Algorithmic Finance 1 (2011) 157–181
DOI 10.3233/AF-2011-012
IOS Press

Algorithmic trading in the Iowa electronic markets

James E. Schmitz
Barclays Capital Inc., New York, NY, USA[1]
E-mail: james.schmitz@gmail.com

Abstract. The Iowa Electronic Markets are small, real-money financial markets designed to aggregate information about future events. The market microstructure of these markets is studied and a market making model is developed to provide liquidity for one set of securities offered by this exchange. A computer program was created to employ the market making model and profit from the market's inefficiencies. Using invested capital, the system traded 34% of the total market volume and achieved a Sharpe ratio of 9.9. This paper reveals the details of how this algorithmic trader worked to show how it functioned and the value it added to the Iowa Electronic Markets.

Keywords: Prediction markets, algorithmic trading, market making, Iowa electronic markets, market microstructure.

1. Introduction

The Iowa Electronic Markets (IEM) are small, real-money financial markets run by the Henry B. Tippie College of Business at the University of Iowa. The values of the securities that trade in these markets are linked to the outcome of real world events, such as political elections. The purpose of these markets is to serve as an educational tool for students to learn about financial markets and how information is priced into security prices. Researchers also use these markets to study market efficiency and the ability of market prices to aggregate the collective information of its participants.

One of the most popular markets the IEM offered in 2008 was a market based on the winner of the U.S. Presidential election. Both in theory and in reality, the prices of the securities in this market rose and fell based on the aggregate opinion of the market participants. It is this market that is the focus of this paper.

A computer program was created to function as a market maker by trading in this market and providing liquidity to other market participants. This market maker can be called an "electronic liquidity provider" or an "algorithmic trader". Computer programs that profitably trade securities without human interaction are mysterious to many because those involved in their creation and operation are reluctant to share their proprietary information. Algorithmic trading systems can be very lucrative, so there is a rational aversion to sharing information about how they work. Algorithmic trading can also be controversial, and the benefit it provides to financial markets and human traders is the source of much debate.

Due to the relatively small size and academic nature of the IEM, there are fewer reasons to keep the information about how this algorithmic trader worked confidential. All of the details about how this system functioned are disclosed in this paper in order to promote a more informed discussion about the benefits of algorithmic trading and its impact on financial markets. This research paper shares the quantitative methodology the system used to make trading decisions, examines the trading results, and discusses the influence it had on the efficiency of this market.

[1]Disclaimer: The views I express in this research paper are my own and were developed before starting at my current position and do not necessarily reflect those of my current employer or my fellow employees.

2. Literature review

The Iowa Electronic Markets are the focus of extensive research. Berg et al. (2008) studied the predictive ability of these financial markets and their accuracy at predicting the outcome of the vote-share of political elections. They characterized the IEM as a unique market that is different from the highly complex financial markets modern society depends on and the stylized markets created by researchers in a laboratory. The securities in the IEM have finite horizons upon which the terminal value is revealed, enabling researchers to measure how well the final outcome was predicted by market prices. Berg et al. found that the average poll error for a group of 15 elections was 1.91%. The average market error for the same elections was 1.49% or 1.58%, using the market price as of midnight on election eve or the volume weighted average price of all trades in the week before the election. Forsythe et al. (1992) investigated the relationship between market prices and election poll results and found no evidence that market price changes were preceded by changes in poll results. Instead, they argued that traders are able to independently interpret the meaning of an event and react quickly with market trades without depending on opinion polls for guidance.

IEM researchers can survey trader preferences and biases to understand how they influence trader behavior. Forsythe et al. also investigated the judgment biases and preferences of individual traders by conducting telephone and mail surveys of all market participants. The two biases studied were the "assimilation-contrast effect," causing a trader's preference for an outcome to influence their interpretation of information concerning the likelihood of that outcome, and the "false consensus effect," causing traders to overestimate the portion of the voter population that shares their views. They found significant evidence of both judgment biases, but argued that the market performed well in spite of the prevalence of biased traders because of the actions of subset of "marginal traders." These marginal traders differed from "average traders" in that they were less prone to judgment biases and traded the securities at prices very close to the equilibrium price. They acted as arbitrageurs and profited by trading with the biased supporters of both candidates. The marginal traders earned higher median returns of 9.6% as a result, compared to 0.0% for the non-marginal traders. The researchers found that the marginal traders were able to do this by more correctly identifying when important events happened and when they did not. They showed that if market prices were rising (falling), the marginal traders were buying (selling) if there was important news events and selling (buying) if there was not.

Another analysis of IEM prices and market efficiency was conducted by Berg and Rietz (2010). These researchers studied prediction markets with contracts dependent on the monthly prices of Microsoft stock or financial indices. They employed a logit model to estimate the true probabilities of different outcomes and compared them to the market prices. They found that the market prices were efficient for short time horizons, but for long horizons there was an overconfidence bias causing the market to over-price the most likely outcomes and under-price the least likely outcomes. This contrasts with studies of racetrack betting, where bettors overpay for the least likely outcomes. The researchers speculated that this is because of the ability of IEM traders to short sell securities, which is not found in racetrack betting. Berg and Rietz then backtested a trading strategy that buys the under-priced securities and short sells the over-priced securities. They found that they could achieve monthly returns in the range of 0.31% and 1.2%, depending on the investor's level of risk aversion. They also discovered that they could achieve higher returns by building a model that factors in the current price of the stock or index. A trader employing this model could have achieved monthly returns in the range of 0.8% to 2.31%. Since including the additional information improved the strategy's returns, the prediction markets could not have been fully efficient.

An analysis of market efficiency of the online betting exchange operated by tradesports.com was conducted by Borghesi (2007). The researcher studied bets placed on the outcome of National Football League (NFL) sporting events. Borghesi found that before the start of a game, average contract prices did not differ significantly from asset values, but after the start of a game and before contract expiration, contract prices were significantly overpriced for lower priced securities. This contrasts the results of Berg and Rietz (2010) but is consistent with studies of racetrack betting. Borghesi suggested the mispricings were caused by the relatively large information shocks that occur after scoring events during the game and a reluctance or difficulty in selling or shorting contracts. Borghesi found that the most significant mispricings took place immediately after scoring events and proposed several trading strategies to exploit the inefficiencies.

A second analysis of market efficiency of the betting exchange tradesports.com was conducted by Tetlock (2008). The researcher studied the relationship between liquidity and prediction accuracy of securities linked to the outcome of financial and sporting events. Tetlock found that increased liquidity did not reduce deviations between security prices and event outcomes, and that in some cases increased liquidity was actually associated with increased price deviations. Liquidity was measured using bid-ask spread size, market depth and cumulative trading volume. Similar to the results of Borghesi (2007), Tetlock found that low priced securities were overpriced and high priced securities were underpriced. The researcher suggested two reasons for these characteristics about this market. The first is that in this market, limit order traders had the tendency to buy the lower priced securities and sell the higher priced securities. These limit orders were filled by market orders, so the market orders also had different buy-sell imbalances for the low and high priced securities. Tetlock also observed that during periods of high information, limit orders had negative expected returns. This suggests that in this market the limit order traders were naively providing liquidity to the market, which could have slowed down market prices' response to new information.

Gil and Levitt (2007) studied a different online betting exchange operated by intrade.com and also found market inefficiencies. They analyzed bets placed on the outcomes of 2002 World Cup matches and found that although market prices increased rapidly after goals were scored, the prices continued to increase for the next 10-15 minutes. Gil and Levitt also found that on average there were negative returns for bets placed on the match favorites and that there was an advantage to market participants who were willing and able to be short sellers. In addition, the researchers examined the behavior of market makers in these markets. They found that on average the market makers earned negative returns. About 30% of the market makers' total losses came from trades made in the minute before and after a goal was scored, suggesting that the other market participants were able to take advantage of new information before market makers were able to change their bid and ask prices.

Berg and Rietz (2006) discussed the Iowa Electronic Markets and the presence of computer programs employing algorithmic trading strategies. They disclosed that they first became aware of the presence of algorithmic traders in the 2000 and 2004 elections, but acknowledged that there may have been earlier algorithmic traders that were unknown to them. In the 2004 election there was one algorithmic trader that was involved in 21% of the total units traded in the winner-takes-all market. They wrote that some computer programs seemed to be taking advantage of arbitrage opportunities and made positive profits. Others seemed to be using price-movement strategies. Berg and Rietz also talked about news events and the speed at which market prices reacted to relevant information. They presented anecdotal evidence of trader reactions before official announcements of key campaign events. Additionally, Berg and Rietz (2006) discussed large orders and market impact. Like more complex markets, prices were impacted by large orders but the effects diminished rather quickly. The markets were robust to manipulation. If a trader did attempt to manipulate prices, the price would revert to the pre-manipulation levels. Since the market prices seem to accurately predict future events, it does not seem that they are being successfully manipulated. The authors pointed out that this does not rule out the possibility of short term manipulation.

The potential for market manipulation in the Iowa Electronic Markets and other political markets was researched by Rhode and Strumpf (2008). In addition to the IEM, they studied tradesports.com and organized betting that took place on Wall Street between the years 1880 and 1944. In each case Rhode and Strumpf found that the political markets were robust to manipulation and that prices could not be distorted beyond short time periods. If prediction markets are vulnerable to manipulation it would challenge the assertion that they have significant predictive ability. Politically motivated market participants have an incentive to manipulate the prices of political markets if the prices can also influence voters and the outcome of an election. Historical records indicate that during the years the political betting markets were active on Wall Street, claims of price manipulation and staged bets were common. Nevertheless, the markets still had high predictive ability and any public claims of market manipulation were not followed by large permanent changes in prices. A more recent claim of price manipulation took place before the 2004 election in the tradesports.com presidential market. In that market the contract for one candidate sharply declined on two occasions. Afterwards it was rumored that one large investor was responsible for the price moves and was attempting to manipulate the price and influence the outcome of the election. It is unknown if these allegations are true, but in any case market prices had returned to their pre-decline levels within an hour and any manipulation that may have taken place had failed. Finally, Rhode and Strumpf experimented in the IEM before the 2000 U.S. Presidential election by making large random trades in an attempt to manipulate prices. They found that after each attempted manipulation, prices would revert to the pre-manipulation levels in a few hours. As a result, they determined that the IEM is robust to manipulation.

Oliven and Rietz (2004) discussed IEM market structure and investigated traders' efficient and inefficient trading practices. As will be explained in section 3.1, the market is structured in such a way that there are always two ways for a price-taking trader to achieve a desired position. For example, to buy a security, the trader can either buy the desired security, or they can buy all the securities in a bundled transaction and then sell the undesired securities. Either way the end position is the same, but the execution costs can be different because of small differences in prices. An efficient trader will always rationally choose to minimize execution costs. In this situation, Oliven and Rietz described a failure to make the rational decision a "price-taking violation of individual rationality." Traders who submit limit orders can also behave irrationally. If a trader sends a limit order that presents a potential arbitrage opportunity to other traders, that trader is committing a "market-making violation of individual rationality." These violations of rationality create opportunities for more efficient traders to make excess profits.

Brahma et al. (2010) explored inventory-based and information-based market making models for prediction markets. They studied an inventory based market making model developed by Hanson (2007) that is commonly used in prediction markets. They developed a new Bayesian information-based market making model and compared the two models in simulated trading and a laboratory market with student traders. The researchers found that in simulations the Bayesian market maker was superior to the inventory based model both in profitability and average error from the simulated true value but not for maximum drawdown. They developed a unique way to create an experimental market in a laboratory to test the two market makers simultaneously and got similar results as the simulations. Brahma et al. pointed out that although a liquidity providing market maker risks losing to better informed or smarter traders, smart market making algorithms may be able to exploit human trader errors or overconfidence. The algorithmic trading system developed in this paper attempts to do just that.

Das (2008) studied market makers and their effects on markets and price movements. He analyzed this from the perspective of how to design a market making algorithm and how market makers influence price dynamics. He built a market making model based on the work of Glosten and Milgrom (1985) that tries to estimate the true value of the security from the information conveyed in the trades and sets prices based on this estimate. In simulated trading, he showed that competition among market makers leads to faster price discovery and smaller spreads. Das argued that the presence of market makers improves market quality.

3. Mechanics of the Iowa electronic markets

3.1. Market prices

This financial market had two securities: one representing the Democratic candidate (DEM), Barack Obama, and the other representing the Republican candidate (REP), John McCain. There were no explicit transaction costs from trading these securities in this market. At maturity, each security pays $1 if the security's candidate wins and $0 if the candidate loses. For these securities, a win is when the candidate captures more votes than his opponent as reported by The New York Times three days after the election and is not based on the electoral college. Since there can be only one winner, the total payout of both securities at maturity is guaranteed to be equal to $1. To avoid arbitrage, the value of the two securities together must also be equal to $1 for all times before maturity, ignoring any discount functions. In addition, the price movements must be inversely correlated because when the value of one security goes up, the other must go down. Therefore, although there are two securities, this market has only one risk factor and one real asset. Buying one security is the same as selling the other. In this paper, when the term "securities" is used, it is referring to the two securities representing the two candidates, and when the term "asset" or "market asset" is used, it is referring to the one true asset that is traded in this market. A long position in the market asset is defined as a long position in the Democratic candidate and a short position in the Republican candidate, and a short position in the market asset is the opposite.

The IEM provides three pieces of price information for each security; the bid price, the ask price, and the last trade price. The minimum tick size for all prices is one tenth of a penny. Unlike many other financial markets, last trade size, bid size, and ask size are not disclosed. This poses unique challenges to an electronic market maker and to any researcher studying the microstructure of this market.

Since there are two securities that are by definition inversely correlated, it simplifies things to invert the price of one of them so it is comparable to the other. For example, if these are the market prices:

Security	Bid	Ask
DEM	0.572	0.602
REP	0.424	0.450

Subtracting the REP bid-ask prices from 1 yields the bid-ask prices for REP_{inv}:

Security	Bid	Ask
DEM	0.572	0.602
REP_{inv}	0.550	0.576

REP_{inv} has the same market exposure as DEM. Notice that the REP bid price determines the REP_{inv} ask price, and the REP ask price determines the REP_{inv} bid price.

An important feature of this market is the exchange allows market participants to buy and sell "bundles" of securities with the exchange itself for a fair price of $1. A bundle is a set of securities representing all of the available securities in the market; in this case, the DEM and REP securities. This is much like a creation or redemption of an ETF in the stock market, but without a fee. This allows a trader two alternatives to trade into a position: they can either buy the one security they want or they can exchange their money for bundles and then sell the other undesired security. The end result will be the same.

Since there are two mechanisms for trading the same asset, the market asset's true bid and ask prices must be identified. The asset's bid price is the maximum of the DEM and REP_{inv} bid prices and the asset's offer price is the minimum of the DEM and REP_{inv} offer prices:

Security	Bid	Ask
DEM	**0.572**	0.602
REP_{inv}	0.550	**0.576**

If a trader wanted to trade in this market it would be efficient and rational for them to only trade using the asset's true bid and ask prices. These two prices make up what can be considered to be the "good spread" and the other two prices make up the "bad spread."

Spreads	Bid	Ask
Good spread	0.572	0.576
Bad spread	0.550	0.602

It would be inefficient for a trader to use a market order to buy the market asset on the bad offer or sell the market asset at the bad bid. Nevertheless, this happened frequently. Doing so would be considered a "price-taking violation of individual rationality," as discussed in Oliven and Rietz (2004). The researchers found that in the 1992 Presidential election, 37.7% of market orders made this error. This is consistent with estimates of the frequency of these violations taken from the algorithmic trader's own trading data.

For the market making efforts described in this paper, the current value of the market asset is defined as the midquote of the good spread. This was considered to be the market asset's efficient price. In the above example, the efficient price is 0.574. This price can be used to value the DEM and REP securities. The value of the DEM security is 0.574 and the value of the REP security is 0.426. In both cases this happens to be closer to the bid side than the ask side of the displayed DEM and REP security spreads. This method is more accurate than simply looking at the midquote of each individual security. Since there is only one asset in this market, the order books of both securities should always be considered together.

3.2. *Market activity*

Figure 1 shows several plots of market activity during the time of this research project. The upper subplot graphs the daily price movements of the market asset's efficient price, as defined in section 3.1. The second subplot shows the daily percent volatility of the market asset price, calculated with a 7 day rolling window and measured in 15 minute intervals.

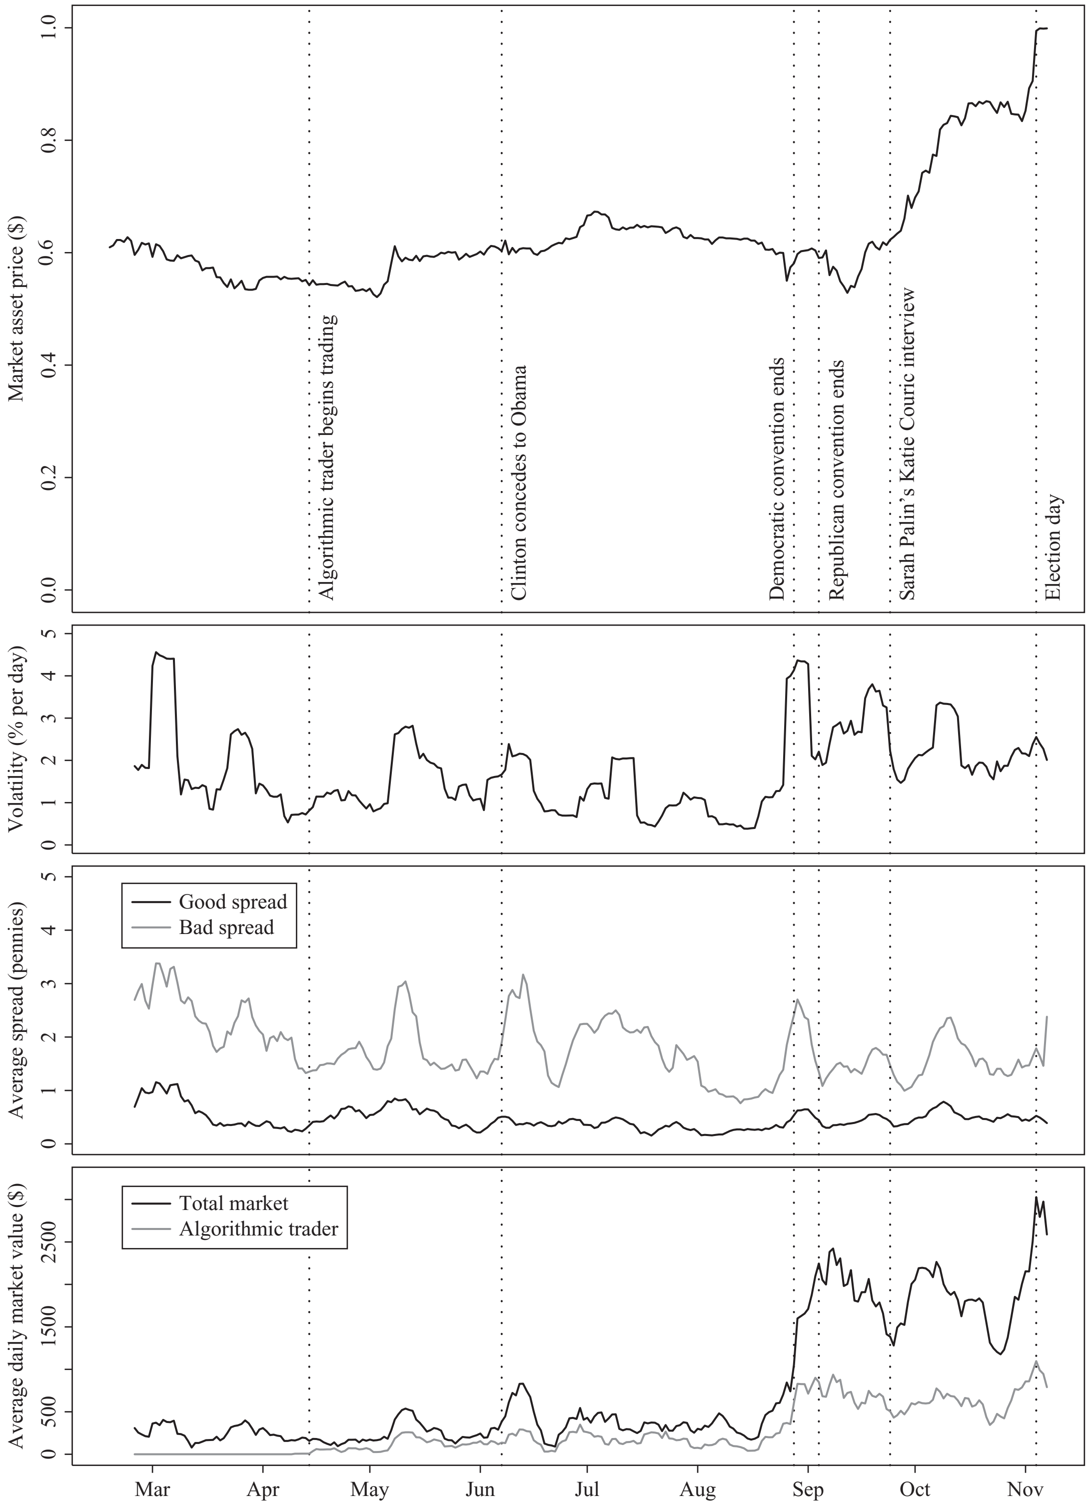

Fig. 1. Four plots of market activity. The upper subplot shows the market asset's price movements with significant political events annotated. The other three plots show the daily percent volatility, average good and bad spread sizes, and average daily market activity in dollars, all calculated with a 7 day rolling window.

The third shows the time weighted average good and bad spread sizes calculated with a 7 day rolling window. The last subplot shows the 7 day moving average of market activity, measured in dollars. A few significant political events are annotated on the charts to show the effect they had on price, volatility and volume.

It should be noted that the "Winner Takes All" market was not the only political financial market offered by the Iowa Electronic Markets. Two others are political markets based on the Democratic and Republican nominations. The securities in those markets reached maturity when their respective party's conventions ended, and at that time money flowed into the accounts of traders who held the securities representing Obama and McCain. Some of that capital was reinvested in "Winner Takes All" securities, explaining the surge in market activity when the conventions ended. This may also partially explain the price fluctuations around the time of the conventions. Volatility also increased noticeably around the time of the conventions. This makes sense as there was more information about the candidates being revealed every day as the time remaining before the election decreased.

The trading system began collecting data in late February while the model and the infrastructure to implement this project were being developed. The system was operating in a limited test mode by April and was fully operational by May. Table 2 shows that before the conventions, the algorithmic trader was 44% of the market volume. After the conventions the market and the system became much more active but its percent of market share declined. This happened partially because of capital limitations and position limits built into the algorithmic trader.

3.3. *Arbitrage*

In this market it is possible for arbitrage opportunities to arise. One way this can happen is when a trader commits what Oliven and Rietz deemed a "market-making violation of individual rationality." This occurs when a trader adds a limit order to one of the order books that causes the good spread to become crossed. A crossed market is when the bid price exceeds the ask price. For example, in this market:

Security	Bid	Ask
DEM	0.572	0.602
REP	0.424	0.450

A bid of 0.580 on the DEM security would result in these spreads:

Spreads	Bid	Ask
Good spread	0.580	0.576
Bad spread	0.550	0.602

Notice the good spread is now crossed.

A more rational trader that intended to buy the DEM security could have bought a bundle from the exchange and sold the REP security at the bid price of 0.424, leaving a long position in the DEM security. The trader could have saved 0.004 per share and gotten the trade filled immediately. But if they don't realize this possibility, the displayed market will become:

Security	Bid	Ask
DEM	0.580	0.602
REP	0.424	0.450

The arbitrage opportunity is to sell both the DEM and REP securities at their respective bid prices and then buy bundles by trading with the exchange. This is profitable because the pair of securities can be sold at a price of 1.004 and bought back from the exchange with a bundled transaction for 1 dollar. The exchange has a special market order for buying and selling bundles using the sum of the displayed bid or ask prices that ensures equal numbers of each security are traded. After selling bundles using one of these market orders, the new displayed market could hypothetically become:

Security	Bid	Ask
DEM	0.580	0.602
REP	0.416	0.450

For this example, assume that the size of the REP 0.424 bid order is smaller than the DEM 0.580 bid order. When both bids are hit by the bundle market order, the REP bid of 0.424 is eliminated but the DEM bid remains. Since the previous REP bid is now gone, the new REP bid is the next order in its order book.

It may seem trivial but it is actually important for the efficiency of this market to have at least one market participant who monitors the prices for this situation and acts accordingly. If this didn't happen the the two security prices would not maintain the proper inverse price relationship. Also, the method of determining the market asset's price discussed in this paper is undefined when the good spread is crossed. The arbitrage opportunity would have to be eliminated to determine the midquote and the market maker's view of the market asset value.

A second arbitrage opportunity involves providing liquidity to the bad spread and taking liquidity from the good spread. For example, in the current state of the example market:

Security	Bid	Ask
DEM	0.580	0.602
REP	0.416	0.450

The good and bad spreads are:

Spreads	Bid	Ask
Good spread	0.580	0.584
Bad spread	0.550	0.602

A trader could place a buy order for the market asset for a price of anything less than 0.580 and a sell order for the market asset for any price more than 0.584. This would correspond to a sell order for the REP security at a price above 0.420 and a sell order for the DEM security at a price above 0.584. If either of the trader's sell orders were lifted, they can immediately sell the other security at its bid price to eliminate the risk of the position. One flaw with this strategy is that since the order sizes are not visible in this market, the trader wouldn't know how many shares they would be able to sell by hitting the bids. This would force the trader to estimate the order book sizes. Sometimes the estimates would be too large, resulting in a position that could not be liquidated profitably. Because of this problem, this strategy doesn't work as well as it would in another market with visible order sizes.

3.4. Market gaming

The absence of displayed order sizes in this market creates the potential for traders to game this market or behave in a way that frustrated the system's ability to determine the efficient price of the market asset. The problem is a limit order for one share placed in the bid or ask order books appears to other market participants to be identical to a limit order for a larger size. Although it is possible that the traders who placed one share orders had little capital and only wanted to trade one share, one can also speculate that the traders were attempting to alter other market participants' perception of current prices with small orders. It is clear that traders placed one share orders with some frequency because sometimes those market participants would also inadvertently create an arbitrage opportunity by committing a market-making violation of individual rationality. When the algorithmic trader tried to exploit the resulting arbitrage opportunity, it would trade only one bundle. About 10% of the system's 1,267 arbitrage attempts resulted in a trade of only one bundle, indicating that the smaller order for the two securities was for one share.

It is unclear what traders' motivations are for doing this. One possible explanation is that it is a trader's attempt to draw more favorably priced orders to the opposite side of the order book they intend to trade. For example, if a trader wanted to sell the DEM security but thought the spread was too large, they could narrow the spread with a buy order for one share and hope that this influences another trader's view of what the current market price is for that security. The second trader may then place a larger buy limit order at a price that is equal to or better than the one share order. Once that happens, the first trader could remove their one share order and sell the DEM security at the improved price.

This is a different kind of market manipulation than what was discussed in Berg and Rietz (2006), but it is still important. Clearly this behavior can only have a short term effect and cannot impact the market's ability to aggregate information or predict the outcome of an election. Nevertheless, it does have an effect on the short term microstructure of this market. Even if it is not considered manipulation of market prices, it can manipulate the actions of other market participants, particularly the actions of the electronic market maker described in this paper.

4. Market making

Consider a market maker that is willing to alter its own portfolio away from its current holdings in exchange for some compensation for bearing the risk of a new portfolio. The compensation comes from the size of the spread earned on each trade, which is the difference between the execution price and the efficient price of an asset. The risk comes from the potential of the altered portfolio to decline in value relative to the desired portfolio.

When an uninformed market maker trades with potentially informed traders, the compensation the market maker earns on each trade will be partially diminished by an adverse price movement from the market impact of the counter-party's trade. The market impact is a result of the information content of the trade, or in other words, the counter-party's informed view of the correct value of the security. A market maker needs to demand extra compensation to make up for the adverse selection associated with the information asymmetry between itself and liquidity demanders. It should be noted that a computer program functioning as a market maker does not necessarily have to be uninformed. A computer program could be created to get information by monitoring prices on other betting websites such as intrade.com or by parsing news headlines. The computer program discussed in this paper, however, will be uninformed. It will know nothing about the meaning of the securities it is trading or the election. Most (if not all) of the traders the market maker trades with will be humans who can interpret the events of the political process associated with an election. Many of those traders will be biased and make mistakes as discussed in Oliven and Rietz (2004), but in aggregate they will be more informed than a computer program and will be able to predict security price movements in a way a computer program cannot.

The market maker's risk of a new position is a function of the volatility of its undesired holdings and the expected time for that position to be offset by a later trade. If a market maker were to provide liquidity to a trader by buying a certain amount of an asset, it can expect that some time later it will sell the same asset to another trader, offsetting the risk of the earlier buy trade. During the time the market maker is waiting for the offsetting trade, it faces the risk that the price will move. The unrealized profit earned from the spread on the first trade can potentially be reduced or turn into a loss. The market maker needs to find an optimal balance between the profit earned from the spread when providing liquidity and its aversion to the risk of the undesired portfolio holdings.

To develop the market making model, consider an exponential utility function to evaluate changes in wealth, W:

$$\mathbf{U}(W) = -e^{-\lambda W}, \quad \lambda > 0 \tag{1}$$

Assume that W is distributed normally with a mean of μ and a variance of σ^2. The expectation of $\mathbf{U}(W)$ is then given by:

$$\mathbf{EU}(W) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} -e^{-\lambda W} e^{-\frac{(W-u)^2}{2\sigma^2}} dW \tag{2}$$

By re-arranging the terms and employing calculus, one can reduce equation 2 to:

$$\mathbf{EU}(W) = -e^{-\lambda(\mu - \frac{\lambda\sigma^2}{2})} \tag{3}$$

Maximizing the expected utility of changes in wealth W is then equivalent to maximizing the expression

$$\mu - \frac{\lambda\sigma^2}{2} \tag{4}$$

In other words, the exponential utility function is equivalent to the mean variance utility function when changes in wealth are normally distributed. We can then write the market making model's utility function as:

$$\mathbf{U}(W) = \mathbf{E}[W] - \frac{\lambda}{2}\mathbf{Var}[W] \tag{5}$$

This equation can be used to evaluate the utility of the change in wealth W from one trade of X shares earning a spread of s per share. This utility function was also used by Berg and Rietz (2010) to develop a trading strategy to exploit the overconfidence bias in IEM prices. Use of this utility function is justifiable in the case where the distribution of W is normally distributed and constant absolute risk aversion (CARA) is desired. Without any prior data to study, it is reasonable to assume that the distribution of W will be normal. It is desirable for the market maker to exhibit CARA so it does not change its affinity for risk when it has a long or short position in an asset and the asset price, and therefore the market maker's net wealth, changes.

The expected increase in wealth from the trade is $\mathbf{E}[W]$ and the variance of the wealth increase is $\mathbf{Var}[W]$. The constant value λ is the risk aversion parameter. By adjusting λ, the market maker can control how much it demands to be compensated for accepting a given level of risk. This utility function is useful for evaluating the trade-off between the expectation of profits from providing liquidity to other market participants and aversion to the risk of undesired positions.

The expected increase in wealth comes from the spread earned on each share traded, minus the market impact from the information content of the new trade. The market impact of the new trade will also affect the value of the current undesired position, either positively or negatively. For example, if an informed trader sells an asset to a market maker, the market impact of the trade will cause the price to decline slightly. If the market maker had a long position before the trade, that market impact will have a negative impact on that position. If instead the market maker had a short position, the decline in price will have a positive impact on the value of the position.

The market maker can estimate the expected increase in wealth when it buys a security on its bid price or sells at its ask price:

$$\mathbf{E}[W] = s|X| - G(Z)(X + A - D) \tag{6}$$

In this equation, Z is the total size of the counter-party's trade, which may be partly or completely filled by the market maker. A is the market maker's current position before the trade and D is the desired position, making the difference $A - D$ the current undesired position. X is the potential alteration to the market maker's portfolio, which will be positive (negative) when the market maker is buying (selling). The market maker's per-share compensation for altering its portfolio is the spread s and the total compensation $s|X|$ is positive when buying or selling. The function $G(Z)$ is a function modeling the permanent impact of the trade. Defining P_{post} as the market asset's efficient price sufficiently far enough into the future after the trade that the temporary impact has completely decayed and P as the price immediately before the trade, the function $G(Z)$ becomes an estimate of the price change $P_{post} - P$:

$$G(Z) = P_{post} - P \tag{7}$$

According to Huberman and Stanzl (2004), the function $G(Z)$ must be linear in trade size to enforce a "no arbitrage" condition on the market. Any non-linear function implies market participants can construct a trading pattern that buys and sells securities and earns risk-free profits. The linear market impact model is

$$\frac{P_{post} - P}{P} = \gamma Z \tag{8}$$

Market impact is then a linear function of the constant γ. Defining $Y = Z - X$ as the estimated size of the trade filled by other market participants at the same price, the market impact function $G(Z)$ becomes

$$G(Z) = \gamma P(X + Y) \tag{9}$$

and the estimate of the expected profit from a buy or sell trade is:

$$\mathbf{E}[W] = s|X| - \gamma P(X + Y)(X + A - D) \tag{10}$$

The variance of the expected increase in wealth comes from the variance of changes in the undesired position value, including the new position X. The market making model can estimate the variance of net wealth as:

$$\mathbf{Var}[W] = (\sigma\sqrt{t}P(X + A - D))^2 \tag{11}$$

The variable P is the efficient price of the market asset, so the quantity $P(A - D + X)$ is the value of the new undesired position measured in dollars. Notice that X can increase or decrease the variance of the market maker's position, depending on the value of $A - D$. The parameter σ is the estimated one-day volatility of the market asset and t is the expected time in days the market maker will hold the position before receiving an offsetting trade. Volatility scales by the square root of time.

Substituting Equations 10 and 11 into equation 5 yields:

$$\mathbf{U}(W) = sX - \gamma P(X + Y)(A - D + X) - \frac{\lambda}{2}(\sigma\sqrt{t}P(A - D + X))^2 \tag{12}$$

To find the optimal number of shares the market maker would be willing to buy or sell at a price earning a spread of s, simply calculate the first derivative of $\mathbf{U}(W)$, set the result equal to zero, and solve for X. The optimal number of shares for the market maker to be willing to buy on the bid or sell at the offer becomes:

$$X_{bid} = \max\left[0, \frac{s - \gamma P(Y + A - D) - \lambda\sigma^2 tP^2(A - D)}{2\gamma P + \lambda\sigma^2 tP^2}\right] \tag{13}$$

$$X_{ask} = \max\left[0, \frac{s - \gamma P(Y - A + D) + \lambda\sigma^2 tP^2(A - D)}{2\gamma P + \lambda\sigma^2 tP^2}\right] \tag{14}$$

Here, the result is enclosed in the $\max$ function to prevent the market maker from attempting to trade a negative number of shares on the bid or at the ask. The model can yield a negative number when $|A - D|$ is large.

To make use of the model, the market maker needs to estimate parameter values from market data and choose a value for the risk aversion parameter λ. Estimating σ is straightforward. Due to the lack of transparency regarding trade sizes, the t, Y and γ parameters can only be approximated in advance by observing average market volumes, asset prices and the average number of trades per day. Later, after the market maker has been operating for a reasonable period of time, the parameters can be calibrated using the market maker's own trading history. The risk aversion parameter λ is more difficult because there is no "correct" value that can be calculated. It can only be set based on risk preferences and tolerance for the volatility of the algorithmic trader's net wealth.

It would be prudent for the market maker to continuously re-estimate the parameters to adapt to changing market conditions. Table 1 shows the range of parameter values used by the market maker during its participation in this market. The market maker also needs to determine a mechanism for determining the value of the efficient asset price P. This is calculated as the midquote of the good spread.

The model can be explored by fixing the parameter values and adjusting one parameter at a time to gain intuition about how the model works. If the market maker's valuation of the efficient price P is fixed to 0.5, the market maker would be willing to buy or sell the asset at prices below or above 0.5, respectively. The spread s earned on the trade is then the difference between any price and 0.5. The number of shares the market maker would be willing to buy or sell at any price can be calculated using formulas 13 and 14.

Figure 2 plots the size of the limit orders the market maker would place in the IEM order books to buy or sell the market asset at different prices. This plot can be interpreted in two ways. For a given price, the plot shows the number of shares the market maker is willing to trade at that price. Alternatively, for given limit order size, the plot shows how much compensation earned from the spread the market maker would require if a limit order of that size were to be filled and alter its asset holdings.

Table 1
Market Making Parameters

Parameter	**Min Value**	**Max Value**
σ	1% *per day*	3% *per day*
t	0.083 *days*	0.25 *days*
Y	75 *shares*	150 *shares*
γ	2.5e-5 $shares^{-1}$	5.0e-5 $shares^{-1}$
λ	6 *day* / $	10 *day* / $

In Figure 2, notice the market maker is willing to trade in larger sizes when the profit opportunity is larger. When its current undesired position $A-D$ equals 0, the size of its orders are equal for equal sized spreads earned from either buying or selling. This is not the case when the market maker has a long position in the asset. It becomes a more aggressive seller and less aggressive buyer, as expressed by the larger sell orders and smaller buy orders for the same prices. This is because the market maker's risk aversion makes it reluctant to further increase its position but more willing to reduce its position. It is demanding more compensation for a trade that moves its portfolio further away from its desired portfolio but is willing to accept less compensation for a trade that moves its portfolio closer to its desired portfolio.

Figure 3 fixes $A-D$ to 0 but adjusts the risk aversion parameter λ. This plot shows that an increase in risk aversion causes the market maker to be less willing to buy or sell the market asset at the same prices. The market maker is demanding more compensation to accept the same level of risk. Increasing the volatility parameter σ has a similar effect, as shown in Figure 4. Figure 5 shows that increasing t has a similar effect as well. A larger t value means the market maker expects to hold on to the undesired risky position for a longer period, and its risk aversion compels it to demand more compensation for doing so.

Adjusting the market impact parameter γ has a much different effect on the market maker's trading behavior. As illustrated in Figure 6, when γ increases, the range of prices at which the market maker is not willing to trade a single share increases. This may be viewed as the minimum spread at which the market maker is willing to accept for any trade as compensation for the information asymmetry between itself and the more informed market participants. It is important that the model has this feature. Otherwise, it would be willing to trade a small number of shares at prices that earned minimal spreads. Those trades would frequently result in losses because the market impact of the counter-party's trade would prevent the market maker from later receiving an offsetting trade that results in a profit.

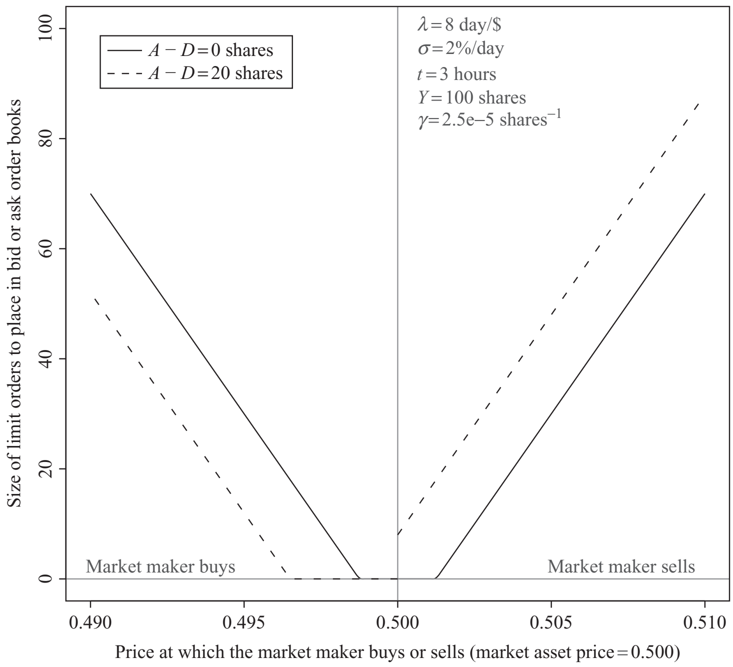

Fig. 2. Plot of the number of shares the market maker would be willing to trade at different prices when the undesired position $A - D$ is adjusted, assuming the market asset value is 0.5. Below an asset price of 0.5 the chart shows the number of shares placed on the bid and above 0.5 the chart shows the number of shares placed at the offer.

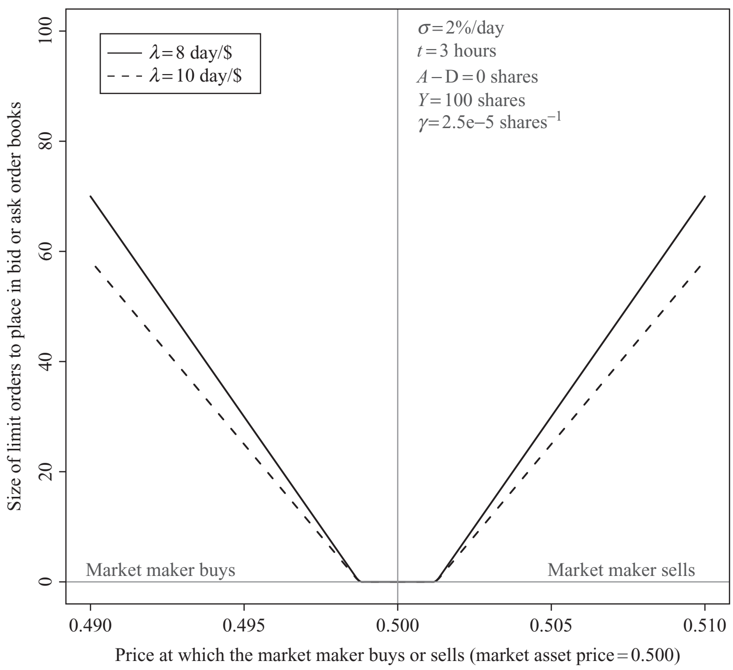

Fig. 3. Plot of the number of shares the market maker would be willing to trade at different prices when the risk aversion parameter λ is adjusted, assuming the market asset value is 0.5. Below an asset price of 0.5 the chart shows the number of shares placed on the bid and above 0.5 the chart shows the number of shares placed at the offer.

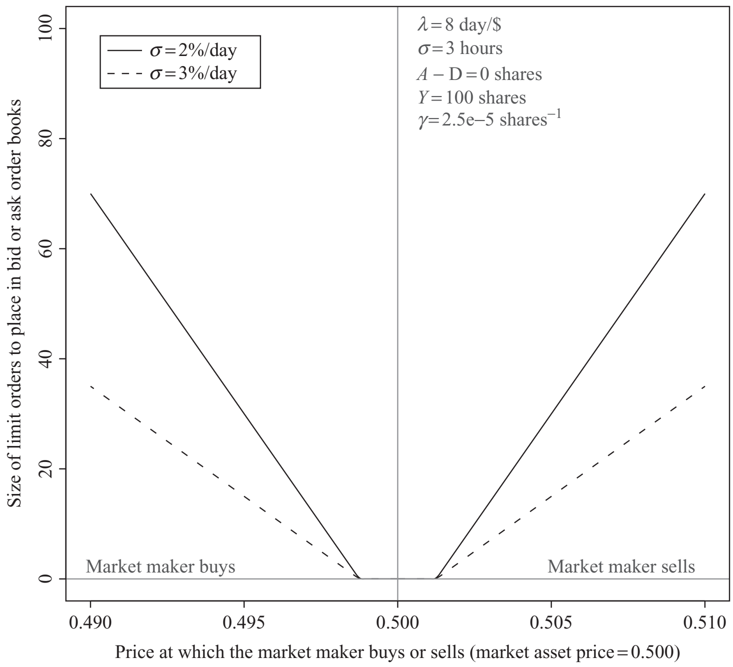

Fig. 4. Plot of the number of shares the market maker would be willing to trade at different prices when the volatility parameter σ is adjusted, assuming the market asset value is 0.5. Below an asset price of 0.5 the chart shows the number of shares placed on the bid and above 0.5 the chart shows the number of shares placed at the offer.

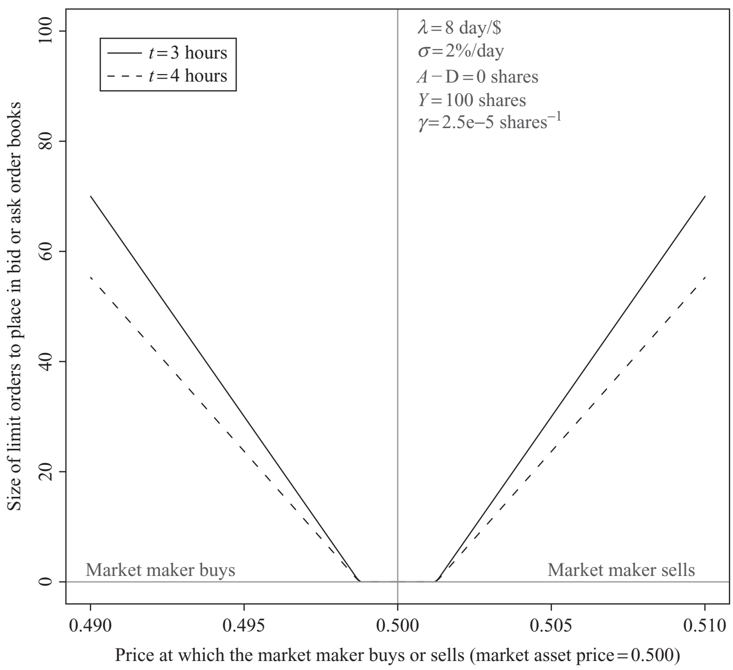

Fig. 5. Plot of the number of shares the market maker would be willing to trade at different prices when the holding period parameter t is adjusted, assuming the market asset value is 0.5. Below an asset price of 0.5 the chart shows the number of shares placed on the bid and above 0.5 the chart shows the number of shares placed at the offer.

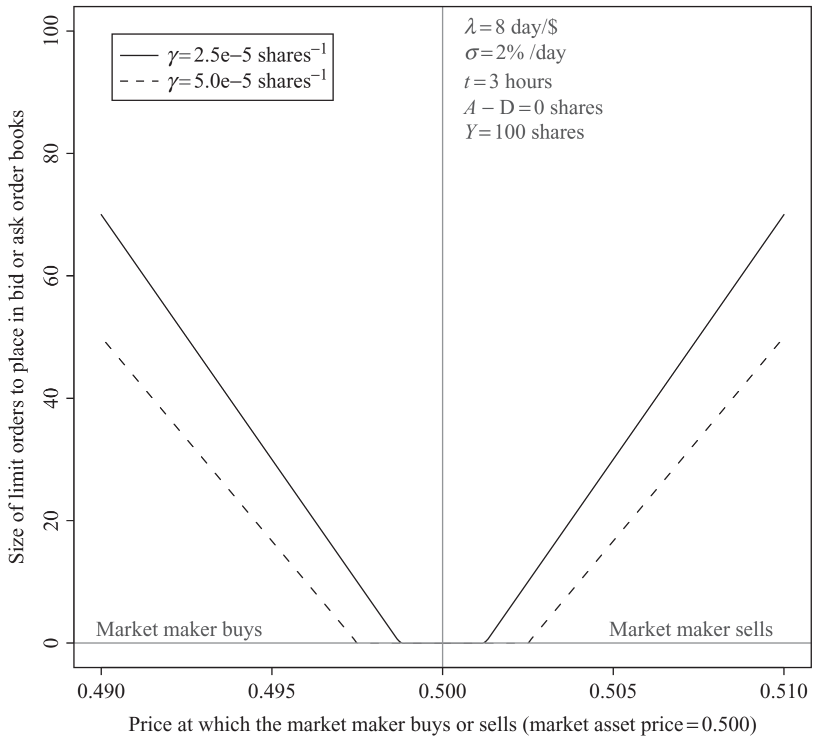

Fig. 6. Plot of the number of shares the market maker would be willing to trade at different prices when the market impact parameter γ is adjusted, assuming the market asset value is 0.5. Below an asset price of 0.5 the chart shows the number of shares placed on the bid and above 0.5 the chart shows the number of shares placed at the offer.

5. Implementation

A computer program was built that had the ability to continuously monitor price and position data in the Iowa Electronic Markets and send market orders and limit orders to the exchange. This computer program had the ability to interact with the same website the human traders use as their trading interface by sending and receiving HTTP requests. Using this interface, the computer program could exploit arbitrage opportunities and employ the market making model developed in this paper to provide liquidity to other market participants.

The market making algorithm was implemented to observe the bid and ask prices set by other market participants every 15 seconds and determine how many shares it would be willing to trade at those same prices using Equations 13 and 14. It only placed limit orders at prices equal to the displayed bid and ask prices. By doing this, the algorithm was consistently "joining" other market participants on the bid or the ask and was not narrowing the bid-ask spread. This is different from a market maker that will "step in front" of other orders and narrow the bid-ask spread by placing buy (sell) limit orders at prices that are higher (lower) than those placed by other market participants. Since this market making algorithm was dependent on other market participants to determine the efficient price and the market does not present order book information beyond the best bid or best ask prices, stepping in front of other orders would obfuscate the algorithm's view of the efficient price and make it more difficult for it to function.

Sometimes the bid and ask prices of the good spread would be equal, making the market "locked." In a locked market the bid price and ask price of the good spread are both equal to the midquote price, and therefore the market making algorithm's view of the efficient price. If at that time the market maker had a long or short position in the market asset and was willing to trade at the good spread's bid or ask prices, it was essentially willing to earn a spread of zero on a trade. It is clear from Figure 2 that this could happen. When that happened, the system could also send a market order to trade the other security immediately instead of waiting for another market participant to fill its order. The algorithmic trader was programmed to do this in a locked market to try to exhaust a long or short position. This is the only time it would use market orders.

Finally, the implementation of the market making algorithm included position limits to protect it from suffering too great of a loss if the price were to move quickly in one direction. In this market the price was guaranteed to reach \$0 or \$1 by the market's termination date so there was always the possibility of a large loss from an adverse price move. The position limits also made sure the system had some un-utilized capital available to take advantage of arbitrage opportunities as they became available.

6. Results

The algorithmic trader started with an initial investment of \$150 and earned a profit[2] of \$547.83. Almost every week (and day) was profitable, with an average daily return of 0.76%. There was one serious drawdown on the day before the election, November 3rd, when the system lost \$44.84 due to a technical problem with the implementation. No other day had a loss of more than \$0.50. The annualized Sharpe ratio over the full period of operation was 9.9.

Table 2 shows the trading activity and market share measured in dollars, shares and percent. Most likely the trading system was the most active market participant during the duration of this research project. The amount of activity is reported for the entire time the system was in operation and broken down into relevant periods. The period before the first party convention is separated from the period during and after the conventions because of the large increase in market volumes. The period right before the election is also segmented because of the unique price movements that occurred at that time.

The algorithmic trader earned \$241.27 from its arbitrage strategy and \$293.92 from its market making strategy. It also earned \$12.64 from discretionary positions. The system's ability to take discretionary positions (by setting parameter D in Equations 13 and 14) was not used until the day before the election. For research purposes it was considered desirable for the algorithmic trader to operate with as little human intervention as possible. On the other hand, it was thought that the system might suffer a large loss at election time due to an adverse price movement or that it might not be able to function correctly. Due to this concern, the discretionary position functionality was used to protect its profits with a long position in the market asset. Surprisingly, this was largely unnecessary. Not only did the system earn a profit on election day, it was also able to earn a profit providing two sided markets on the day after election day as well, when the winner of the election became common knowledge.

The upper plot in Figure 7 shows the cumulative profits of the market making and arbitrage strategies. It also shows the small profit from discretionary positions right before the election. Notice in the upper chart that there was a profit surge around the time of the party conventions. Most of this was from the

[2]The invested capital and all the profits were donated to charity.

Table 2

Trading Volumes and Market Share by time period. Dollars traded and shares traded are the the total notional or shares traded in each time period. This table excludes bundled transactions to create or redeem securities for \$1.

	Algorithmic Trader	Total Market	% Market Share
Full Period 4/15-11/7			
Dollars Traded (\$)	70,376.04	187,721.49	37%
Shares Traded	140,911	418,008	34%
Pre-Convention 4/15-8/24			
Dollars Traded (\$)	19,571.27	43,692.85	45%
Shares Traded	38,703	88,220	44%
Post-Convention 8/25-10/20			
Dollars Traded (\$)	37,635.20	106,489.07	35%
Shares Traded	74,916	228,846	33%
Election Time 10/21-11/4			
Dollars Traded (\$)	12,363.90	33,137.59	37%
Shares Traded	26,225	94,386	28%

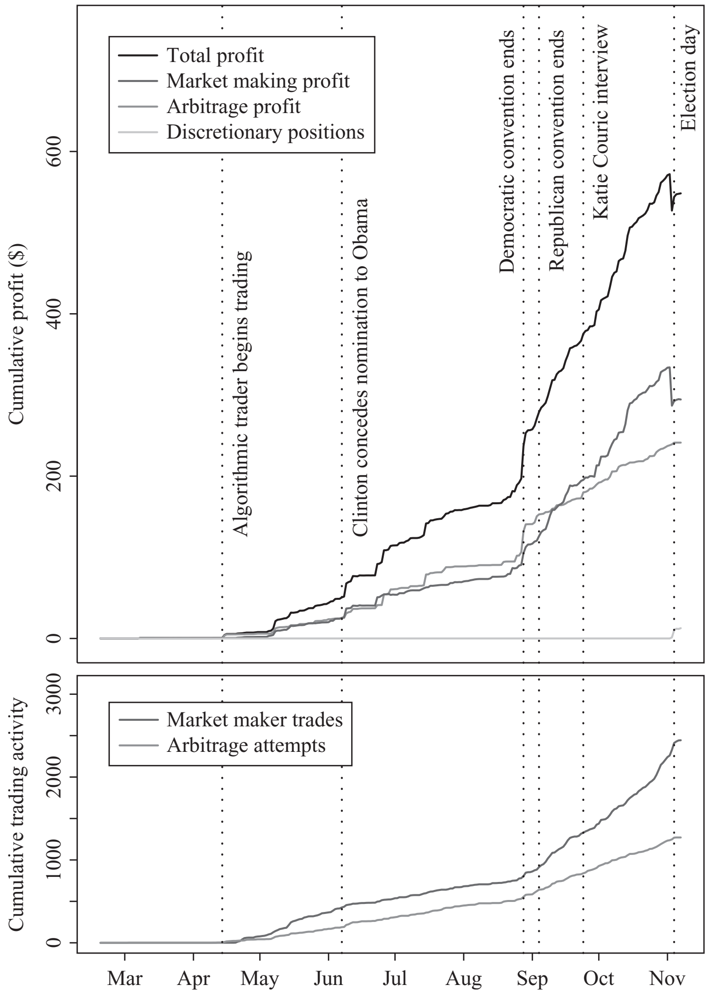

Fig. 7. Plots of trading activity. The upper subplot shows the cumulative profits from discretionary positions and the market making and arbitrage strategies. The lower subplot shows the cumulative number of arbitrage attempts and market making trades.

arbitrage strategy. On one day during the Democratic convention the algorithmic trader earned over $25 from arbitrage alone. There was also a sustained increase in the rate at which profits were earned around the time of the party conventions, which corresponds to the increase in market activity shown in Figure 1.

The technical problem the computer program suffered on the day before the election affected its ability to correctly communicate limit prices to the exchange through the web interface when the price of the REP security dropped below $0.10. The system began trading at prices that were different from what it intended, resulting in a large loss of capital. The market impact of these trades affected the security prices, which recovered quickly once the problem was corrected.

The lower plot in Figure 7 shows the cumulative activity of the two strategies. The system made a total of 1267 arbitrage attempts, 1264 of which were profitable. Three attempts resulted in a small loss because one leg of the arbitrage disappeared before the market order was executed. This could have been because of random market activity or from an electronic competitor for arbitrage opportunities. The market making strategy made a total of 2443 trades and provided liquidity with 76.5% of its volume. The remainder of its trades were market orders to exhaust undesired positions when the good spread was locked.

The daily performance statistics are presented in Table 3. Results for the full period of operation and relevant sub-periods are shown. The total performance statistics of all of the system’s activity is in the first

Table 3

Daily performance results by time period. "All Returns" results are actual trading results from discretionary positions and the market making and arbitrage strategies. "Buy & Hold Benchmark" results are simulated returns assuming capital was only invested in the DEM security instead of being used for trading strategies. Other results are calculated returns for each strategy with the profits from other strategies removed.

	All Returns	Market Making & Arbitrage	Market Making Only	Arbitrage Only	Buy & Hold Benchmark
Full Period 4/15-11/7					
Average	0.76%	0.75%	0.53%	0.47%	0.31%
Std Dev.	1.46%	1.46%	1.24%	1.10%	1.94%
Min	−6.22%	−6.39%	−9.76%	0.00%	−8.18%
Median	0.30%	0.29%	0.22%	0.09%	0.08%
Max	11.60%	11.60%	6.06%	10.01%	9.89%
Sharpe*	9.91	9.80	8.19	8.15	3.11
Pre-Convention 4/15-8/24					
Average	0.61%	0.61%	0.35%	0.38%	0.08%
Std Dev.	1.26%	1.26%	0.84%	0.98%	1.24%
Min	−0.14%	−0.14%	−0.72%	0.00%	−3.78%
Median	0.18%	0.18%	0.09%	0.04%	0.00%
Max	8.07%	8.07%	6.05%	6.85%	5.83%
Sharpe*	9.21	9.21	7.90	7.36	1.28
Post-Convention 8/25-10/20					
Average	1.26%	1.26%	1.15%	0.73%	0.69%
Std Dev.	1.67%	1.67%	1.26%	1.45%	2.75%
Min	0.00%	0.00%	−0.23%	0.00%	−8.18%
Median	0.75%	0.75%	0.78%	0.27%	0.58%
Max	11.60%	11.60%	6.06%	10.01%	6.47%
Sharpe*	14.41	14.41	17.39	9.59	4.77
Election Time 10/21-11/4					
Average	0.26%	0.16%	−0.09%	0.40%	0.98%
Std Dev.	1.91%	1.87%	2.73%	0.32%	3.02%
Min	−6.22%	−6.39%	−9.76%	0.01%	−2.48%
Median	0.41%	0.41%	0.38%	0.40%	0.00%
Max	2.48%	1.64%	2.10%	1.17%	9.89%
Sharpe*	2.58	1.62	−0.64	24.42	6.18

* The Sharpe ratio is annualized assuming 365 days a year.

column. The second column shows the performance statistics with the discretionary positions removed. The "buy and hold benchmark" shows simulated returns assuming the capital was only invested in the DEM security[3] instead of being used for the algorithmic trading strategies. The other columns show the performance of either the market making or arbitrage strategies with the profits from all other activity removed.

It is clear from this table that the algorithmic trader had higher returns and a higher Sharpe ratio than a trader with advance knowledge of the election outcome would have been able to achieve. It would have been very difficult for a human trader focused on campaign news to achieve superior trading results. Employing this algorithmic trading strategy is less risky and has higher investment returns than attempting to interpret political news events and predict the outcome of the election.

Table 3 also shows the increase in profitability in the post-convention period. The Sharpe ratio increased from 9.2 before the conventions to 14.4 after. Much of the increase is due to an increase in spread profits, to be discussed in more detail later. The table also shows much poorer results for the two week period right before the election. This is largely due to the large

[3] The Democratic candidate won the 2008 Presidential election.

Table 4

Daily residual returns by time period. Residual returns calculated with the simulated "Buy & Hold" strategy as the benchmark. "All Returns" results are actual trading results from discretionary positions and the market making and arbitrage strategies. Other results are calculated returns for each strategy with the profits from other strategies removed.

	All Residual Returns	Market Making & Arbitrage	Market Making Only	Arbitrage Only
Full Period 4/15-11/7				
Average	0.44%	0.43%	0.22%	0.16%
Standard Deviation	2.38%	2.40%	2.33%	2.17%
Min	−7.62%	−8.73%	−11.16%	−9.34%
Median	0.26%	0.26%	0.17%	0.08%
Max	11.86%	11.86%	9.84%	8.79%
Information Ratio*	3.54	3.43	1.79	1.37
Pre-Convention 4/15-8/24				
Average	0.53%	0.53%	0.26%	0.29%
Standard Deviation	1.72%	1.72%	1.53%	1.50%
Min	−4.21%	−4.21%	−5.33%	−4.66%
Median	0.32%	0.32%	0.17%	0.09%
Max	11.86%	11.86%	9.84%	6.97%
Information Ratio*	5.85	5.85	3.29	3.74
Post-Convention 8/25-10/20				
Average	0.57%	0.57%	0.47%	0.04%
Standard Deviation	3.27%	3.27%	3.14%	3.08%
Min	−5.86%	−5.86%	−6.32%	−6.45%
Median	0.19%	0.19%	0.19%	−0.04%
Max	10.39%	10.39%	9.43%	8.79%
Information Ratio*	3.34	3.34	2.83	0.27
Election Time 10/21-11/4				
Average	−0.72%	−0.82%	−1.07%	−0.57%
Standard Deviation	3.34%	3.56%	4.08%	3.03%
Min	−7.62%	−8.73%	−11.16%	−9.34%
Median	0.34%	0.34%	0.23%	0.16%
Max	3.12%	3.12%	2.96%	3.07%
Information Ratio*	−4.10	−4.38	−4.99	−3.60

* The Information Ratio is annualized assuming 365 days a year.

loss from the technical problem on the day before the election.

Residual performance statistics using the simulated "buy and hold" returns as the benchmark are presented in Table 4. This table shows statistics for the returns in excess of the benchmark returns for each day. Interestingly, the information ratio (IR)[4] for the pre-convention period is higher than the IR for the post-convention period. This is because the Sharpe ratio for the buy & hold benchmark is higher for the post-convention period, as shown in Table 3. The IR for the two week period before the election is negative because of the technical problem.

The market making profits were further analyzed using an approach similar to the analysis done by Menkveld (2011). The net profit π in any time interval is the average net cash flow, assuming it starts and ends with no position:

$$\pi = \sum_{t=1}^{T} -\Delta n_t P_t \tag{15}$$

In this equation, P_t is the transaction price and n_t is the net inventory position at time t. The equation can be rewritten to decompose the market

[4] The information ratio is calculated in the same manner as the Sharpe ratio except it uses the residual returns instead of the actual returns. The residual returns are the actual returns minus the benchmark returns.

making algorithm's revenue into a spread profit and a positioning profit:

$$\pi = \sum_{t=1}^{T} n_{t-1}\Delta p_t + |\Delta n_t| s_t \tag{16}$$

where p_t is the midquote price and s_t is the effective half spread $|P_t - p_t|$. The positioning profit is then $n_{t-1}\Delta p_t$ and the spread profit is $|\Delta n_t| s_t$. In other words, the spread profit is the profit the market making algorithm earns from buying at prices less than the efficient price and selling at prices higher than the efficient price. The positioning profit is the profit the algorithm earns from favorable price movements of its current asset holdings.

Figure 8 shows the cumulative decomposed market making profits. The figure shows that the cumulative positioning profit was negative. In fact, the daily positioning profit was positive on only 16% of the trading days. In the three week period before the election, the daily positioning profit was positive on only one day. This was expected and is the result of the information asymmetry between the algorithmic trader and the human market participants. The system was continuously holding losing positions but was able to do so profitably because the spread profits exceeded the positioning losses. It was able to accomplish this by not holding any position long enough to (on average) allow the positioning loss to exceed the profit earned on the spread.

Table 5 shows statistics on the daily decomposed market making profits, broken down by time period. The average positioning profit was negative but the maximum positioning profit was positive, indicating that sometimes the market making algorithm profited from price movements. The average spread profit was positive and exceeded the average positioning loss in every period except for the period right before the election. The spread profit was non-negative for every day except for the day before the election when

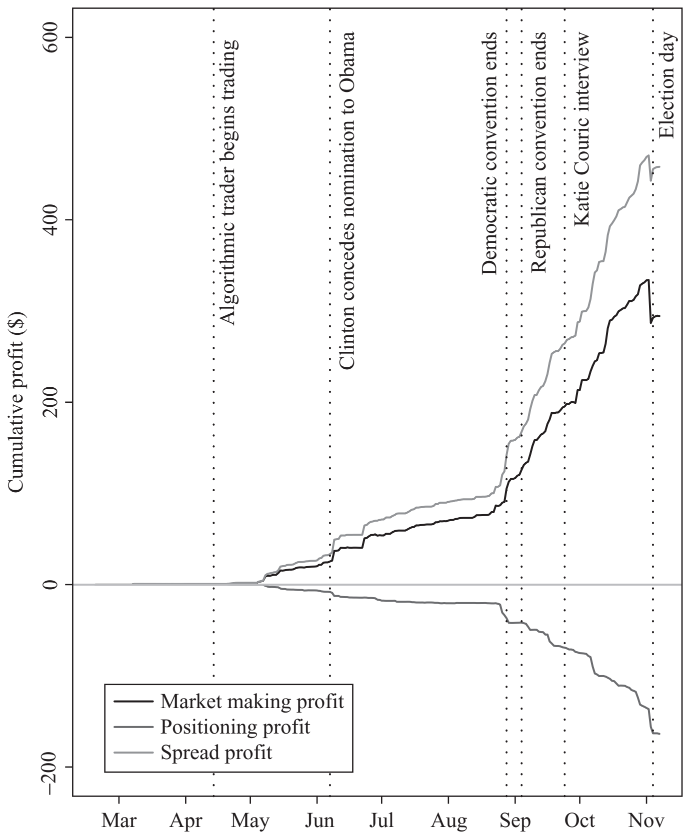

Fig. 8. Decomposition of the cumulative profits from the market making strategy. The positioning profit is the profit or loss from being long or short the market asset when the price changed. The spread profit is the profit from earning the spread. The market making profit is the sum of the positioning profit and the spread profit.

Table 5
Daily decomposition of the profits from the market making strategy by time period. The positioning profit is the profit or loss from being long or short the market asset when the price changed. The spread profit is the profit from earning the spread. The total profit is the sum of the positioning profit and the spread profit.

	Positioning Profit ($)	Spread Profit ($)	Total Profit ($)
Full Period 4/15-11/7			
Min	−19.405	−27.773	−47.178
Median	−0.080	0.695	0.417
Max	0.348	18.253	17.005
Average	−0.790	2.210	1.420
Standard Deviation	2.058	4.363	4.469
Pre-Convention 4/15-8/24			
Min	−3.494	−0.003	−1.470
Median	−0.025	0.273	0.181
Max	0.348	14.130	10.636
Average	−0.158	0.808	0.650
Standard Deviation	0.406	1.782	1.558
Post-Convention 8/25-10/20			
Min	−10.383	0.000	−0.767
Median	−0.700	3.622	2.447
Max	0.197	18.253	17.005
Average	−1.573	5.342	3.769
Standard Deviation	2.286	4.940	4.060
Election Time 10/21-11/4			
Min	−19.405	−27.773	−47.178
Median	−1.306	3.186	1.825
Max	−0.112	17.495	9.832
Average	−3.494	2.921	−0.573
Standard Deviation	5.070	9.662	13.139

the system suffered a technical problem. Also notice that the increased profitability in the post-convention period is largely due to the increased spread profits from market making.

Table 6 shows statistics on the market making algorithm's volume. Trades where the system provided liquidity and took liquidity are analyzed separately. The trades where the algorithmic trader took liquidity are the trades that took place when the good spread was locked and the system used a market order to exhaust an undesired position. Trades that provided liquidity are broken down further to show the volume that traded on the good and bad spreads. The volume that traded on the bad spread exceeds the volume that traded on the good spread, indicating that the majority of the spread profit in Table 5 was from traders that committed price taking violations of rationality and traded on the bad spread. Refer to Figure 1, which shows that the average size of the bad spread is about 1.6 pennies, whereas the average size of the good spread is about 0.5 pennies. It is clear from Equations 13 and 14 that the market making algorithm would be willing to trade in larger sizes when the spread size increased, so it makes sense that the system would trade the majority of its volume on the bad spread to maximize its profits.

Table 7 shows statistics on the daily arbitrage profits. As expected, the profits were positive every day, with profitability increasing during and after the party conventions when market volume increased.

7. Discussion

7.1. Market structure

The structure of the Iowa Electronic Markets helped shape the profit opportunities available to the algorithmic trading application developed in this paper. Since this market had two securities but only

Table 6
Market making strategy's trading volume decomposed by time period. The Good Spread and Bad Spread Volumes are for trades where the market making algorithm traded with limit orders, providing liquidity on either the good or bad spreads. The Provide Volume is for all trades where the algorithm provided liquidity. The Take Volume is for all trades where the system traded with market orders, taking liquidity. This table excludes the 2038 shares traded as a result of the technical problem that occurred on the day before the election.

	Good Spread Volume	Bad Spread Volume	Provide Volume	Take Volume
	Full Period 4/15-11/7			
Total Shares	9,834	34,585	44,419	13,650
Avg. Shares Per Day	47.51	167.08	214.58	65.94
Percent	16.94%	59.56%	76.49%	23.51%
Trade Count	424	1,452	1,876	567
Avg. Trade Size	23.19	23.82	23.68	24.07
	Pre-Convention 4/15-8/24			
Total Shares	1,284	7,390	8,674	3,215
Avg. Shares Per Day	9.73	55.98	65.71	24.36
Percent	10.80%	62.16%	72.96%	27.04%
Trade Count	91	440	531	219
Avg. Trade Size	14.11	16.80	16.34	14.68
	Post-Convention 8/25-10/20			
Total Shares	4,991	19,999	24,990	7,424
Avg. Shares Per Day	87.56	350.86	438.42	130.25
Percent	15.40%	61.70%	77.10%	22.90%
Trade Count	170	697	867	231
Avg. Trade Size	29.36	28.69	28.82	32.14
	Election Time 10/21-11/4			
Total Shares	3,559	6,821	10,380	2,429
Avg. Shares Per Day	237.27	454.73	692.00	161.93
Percent	27.79%	53.25%	81.04%	18.96%
Trade Count	163	297	460	102
Avg. Trade Size	21.83	22.97	22.57	23.81

one real asset, it was possible for traders to trade inefficiently and commit market-making and price-taking violations of rationality. If there was only one security this would not be possible. The algorithmic trader profited from these violations of rationality. All of the algorithmic trader's arbitrage profits were the result of traders who committed market making violations of rationality. Table 6 shows that the majority of the market maker's trades that provided liquidity to other traders were on the bad spread. This means that those liquidity takers were committing price-taking violations of rationality. If they had traded more efficiently on the good spread the market maker's profits and volume would have been lower.

It is easy to criticize the traders who committed violations of rationality, but the structure of the market actually made it difficult to be rational. Consider the case of a trader with $100 who wants to invest the full amount in the DEM security. If these are the market prices:

Security	Bid	Ask
DEM	0.572	0.578
REP	0.424	0.450

The rational trading decision is to buy bundles and sell the REP security, not buy the DEM security. The good spread is constructed with the DEM and REP bid prices, so buying the DEM security on the offer price would be a violation of rationality.

Imagine that the trader buys 100 bundles for $100 and sells 100 REP securities at the bid price of 0.424. The trader now has 100 DEM securities and $42.40. The trader has not yet invested the full amount in the DEM security so the trader must now buy 42 bundles and continue to sell REP securities at the bid price of 0.424. But this time, imagine that the trader can only

Table 7
Daily Arbitrage profit by time period. The arbitrage profit is the risk-free profit earned from trading pairs of securities at prices above or below $1.

	Arbitrage Profit($)
Full Period 4/15-11/7	
Min	0.000
Median	0.195
Max	25.651
Average	1.166
Standard Deviation	2.617
Pre-Convention 4/15-8/24	
Min	0.000
Median	0.064
Max	13.377
Average	0.718
Standard Deviation	1.913
Post-Convention 8/25-10/20	
Min	0.000
Median	0.872
Max	25.651
Average	2.167
Standard Deviation	3.831
Election Time 10/21-11/4	
Min	0.035
Median	1.525
Max	4.379
Average	1.527
Standard Deviation	1.181

sell 20 REP securities before the REP 0.424 bid order is depleted. The trader now has 142 shares of the DEM security, 22 shares of the REP security, and $8.88. If these are the new market prices:

Security	Bid	Ask
DEM	0.572	0.578
REP	0.418	0.450

The trader would be committing a price-taking violation of rationality by continuing to sell the REP security. Because of the market impact of the previous trade, the good spread is now constructed with the DEM bid and offer prices. In this situation the rational trader should proceed by first selling 22 bundles and then buying the DEM security.

After selling the bundles for $22 the trader now has 120 shares of the DEM security and $30.88. If the trader is able to buy 53 DEM securities on the offer price of 0.578, the trader would have achieved the desired position after a total of 6 transactions. More transactions would be required if 53 DEM securities were not available for sale at that price.

Compare the trader described above with a trader who willingly commits a price-taking violation of rationality and is able to buy 173 DEM securities on the offer price of 0.578. The less rational trader made one transaction and had execution costs that are only $0.24 higher than the rational trader. Small improvements in execution costs may not motivate all traders to trade as efficiently as they can.

The above scenario is realistic, particularly for traders with larger accounts who are more likely to not find sufficient liquidity and impact prices. It is easy to understand why a trader would accidentally or intentionally commit a violation of rationality in this market.

A structural change to the IEM that would eliminate these problems is the addition of phantom orders. With phantom orders, limit orders for the DEM and REP securities would be mirrored in each other's order books. If a trader attempted to trade with a phantom order in one order book, the exchange would automatically execute a bundled transaction and trade the other security. The concept of a good spread and a bad spread would be eliminated because the spreads for each security would be equivalent. This would simplify the number of transactions needed to short securities and trade without violations of rationality. In addition, this would lower traders' execution costs.

7.2. *Comparison with other markets*

Although the algorithmic trader was successful and profitable, it is questionable how well the market making model described in Equations 13 and 14 would work in a more developed financial market like the modern equity or bond markets. The basic concept would still be valid but most likely it would need to be modified to be applied successfully. An electronic market maker would likely need to trade many securities simultaneously and manage risk across the entire portfolio. The risk model in equation 11 could be replaced with a model that estimates the risk of a group of securities. Two possibilities are a factor model or a covariance matrix based model. The market impact model in equation 9 could be replaced with a different statistical model that predicts short term movements of the efficient price. Also, a market maker would need to have more sophisticated logic to handle order placement across multiple trading exchanges and make decisions about if it should place limit orders at prices equal to the best bid or offer, at a price better or worse than the current best bid or offer, or not at all. It would be important to minimize the amount of trading done with informed counterparties.

Another important point is that the counterparties a market maker would be trading with in the modern equity or bond markets would be more sophisticated that the traders found in the IEM. Often counterparties would be other algorithms looking to accumulate large positions over an extended period of time. The market maker could experience larger positioning losses before it could find another counterparty to trade with to offset its position. This is very different from the counterparties found in the IEM, who would typically trade in short bursts and would make buy and sell decisions that were largely uncorrelated with the trading decisions of previous market participants. Also, the algorithms would trade more efficiently and would rarely make the equivalent of any "violations of individual rationality."

The arbitrage strategy described in section 3.3 could work in any modern financial market that has multiple securities with identical risk exposure. Instead of either buying or selling pairs of securities in bundled transactions, a very fast trading system could buy and sell a pair of securities, netting a small profit without any risk exposure. If a pair of securities with identical risk exposure could be found that had different sized spreads, an electronic trader could employ the second arbitrage strategy described in section 3.3. This would work because the sizes of the orders in the order books would be visible. It would be necessary to have a very fast trading system because most modern financial markets already have many electronic traders employing strategies similar to the arbitrage strategies described herein. A new electronic trader would either need to be faster than its competitors or more creative about how the strategies are applied.

The arbitrage strategy could also be employed on baskets of securities instead of two individual securities. As long as the basket of securities had the correct risk exposure, the strategy would still be profitable. This concept is the foundation of index arbitrage.

7.3. Benefits of an electronic market maker

Another important question is if the Iowa Electronic Markets benefited from the presence of an electronic market maker. This can be evaluated using the same approach used in the literature to examine market efficiency and average or marginal traders.

Oliven and Rietz (2004) discussed the concepts of price taking and market making violations of rationality. By construction, the algorithmic trader described in this paper never committed either of these violations. It is reasonable to assert that the market benefited from a market participant that traded a significant amount of volume and never made either of these mistakes. The computer program was also incapable of judgment biases such as the "assimilation-contrast effect" or the "false-consensus effect" described in Forsythe et al. (1992). It also could not have an overconfidence bias, as discussed in Berg and Rietz (2010).

Oliven and Rietz (2004) argued that markets can be efficient despite the presence of biased market participants because a few unbiased arbitrageurs would eliminate the effects of the biased traders. Forsythe et al. (1992) suggested that there were some traders who did not have any judgment biases and profited from acting as arbitrageurs. These unbiased traders would sell a security to the biased supporters of the associated candidate and would buy the security from biased supporters of the other candidate. This is precisely what the algorithmic trader described in this research paper did. The computer program did not know which candidate was going to win the election or which traders held unbiased views about the future outcome of the election. The only thing that mattered was that supporters of both candidates placed market orders for their candidate's securities at around the same time. The market making algorithm would arbitrage one biased trader's views against another, and make a profit as a result.

Forsythe et al. also discussed news events and argued that marginal traders would be able to correctly assess the relevance or irrelevance of news events and would consistently buy or sell securities appropriately. A marginal trader would recognize relevant news and profit from the foreseen price changes. This is one criterion where the algorithmic trader cannot be considered marginal. Because this system could not interpret news, it was never able to foresee price changes and failed to achieve a positioning profit on 84% of the trading days and all but one of the trading days in the three weeks before the election. The system was never able to identify when important news events happened and when they did not. This was not a problem. Since there were always biased traders who also failed to correctly interpret news events, the market making algorithm was able to earn a spread profit that exceeded the positioning loss the majority of the time.

Another way the algorithmic trader cannot be considered to be a marginal trader and did not aid in the efficiency of the market has to do with the technical problem that affected the system on the day before the election. The technical problem is a frustrating blemish on an otherwise successful venture. When evaluating the market making model in Equations 13 and 14, the data for that day can be excluded[5] because that data does not represent what the model can achieve. When evaluating the presence of algorithmic traders in the Iowa Electronic Markets, this event cannot be ignored. Technical problems are part of the reality of algorithmic trading and it is only fitting that this problem happened here. The possibility or actuality of an electronic trader behaving irrationally or nonsensically is one way electronic traders are not marginal traders and negatively affect the market's ability to aggregate information and predict future events.

Unfortunately this problem happened on the day before the election, a day some researchers use to derive the market's ultimate decision of what the election outcome will be. Thankfully it appears that market prices recovered very quickly. This rapid recovery is a testament to the robustness of the Iowa Electronic Markets and its ability to compensate for a mistake of one temporarily irrational trader.

7.4. Future research

New securities linked to political events, including the 2012 U.S. Presidential election, are constantly being made available for trading in the Iowa Electronic Markets. The availability of profit opportunities and the information contained in this research paper means it is likely there will be more competition for the potential profits. This competition will make the market more efficient and a better aggregator of information. An increase in electronic market making and more sophisticated market making strategies will result in more liquidity, smaller spreads, and less market impact for trades. More arbitrageurs will also create more competition for arbitrage opportunities. As a result, failed arbitrage attempts will become a more frequent occurrence.

More competition for trading profits in the Iowa Electronic Markets means employing the strategies documented in this paper may not be as successful in the 2012 election market as they were in the 2008 election market. It will be more challenging to achieve the same profit and Sharpe ratio, but the Iowa Electronic Markets will benefit from this increased competition. It is this process of competition through technological and quantitative innovation that causes markets to evolve and continue to become more efficient.

[5]Except for Table 6, every table and figure in this paper analyzing the algorithmic trader's profitability included the data for that day. Table 6 analyzes the market making strategy's trading volume with the erroneous trades removed.

8. Acknowledgments

I would like to thank Ritesh Bansal, Jacek Janczewski, Petter Kolm, Lee Maclin, Philip Maymin, Jeff Miller, Georgios Petropoulos, Wolfgang Ruess and William White for helpful comments and suggestions. I would also like to thank an anonymous referee for their insight and meaningful comments. Any errors are strictly my own.

Finally, I would like to thank Jim Gatheral, Lee Maclin, and Nassim Taleb, whose love of mathematics and finance inspired me to pursue this research project.

References

Berg, J.E., Rietz, T.A., 2006. The Iowa electronic markets: stylized facts and open issues. In: Hahn, R., Tetlock, P. (Eds.), Information Markets: A New Way of Making Decisions. AEI/Brookings Center for Regulatory Studies, pp. 142–169. http://tippie.uiowa.edu/iem/research/abstracts/bergrietz_2006.html, http://www.bibsonomy.org/bibtex/2be8d70fcfb118a63cc71fc1815869826/cgz.

Berg, J.E., Rietz, T.A., 2010. Longshots, overconfidence and efficiency on the Iowa electronic markets. Working Paper, Tippie College of Business, University of Iowa. http://tippie.uiowa.edu/iem/research/abstracts/bergrietz_2002.html, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1645062.

Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A., 2008. Results from a dozen years of election futures markets research. In: Plott, C.R., Smith, V.L. (Eds.), The Handbook of Experimental Economics Results. Elsevier Science, Amsterdam.

forthcoming. http://tippie.uiowa.edu/iem/research/abstracts/bergforsythenelsonrietz_2008.html, http://econpapers.repec.org/bookchap/eeeexpchp/5-80.htm.

Borghesi, R., 2007. Price biases in a prediction market: NFL contracts on tradesports. J. Prediction Markets. 1 (3), 233–253(21). University of Buckingham Press. http://www.ingentaconnect.com/content/ubpl/jpm/2007/00000001/00000003/art00005.

Brahma, A., Das, S., Magdon-Ismail, M., September 2010. Comparing prediction market structures, with an application to market making. *ArXiv e-prints*.

Das, S., 2008. The effects of market-making on price dynamics. In: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems – Volume 2. International Foundation for Autonomous Agents and Multiagent Systems. Estoril, Portugal.

Forsythe, R., Nelson, F., Neumann, G.R., Wright, J., 1992. Anatomy of an experimental political stock market. Am. Econ. Rev. 82 (5), 1142–1161. American Economic Association. http://www.jstor.org/discover/10.2307/2117471?uid=3739832&uid=2129&uid=2&uid=70&uid=4&uid=3739256&sid=21100835301151.

Gil, R., Levitt, S.D., 2007. Testing the efficiency of markets in the 2002 World Cup. J. Prediction Markets. 1, 255–270. http://ideas.repec.org/a/buc/jpredm/v1y2007i3p255-270.html.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. J. Financ. Econ. 14 (1), 71–100. Elsevier, http://ideas.repec.org/a/eee/jfinec/v14y1985i1p71-100.html.

Hanson, R., 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. J. Prediction Markets. 1 (1), 3–15. http://ideas.repec.org/a/buc/jpredm/v1y2007i1p3-15.html.

Huberman, G., Stanzl, W., 2004. Price manipulation and quasi-arbitrage. Econometrica. 72 (4), 1247–1275. The Econometric Society. http://www.jstor.org/discover/10.2307/3598784?uid=3739832&uid=2129&uid=2&uid=70&uid=4&uid=3739256&sid=21100835301151.

Menkveld, A.J., 2011. High frequency trading and the *New-Market* makers. EFA 2011 Paper.

Oliven, K., Rietz, T.R., 2004. Suckers are born but markets are made: individual rationality, arbitrage, and market efficiency on an electronic futures market. Manag. Sci. 50 (3), 336–351. INFORMS Institute for Operations Research and the Management Sciences (INFORMS), Linthicum, Maryland, USA©2004. http://dl.acm.org/citation.cfm?id=1014288.

Rhode, P.W., Strumpf, K., 2008. Manipulating political stock markets: a field experiment and a century of observational data. http://www.unc.edu/∼cigar/strumab.htm.

Tetlock, P.C., May 1, 2008. Liquidity and prediction market efficiency. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=929916.