Trading Strategies and Market Microstructure: Evidence from a Prediction Market*

David M. Rothschild† Rajiv Sethi‡

November 22, 2015

Abstract

We examine transaction-level data from Intrade's 2012 presidential winner market for the entire two-year period for which trading occurred. The data allow us to compute key statistics, including volume, transactions, aggression, directional exposure, holding duration, margin, and profit for each of 6,300 unique trader accounts. We identify a diverse set of trading strategies that constitute a rich market ecology. These range from arbitrage-based strategies with low and fleeting directional exposure to strategies involving large accumulated positions in one of the two major party candidates. Most traders who make directional bets do so consistently in a single direction, unlike the information traders in some canonical models of market microstructure. We present evidence suggestive of manipulation by a single large trader, and consider the possible motives for such behavior. Broader implications for the interpretation of prices in financial markets and the theory of market microstructure are drawn.

*This project was initiated while Sethi was a visiting researcher at Microsoft Research, New York City. We thank Ron Bernstein for providing us with transaction level data and Suxin Shen for exceptional research assistance. No user identities were disclosed to us by Intrade.

†Microsoft Research, New York City

‡*Corresponding Author*: Department of Economics, Barnard College, Columbia University and the Santa Fe Institute. Phone 212-854-5140. Email: rs328@columbia.edu.

1 Introduction

A central concern in the economics of finance is the process by means of which asset prices adjust in response to new information about the size and certainty of promised income streams. In order for prices to respond to news, there must exist a set of market participants who find it profitable to acquire information and trade on this basis. But this raises the question of how such informed traders are able to find willing counterparties.

Two broad approaches this question have been taken in the theoretical literature. In the classic models of Copeland and Galai (1983), Kyle (1985), and Glosten and Milgrom (1985), market-makers operating under competitive conditions are willing to trade with all counterparties, including those with superior information, provided that there also exist some traders with pressing liquidity needs. The losses incurred by this latter set of individuals allow for positive returns to information traders, while intermediaries break even. In these models market participants may have heterogeneous information, but form beliefs based on a common prior. An alternative approach to trading is based on the hypothesis that market participants have heterogeneous priors or mental models, and hence differ in their interpretations of public information (Miller, 1977; Harrison and Kreps, 1978). In this framework trading occurs on the arrival of new information because individuals, even if symmetrically informed, agree to disagree about its payoff implications.[1]

These models (as well as many others descended from them) are based on a partition of the trading population into multiple groups, each using a distinctive strategy. In principle, one could distinguish among theories by examining the empirical distribution of trading strategies in an asset market. Deducing the nature of a strategy from a sequence of trades is not an easy task, however, especially when large numbers of individuals are interacting simultaneously with each other and responding rapidly to changes in market states and the arrival of information. At a minimum, one needs transaction level data in which each trade can be associated with a buyer and seller account.

Recent work by Kirilenko et al. (2011) has used transaction-level data from the S&P 500 E-Mini futures market to partition accounts into a small set of groups, thus mapping out an "ecosystem" in which different categories of traders "occupy quite distinct, albeit overlapping, positions." Their concern was primarily with the behavior of high frequency traders both before and during the flash crash of May 6, 2010, especially in relation to liquidity provision. Subsequent work by Baron et al. (2012) has used similar data over a longer time frame to compute the risk-adjusted profits of high frequency traders.

In this paper we examine transaction-level data from a very different source, Intrade's prediction

[1]See also Harris and Raviv (1993), Kandel and Pearson (1995), Scheinkman and Xiong (2003), and Hong et al. (2006). Equilibrium models with common priors and private information, such as Hellwig (1980) and Grossman and Stiglitz (1980), explore the question of whether prices fully aggregate all dispersed information, but do not consider the process of trading.

market for the 2012 US presidential election. With wagers topping $230 million on the 2012 US election, this was the largest exchange in both volume and transactions for this event.[2] The first trade in this market occurred on November 16, 2010, and the last was executed on November 7, 2012, a few hours after all polls had closed. When the market opened it was not clear who the eventual nominees would be, and a total of 22 different contracts were listed.[3] Trading was initially active in all contracts, but as it became clear that the election would be contested by Obama and Romney, volume in the remaining contracts diminished substantially.

Contracts on prediction markets such as Intrade are structured as binary options that pay a fixed amount if the referenced candidate is elected, and nothing otherwise. We obtained data on each transaction in this market for the entire two years of its existence, including price, quantity, time of trade, and aggressor side, as well as a unique identifier for each account that allowed us to trace the evolution of trader positions and profits. There were 6,300 unique accounts, almost 300,000 total transactions, and close to 13 million contracts traded in total. No identities could be deduced from this data, but it is possible to make inferences about strategies from the pattern of trades.

The data allows us to compute volume, transactions, aggression, directional exposure, holding duration, margin, and profit for each of the trader accounts. We find that a relatively small number of highly active traders dominate total volume and transactions. Most traders either take liquidity consistently or provide it consistently, although the most active traders have intermediate levels of aggression. The vast majority of traders hold their positions to expiration but a small number of high volume traders have extremely short holding periods, closing out positions almost as soon as they have entered them.

Based on a range of computed characteristics, we partition the full set of accounts into six categories that we associate with distinct trading strategies. One of these is arbitrage, which seeks to exploit temporary inconsistencies between the prices of the two contracts, and tends to be algorithmically implemented. The others are unidirectional or exhibit varying levels (extreme, high, moderate and low) levels of bias. A unidirectional trader's portfolio is always long for one of the two major candidates, while a low bias trader's portfolio fluctuates between long in one candidate and long in another. Low or moderate bias traders that are not arbitrageurs most closely resemble the information traders in standard models.

Having developed a taxonomy based on the behavior of traders we turn to their beliefs and

[2] Intrade closed the accounts of all US residents in December 2012 in the wake of a suit brought by the CFTC alleging that the company "solicited and permitted" US persons to trade commodity options without being a registered exchange. It ceased all trading activity in March 2013 and remains closed.

[3] All but one of the contracts referenced a particular potential candidate (Obama, Clinton, Biden, Romney, Palin, Thune, Gingrich, Pawlenty, Daniels, Huckabee, Paul, Trump, Bachmann, Huntsman, Barbour, Cain, Santorum, Perry, Christie, Bloomberg, and Johnson). The last contract referenced Other, and would have paid out if none of the listed candidates had won.

motivations. Traders could be engaging in cross-market arbitrage, hedging exposures in more conventional asset markets, attempting to manipulate prices, or directly speculating on the electoral outcome (the canonical reason for trading). Contracts very similar to those on Intrade were also offered on Betfair, the world's largest prediction market, and a persistent price disparity between the two created incentives for arbitrage across exchanges. This can only account for unidirectional strategies in one direction, however, since the Romney contract was consistently overpriced on Intrade relative to Betfair. Hedging may have been a motive, given that prior research has linked electoral outcomes to movements in both broad market indexes as well as particular baskets of securities (Knight, 2006; Snowberg et al., 2007). We show, however, that at least as far as market indexes are concerned, prior correlations appear not to have been relevant to the 2012 election. This also suggests that price manipulation for financial gain, by temporarily affecting the S&P futures market for instance, would not have been effective.

It does appear, however, that price manipulation by a single trader who accumulated a large directional position on Romney may have been a factor. This trader accounted for more than one-third of all bets placed on Romney to win, and about one-fifth of all bets placed on Obama to lose.[4] These positions were accumulated by placing large bids for Romney and large offers for Obama that effectively created a firewall, preventing prices from moving in response to incoming information. This resulted in remarkable stability in Intrade prices for several hours on Election Day, and at other critical moments of the campaign, even as prices on Betfair were moving sharply. On Election Day, these orders were removed just as voting ended in Colorado, the last swing state to close its polls. The effect was a sharp price movement and immediate convergence to the Betfair prices. Financial gain though correlated changes in stock prices seems an unlikely motivation for this activity, since these appear to have broken down in 2012. More plausibly, this trader could have been attempting to manipulate beliefs about the odds of victory in an attempt to boost fundraising, campaign morale, and turnout.

Manipulation aside, one of our most striking findings is that 87% of traders, accounting for 32% of volume, never change the direction of their exposure even once. A further 37% of volume comes from 7% of traders who are strongly biased in one direction or the other. A handful of arbitrageurs account for another 16% of volume, leaving just 6% of accounts and 15% of volume associated with individuals who are unbiased in the sense that they are willing to take directional positions on either side of the market. This suggests that information finds its way into prices largely through the activities of traders who are biased in one direction or another, and differ not only with respect to their private information but also with respect to their *interpretations* of *public* information. In this sense our findings are supportive of models of trading based on heterogeneous prior beliefs, such as Miller (1977) and Harrison and Kreps (1978). Fundamental belief heterogeneity of this kind is an increasingly common assumption in models of speculation, given the difficulty of accounting for trade under the common prior assumption.[5] We argue here that heterogeneous prior beliefs

[4]The trader held the vast majority of these contracts to expiration, but closed out a small number on earlier dates.

[5]See Milgrom and Stokey (1982) for a classic no-trade theorem, and Harris and Raviv (1993), Kandel and Pearson

imply a certain distribution of directional exposure across traders in binary options markets, and that this prediction is consistent with the patterns observed in our data.

2 Data

The data come from Intrade's presidential election market, and include a record for each transaction completed during the two year period over which the market was open.

The contracts were structured as cash-or-nothing binary options; for instance, the Obama contract paid \$10 contingent on his winning the presidential election (and nothing otherwise), while the Romney contract paid \$10 contingent on a Romney victory. These contracts could be bought and sold just like stocks or futures contracts in financial markets. As in other electronic exchanges, the mechanics of trading involved the use of a continuous double auction. That is, traders could place a bid to buy or an offer to sell a specified number of contracts at a specified price. If an incoming order was compatible with an earlier order that had yet to trade, a transaction would occur. If no such compatible order existed, then the new order would be recorded and made available to trade against future orders. This list of orders that have been received but await a compatible counterparty is referred to as an order book. Whenever a trade occurs, it involves a new, incoming order and an older one that is resting in the order book. The former is said to be *aggressive* or liquidity-taking, and the latter *passive* or liquidity-providing. A new order that results in a trade is said to be *marketable*.

Each observation in our data identifies the contract traded, time of trade, price, quantity, buyer and seller accounts, and the aggressor side. Accounts are identified only by a unique number so that trader anonymity is fully protected. The timestamp on each trade allows us to compute average holding periods for each trader in a manner described below. By tracking trader portfolios over the period, we can also determine changes in amount of money risked, as well as the aggregate profit or loss. The exchange required traders to have cash deposits that were enough to cover their worst-case loss. This amount is referred to as the trader's *margin*.

2.1 Volume

There were about 287,000 distinct transactions over this two year period, involving 12.9 million contracts and 6,300 unique accounts.[6] The average transaction was for about 45 contracts, while the largest single trade involved ten thousand.

(1995), Scheinkman and Xiong (2003), Hong et al. (2006), Geanakoplos (2010) and Simsek (2013) for models of speculation based on heterogeneous priors. When prior beliefs are unobservable, then private information need not be fully revealed even if posteriors are common knowledge (Sethi and Yildiz, 2012).

[6]In all, 7.6 million traded contracts referenced either Obama or Romney, since volume in the remaining contracts dried up once the decisive primaries were over.

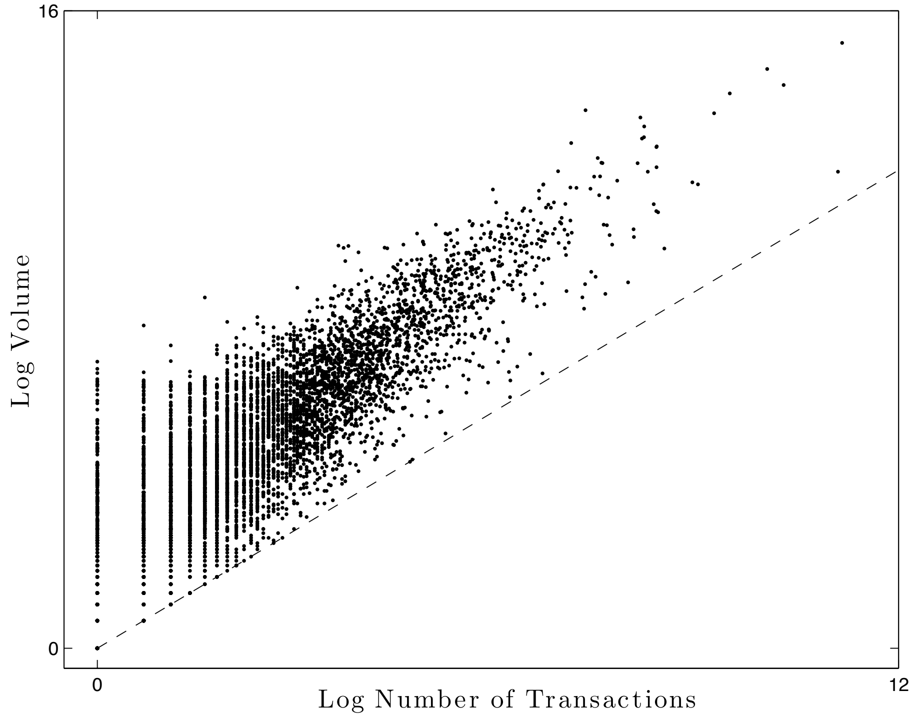

Figure 1: Distribution of Transactions and Volume

The distribution of transactions and volume (both in logs) is shown in Figure 1, where each point corresponds to a distinct trader. Volume refers to the total number of contracts traded, which must be at least as great as the number of transactions. Hence all points lie on or above the line defined by the equality of the two. The positive association between volume and trading frequency is clear from the figure, but there is also considerable variation across traders in the size of the average trade.[7] Many traders engaged in just a handful of transactions over this period: 55% had less than ten trades, and 12% had just one. But some of these had fairly high volume. Of the traders with just a single transaction, for instance, 78 had volume exceeding a hundred contracts and three had volume exceeding a thousand.

More generally, both volume and total transactions are dominated by a relatively small number of traders. The single largest trader was responsible for more than 15% of total volume, with close to four million contracts bought or sold in about 70,000 distinct transactions.[8] This was also the most frequent trader, accounting for more than 12% of total observations. The largest 63 traders

[7]The number of transactions is an upper bound for the number of marketable orders, since a single marketable order that trades against more than one resting order will appear as multiple transactions in our data.

[8]Since two parties are jointly responsible for trade in a single contract, aggregating volume across all traders results in double the number of contracts traded. The most a single trader can contribute to this total is one-half. The same applies to total transactions.

(just 1% of the trading population) were responsible for 67% of volume, while the 63 most frequent traders accounted for 60% of transactions.

Table 1 contains information about ten selected traders. The first two columns contain volume (the total number of contracts traded) and trades (the total number of distinct transactions). Each transaction here has a unique price and counterparty, so a single large marketable order can give rise to multiple transactions. The third column, scope, contains the number of distinct candidate contracts traded, and ranges between 1 and 22. Aggression is the share of total volume that resulted from a marketable order placed by the trader in question; these are orders that trade against resting orders and remove liquidity. Direction is defined more precisely below, and refers to the extent to which exposure increasing trades were bets on Obama (coded as positive direction) or Romney (coded as negative direction). This measure ranges between $+1$ for traders who were always betting on Obama (or against Romney), to -1 for those who were doing the opposite. Holding is the median amount of time, in seconds, that elapses before an exposure increasing trade is closed out or reversed, and duration is a normalized measure of holding; the computation of both is described in detail below. Margin is the largest amount of money that the trader has at risk at a single point in time over the entire cycle; this is the most that the trader could conceivably have lost. The last column is just the realized profit or loss.

	Volume	**Trades**	**Scope**	**Aggression**	**Direction**	**Holding**	**Duration**	**Margin**	**Profit**
A	3,961,242	69,977	22	0.77	0.19	0	0.00	\$9,877	\$61,871
B	2,062,908	22,738	2	0.27	-1.00	520,428	1.00	\$6,882,186	-\$6,882,186
C	1,380,406	29,134	22	0.31	-0.49	2,491	0.00	\$737	\$11,921
D	321,818	1,207	2	0.83	0.92	51,470	1.00	\$2,099,441	\$867,059
E	174,712	4,340	22	0.79	-0.13	7	0.00	\$415	\$1,058
F	156,413	65,652	22	0.72	0.38	0	0.00	\$1,375	\$2,147
G	138,264	1,707	9	0.39	1.00	1,802,885	1.00	\$535,018	\$318,975
H	72,563	392	2	0.73	1.00	149,180	1.00	\$479,896	\$233,414
I	68,416	858	2	0.37	-1.00	664,333	1.00	\$121,609	-\$121,609
J	44,195	350	1	0.64	-1.00	152,360	1.00	\$149,998	-\$149,998

Table 1: Characteristics of Selected Traders

The traders in Table 1 are those who traded at least 10,000 contracts referencing the two main candidates (Obama and Romney) and were among the top three overall with respect to one or more of the following criteria: volume, trades, frequency, margin, profit, or loss (here high frequency corresponds to low duration). They are ordered by total volume.[9] Traders A and F both had zero median holding period. That is, more than half of all exposure increasing trades were closed out within the same second. In fact, for trader A, more than half were reversed within the same *millisecond*—these were mutually offsetting trades with exactly the same timestamp. Such arbitrage strategies can only be implemented algorithmically. Trader A was also the largest by

[9]Many accounts had greater volume and transactions than some on this list. For instance, trader J was 56th on volume, 33rd on volume in the two main contracts, 188th on trades, 19th on margin risked, and 2nd on total loss.

volume and transactions overall, and trader F had the second highest number of trades. Trader D was the most profitable, while B incurred the greatest loss.

The two largest traders adopted completely different strategies, with one of them building up large directional exposure and holding almost all contracts to expiration, while the other engaged in arbitrage, selling all 22 contracts almost simultaneously, with virtually no sustained directional exposure. First consider trader B, who traded only contracts referencing the two major party candidates. The evolution of this trader's positions over the 120 day period leading up to the election are shown in Figure 2. He consistently bought the Romney contract and sold the Obama contract, building a large directional position over time. This required increasing amounts of margin and a substantial exposure (close to 7 million dollars) to a Romney loss on Election Day.

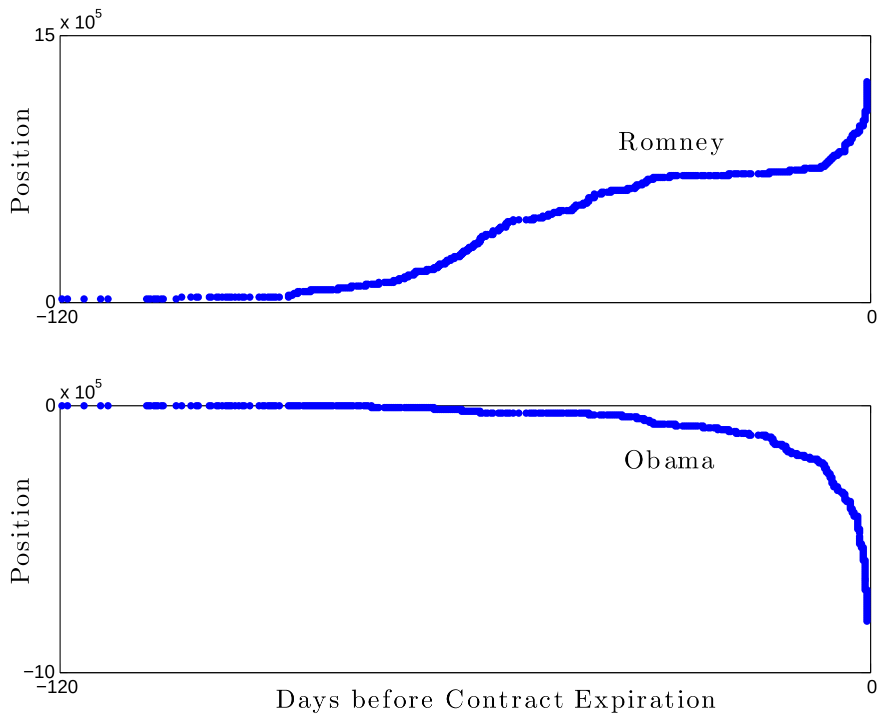

Figure 2: Evolution of Positions for Trader B

Next consider trader A, who traded each of the 22 listed contracts. The evolution of this trader's positions in the two major party nominees over the 120 day period leading up to the election are shown in Figure 3. He accumulated large and virtually identical short positions in these two contracts (and in all other contracts) over time. That is, he bet on *both* Obama and Romney to lose, at prices that ensured a certain profit (since one of the two was guaranteed to lose). This requires a minimal amount of cash margin. The trader quickly covered directional positions in one contract with exactly offsetting positions in the complementary contracts. This behavior is strongly suggestive of an arbitrage based strategy, a point to which we return below.

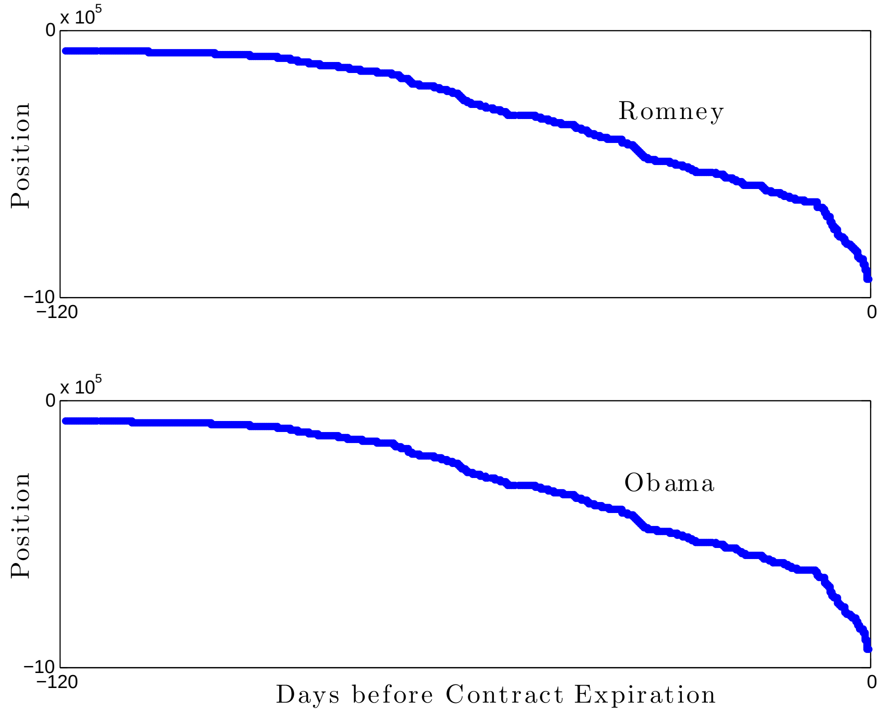

Figure 3: Evolution of Positions for Trader A

3 Characterizing Strategies

While the strategies of the two largest traders are quite easy to characterize based on visual inspection, this is not the case for the majority of market participants. In order to describe the ecology of strategies in the market as a whole, we compute several measures for each trader using the transaction level data: aggression, direction, exposure, duration, margin and profits. Some of these measures require that we restrict attention only to traders who bought or sold at least one contract in one of the two major party nominees. There are 5,906 such traders, constituting 94% of the total trading population, together accounting for 97% of total volume across all contracts.

3.1 Aggression

Each order has two transacting parties, one of which is passive and provides liquidity, while the other is aggressive and takes liquidity. The aggressive party is the one who initiates the trade by placing an order that is immediately marketable against a resting limit order previously placed by the passive party. For each trader we can compute an index of aggression, defined simply as the proportion of this trader's total volume (in the two major contracts) that was initiated by an aggressive order by the trader. There are substantial variations across traders in the degree of aggression: 5% are always passive, 37% are always aggressive, and the remainder lie somewhere in between, sometimes taking and sometimes providing liquidity.

Common arbitrage based strategies gives rise to intermediate values of the aggression measure by design. For instance, consider the strategy of placing an offer to sell each contract at a price that, in conjunction with the current bids in the complementary contracts, yields an arbitrage profit. If an offer is met the trader immediately covers by selling all other contracts at the prevailing bid prices. The first trade is passive, while the remainder are aggressive. When only two contracts are involved the resulting aggression measure is exactly one-half. If, in addition to this, the strategy sells all contracts aggressively when the sum of the bid prices exceeds the amount to be paid to the winning contract, the overall aggression measure will lie between one-half and 1. The three traders in Table 1 who had the shortest median holding period—between 0 and 7 seconds—all had aggression measures consistent with such strategies. In contrast, trader C has a short median holding period but is mostly passive, suggesting a more traditional market making strategy of placing bids and asks simultaneously for multiple contracts and waiting for these to trade.

3.2 Direction

Restricting attention only to trades in contracts that reference one of the two major party nominees, we assign the number $+1$ to a contract that increases exposure to an Obama loss, and -1 to a contract that increases exposure to a Romney loss. These numbers are then averaged over all exposure-increasing contracts to obtain an index of directional exposure.

A trader with direction $+1$ is never betting against Obama. Such a trader may be buying and selling Obama contracts repeatedly, but all sales are exposure reducing while all purchases are exposure increasing. Similarly, an index of -1 indicates a trader who is never betting against Romney. Values of direction close to zero correspond to traders who switch exposure from Obama to Romney or vice versa at some point, perhaps repeatedly, and whose exposure increasing trades do not systematically favor one candidate or the other.

In constructing a taxonomy of trading strategies, we will be more interested in the size rather than the sign of direction. Accordingly, we define the *bias* of a trading strategy as the absolute value of its direction. We see from Table 1 that the three traders with the largest loss all had bias equal to 1, as did two of the there traders with the largest gain. Trader D, who had the largest profit in the data, had bias of 0.92. These traders also had duration equal to 1. In contrast, the four traders with bias below 0.5 also had duration close to zero. These ten traders either have high frequency and low bias, or low frequency and high bias.

Very high levels of bias are common among traders and account for a sizable portion of volume. More than 89% of traders, responsible for over 58% of volume in the two main contracts, have bias exceeding 0.9. In fact, 87% of traders, accounting for 32% total volume, have bias precisely equal to 1. This latter group, which includes traders G, H, I and J in Table 1, do not change the direction of their exposure even once.

3.3 Holding Duration

For any given trader we have a list of all trades and the times at which they occur. Let (q_i, t_i) denote the quantity traded and timestamp for the ith transaction involving the trader. Restricting attention only to trades in one of the two major party nominees, we classify a trade as long Obama if it involves either a purchase of the Obama contract or a sale of the Romney contract. Otherwise the trade is long Romney. At any given time, a trader's net position is either long Obama, or long Romney, or zero. Each trade either increases the trader's directional exposure or reduces it.

Starting with the first trade, which clearly increases exposure, we identify the earliest time t_j at which exposure is reduced. If the number of contracts in the first trade is q_1 while the number in the direction-reversing trade is q_j, we compute $q = \min\{q_1, q_j\}$ and assign the time difference $t_j - t_1$ to precisely q contracts from the first trade. We then subtract q from both q_1 and q_j, resulting in a new series of trades and times to which the same procedure can be applied. This is iterated until all trades have been reversed, with any residual position reversed at the time of contract expiry. This procedure yields a holding period for each contract purchased or sold that increases directional exposure. Trades that reduce directional exposure simply close out or cover a position previously entered. One measure of duration is the *median* holding period for all exposure increasing trades.

One problem with this measure is that for trades with the same timestamp, the order in which they appear in the data affects the computation of the measure. This is easy to deal with, however, by treating mutually offsetting trades with the same timestamp separately, and assigning them zero holding period. The holding period for other exposure increasing trades may then be computed after removing these from the data. We adopt this procedure, but after first rounding the timestamps to the nearest second (the original data had timestamps in milliseconds). That is, we treat transactions that occur within the same second as if they occur simultaneously.

As in the case of aggression, there is wide variation in holding periods. Three traders had median holding precisely equal to zero: more than half of exposure increasing trades were reversed within the same second. Two of these (A and F) are shown in Table 1; the third traded just fourteen main market contracts and is omitted from the table. Trader G in Table 1 had median holding about three weeks; across all traders the longest holding period was almost two years.

For traders who never reduce their exposure to a particular candidate, the median holding period just reflects the median time before contract expiration at which they enter their positions. This can vary from a few hours (if they were especially active on election day) to two years. In order to obtain a normalized measure of duration, we divide the median holding period by the median time to contract expiration of their trades. This is constrained to lie between 0 and 1, and equals 1 for those who never reduce their exposure to a candidate. This normalized measure of duration is reported in Table 1 for the ten selected traders. Four of these have very short duration, exiting positions almost as soon as they enter them; the remaining six hold their median trade essentially to expiry.

Of those 5906 traders who have at least one major market trade, a clear majority—64% in all, accounting for almost 20% of volume in the two main contracts—have duration precisely equal to 1. These individuals bet only in one direction and hold the median contract to expiration. But significant volume comes from a small number of high frequency traders. Just two traders (A and F in Table 1) with zero median holding account for almost 14% of volume in the two main contracts.

3.4 Margin

Intrade required traders to post cash margin equal to their worst-case loss, and this amount was frozen in trader accounts as soon as a transaction was made. Hence, for example, the sale of the Obama contact for \$6.00 would result in the freezing of \$4.00 in a trader's account, while \$6.00 would be frozen in the account of his counterparty. These amounts are released when positions are exited, net of any profits or losses made on the trade.

Importantly, the contracts are margin-linked for sales: a trader who sold Obama at \$6.00 and (at some later date) sold Romney at \$3.50 would have \$4.00 frozen on the initial trade and \$3.50 *released* on the second, since the worst case loss (under the assumption that only one of the two candidates could win) was now \$0.50. And if the latter contract were sold for \$4.50 instead, the entire margin from the first trade would be released, since no loss on the combined position is possible.

The eighth column of Table 1 shows the *maximum* amount of margin frozen in each of the trader accounts over the course of the entire cycle. This is the largest amount of cash that each trader had at risk. Particularly striking is the amount for trader C, who managed to participate in over 29,000 trades using just \$737 in cash. Just prior to contract expiration, this trader had short positions exceeding 73,000 contracts in *both* Obama and Romney, along with significant short positions in all other contracts, and would have netted a profit of at least \$11,921 regardless of the electoral outcome. The four high frequency traders in the table all had very low levels of margin relative to the volume of contracts traded, and all netted a profit well in excess of the margin posted.

3.5 Profits

The computation of profits for each account is a straightforward matter, since we know the prices at which all positions are entered and exited, and the terminal value of all contacts held to expiration. The last column of Table 1 reports realized profits for the ten selected traders, which include the three most profitable and the three with the biggest loss. Trader B lost close to \$7 million on an accumulated position that was heavily short Obama and long Romney, and this was by far the largest loss recorded in the data. Trader D bet heavily on an Obama victory and had the highest profit recorded in the data. These were both risky bets that could have gone either way. In contrast, Traders A, C, E, and F used arbitrage-based strategies with minimal risk exposure to book small profits on each of a large number of transactions.

4 Trading Strategies

Trading strategies can differ along multiple dimensions. The degree to which traders take or provide liquidity varies widely, as does the extent to which they build up large directional positions rather than entering and exiting positions repeatedly. Holding periods also vary widely, as does the willingness to switch positions from one direction to another. Some traders are clearly engaged in arbitrage across the two contracts, building up large negative positions in both, with a very small amount of margin frozen relative to position size.

These two strategies discussed above (for traders A and B respectively) occupy distinct niches in a rich trading ecology that we now attempt to characterize in detail. The goal is to construct a taxonomy of strategies, based on bias and holding period. Each account can be placed in one of six different categories, based on the following definitions:[10]

1. **Arbitrage**: median holding period less than 10 minutes (traders A, E, and F)
2. **Unidirectional**: bias equal to 1, non-arb (traders G, H, I, and J)
3. **Extreme Bias**: bias above 0.9 but below 1, non-arb (traders B and D)
4. **High Bias**: bias between 0.5 and 0.9, non-arb
5. **Moderate Bias**: bias between 0.25 and 0.5, non-arb (trader C)
6. **Low Bias**: bias less than 0.25, non-arb

The incidence of each strategy in the trading population, the volume attributed to each, and various characteristics (aggression, duration and bias) are shown in Table 2.

The unidirectional strategy accounts for 87% of traders and almost a third of total volume. These traders never switch the direction of their exposure. The second category is composed of traders with extreme levels of bias (exceeding 0.9). These traders are almost always either long

[10]The qualifier "non-arb" in the definition of all non-monotone strategies is necessary to prevent double counting. Although the 40 arbitrage strategies all have very low duration (none exceeding 0.02), they have widely varying biases ranging from precisely zero to precisely 1.

Strategy	Traders	%	Volume	%	Aggression	Duration	Bias
Unidirectional	5,118	87%	4,901,262	32%	0.65	0.75	1.00
Extreme Bias	136	2%	3,987,006	26%	0.38	0.65	0.97
High Bias	272	5%	1,699,355	11%	0.42	0.22	0.71
Moderate Bias	173	3%	1,293,289	9%	0.41	0.06	0.43
Low Bias	167	3%	926,702	6%	0.40	0.10	0.13
Arbitrage	40	1%	2,368,380	16%	0.73	0.00	0.20
Total	5,906	100%	15,175,994	100%			

Table 2: A Taxonomy of Trading Strategies

Obama or long Romney, and the mean level of bias in this group is 0.97.[11] There are few traders in this group but they include some very large ones (including B and D in Table 1) and together account for 26% of volume. The high and moderate bias categories account for a further 11% and 9% of volume respectively, and those with low bias account for 6%.

Arbitrageurs also have low bias on average, but are distinguished by their very short duration. Less than 1% of traders fall into this category but they account for 16% of total volume. As noted above, a substantial portion of this volume comes from traders A and F in Table 1. These two traders (and a handful of others in the data) are similar to the high-frequency traders in Kirilenko et al. (2011) in that their strategies are algorithmically implemented and characterized by short duration and limited directional exposure. Such traders account for about one-third of total volume in the S&P futures market.

Recall that a bias exceeding 0.5 implies that three-quarters of all exposure increasing trades are in one direction. These traders do switch direction, but strongly favor one direction over the other. Taken together, therefore, strategies that restrict their exposure largely to one side of the contract are responsible for 69% of total volume. A further 16% is associated with arbitrage, leaving just 15% assigned to strategies with moderate or low bias, whose exposure increasing trades do not favor either side of the contract more than three-fourths of the time.

Some features of the ecological composition of the market may be seen in Figure 5, which plots duration and bias for all 5,906 strategies that were active in one or both of the main candidate markets. The area of each circle is proportional to the trader's (main market) volume, and the arbitrage based strategies are shown in the darker green. All unidirectional strategies (with bias equal to one) are clustered along the top of the figure. The two largest circles correspond to traders B (duration and bias both close to 1, main market volume 2.06m) and A (zero duration, low bias, main market volume 1.94m) respectively.

It is clear from the figure that, for the most part, strategies are either extremely biased or have very low duration (or both). Strategies with long holding periods tend to be extremely biased, but even short duration strategies are often highly biased. There are many strategies with duration

[11] The levels of aggression, duration and bias within groups are computed by taking the (main market) volume-weighted average of these magnitudes.

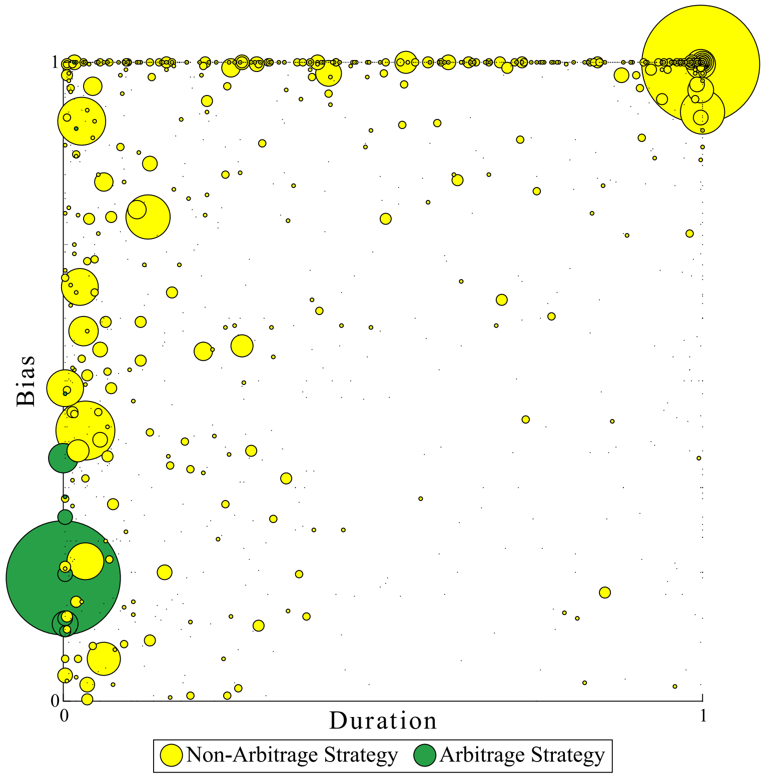

Figure 4: Duration and Bias for Traders Active in Major Markets

approximately zero that had holding periods in excess of ten minutes, and were therefore not classified as arbitrageurs. Some of these are likely to have been using market making strategies, such as posting bids and offers across multiple contracts, or arbitrage strategies that were manually rather than algorithmically executed. Trader C, for example, has near-zero duration and a very low ratio of margin to profits, indicative of an arbitrage or market making strategy. But the 42 minute median holding period makes it possible that the trader was executing an information based strategy with very quick turnover.

The most striking feature of the data is the relatively small number of unbiased traders with moderate or long holding periods and sizable volume. Most large traders were either heavily biased (like B and D) or engaged in arbitrage or market making strategies (like A and F). This has implications for our understanding of how prices come to reveal information. Polling data was appearing with rapid frequency over this period, and campaign related news dominated headlines. If all this information was finding its way quickly into prices, it must have been at least in part through the actions of *biased* traders.

A plausible channel through which these biases could arise is *wishful thinking*, or the tendency of people to "bend their expectations to coincide with their preferences" (Granberg and Brent, 1983). Uhlaner and Grofman (1986) found strong evidence of this in the 1980 Presidential race between Carter and Reagan, with more then 80 percent of each candidate's supporters expecting their favored candidate to win. Even in elections that are viewed by many as foregone conclusions, wishful thinking is prevalent: a quarter of those intending to vote for McGovern in 1972 and more than 30 percent of Goldwater voters in 1964 thought that their favored candidate would win; in contrast, the proportion of Nixon supporters in 1972 and Johnson supporters in 1964 who believed that their favored candidate would win was close to 100 percent (Granberg and Brent, 1983).

5 Beliefs and Objectives

Our taxonomy of trading strategies was based on *what* traders were doing, and we now turn to the question of *why*. For the arbitrage based strategies the answer is straightforward: the link between the two contracts presents an opportunity for risk-free profit that can be grasped with a modest capital outlay and some programming skill. Neither beliefs about the electoral outcome nor complicated objectives need come into play.

The beliefs and motivations underlying the other strategies are harder to deduce. Traders could be engaging in arbitrage across exchanges, speculating based on subjective beliefs about the electoral outcome, hedging exposures in other markets, or attempting to manipulate prices to further political or financial goals. We consider each of these possibilities in turn.

5.1 Cross-Market Arbitrage

The arbitrage strategies identified in Table 2 are based on exploiting inconsistencies in the prices of the two contracts on the *same* exchange. But similar contracts were available on other exchanges, allowing for cross-market arbitrage. For such transactions we would only observe one leg of the activity, making it appear that the trader was building a directional position when in fact his aggregate portfolio was characterized by negligible bias.

One exchange that offered virtually identical contracts to those on Intrade was Betfair, and there was a striking and persistent disparity in prices across the two trading venues over this period. For the last few months of the election cycle, the headline Obama contract was trading at a higher price on Betfair relative to Intrade, with the disparity fluctuating between 5 and 10 percentage points (Rothschild and Pennock, 2014). Even accounting for the fact that Betfair takes 2-5 percent of winnings, depending on account size, this represented an arbitrage opportunity.[12]

Exploiting this disparity was considerably more complicated and costly than the within-exchange arbitrage described earlier. Since positions across the two exchanges were not margin

[12] A trader who bought the Obama contract on Intrade and sold it on Betfair would have avoided this transactions cost, since all winnings would have been on Intrade. But this could not have been known in advance of the election result.

linked, building a directional position on one exchange with offsetting positions in the other required substantial margin to be posted in both. In addition, Betfair contracts were denominated in British pounds, subjecting traders to exchange rate risk or requiring them to hedge this using currency futures, in addition to the cost of currency conversion. Most importantly, since the disparity was always in the same direction, such cross-market arbitrage would result in a long Obama position on Intrade, and cannot therefore account for monotone and unidirectional positions that were long Romney. That is, it cannot account for the behavior of trader B in Table 1, nor for a substantial portion of other unidirectional traders. Nevertheless, we cannot rule out the possibility that some traders with unidirectional positions that were long Obama had offsetting positions on Betfair.

5.2 Speculation

The standard account of financial market behavior in economics is based on preferences over lotteries, represented by the maximization of expected utility. This model allows us to make some inferences about the distribution of trader beliefs.

At a given point in time, consider a trader with portfolio (y, z), where y is cash and z is the (possibly negative) number of Obama contracts owned, measured in units of face value. That is, z is the amount of cash that the trader is contracted to receive in the event of an Obama victory. Let π denote the current market price of such contracts, and p the trader's belief in the likelihood of an Obama victory.[13] If this trader were to adjust his portfolio by buying d units at price π, his expected payoff would be

$$pu(y - \pi d + z + d) + (1 - p)u(y - \pi d).$$

In order for the trader to hold the portfolio (y, z), the demand $d = 0$ must be optimal. The necessary first-order condition for this is given by

$$p(1 - \pi)u'(y + z) = (1 - p)\pi u'(y).$$

For traders with $u'' < 0$, $z > 0$ can be optimal if and only if $p > \pi$. That is, a risk-averse trader will be long Obama if and only if his subjective belief in an Obama victory exceeds the market price. Less obviously, the same is also true of risk-seeking traders, since these markets allow for positions in either direction. A risk-seeking trader who was long Obama but believed that the price was higher than the subjective probability of victory could switch direction, securing an increase in expected return without a reduction in risk-exposure. This is clearly not possible in markets for lottery tickets since these assets cannot be sold short. But unlike conventional lotteries, prediction markets do not force individuals to accept negative expected returns in order to increase risk exposure.

[13] Here we are neglecting the bid-ask spread, assuming that small amounts can be purchased or sold at price π. The spread over this period was negligible, amounting to just a penny or two per $10 of face value.

The implication of this reasoning is that under the standard expected utility hypothesis, *unidirectional traders must believe that the market price is wrong in precisely the same direction throughout the observational window.* Those with direction $+1$ must consider the Obama price to be consistently too low relative to their subjective belief, while those with direction -1 must think the opposite.

Biased traders also believe that the market is systematically mispricing the Obama contract in the same way for much of the time. Traders with extreme bias are almost unidirectional; the volume weighted average bias in this group is 0.97 (see Table 2). Those with high bias do occasionally change direction, but the average bias within this group is 0.71.

Recall that taken together, the unidirectional, extreme, and high bias categories account for 94% of traders and almost 70% of total volume. It is entirely possible that many of these individuals trade in response to new information, but since they are strongly biased in two opposing directions, it is very difficult to reconcile their behavior as a group with the common prior hypothesis. That is, to the extent that they are trading on news, they must be doing so with very different interpretations of information.[14]

The moderate and low bias traders most closely resemble the information traders in standard model of financial market microstructure (Glosten and Milgrom, 1985; Kyle, 1985). However, they are responsible for just 6% of accounts and 15% of volume, and their counterparties on most trades are very likely to be biased. It is the interaction between the various groups of traders, and especially interactions among traders with strong but opposing directional biases, that govern the behavior of asset prices in this market.

5.3 Hedging

Electoral outcomes are known to affect a variety of asset prices, including broad indexes as well as specific sectors. In principle, individuals with such exposures can hedge their risk by taking offsetting position in prediction markets. This can give rise to positions that appear to conflict with subjective beliefs about the outcome.

As an example, consider the findings of Snowberg et al. (2007) regarding the race between Bush and Kerry in 2004. Using the fact that a set of flawed exit polls were leaked at about 3pm on Election Day, the authors were able to identify market beliefs about the asset price implications of the electoral outcome. The exit polls pointed to a Bush defeat, and appear to have been widely

[14] We can also imagine that some some traders gain positive utility from owning either Obama to win or Romney to win. This is observationally equivalent to the case of extreme beliefs in one direction of the other: such traders act as if they are highly optimistic about the prospects of victory for their preferred candidate. Whether or not partisan leanings correlate with the beliefs inferred from positions is an interesting question, but one that cannot be addressed with our data (since we observe only market behavior).

believed. The price of the Bush contract on TradeSports (a precursor to Intrade) fell sharply from about 55 cents to 30 cents per dollar of face value upon release of the polls, remained close to this depressed level for several hours, and then rose sharply to 95 cents shortly before the election was called. This movement was mirrored in the price of the S&P futures contract with the closest settlement date, which dropped by 1% on release of the flawed polls, remained down until the error became apparent, and rose by 1.5% in concert with the rise in the price of the Bush contract.[15]

The finding that electoral outcomes have predictable asset price effects raises the possibility that individuals were hedging positions in conventional financial markets with partially offsetting positions on Intrade. This could account for monotone strategies with long duration in either direction. However, we provide some evidence below that this would not have been an effective hedge during this particular electoral cycle, at least with respect to a broad based index such as the S&P 500.

The link between electoral outcomes and asset prices also creates an incentive for the deliberate manipulation of prediction market prices. For instance, if a large trader could temporarily elevate or depress a price on an election contract, resulting in corresponding temporary changes in the prices of other assets, opportunities for substantial windfall gains could arise. We next explore this possibility, as well as other incentives for manipulation.

6 Manipulation

Although the main electoral markets on Intrade had high volume, broad participation, and tight spreads over this period, the amounts wagered were small relative to those in stock and bond markets, and also relative to the costs of contemporary electoral campaigns. This raises the question of whether price manipulation was attempted on this market, and if so, for what purpose.

At about 3:30pm on Election Day, the order books for the Obama and Romney contracts were discovered to have a very unusual structure.[16] There were bids to buy more than 40,000 contracts in the Romney book at prices between 28 and 30, an order of magnitude greater than the number of offers to sell. The Obama order book was similarly skewed, with offers to sell more than 35,000 contracts at prices between 70 and 71, vastly great than the number of bids for the contract. Note that prices are quoted as a percentage of face value, which on Intrade was $10 per contract. Hence the margin frozen for this set of orders alone was about $200,000. The left side of Figure 5 shows a snapshot of the quantity of open orders, at the top three prices, to buy and sell Obama contracts, taken every three minutes, from noon ET onward on Election Day.

[15] Along similar lines Knight (2006) found that increases in the likelihood of a Bush victory over Gore in 2000, as measured by prices in the Iowa Electronic Markets, were associated with a substantial return differential between Bush-favored sectors (such as tobacco) and Gore-favored sectors (such as alternative energy).

[16] See pic.twitter.com/EHvm1DGl and http://t.co/kJJeFkqJ for screenshots of order books posted on the afternoon of election day.

This order book structure is highly unusual at a time when information is coming in thick and fast, because the party with resting orders stands to lose a great deal in transactions with more informed traders.[17] Although the evidence is not conclusive, it appears as if someone was trying to put a floor of about 30 on the Romney price and a corresponding ceiling of 70 on the Obama price. If so, the strategy was successful: the floor and ceiling both held firm for several hours until the firewall finally collapsed at around 9pm ET. While prices on Betfair were fluctuating freely in response to incoming information, and broadly moving in the Obama direction, those on Intrade were remarkably stable despite heavy volume.

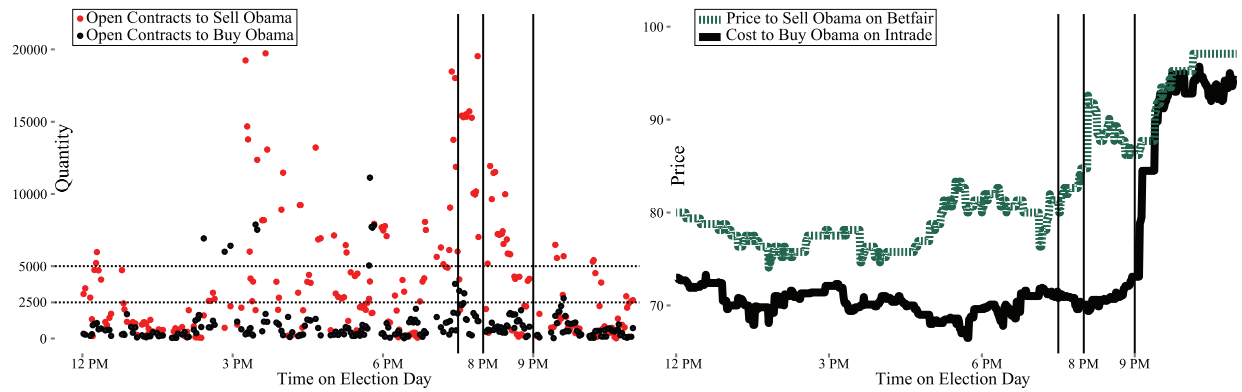

Figure 5: Betfair and Intrade on Election Day

Starting at noon ET on Election Day, the right side of Figure 5 shows prices on Betfair and Intrade for Obama to win. The Betfair premium on the Obama contract was relatively modest at the start of the day but the gap began to widen in the afternoon. At 7:30pm, when the Betfair price began another upward movement, there were over 15,000 contracts for sale at a price of about 70 in just the top three positions of the Obama order book on Intrade. Breaking though this barrier would have required over a hundred thousand dollars, and since there was no way to transfer funds to the exchange at short notice, this was enough to maintain the price ceiling.

Since most of the limit orders that held the floor and ceiling in place were eventually met, we are able to confirm from the data that they were placed by Trader B. This trader spent over \$375,000 to hold prices in place during the 90 minute interval starting at 7:30pm. At 9pm he pulled out of the market for a while, and the price of the Obama contract shot up quickly to reach the Betfair level. His return about an hour later, this time in support of a much lower Romney price, caused a wedge between the two markets to open up again. Whatever his motives may have been, there can be little doubt that his activity had a price impact.

Trader B was responsible for one-third of the total money on Romney over the two week period leading up to the election, and about a quarter over the entire cycle. The result was a

[17] The extreme order book asymmetry that was apparent on Election Day was also in evidence in earlier periods of high information flow, for instance during Hurricane Sandy. This was verified using data collected in real time, every three minutes, on the prices and quantities in the first three positions of the order books for the two contracts.

loss of close to seven million dollars. What could possibly have motivated this activity? Given that the trader bet on Romney and not Obama, we can rule out cross-market arbitrage with Betfair as a motivation. This leaves three possibilities: (i) the trader was convinced that Romney was underpriced throughout the period and was expressing a price view, (ii) he was hedging an exposure held elsewhere, or (iii) he was attempting to distort prices in the market for some purpose.

A trader believing that the Romney contract on Intrade was underpriced should have bought the contract on Betfair instead, where it was consistently cheaper. Nevertheless, given the additional difficulties and costs of using Betfair, especially for US based traders, we cannot rule out the possibility that this trader was simply expressing a price view.

What about hedging? If the findings of Snowberg et al. (2007) applied to the 2012 election, then a long Romney position would hedge a short position in S&P futures. In order to explore the relevance of this possibility, we examined movements in the index during the first debate, in which Obama's performance was widely acknowledged to have been disastrous. The price of the Obama contract stood at about 71 cents per dollar of face value at the start of the debate, tumbled to 66 during the event, and ended the day at around 65 despite a short-lived recovery in the interim. These price movements are shown in Figure 6.

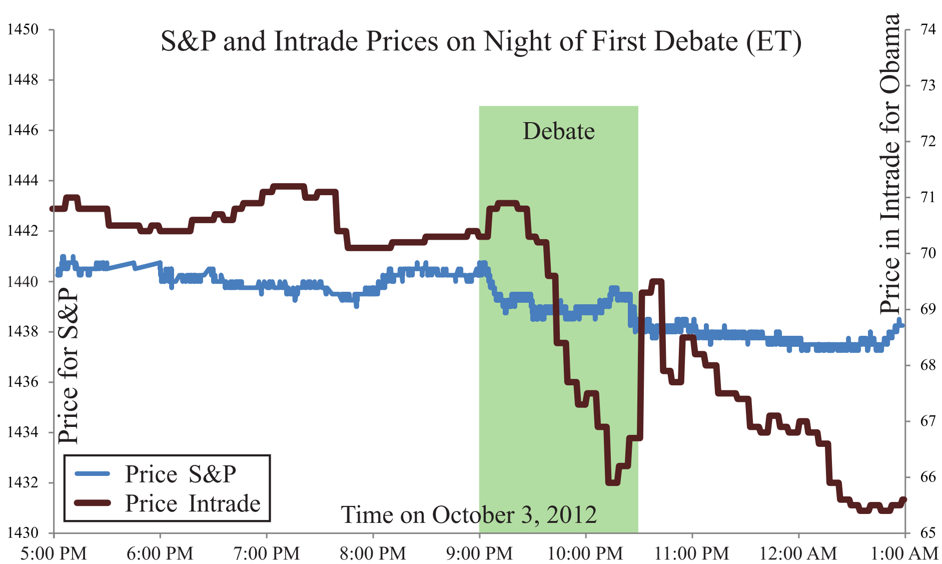

Figure 6: Responses of Intrade and S&P futures to the First Debate

The figure also shows movements in the price of the S&P 500 futures contract over this period. If anything, the index *fell* modestly on the basis of Obama's performance, which is the opposite of what one would expect based on the analysis by Snowberg et al. (2007) of the 2004 election. A similar effect occurred on Election Day itself, with a modest rise in the index occurring as the election was called.[18] Based on these admittedly crude observations, it appears that the use of Intrade to hedge a market position based on correlations in prior elections would have been ineffective at best, and possibly counterproductive. For similar reasons, forcing up the Romney price on Intrade in the hope that this would provide a temporarily boost to stock prices in the aggregate would have resulted in losses rather than windfall gains. For these reasons, we consider the hedging of market risk and the manipulation of Intrade for financial gain to be unlikely explanations for the behavior of this trader.[19]

This still leaves open the possibility of market manipulation for political purposes (Thompson, 2012). The trading losses, while hardly trivial, pale in comparison with the cost of contemporary political campaigns.[20] It is conceivable that the trader could have been attempting to manipulate beliefs about the odds of victory in an attempt to influence fundraising, campaign morale, voter preferences, and turnout.

The viability of a campaign is clearly an important consideration for potential donors.[21] Mutz (1995) distinguishes between "loyalty-based" and "hesitancy-based" giving, where the former motive may be triggered when one's preferred candidate is slipping in the polls and is in need of a boost, while the latter is contingent on a demonstrated likelihood of prevailing in the election. Hesitancy-based giving may be especially relevant in the wake of the Supreme Court's 2010 *Citizens United* ruling, allowing unlimited contributions to political action committees that are formally separate from but informally closely tied to individual campaigns. The manipulation of prediction market prices costs money, since bets must be placed at unfavorable odds, but could pay for itself several times over if it maintains flows to campaigns or super PACs.

There is also the possibility that beliefs about the likelihood of victory can have a direct effect on voter preferences. While the existence of such bandwagon effects remains a matter of dispute, there is some recent evidence that perceived momentum for a party can result in increased support. In a survey experiment involving a very large sample (over 23,000) of the Dutch electorate, van der Meer et al. (2015) find that while polling levels do not significantly affect vote intentions, polling momentum does. Here subjects were exposed to identical (and genuine) polling data but it was framed differently for different treatment groups. Positive framing, using a more recent reference point to suggest momentum, was found to have a positive and statistically significant impact on vote intentions. In the US context, Bartels (1985) provides evidence of a bandwagon effect for Bush in the 1980 nomination contest, and Skalaban (1988) claims an effect favoring Reagan in the

[18] One possible reason for the modestly positive market impact of the Obama victory was the belief that monetary policy would be tighter under a Romney administration (Popper, 2012).

[19] We cannot, of course, rule out the possibility that hedging more specific exposures or manipulating a narrow set of securities would have been effective.

[20] Aggregate expenses for the 2012 presidential cycle were over $2.6 billion (http://www.opensecrets.org/bigpicture).

[21] The evidence on the sensitivity of contributions to polling data comes mostly from primaries; see Mutz (1995) for the 1988 Democratic presidential primary, Fuchs et al. (2000) for the 1989 New York Democratic mayoral primary, and Adkins and Dowdle (2002) on pre-primary fundraising over the 1980-2000 period.

general election of that year, but neither study is experimental.[22]

Anecdotal evidence certainly suggests that campaigns consider positive momentum to be valuable. For instance, a poll showing a substantial tightening of the Bush-Clinton race close to the 1992 election was cited as evidence for momentum by the trailing Bush campaign, but dismissed as an outlier by the Clinton campaign (Ansolabehere and Iyengar, 1994). To the extent that such bandwagon effects exist, or are believed to exist, there are incentives for partisans to invest in the manipulation of prediction market prices. Based on the pattern of transactions in our data, it is likely that at least one well-funded individual responded to these incentives in the run up to the 2012 election.

Attempts to manipulate prediction market prices have historically been futile (Rhode and Strumpf, 2008), but this may have been an instance when a modest distortion was successfully sustained over several days, with more substantial effects over the last few hours. Our findings suggest that with a moderate amount of capital and a strategy of placing large limit orders at plausible prices, prediction markets with relatively high liquidity and broad participation can be manipulated for a time. However, in order for manipulation to be successful in meeting broader financial or political goals, it is important that it remain undetected in real time. It is doubtful that any such broader goal was met. The possibility of price manipulation on Intrade was salient as the election approached, and the availability of multiple trading venues made it easy to spot outliers. Social media can serve as an amplifier of misinformation at times, but it can also serve as an antidote. The rapid spread of information through dense online networks can prevent manipulation from remaining undetected for long, and thereby undermine any broader goals that manipulation is intended to achieve.

Historically, prediction markets have generated forecasts that compete very effectively with those of the best pollsters.[23] This was also the case in the 2012 cycle, attempts at manipulation notwithstanding. Forecasts based on price averages across multiple trading venues, with corrections for the favorite-longshot bias that is known to arise in sports betting and related markets, were especially accurate (Rothschild, 2014). Nevertheless, it is worth knowing that a highly visible market that drove many a media narrative could be manipulated at a cost less than that of a primetime television commercial.

[22]For issue based polling, bandwagon effects have been clearly and convincingly demonstrated using experiments; see Marsh (1985) and Nadeau et al. (1993) on abortion, and Rothschild and Malhotra (2014) on free trade agreements. Beliefs can also affect vote intentions for tactical reasons in multi-candidate contexts; see for instance Blais et al. (2006) on the 1988 Canadian election.

[23]See Berg et al. (2008) on the performance of 49 markets in 13 countries, Wolfers and Leigh (2002) and Leigh and Wolfers (2006) on two successive Australian elections, and Rothschild (2009) and Rothschild (2014) on the US presidential elections of 2008 and 2012.

7 Discussion

In this paper we have taken a close look at transaction level data in an asset market that has recently had broad cultural and political importance. Doing so allowed us to characterize a rich ecology of trading strategies that is dominated by traders who seldom, if ever, change directions. This has implications for theoretical models of market microstructure. It appears that the trading process is driven, in large measure, by individuals with different *interpretations* of *public* information. The results of opinion polls or the words spoken in debate are not facts in dispute, but there can be considerable disagreement about their meaning and importance.

A strength of prediction markets is that prices reflect not just public information from polls and models but also scattered private information held by traders. Empirical studies of the performance of such markets have consistently found them to be at least as accurate as opinion polls, and much quicker to incorporate new information. Nevertheless, prediction markets also have a number of characteristics that could make them vulnerable to failure. Traders are predominantly young and male, and include many who are not eligible to vote in the elections on which they are betting. Sophisticated traders probably make allowances for this fact, but even so, it is unlikely that their beliefs are a representative sample of the population at large. Second, the bets placed by individuals depend on their respective budgets, and these can—and indeed, do—vary dramatically across traders. Prices skew towards the beliefs of those who are best funded, who need not be the best informed, though this effect is somewhat mitigated by the fact that budgets evolve over time based on prior trading performance. Third, traders are not immune from the wishful thinking that has been amply documented among voters in elections. To some degree the biases of traders with opposing preferences will tend to cancel each other out, but it is not clear why this should result in unbiased forecasts in the aggregate.

In addition, there is a paradox at the heart of the prediction market model. The more accurate market forecasts are perceived by the general public to be, the greater is the incentive for deep-pocketed partisans to try to manipulate prices to alter perceptions of the state of the race. They may do so to sustain campaign contributions, maintain morale among supporters, boost turnout, and influence voter preferences over candidates. The relatively small size of these markets makes manipulation relatively inexpensive, especially in comparison with the aggregate expenditures involved in a national campaign.

Given the mismatch between the characteristics of traders and the electorate, the unbalanced budget constraints, the propensity for wishful thinking, and the incentives for manipulation, the high level of prediction market forecasting accuracy is something of a mystery. It has been argued by Berg et al. (2008) that the traders who really matter, the ones who post prices against which others are likely to trade, are not vulnerable to wishful thinking. But in the case examined here, most orders to sell Romney contracts or buy Obama contracts traded against orders posted by a single large trader who appears to having been attempting to manipulate beliefs.

Kets et al. (2014) suggest a possible explanation for the effectiveness of the wisdom of crowds even in the presence of wishful thinking. In their model beliefs are arbitrarily given, but the wealth of individuals evolves over time as a result of trading profits and losses as events are repeatedly realized. In the long run the distribution of wealth is such that the market clearing price matches the objective likelihood of the events in question. This model is not directly applicable to the case of elections, each of which is a unique event. Nevertheless, it suggests that endogenous transfers of wealth may be a key element in accounting for the impressive forecasting accuracy of prediction markets.

A promising direction for theoretical research that is suggested by our findings is the development of models that allow for public information to be rapidly and reliably reflected in prices in the presence of heterogeneous interpretations of this information. Ideally, the model should allow for responsive prices even with unidirectional traders. One possibility is to consider trading among partisans who differ systematically in their interpretations of public signals, such that news favorable to their side is considered to be highly significant, while news unfavorable to their side is viewed as relatively unimportant. To illustrate, consider a contract that pays if and only event A occurs. Let A–partisans denote those who are predisposed to believe that A will occur, while B–partisans are predisposed to believe the opposite. The appearance of news that makes A appear more likely will then cause A–partisans to buy vigorously, driving up the price to levels at which B–partisans will wish to sell. The news does lead to a higher price for the asset, but with B–partisans increasing their short positions. Similarly, news that makes A appear less likely will cause a price decline, but with A–partisans increasing their long positions at prices that they consider too low. Prices do absorb information in both cases, but without any trader changing direction.

Such dynamics are most likely to arise in electoral prediction markets and in sports betting, but similar effects may well exist also for common stock. Especially in the case of consumer durables, attachment to products and the companies that make them is widespread. It would not be surprising if one were to find Apple or Samsung partisans among investors, just as one finds them among consumers. A more thorough development of this idea is beyond the scope of the present empirical exercise, but seems worth pursuing in future research.

References

Randall E. Adkins and Andrews J. Dowdle. The money primary: What influences the outcome of preprimary presidential nomination fundraising? Presidential Studies Quarterly, 32(2):256–275, 2002.

Stephen Ansolabehere and Shanto Iyengar. Of horseshoes and horse races: Experimental studies of the impact of poll results on electoral behavior. Political Communication, 4(413-430), 1994.

Matthew Baron, Jonathan Brogaard, and Andrei Kirilenko. The trading profits of high frequency traders. Unpublished Manuscript, 2012.

Larry M. Bartels. Expectations and preferences in presidential nominating campaigns. American Political Science Review, 79(3):804–815, 1985.

Joyce Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz. Results from a dozen years of election futures markets research. In Handbook of Experimental Economics Results, Volume 1, 2008.

André Blais. To Vote Or Not to Vote? The Merits and Limits of Rational Choice Theory. University of Pittsburgh Press, Pittsburgh, PA, 2000.

André Blais, Elisabeth Gidengil, and Neil Nevitte. Do polls influence the vote? In Richard G. C. Johnston and Henry E. Brady, editors, Capturing Campaign Effects, pages 263–279. University of Michigan Press, 2006.

Thomas E. Copeland and Dan Galai. Information effects on the bid-ask spread. Journal of Finance, 38(5):1457–1469, 1983.

Morton Deutsch and Harold B Gerard. A study of normative and informational social influences upon individual judgment. Journal of Abnormal and Social Psychology, 51(3):629–636, 1955.

Ray C. Fair. Predicting presidential elections and other things. Stanford University Press, second edition, 2011.

Ester R. Fuchs, E. Scott Adler, and Lincoln A. Mitchell. Win, place, show: Public opinion polls and campaign contributions in a New York City election. Urban Affairs Review, 35(4):479–501, 2000.

John Geanakoplos. The leverage cycle. In NBER Macroeconomics Annual 2009. University of Chicago Press, 2010.

Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics, 14:71–100, 1985.

Donald Granberg and Edward Brent. When prophecy bends: The preference–expectation link in us presidential elections, 1952–1980. Journal of Personality and Social Psychology, 45(3):477, 1983.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. American Economic Review, 70(3):393–408, 1980.

Milton Harris and Artur Raviv. Differences of opinion make a horse race. Review of Financial Studies, 6(3):473–506, 1993.

J. Michael Harrison and David M. Kreps. Speculative investor behavior in a stock market with heterogeneous expectations. Quarterly Journal of Economics, 92(2):323–336, 1978.

Martin F. Hellwig. On the aggregation of information in competitive markets. Journal of Economic Theory, 22:477–498, 1980.

Harrison Hong, Jose Scheinkman, and Wei Xiong. Asset float and speculative bubbles. Journal of Finance, 61(3):1073–1117, 2006.

Eugene Kandel and Neil D Pearson. Differential interpretation of public signals and trade in speculative markets. Journal of Political Economy, 103(4):831–872, 1995.

Willemien Kets, David M. Pennock, Rajiv Sethi, and Nisarg Shah. Betting strategies, market selection, and the wisdom of crowds. In Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.

Andrei A. Kirilenko, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. The flash crash: The impact of high frequency trading on an electronic market. Available at SSRN: http://ssrn.com/abstract=1686004, 2011.

Brian Knight. Are policy platforms capitalized into equity prices? evidence from the bush/gore 2000 presidential election. Journal of Public Economics, 90(4–5):751–773, 2006.

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

Andrew Leigh and Justin Wolfers. Competing approaches to forecasting elections: Economic models, opinion polling and prediction markets. Economic Record, 82(258):325–340, 2006.

Catherine Marsh. Back on the bandwagon: The effect of opinion polls on public opinion. British Journal of Political Science, 15(1):51–74, 1985.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. Journal of Economic Theory, 26(1):17–27, 1982.

Edward M Miller. Risk, uncertainty, and divergence of opinion. Journal of Finance, 32(4):1151–1168, 1977.

Diana C Mutz. Effects of horse-race coverage on campaign coffers: Strategic contributing in presidential primaries. Journal of Politics, 57(4):1015–1042, 1995.

Richard Nadeau, Edouard Cloutier, and J-H. Guay. New evidence about the existence of a bandwagon effect in the opinion formation process. International Political Science Review, 14(2): 203–213, 1993.

Thomas R Palfrey and Howard Rosenthal. Voter participation and strategic uncertainty. American Political Science Review, 79(1):62–78, 1985.

Nathaniel Popper. Which changed first, the polls or the markets? New York Times, October 16 2012.

Paul W. Rhode and Koleman S. Strumpf. Manipulating political stock markets: A field experiment and a century of observational data. Working Paper, University of Arizona, 2008.

William H Riker and Peter C Ordeshook. A theory of the calculus of voting. American Political Science Review, 62(1):25–42, 1968.

David Rothschild. Forecasting elections comparing prediction markets, polls, and their biases. Public Opinion Quarterly, 73(5):895–916, 2009.

David Rothschild. Combining forecasts for elections: Accurate, relevant, and timely. International Journal of Forecasting, 2014.

David Rothschild and Neil Malhotra. Are public opinion polls self-fulfilling prophecies? Research & Politics, 1(2):1–10, 2014.

David Rothschild and David M Pennock. The extent of price misalignment in prediction markets. Algorithmic Finance, 3(1):3–20, 2014.

Jose Scheinkman and Wei Xiong. Overconfidence and speculative bubbles. Journal of Political Economy, 111(6):1183–1219, 2003.

Kay Lehman Schlozman and John T Tierney. Organized Interests and American Democracy. Harper & Row, New York, 1986.

Rajiv Sethi and Muhamet Yildiz. Public disagreement. American Economic Journal: Microeconomics, 4(3):57–95, 2012.

Alp Simsek. Belief disagreements and collateral constraints. Econometrica, 81(1):1–53, 2013.

Andrew Skalaban. Do the polls affect elections? some 1980 evidence. Political Behavior, 10(2): 136–150, 1988.

Erik Snowberg, Justin Wolfers, and Eric Zitzewitz. Partisan impacts on the economy: evidence from prediction markets and close elections. Quarterly Journal of Economics, 122(2):807–829, 2007.

Derek Thompson. Should presidential campaigns spend more money manipulating intrade? The Atlantic, October 23 2012.

Carole Jean Uhlaner and Bernard Grofman. The race may be close but my horse is going to win: Wish fulfillment in the 1980 presidential election. Political Behavior, 8(2):101–129, 1986.

Tom W. G. van der Meer, Armen Hakhverdian, and Loes Aaldering. Off the fence, onto the bandwagon? a large-scale survey experiment on effect of real-life poll outcomes on subsequent vote intentions. International Journal of Public Opinion Research, 2015.

Justin Wolfers and Andrew Leigh. Three tools for forecasting federal elections: Lessons from 2001. Australian Journal of Political Science, 37(2):223–240, 2002.

Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. Working paper No. 12200, National Bureau of Economic Research, 2006.