief $b_{n,e} \in [0, 1]$ that corresponds to the probability that the event e will occur with an outcome = 1. Existing models of prediction markets^{7,8,18} have empirically verified that this belief value is drawn from a Beta distribution. Following these studies, we have assumed in our prediction market that a trading agent n 's initial belief about the outcome of an event being 1 is drawn independently from a continuous beta distribution $Beta(\alpha, \beta)$, where α and β are positive shape parameters that determine the shape and skew of the beta distribution's curve. There are two cases when the belief of an agent is dynamically updated. First case is when the market price changes and the second case is when a new information about the event becomes available to the trading agent. For the first case, the belief update function used by our trading agents is based on Gjerstad's prediction market model⁷ where, during each day, every time a market price of the security for an event e changes, trading agent n updates its beliefs as a weighted sum of the event e 's observed market

prices during that day and its own beliefs using the belief update Equation 1.

$$b_{n,e} \leftarrow \mu_{n,e,d} \cdot b_{n,e} + (1 - \mu_{n,e,d}) \cdot \pi_{e,d}, \quad (1)$$

where $\pi_{e,d}$ denotes the current market price of the security for event e during day d , $b_{n,e}$ is the belief of agent n for event e , and $\mu_{n,e,d} \in (0, 1)$ is agent n 's belief weight factor⁹ that represents the confidence of agent n in its current belief value about the outcome of an event e . At the end of each day, agent n finds out whether its buy or sell order for that day was executed or not and updates its belief weight factor using the equation:

$$\mu_{n,e,d} = \begin{cases} U[0.5, 1] , & \text{if buy/sell order executed;} \\ U[0, 0.5] , & \text{if buy/sell order wasn't executed.} \end{cases}$$

This ensures that if the buy/sell order was executed (not executed) an agent gives a higher (lower) preference to its own previous belief value while updating its belief using Equation 1. The second case of the belief update is discussed in Section 3.1.3 after we introduce the information modeling in our prediction market.

3.1.2 Dynamic Information Modeling in Prediction Markets

Empirical data from real prediction markets such as the Iowa Electronic Marketplace¹ show that the belief value of trading agents for an event e is significantly affected by new information that arrives in the market. Therefore, it makes sense to incorporate the effect of new information while updating agents' belief values. However, previous prediction market models and Equation 1 do not consider the effect of new event information on the belief values of a trading agent. To address this deficit, we have introduced a second step in the update of belief values that uses new parameters in our prediction market, as described below:

- a. *Information impact parameter* $\zeta_{i,e}$: In a prediction market, the price of the security related to an event is changed by the traders every time new information about the event becomes available. Also, different information about an event result in

different degrees of price change of the security²⁰. We have incorporated the effect of new information on the prices in the prediction market using an information impact parameter $\zeta_{i,e} \in [-0.5, 0.5]$ for information i related to event e . Larger values of $\zeta_{i,e}$ indicate high-impact information or unexpected news related to an event, which significantly affects the price of security related to the event. Smaller values of $\zeta_{i,e}$ indicate a normal-impact or low-impact information. These correspond to expected news about an event and don't affect prices significantly.

- b. *Information reliability parameter $r_{n,e}$* : Previous research on online financial markets²¹ have shown that different people rely on the new information in the market to different degrees and that an information reliability parameter can be used to represent the level of agent n 's reliance on information regarding event e . Following a study on the relation between trust and risk reported by Molm *et al.*²², we consider that the risk coefficient of an agent is affected by the reliability it associates with new information related to events. To model this, we have introduced an information reliability parameter, $r_{n,e}$ of an agent n for an event e . If a trading agent is a risk-taker, then its reliability parameter is high ($r_{n,e} \in U[0.5, 1.0]$), if it is risk-neutral, then $r_{n,e} = 0.5$, and if it is risk-averse, then its reliability parameter is low ($r_{n,e} \in U[0, 0.5]$). Thus, new information about an event e is perceived by the trading agents with different amounts of information reliability. Human traders in prediction markets have also been reported to dynamically update their reliability on event information, albeit by a small amount²³. To model this behavior, the trading agents in our prediction market dynamically update the reliability parameter at the end of each day using the following equation:

$$r_{n,e,d} = \begin{cases} U[0, 0.5] & , \text{ if } \theta_{n,e,d} > 0; \\ 0.5 & , \text{ if } \theta_{n,e,d} = 0; \\ U[0.5, 1.0] & , \text{ if } \theta_{n,e,d} < 0. \end{cases}$$

- c. *Information availability parameter $a_{n,e}$* : In real prediction markets, different traders have access to different amounts of information related to events. Traders with

access to a larger amount of information related to events are able to make informed decisions about the price of the securities related to events. We have modeled different levels of availability of information among the trading agents using an information availability parameter $a_{n,e} \in \{0, 0.5, 1\}$. $a_{n,e} = 1$ indicates that agent n has access to all available information about event e in the market, while $a_{n,e} = 0.5$ and $a_{n,e} = 0$ indicate limited and no access to information about event e respectively. Although the availability of information to human traders follows a more complex model, we have restricted the information availability parameter to three discrete values to get easily analyzable results of the effect of information availability on the trading agents, without loss of generality.

3.1.3 Agent Belief Revision based on new information

On each day in the prediction market, a trading agent has to make a decision about the securities it wants to purchase, sell or hold. The belief value about the outcome of the events related to the securities plays a pivotal role in making this decision. Therefore, every agent has to update its belief values for each event it is interested in based on the new information in the market related to those events. Because new information about events can arrive into the market from various sources with different reliabilities, an agent should adjust the impact of the information with the corresponding information reliability parameter. The reliability-weighted information impact parameter of an event then increases the belief value about the outcome of that event if its positive, and decreases the belief value if it is negative. Finally, to prevent rapid changes to the belief value, belief updates should also include the current belief value of an event. The belief update equation based on these factors that is used by a trading agent in our prediction market is given by:

$$b_{n,e} \leftarrow r_{n,e,d} \cdot (b_{n,e} - \bar{\zeta}_{i,e}^- \cdot \ln b_{n,e}) + (1 - r_{n,e,d}) \cdot b_{n,e} \quad (2)$$

where $\bar{\zeta}_{i,e} = \sum_i \zeta_{i,e}$ denotes the average value of the information impact parameter for all information related to event e by agent n , and $r_{n,e,d}$ denotes the information

reliability parameter agent n has for event e on day d . A sample surface plot that




Fig. 1 A variation of the new belief and the current belief with the information impact factor (a), with the information reliability (b).

illustrates the relationship between the new belief value and the current belief value and the information impact factor, $\bar{\zeta}_{i,e}$ in Equation 2 is given in the graph (a) of Figure 1 with $r_{n,e} = 0.5$. As shown in the graph a positive (negative) value of the information impact factor, increases (decreases) the new belief value from its previous value, while a zero value of the information impact factor does not change the belief value. The graph (b) in Figure 1 shows the relationship between the new belief value, the current belief value, and the information reliability parameter with $\bar{\zeta}_{i,e} = 0.2$. We can see that the agent with higher information reliability accounts more for the information impact, while the agent with smaller information reliability places more weight on its past belief.

3.1.4 Utility Maximization and Optimum Trading Quantity Calculation

The main problem facing a trading agent in a prediction market is to determine what quantities of each security to buy and sell on each day at the trading price of the security so that it can maximize its own utility or monetary gain. To achieve this, a trading agent dynamically incorporates new information about the event related to

each security and dynamically obtains updates of the current market price of each security to estimate the optimum quantity of the security it should buy or sell. When the prediction market starts on day d_0 , each trading agent n is provided with an initial wealth of $w_{n,d_0} > 0$. We consider each trading agent in our prediction market as a utility maximizer that buys or sells securities to maximize its wealth on each day. The instantaneous utility of agent n on day $d - 1$ is updated by the following equation:

$$\dot{u}_{n,d} = \sum_{e \in E} q_{e,d}^{sell} \cdot \pi_{e,d}^{sell} - \sum_{e \in E} q_{e,d}^{buy} \cdot \pi_{e,d}^{buy} \quad (3)$$

where, $q_{e,d}^{sell}$ is the number of securities related to an event e sold by agent n on day d , $q_{e,d}^{buy}$ is the number of securities related to an event e bought by agent n on day d , and $\pi_{e,d}^{sell}$ and $\pi_{e,d}^{buy}$ are respectively the ask and bid price of the security for event e during day d .

Equation 3 maintains the instantaneous utility of the agent n on day d depending on the transactions it has made on that day, however it does not account for its risk type. Previous research^{7,9} has shown that the correct behavior of prediction market models can be obtained only if the risk-taking and risk-averse behaviors of the human traders are considered while trading securities in prediction markets. Therefore, in our prediction market we adopt a constant relative risk averse (CRRA) utility function $\tilde{u}_{n,d}$ for agent n with a relative risk aversion coefficient. CRRA utility functions have been widely used to model risk behaviors. Relative risk aversion coefficient, θ_n , is used to classify trading agent n 's risk levels as follows. If $\theta_n > 0$, the agent n is risk-averse, if $\theta_n = 0$, the agent n is risk-neutral, and if $\theta_n < 0$, the agent n is risk-seeking. The trading agents' risk coefficients are normally distributed in our simulations. Following Gjerstad's trading agent utility model⁷, during each day a trading agent uses its instantaneous utility and its risk-taking coefficient to calculate its modified instantaneous utility for that day, using Equation 4.

$$\tilde{u}_{n,d}(\dot{u}_{n,d}, \theta_n) = \begin{cases} \frac{\dot{u}_{n,d}^{1-\theta_n}}{1-\theta_n} & , \text{ if } \theta_n \neq 1; \\ \ln(\dot{u}_{n,d}) & , \text{ if } \theta_n = 1. \end{cases} \quad (4)$$

The utility value for an agent can fluctuate considerably in successive days based on its trading pattern, and the price of securities in the market. Large fluctuations in the utility can in turn delay the convergence to equilibrium because it can cause large fluctuations in the number of securities being bought or sold in the market. To prevent these fluctuations we have used a daily weighted utility value for an agent's utility given by:

$$u_{n,d}(\tilde{u}_{n,d}) = \sum_{h=d}^{d-H} \delta_h \cdot \tilde{u}_{n,d}, \quad (5)$$

where δ_h is a discount factor that ensures that more recent utility levels have a greater influence on the current weighted utility, $\sum_{h=d}^{d-H} \delta_h = 1$, and H is the number of days considered from the past.

The final objective of a rational trading agent is to maximize the utility it receives from purchasing and selling securities in the prediction market. To achieve this objective, a trading agent has to calculate the optimum quantity of securities it should trade so that the resulting utility is maximized. Let $Q_{n,e,d}$ denote the optimum quantity of securities related to event e for trading agent n during day d . The trading agent's decision is to choose quantity to buy or sell for each security, so that the expected utility is maximized⁷. We have used a similar model in our paper with the expected utility given by Equation 6.

$$\begin{aligned} E_e(u_{n,d}(\theta_n, Q_{n,e,d})) &= \\ &= \frac{1}{1-\theta_n} b_{n,e,d} (u_{n,d} + Q_{n,e,d} \cdot (1-\pi_{e,d}))^{1-\theta_n} + \\ & (1-b_{n,e,d}) (u_{n,d} - Q_{n,e,d} \cdot \pi_{e,d})^{1-\theta_n}, \text{ if } \theta_n \neq 1 \\ &= b_{n,e,d} \cdot \ln(u_{n,d} + Q_{n,e,d} \cdot (1-\pi_{e,d})) + \\ & (1-b_{n,e,d}) \ln(u_{n,d} - Q_{n,e,d} \cdot \pi_{e,d}), \text{ if } \theta_n = 1 \end{aligned} \quad (6)$$

where $\pi_{e,d}$ is the current market price of the security for an event e on day d .

The optimum quantity $Q_{n,e,d}$ is obtained by taking the first-order derivative of the Equation 6, setting the resulting equation equal to 0, and solving it to obtain the following solution:

$$Q_{n,e,d}(\pi_{e,d}, b_{n,e,d}, \theta_n, u_{n,d}) = \frac{((1 - \pi_{e,d})^{\frac{1}{\theta_n}} b_{n,e,d}^{\frac{1}{\theta_n}} - (\pi_{e,d})^{\frac{1}{\theta_n}} (1 - b_{n,e,d})^{\frac{1}{\theta_n}}) u_{n,d}}{((1 - \pi_{e,d})(\pi_{e,d})^{\frac{1}{\theta_n}} (1 - b_{n,e,d})^{\frac{1}{\theta_n}} + \pi_{e,d}(1 - \pi_{e,d})^{\frac{1}{\theta_n}} b_{n,e,d}^{\frac{1}{\theta_n}})} \quad (7)$$

If $Q_{n,e,d} > q_{n,e,d}$, that is the quantity required to maximize utility on day d is more than the number of securities of event e that agent n currently has on day d , the agent places a buy order for purchasing $Q_{n,e,d} - q_{n,e,d}$ additional units of the security of event e . On the other hand, if $Q_{n,e,d} < q_{n,e,d}$, the agent places a sell order for selling $q_{n,e,d} - Q_{n,e,d}$ units of the security of event e . The agents set their prices corresponding to their current beliefs; that is, $\pi_{n,e,d}^{buy} = b_{n,e,d}$ if the trader is purchasing securities of event e and $\pi_{n,e,d}^{sell} = b_{n,e,d}$ if the trader is selling securities of event e during day d .

Figure 2 gives a flowchart of the algorithm used by a trading agent in our agent-based prediction market. Each trading agent n enters the prediction market with some initial wealth, a belief about the outcome of each event that is initially drawn from a Beta distribution, a risk coefficient, and, information reliability and information availability parameter values. First, the agent checks its information availability parameter $a_{n,e}$ and updates its beliefs about an event e (using Equation 2) with a probability corresponding to the value of $a_{n,e}$. Throughout the day, the trading agent n checks whether the market price for the security associated with an event e has been updated. When the market price is changed, the trading agent updates its beliefs using Equation 1. At the end of each day, the trading agent n updates its utility using Equation 5 and calculates $Q_{e,d}$, the number of units of security corresponding to event e it needs to maximize its utility. Based on $Q_{e,d}$ value it then makes a decision to buy, sell, or to hold a security associated with an event e and submits an order to the prediction market. If the order is accepted, the trading agent n updates its weighted utility using Equation 3. Finally, the trading agent's information reliability parameter is updated and the new day starts.

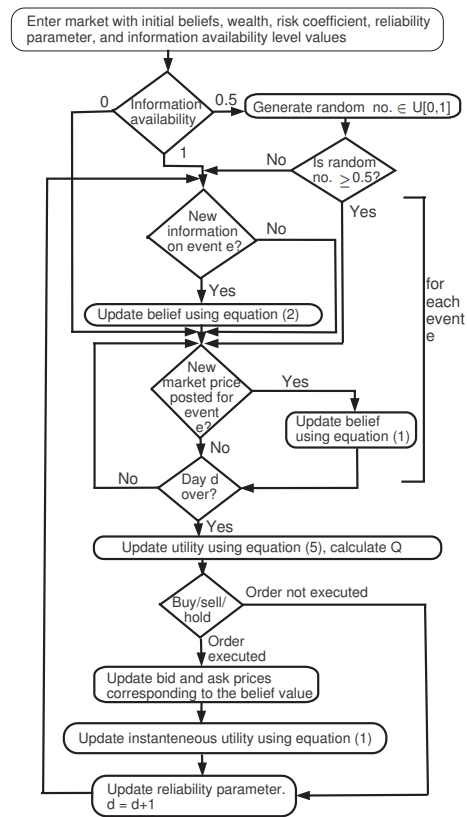


Fig. 2 A flowchart showing the operation of the trading agents in the prediction market.

4 Empirical Evaluation

We have verified the performance of our prediction market through several simulations while varying different market, event and agent related parameters. The first objective of our simulations is to clarify that our prediction market behaves in a manner similar to actual prediction markets with human traders. Secondly, we attempt to understand the dynamic behavior of trading agents' prices and utilities under different values of the information related parameters. For most of our simulations, we used data sets from the movie markets on IEM. A security in a movie market on IEM corresponds to a boolean event - whether the box office collections of a certain Hollywood movie will reach a predetermined dollar amount. The movie market securities are particularly suitable for

Price History

Movie_MINC

These markets are open continuously. Low/High/Average prices refer to prices beginning at 12:00:01am and ending at 11:59:59pm. LastPrice refers the price of the last trade before midnight. Today's prices are not included in this list.

Date	Contract	Units	\$Volume	LowPrice	HighPrice	AvgPrice	LastPrice
11/01/01	MINC180H	1	0.150	0.150	0.150	0.150	0.150
11/02/01	MINC180H	1	0.150	0.150	0.193	0.150	0.150
11/03/01	MINC180H	0	0.000	0.000	0.000		0.150
11/04/01	MINC180H	40	15.790	0.154	1.000	0.395	0.500
11/07/01	MINC180H	12	7.277	0.567	0.645	0.606	0.570
11/08/01	MINC180H	24	15.152	0.567	0.669	0.631	0.601
11/09/01	MINC180H	9	5.580	0.610	0.670	0.620	0.610
11/10/01	MINC180H	3	1.830	0.000	0.000	0.610	0.610
11/11/01	MINC180H	16	12.420	0.600	0.800	0.776	0.800
11/12/01	MINC180H	15	13.890	0.750	0.989	0.926	0.984
11/13/01	MINC180H	13	11.335	0.854	0.950	0.872	0.899
11/14/01	MINC180H	20	17.801	0.850	0.940	0.890	0.929
11/15/01	MINC180H	22	20.110	0.880	0.929	0.914	0.924
11/16/01	MINC180H	38	34.627	0.880	0.929	0.911	0.925
11/17/01	MINC180H	13	11.597	0.880	0.920	0.892	0.880
11/18/01	MINC180H	4	3.640	0.880	0.920	0.910	0.920
11/19/01	MINC180H	16	14.635	0.870	0.960	0.915	0.960
11/20/01	MINC180H	24	22.288	0.880	0.970	0.929	0.900
11/21/01	MINC180H	0	0.000	0.000	0.000		0.900
11/22/01	MINC180H	51	48.360	0.929	0.960	0.948	0.960
11/23/01	MINC180H	7	6.760	0.950	1.000	0.966	0.950
11/24/01	MINC180H	1	0.995	0.995	0.995	0.995	0.995
11/25/01	MINC180H	6	5.985	0.993	1.000	0.998	1.000
11/26/01	MINC180H	17	16.718	0.953	1.000	0.983	1.000
11/27/01	MINC180H	4	3.912	0.956	1.000	0.978	1.000
11/28/01	MINC180H	6	5.970	0.990	1.000	0.995	1.000
11/29/01	MINC180H	0	0.000	0.000	0.000		1.000

Fig. 3 A snapshot of the transaction history of the *MINC180H* security corresponding to the Monsters Inc. movie obtained from the Iowa Electronic Marketplace.

our agent-based prediction market because like the events in our market, the events corresponding to the different movies are not correlated with each other. The data set available from the IEM movie market had information on 14 securities consisting of the opening, average and closing prices of each security for each day of trading. A snapshot of one of the securities from this market is shown in Figure 3. Out of the 14 securities, 4 were discarded because they either had very low trading activity, or, a new security with a different outcome (usually a different dollar amount for box collections) related to the same movie was introduced before the original security expired. Most of the securities in the movie market expired after 30 days. One or more new securities were

Parameter	Value
d	2 sec
N	50, 100*, 300, 500, 1000
E	5, 10*, 15, 20, 25
H	10
p_e	$U[0, 1]$
ρ	$U[\frac{1}{30}, \frac{1}{60}]$
$\zeta_{i,e}$	$U[-0.5, 0.5]$
s_e	$U[0, 1]$
$Poisson(\lambda_e)$	$Poisson(\lambda_e = 3^*, 5, 10, 15, 20, 25, 30)$
$r_{n,e}$	$[0, 1]$ depending on θ_n value
w_{n,d_0}	$U[E , 2 E]$
b_{n,e,d_0}	$Beta(1, 1), (2, 2)^*, (7, 3), (3, 7)$
$a_{n,e}$	0, 0.5, 1*
φ_n	0 or 1*
θ_n	$U[-1.0, 1.0]$
μ_{n,d_0}	$U(0, 1)$

Table 2 Parameters used for simulation experiments. Default values are marked by asterisk.

Symbol of the security	Movie Name	Payoff condition to get \$1
<i>Sun50L</i>	Tears of the Sun	$50 \geq \text{BOR} \geq 40$
<i>Cat110L</i>	Cat in the Hat	$\text{BOR} \leq 110$
<i>Die120H</i>	Die Another Day	$\text{BOR} \geq 120$
<i>MINC180H</i>	Monsters Inc.	$\text{BOR} \geq 180$
<i>SH90L</i>	Sleepy Hollow	$90 \geq \text{BOR} \geq 70$
<i>TTI20H</i>	10 Things I Hate About You	$\text{BOR} \geq 20$
<i>MAT45H</i>	The Matrix	$45 \geq \text{BOR} \geq 35$
<i>WINE100H</i>	World is Not Enough	$\text{BOR} \geq 100$
<i>EOS70H</i>	Enemy of the State	$70 \geq \text{BOR} \geq 35$
<i>ISK50L</i>	I Still Know What You Did Last Summer	$\text{BOR} \geq 50$

Table 3 Description of the securities used as the data set in our simulations. BOR stands for the box office receipts of the movie in millions of US Dollars.

introduced into the movie market at an average interval of 30 days. The data from the movie market spans a period of 180 days. Each data set from IEM contains the aggregate price for each security at the end of each day. However, the IEM data does

not contain the individual bid and ask prices from each trader for each event during a day. The aggregation function used to calculate the market price is also not available in the public domain. In our simulations, we have used the daily closing prices from the IEM data set to compensate for the unavailability of the market prices. At the end of each day, trading agents are informed the closing price from the IEM data set. Agents then revise their beliefs for each security using Equation 1, while assuming that the closing price of the security is equal to the security’s market price. Table 3 gives the description of the securities we have used in our simulations. In our agent-based prediction market the values for market related parameters such as the duration of an event e , τ_e , the arrival rate of the events, ρ , and the number of days over which the market runs, D , were set to 30 or 60 days, once in 30 days, and 180 days respectively corresponding to the values of these parameters observed in the movie market. Table 2 gives the list of parameters and their values we have used in our simulations.

4.1 Model Comparison

For our first set of experiments, we compare the performance of our multi-agent prediction market with the model proposed by Gjerstad⁷. The data set for this experiment was taken from the 2008 Democratic Presidential Nomination Market in IEM¹. There are four securities in this market representing four different outcomes, i.e., Hillary Clinton, John Edwards, Barack Obama, or another candidate winning the democratic presidential nomination. This prediction market started with all four securities and ran for 400 days, which was also the duration of all the four securities. Therefore, in this prediction market we set the duration of an event e , τ_e , the arrival rate of the events, ρ , and the number of days over which the market runs, D , to 400 days, 0, and 400 days respectively. Figure 4 shows the difference between the actual market price and the predicted market price while using our agent-based prediction market (a) and the prediction market based on the Gjerstad’s model (b) for the security corresponding to Obama winning the nomination. The error bars show the confidence intervals over 10 runs. As shown in Figure 4, the difference between the predicted and actual




Fig. 4 The comparison of the simulated market prices with real market prices in our agent-based prediction market (a) and the prediction market based on Gjerstad's model (b). The solid line represents perfect prediction of market prices by agents.

prices using our agent-based prediction market is significantly lower (8%) than that reported in Gjerstad's model (35%). The more accurate price prediction in our prediction market can be attributed to our belief update mechanism that considers two additional parameters - the risk coefficient of each agent through the information reliability parameter $r_{n,e}$, as well as the information impact parameter $\zeta_{i,e}$, as compared to Gjerstad's model. These parameters incorporate the information flow across the agents in the market as well as the effect of the information of the agents, and, therefore, enable the agents to make a more informed decision about the belief related to the event in the market.

4.2 Effect of Agent Related Parameters

Effect of information arrival rate. For our next set of experiments under this category, we analyze the effect of different information arrival rates on the belief values of the agents. We consider three events, where each event has a different value of the rate parameter λ_e of the Poisson distribution, that characterizes the rate at which information related to the event arrives. For a Poisson rate parameter λ_e , the rate at which information arrives in the market is proportional to $\frac{1}{\lambda_e}$. A higher value of λ_e means that information related to the event arrives at a slower rate to the market and vice-versa. All other parameters related to the event and the agents are kept at their default




Fig. 5 Effect of the different rates of information arrival on (a) the market price and (b) on the utilities of the trading agents.

values. In prediction market agents' beliefs about an event are expressed through the aggregated market price of the event. We have therefore used the events' market prices to observe the effect of the information arrival rate on the agents' beliefs. To determine suitable values for the information arrival rate, $\frac{1}{\lambda_e}$, we analyzed the instances at which the price of the *SH90L* security in the movie market changed by more than 10%.

We observed that these price fluctuations happened at intervals lying between a lower bound of 3 days, an upper bound of 15 days, and an average of 10 days. Based on this observation, we used the minimum, maximum and intermediate values of the information arrival rate $\frac{1}{\lambda_e}$ at $\frac{1}{15}$, $\frac{1}{3}$ and $\frac{1}{10}$ respectively. We made successive runs of our agent-based prediction market with the information arrival rate set to each of these values of $\frac{1}{\lambda_e}$. Figure 5(a) shows the variation of the market prices for the different values of information arrival rates over the 30-day duration of the *SH90L* security. In Figure 6 we

$\frac{1}{\lambda_e}$	Mean Price	Variance in Price
$\frac{1}{3}$	0.49	0.08
$\frac{1}{10}$	0.52	0.1
$\frac{1}{15}$	0.54	0.12

Fig. 6 Mean and variance of the market prices under different $\frac{1}{\lambda_e}$ values.

have extracted the means and the variances of the market prices corresponding to different values of the information arrival rate, $\frac{1}{\lambda_e}$, from the results shown in Figure 5(a). The correlation coefficient using Pearson correlation between the information arrival rate $\frac{1}{\lambda_e}$ and the variance of the market prices is -0.92 . Based on the high negative

correlation coefficient value, we can infer that lower values of information arrival rate in prediction markets (or, in other words, higher value of λ_e) slightly increases the fluctuation of the security's price and vice-versa. This also agrees intuitively with the empirical evidence observed in real prediction markets - when information about an event arrives less frequently, trading agents (and human traders) can get disoriented about the outcome of the event, and this can lead to their erroneously predicting the probability of the event's outcome during its lifetime. This manifests itself in larger fluctuations in the security's price during the lifetime of the event. On the other hand, more frequently available information keeps trading agents focused on the possible outcome of the event and prevents large fluctuations in the market price.

Based on our analysis of the correlation between information arrival rates and fluctuation of market prices, it would make sense to test if different information arrival rates also impact the utilities obtained by the traders in a prediction market. We used the same experimental settings as in the previous experiment. However, the information arrival rate $\frac{1}{\lambda_e}$ was varied between $\frac{1}{2}$ and $\frac{1}{30}$. The average utility of a trading agent was calculated by averaging the utility over all agents and all days. Figure 5(b) illustrates the relationship between the information arrival rate and the average utility of a trading agent. The correlation coefficient between the information arrival rate and the agents' utility using Pearson's correlation is obtained as -0.98 . The proximity of the coefficient's absolute value to 1 implies that there is almost a linear relationship between information arrival rate and utility. The negative value of the coefficient indicates that a trading agent's utility decreases as information arrival rate increases and vice-versa. The results from this experiment can be analyzed intuitively as well. With higher information arrival rates, information about an event arrives more frequently and results in more up-to-date belief values for the event's security calculated using Equation 2. More up-to-date belief values in turn result in a more accurate value of $Q_{n,e,d}$, the quantity of security that the agent should trade (calculated using Equations 6 and 7) and consequently gives a higher utility to the agent. In summary, this set of experiments show that information arrival rates have a considerable impact on

the belief and utilities of trading agents, and therefore, affect the overall dynamics of a prediction market.

Effect of information impact. In a prediction market, the same information about an event can cause different traders to behave differently based on the traders' perceptions of the impact of that information. In our next set of experiments, we attempt to analyze the effect of different traders having different information impact parameters on the belief values of agents in the market. Our definition of the dependency between agent belief values and the information impact parameters is given in Equation 2, which is reproduced below:

$$b_{n,e,d} = r_{n,e,d} \cdot (b_{n,e,d-1} - \bar{\zeta}_{i,e} \cdot \ln b_{n,e,d-1}) + (1 - r_{n,e,d}) \cdot b_{n,e,d-1}$$

From this equation, we see that a positive(negative) value of the information impact factor causes an increase(decrease) in the belief value of an agent for a security, while a zero value of the impact factor does not change the belief value. This increases the utility the agent receives and results in more buy(sell) orders from the agent (using Equations 6 and 7) before the event expires. To verify the effect of different values of the information impact factor, we consider the effect of a positive-only impact fac-




Fig. 7 A graph showing the effect of the information impact parameter on the market prices.

tor ($\zeta_{i,e} \in U[0.1, 0.5]$), a negative-only impact factor ($\zeta_{i,e} \in U[-0.5, -0.1]$), and the impact factor consisting of mixed positive and negative values ($\zeta_{i,e} \in U[-0.3, 0.3]$).

In prediction markets, trading agents' beliefs about an event are expressed through the aggregated market price of the event. We have therefore used the real closing prices of the *MAT45H* security that has a duration of 25 days from the IEM movie data to analyze the effect of information impact factor on the agents' beliefs. For this set of experiments, each trading agent n is assumed to have fixed information reliability parameter $r_{n,e} = 0.5$ about

Mean of $\zeta_{i,e}$ distribution	Mean Price	Variance in Price
0.3	0.53	0.04
0	0.49	0.07
-0.3	0.45	0.05

Fig. 8 Mean and variance of the market prices under different $\zeta_{i,e}$ values.

each event e . All other parameters are retained at their default values shown in Table 2. In Figure 7, we report the results of the different values of the information impact parameter on the aggregate market prices of the agents over the 25-day duration of the security. Figure 8 shows the means and the variances of the market prices corresponding to different values of the information impact factor from the results shown in Figure 7. The Pearson's correlation coefficient between the means of the information impact factor values and the means of the market prices obtained under the corresponding $\zeta_{i,e}$ values is 0.98. A positive value of the correlation coefficient close to 1 shows that market prices are strongly correlated with the value of the information impact parameter. When the impact factor of the new information about an event has mixed positive and negative values, then it does not have a significant impact on the market prices as shown by the $U[-0.3, 0.3]$ curve in Figure 7.

Effect of information reliability. Our objective for the next set of experiments is to test if different information reliability levels have any effect on trading agents' beliefs. A trading agent expresses its beliefs about the outcome of an event through the ask and bid prices it sets. Therefore, we have one of these prices - the ask price to illustrate the relationship between the information reliability parameter and the agents'

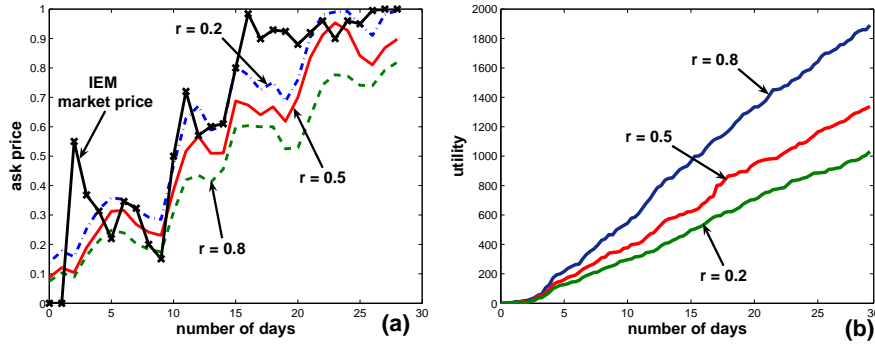


Fig. 9 (a) Ask prices of the traders and (b) the utilities of the traders for different information reliability levels.

beliefs. For this experiment we have used the closing prices from the *MINC180H* security that has a duration of 30 days in the IEM movie market. As mentioned in Section 3.1.1, we assume that an agent’s reliability for new information is directly correlated to its risk coefficient. To model the effect of different risk behaviors of the trading agents on the prediction market, we assume that one-third of the agent population is risk-taking with $r_{n,e} \in U[0.5, 1.0]$, another one-third of the population is risk neutral with $r_{n,e} = 0.5$, while the final one-third of the population is risk-averse with $r_{n,e} \in U[0, 0.5]$. The variation of the ask prices of securities for the different agent populations and the aggregated market price is shown in Figure 9(a). We have also shown the aggregated market price for the *MINC180H* security from the IEM movie market, which does not incorporate the information reliability parameter in the trading agents decision, to contrast the effect of adding the information reliability parameter. We observe that higher values of the information reliability parameter result in lower ask prices, as summarized in Figure 10. The correlation coefficient using Pearson correlation between the mean value of the trading agents’ ask prices and the information reliability parameter is calculated to be -0.99 . This implies a strong negative association indicating that the trading agents with a higher value of the information reliability param-

$r_{n,e}$	Mean	Variance
0.2	0.52	0.13
0.5	0.48	0.09
0.8	0.43	0.06

Fig. 10 Mean and variance of the market prices under different $r_{n,e}$ values.

eter and a lower risk coefficient set lower prices for their ask price. This is because in Equation 2, (belief update equation for agents) when an agent uses a higher value for the information reliability parameter $r_{n,e}$, its updated belief value incorporates the information impact factor $\zeta_{i,e}$ with the previous belief value. Since, the information impact factor also affects the aggregate price for the security in the market, considering this factor keeps the ask price of the agent closer to the market aggregate. The bid prices (not shown here) also follow a similar pattern. Figure 9(b) shows the average utilities of the same trading agent population for the different information reliability values. We observe that when the information is more reliable ($r_{n,e}$ has a higher value closer to 1.0), the utility of the trading agents is higher. This behavior can also be attributed to the fact that if the agent has a high value for its reliability parameter, the agent's belief function updates given in Equation 2 will include new information about an event through the information's impact parameter (ζ). These agents are able outperform by 30 – 45% the agents that rely mostly on the past market buying prices for updating their belief value instead of using the information's reliability.

Effect of information availability. In this set of experiments we set out to test if different information availability levels related to events have any effect on trading agents' beliefs and the agents' utilities. We have once again used agents' ask prices to illustrate the relationship between the information availability and the agents' beliefs because agents express their beliefs by setting their bid and ask prices. By analyzing ask prices, we can also see how different the information availability values lead to different prices being set by agents, which in turn leads to different values in utility. Different levels of availability of information among the trading agents can result in different trading decisions made by the trading agents and therefore in different utility. For example, a trading agent with all the information about an event available with it, can benefit from that information by making more informed trading decisions and get a higher utility. We have used the information availability parameter described in Section 3.1.2, $a_{n,e} = \{0, 0.5, 1\}$, to denote different degrees of information penetration

among the trading agents. We have again reported the results using the closing prices and event's duration from the *MINC180H* security of the IEM movie market. In our simulations, one third of the trading agent population has no access to the information about the event with $a_{n,e} = 0$, another third of the agent population has partial access to the information with $a_{n,e} = 0.5$, while the final third of the population has full access to the information with $a_{n,e} = 1$. The results of our simulations are shown in Figure 11.

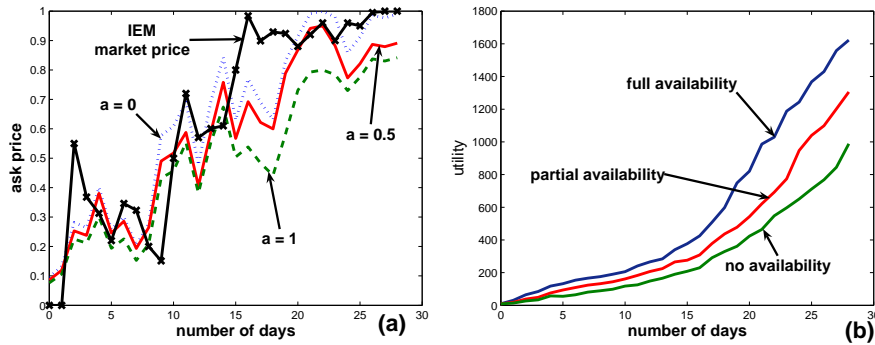


Fig. 11 Effect of the different levels of information availability on (a) the trading agents' ask prices and (b) their utilities.

From these results, we calculated the Pearson's correlation coefficient between the information availability parameter and the market price and found it to be -0.88 . This indicates that there is a strong negative association between the market price and the information availability parameter. This implies that the trading agents with more information (higher value of $a_{n,e}$) about the event are able to set the lowest ask prices for the event's security and purchase the security. Figure 11(b) shows that trading agents with full information availability are able to obtain higher utility than other agents. This behavior can be attributed to the fact that if $a_{n,e} = 1$, the agent's belief function update using Equation 2 is always triggered to include new information about the event e if available, as shown in the flowchart in Figure 2. This results in more up-to-date values of beliefs (using Equation 2) and ensures more accurate value of the quantity of security that the agent should trade (using Equations 6 and 7) and

culminates in higher utility to the trading agent. On the other hand, if $a_{n,e} = 0.5$, the agent’s belief function update using Equation 2 is triggered only 50% of the time, resulting in less utility for agent n . When $a_{n,e} = 0$, the agent n ’s belief function update using Equation 2 is never triggered, resulting in the least utility for agent n . Our result also agrees intuitively with the behavior of real traders in prediction markets - more information about an event enables a trader to make a more informed decision and thereby obtain higher utilities.

Effect of scalability. For our next set of experiments in this section, we have tested the scalability of our multi-agent based prediction market with the number of trading agents and the number of events. We have tested our prediction market with different combinations of 5, 10, 15, 20 and 25 events, and, 50, 100, 300, 500 and 1000 trading agents in the market. Table 4 shows the average utility for a trader at the end of 100 days. We observe that as the number of events in the market increases, the

Traders→/ Events↓	50	100	300	500	1000	Avg. over no. of agents
5	1480.5	1191.9	1071.0	1003.6	893.5	1059.1
10	1560.0	1265.0	1187.5	1089.6	965.2	1183.6
15	1757.16	1456.8	1299.2	1174.6	1006.5	1306.3
20	1791.1	1545.5	1424.5	1287.5	1104.4	1397.8
25	1843.3	1643.2	1534.5	1393.8	1275.2	1511.2
Avg. over no. of events	1631.5	1406.7	1295.0	1153.9	970.9	

Table 4 Utilities of one trader from the simulations with different number of traders and events.

traders’ utilities increase by 30% because they are able to trade in more securities. In a similar manner, when the number of trading agents in the market increases but the number of events is fixed, the average utility of a trader decreases by 40% because of the increased competition between traders in the market. In summary, as the prediction market becomes more populated in the number of events and the number of agents, the utilities of the trading agents increase correspondingly.

Agent Strategy Comparison. For our final set of experiments, we have compared the utility maximizing trading agents described in Section 3.1.1 with zero-

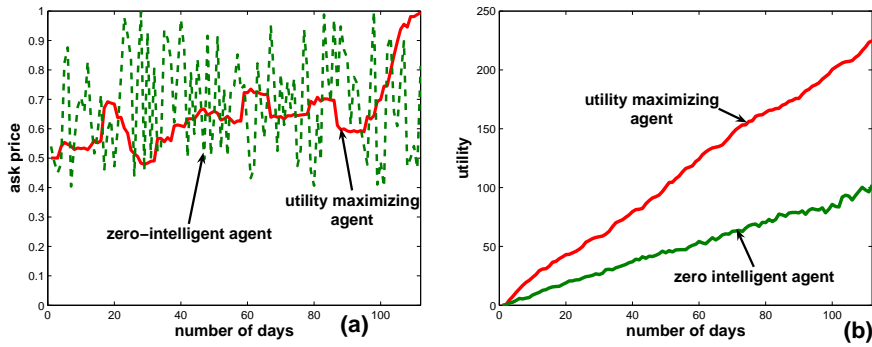


Fig. 12 The comparison of the utility maximizing and the zero-intelligent agent’s (a) ask prices and (b) utilities.

intelligence(ZI) trading agents¹⁸. To the best of our knowledge, ZI agents are the only strategic agents that have been simulated for continuous double auction-based prediction markets. Both our utility maximizing agents and the ZI agents draw the belief about the outcome of an event e on day d from the same distribution $b_{n,e,d}$ that is specified by a Beta distribution. However, ZI agents differ from our utility maximizing agents in two ways. First, instead of using utility maximization criterion in Section 3.1.4 (given in Equations 6 and 7) to determine the price and quantity of securities for trading, they decide to buy or sell the securities with equal probability. Consequently, ZI agents set a bid price at $U[0, b_{n,e,d}]$ or an ask price at $U[b_{n,e,d}, 1]$ for the event e ’s security on day d . Secondly, unlike our utility maximizing agents, ZI agents do not observe market prices, do not maintain a history of past trades, and do not react to any information about the event available in the market.

For this set of simulations, we have used the 2008 presidential winner-takes-all market data from IEM - the *DEM08.WTA* security that pays \$1 if the Democratic Party nominee receives the majority of popular votes cast for the two major parties in the 2008 U.S. Presidential election, \$0 otherwise. We analyzed the price data for the *DEM08.WTA* security manually and calculated the information arrival rate and the impact factor of the new information that was consistent with the price data of this security. Figure 12 shows the comparison of the ask prices and utilities between one

		Aggregate Price	Variance
<i>Sun50L</i>	Our	0.79	0.015
	ZI	0.82	0.03
<i>Cat110L</i>	Our	0.53	0.013
	ZI	0.57	0.041
<i>Die120H</i>	Our	0.89	0.018
	ZI	0.81	0.035
<i>MINC180H</i>	Our	0.71	0.009
	ZI	0.77	0.047
<i>SH90H</i>	Our	0.39	0.012
	ZI	0.51	0.039
<i>TTI20H</i>	Our	0.94	0.019
	ZI	0.89	0.04
<i>MAT45H</i>	Our	0.90	0.016
	ZI	0.85	0.036
<i>WINE100H</i>	Our	0.73	0.013
	ZI	0.79	0.043
<i>EOS100H</i>	Our	0.71	0.014
	ZI	0.74	0.037
<i>ISK50L</i>	Our	0.73	0.015
	ZI	0.78	0.042
		Total Utility	
Total Utility		Our	1407.85
		ZI	600.45

Table 5 Comparison of ZI agent with utility maximizing agent showing the aggregate prices and utilities obtained by these two agents for different securities in the IEM over a period of 100 days.

utility maximizing agent and one ZI agent. In Figure 12 (a) we observe that the variance of the ask prices determined by the utility maximizing agent is 0.01 and the variance of the ask prices set by the ZI agents is 0.08. The lower variance of the ask prices for our agents can be attributed to the fact that the updated belief values for a ZI agent are drawn from a uniform distribution that does not consider any available information about the events in the market or information from their past orders. In contrast, our utility maximizing agent is able to avoid large price fluctuations by using the available market information in the form of past prices, its own risk-taking behavior, and the information impact factor of an event, while setting its price. Therefore, the utility maximizing agent is able to set better predicted prices and obtain a higher utility than the ZI agent by an average of 58.5%. Table 5 displays the aggregate prices of different securities over 100 days obtained with the two agent strategies, the variance of those

prices, and the total utilities obtained by the utility maximizing agent and the ZI agent, for the 10 movie securities described in Table 3. Once again, we observe that the utility maximizing agents outperform the ZI agents for each security.

5 Discussion

The idea of using prediction markets as a form of collective intelligence has been gaining both research and commercial interest over the last few years. Over the last decade, prediction markets have matured from social research tools to commercially used predictors of event outcomes. Therefore, it is important to be able to analyze the dynamic behavior of prediction markets using realistic models of market and trader behavior. As illustrated by the experiments related to the price dynamics reported in Figures 4, 7(a) and 11(a) the prices in our agent-based prediction market are within 0.005 – 0.35% of the actual prices in the IEM prediction market for the same set of events in both markets. This shows that the behavior and dynamics of our agent-based prediction market closely resembles those of a real prediction market. Our simulation results have quantified the effect of information on the price dynamics and utilities of trading agents in a prediction market. suggest that the higher values of information arrival rates prevent large fluctuations in the market price as the trading agents receive new information about the event more frequently. This is also consistent with the observations in real prediction markets, the more information there is about an event in the market. The information impact parameter indicates the positivity, negativity or neutrality of the information about an event. Our results show that the positive impact factor of the new information about an event leads to increased market price and vice versa, provided that the information reliability parameter is constant. This is also intuitive from the real prediction markets - bad (negative) news related to an event causes the probability of an event happening to decrease while positive news about an event leads to an increase of the probability of an event's occurrence. From our simulations we observed that the trading agents with a higher value of the information reliability are able to obtain higher utilities. This allows us to conclude that a higher reliance of the

trading agents on the information in the market allows them to get more news about the event and make more informed decisions. We also observed that the trading agents with more information (higher value of information availability parameter) about the event are able to obtain higher utilities. This result suggests that the trading agents that have access to more information are able to use the available information to obtain higher utilities. This also consistent with the behavior of human traders in real prediction markets, where well-informed traders are able to make better decisions. Finally, we compared the performance of the utility maximizing agents used in our prediction market with ZI agents. Our simulation results showed that an utility maximizing agent is able to set better prices and obtain a higher utilities compared to a ZI agent. We inferred that this is due to the utility maximizing agent using and dynamically updating the information about events available in the market, while ZI agents do not use any such information. Overall, our results suggest that different aspects of the information about events in a prediction market have a significant impact of the prices, utilities and probabilities of the events in the market. These results can be used to obtain a better understanding of traders' behavior in a prediction market in response to information about events.

6 Conclusion and Future Work

In this paper, we have described an multi-agent based prediction market and used it to analyze the behavior of the trading agents in response to various information related parameters in the market. In the future, we plan to investigate the effect of the different utility maximization pricing strategies employed by the trading agents on the prediction market prices and on the utilities of the traders. This will allow us to study, test and compare the behavior of the trading agents in a competitive environment. Prediction markets allow one event's outcome to be contingent on another event's occurrence. Prediction markets can also be used to discover the existence of correlation between different events. In the future, we are also interested in introducing and investigating correlations between events in a prediction market. Finally, we are

interested in exploring another important direction in prediction markets - different revelation mechanisms and how to encourage truthful revelations by the trading agents. We envisage that using multi-agent systems to implement and understand the behavior of prediction markets along the directions outlined in this paper will lead to a better understanding and help to solve the issues and challenges related to prediction markets and distributed information aggregation.

Acknowledgements The authors would like to sincerely thank Professor Justin Wolfers of the University of Pennsylvania for providing us with the data²⁴ which was used for some of the experiments reported in this paper.

References

1. Iowa Electronic Marketplace URL: <http://biz.uiowa.edu/iem>
2. Betfair URL: <http://games.betfair.com>
3. Hollywood Stock Exchange URL: <http://www.hsx.com>
4. Cowgill B, Wolfers J, Zitzewitz E. Using Prediction Markets to Track Information Flows: Evidence from Google. Resource document <http://www.bocowgill.com/GooglePredictionMarketPaper.pdf>. Accessed 14 May 2009.
5. Chen Y. Predicting Uncertain Outcomes Using Information Markets: Trader Behavior and Information Aggregation. *New Mathematics and Natural Computation* 2006; 2(3):1-17.
6. Feigenbaum J, Fortnow L, Pennock D, Sami R. Computation in a Distributed Information Market. *Theoretical Computer Science* 2005; 343(1-2):114-132.
7. Gjerstad S. Risks, Aversions and Beliefs in Predictions Markets. mimeo, U. of Arizona 2005; RePEc:wpa:wuwpmi:0411002.
8. Manski C. Interpreting the Predictions of Prediction Markets. *Economic Letters* 2006; 91(3):425-429.
9. Wolfers J, Zitzewitz E. Interpreting Prediction Markets as Probabilities. NBER Working Paper No. 12200 2006.
10. Tradesports URL: <http://www.tradesports.com>
11. Yootles URL: www.yootles.com
12. Wolfers J, Zitzewitz E. Prediction Markets. *Journal of Economic Perspectives* 2004; 18(2):107-126.
13. Chen Y, Mullen T, Chu C. An In-Depth Analysis of Information Markets with Aggregate Uncertainty. *Electronic Commerce Research* 2006; 6(2):201-22.

-
14. Hanson R. Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation. *Journal of Prediction Markets* 2007; 1(1):3-15.
 15. Hanson R. Combinatorial Information Market Design. *Journal Information Systems Frontiers* 2003; 5(1):107-119.
 16. Dimitrov S, Sami R. Nonmyopic Strategies in Prediction Markets. *Proceedings of EC* 2008; 200-2009.
 17. Das S. The Effects of Market-making on Price Dynamics. *Proceedings of AAMAS* 2008; 887-894.
 18. Othman A. Zero Intelligence Agents in Prediction Markets. *Proceedings of AAMAS* 2008; 879-886.
 19. Maheu J, McCurdy T. News Arrival, Jump Dynamics, and Volatility Components for Individual Stock Returns. *The Journal Of Finance*, 2004; 59(2):755-793.
 20. Corallo E. The “news” effect on some financial variables. *LIUC Papers in Economics* 172 2005.
 21. Kollock P. The Production of Trust in Online Markets. *Advances in Group Processes* 1999; 16(1):99-123.
 22. Molm L, Schaefer D, Collett J. Fragile and Resilient Trust: Risk and Uncertainty in Negotiated and Reciprocal Exchange. *Sociological Theory* 2009; 27(1):1-32.
 23. Ramakrishnan R, Thakor A. Information Reliability and a Theory of Financial Intermediation. *Review of Economic Studies* 1984; 51(3):415-432.
 24. Wolfers J, Zitzewitz E. Five Open Questions for Prediction Markets. In *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, ed. Robert Hahn and Paul Tetlock, AEI-Brookings Joint Center, Washington D.C. 2005.