Title: The Extent of Price Misalignment in Prediction Markets

Abstract: We study misaligned prices for logically related contracts in prediction markets. First, we uncover persistent arbitrage opportunities for risk-neutral institutional investors between identical contracts on different exchanges. Examining the impact of several thousand dollars of transactions on the exchanges themselves in a randomized field trial, we document that price support extends well beyond what is seen in the published order book and that arbitrage opportunities are significantly larger than purely observational measurements indicate. Second, we demonstrate misalignment among identical and logically related contracts listed on the same exchange that cluster around moments of high information flow, when related contracts systemically shut down or fail to respond efficiently. Third, we document bounded rationality in prediction markets; examples include: consistent asymmetry between buying and selling, leaving the average return for selling higher than for buying; and persistent price lags between exchanges. Despite these signs of departure from theoretical optimality, the markets studied function well on balance, considering the sometimes complex and subtle relationships among contracts. Yet, we detail how to improve prediction markets by moving the burden of finding and fixing logical contradictions into the exchange and providing flexible trading interfaces, both of which free traders to focus on providing meaningful information in the form they find most natural.

Corresponding and First Author: David Rothschild, Microsoft Research, 641 6th Ave, 7th Floor, New York, NY 10011, (973) 818-8997, Email: David@ResearchDMR.com, Website: ResearchDMR.com

Second Author: David Pennock, Microsoft Research, Email: DPennock@microsoft.com, Website: DPennock.com

I. Introduction

At 4pm E.T. on October 3, 2012, a trader on Betfair, a U.K. betting exchange, was willing to sell "shares" of Mitt Romney for 19.6 cents on the dollar. A buyer who agreed to the price would give the seller the $0.196 and receive in return a contractual promise from the seller to pay the buyer $1.00 if Mitt Romney won the 2012 U.S. Presidential election. The buyer, if she held the contract until Election Day, would either lose her $0.196 if Obama won or earn $1-$0.196 = $0.804 if Romney won. At the same time, on Irish prediction exchange Intrade, the most aggressive buyers were willing to pay up to 29.8 cents for the same "$1 if Romney wins" contract. Obviously, both exchanges could not be right about the value of the contract. More than informational differences, the contradiction represented free money: a trader could buy contracts on Betfair for $0.196 each and sell identical contracts on Intrade for $0.298, pocketing $0.102 for each contract, minus transaction fees. (If Romney won, the trader would collect on Betfair and lose on Intrade: $0.804 - $0.702 = $0.102; if Obama won, the trader would collect on Intrade and lose on Betfair: $0.298 - $0.196 = $0.102.) How was this price misalignment occurring and what does it mean for markets and exchanges?

A canonical prediction market contract pays $1 if and only if an outcome occurs. We examine activity on two public exchanges offering such binary-payoff contracts for political outcomes: Dublin-based Intrade and London-based Betfair. In practice, the two exchanges, though offering mathematically equivalent binary-payoff gambles, provide users with significantly different trading interfaces. Intrade's interface models a stock market: they list all-or-nothing contracts to buy or sell in a continuous double auction (CDA), though one share pays $10, not $1. On Betfair, traders don't buy "shares" of contracts, but instead state (1) how many British pounds they are willing to "lay" on the outcome and (2) the odds—or amount returned if they're right per pound risked—that they are willing to accept, reported in the decimal format common among European bookmakers. Further, for both exchanges, transaction costs, opportunity cost, currency risk, and counterparty risk dilute the value of any contract.

Even under the weakest form of market efficiency, true arbitrage, or a risk-free profit after fees, should not exist. Some arbitrage occurs between bookies and between bookies and exchanges, largely due to the structural differences between bookies and exchanges; one key distinction is the ability of exchanges to absorb new information instantaneously as bookie prices move more slowly [Franck et al. 2013, Strumpf 2003]. Between exchanges, someone should buy the cheaper contract, raising its price, and sell the higher-priced contract, lowering its price, until the price differential is less than any fees and overhead. Yet between-market price misalignment exists for numerous contracts. For months leading up to the U.S. Presidential election, we documented numerous opportunities to buy contracts in one exchange and sell the same contract in another exchange for more money. We found occurrences in high-liquidity markets, including the major party nomination markets and general election markets, as well as in low-liquidity markets like the state-by-state primaries. We consistently observed executable net profits of between 1 and 5 percent, even after reasonable estimation of the associated transaction costs and risks.

The order book does not capture all of the demand supporting a price misalignment, only the publicly declared limit orders. Thus the order book reflects a lower bound on arbitrage profits: the true opportunity may be several times higher if hidden demand – what we call a *shadow order book* – exists in the form of traders or their programmed agents waiting ready to accept new orders at the margin or to refill the order book as it empties. Not utilizing the order book is the theoretically rational strategy in many situations, where a trader does not want to leave an order that needlessly exposes some of the trader's information or that another trader with newer information can take advantage of it. Some exchanges actually allow explicit reserve or hidden orders, which allow big traders to hide their strategy by placing orders that are only partially revealed on the order book [Frey and Sandas, 2008]. Betfair and Intrade do not explicitly support hidden orders; on these exchanges, the shadow order book takes the form of automated trading bots or human traders "waiting in the wings" ready to pounce quickly as conditions change. Clearing out both order books to fully align prices (modulo transaction fees) would yield the minimum amount of arbitrage gain. For example, suppose Market A has 5 shares for sale at $0.10 and 5 more for sale at $0.12, and Market B has bids to buy 5 shares at $0.15 and 5 shares at $0.13. A trader can buy 10 contracts in Market A for an average of $0.11 and sell them in market B for an average of $0.14. But, what if, once the trader starts selling in Market B at $0.15, new bidders emerge clamoring to buy at $0.14 or $0.15? The profit per share could be closer to $0.04 instead of $0.03, and the total profit could be much higher than the 10x$0.03 = $0.30 profit explicit in the order book. Also note that the contract end date, in our case Election Day, represents only the latest possible time after which traders can reclaim capital; if prices align any time in the interim, traders can close out arbitrage positions, freeing up invested capital.

We used an empirical approach to probe for the existence and extent of a shadow order book, finding that the liquidity for any contract at any price is many times larger than the public bids and asks imply. Thus, we estimate the amount of money that arbitrageurs can extract is an order of magnitude higher than what a purely observational study would indicate. By observing the execution of over three thousand dollars in trades on the two exchanges, we confirmed that the price disparities represented real arbitrage opportunities, and were not a manifestation of inaccurate or delayed price quotes, reasons we could not rule out by simply observing the order books.

We hypothesize three main reasons that arbitrage persisted. First, small individual investors were deterred by the non-negligible risks, especially counterparty risk, combined with higher costs. Second, despite clear arbitrage opportunities, larger or institutional investors concluded that a one-time profit of on the order of tens of thousands of dollars was not enough to warrant their resources. Third, although traders should close the gap for informational reasons anyway, even if no one is profiting from the arbitrage, we find evidence that the gap was held open by a single large trader in one of the exchanges who was not maximizing profit.

We also find price misalignment between contracts listed on the *same* exchange with both mutually exclusive and conditionally related contracts; we attribute this misalignment to inefficient action in related contracts at times of high information flow. We demonstrate where a trader could sell several candidates running for the same election for more than $1 total, yet know that he or she was obliged to pay out only $1 for at most one of the candidates. We found conditional contracts where the price for a candidate to win the presidential election was higher than the price of their party to win, even though the party can run a different candidate but the candidate would not run without their party. For related contracts within markets, both mutually exclusive and conditional contracts, less salient markets became illiquid when the more salient contracts were moving due to increased information flow. In some examples, traders withdrew all bid and ask orders from contingent markets, even if there were safe prices to leave the orders regardless of the outcome of the other event.

We document additional evidence of weak-form inefficiency and bounded rationality. First, between markets, Intrade was leading Betfair during the primary. A lagged Intrade price of 12 hours was better indicator of the Betfair price than the Betfair price 12 hours previous. Second, the sum of all asks in complete sets of mutually exclusive contracts are consistently further from $1 than the sum of bids. Thus, consistent asymmetry between buying and selling across many exchanges, means that, on average, traders will earn higher returns on sell orders than buy orders.

To the extent that prediction markets seek to draw out information from a crowd [Gruca et al. 2007], our findings highlight shortcomings of existing exchange designs and standard industry practice. First, the shadow order book represents a trove of information hidden below the surface; traders, often via programmed robots, are watching and waiting, ready to react to market changes, but researchers cannot regularly capture or understand the degree of hidden price support or the subjective expectations of these waiting traders. Second, the prevalence of arbitrage means that trading surpluses go to uninformed participants who mechanically implement logical propagations, a task that can and should go to a computer. If the goal is to maximize information, prediction markets should reserve the highest reward for the most efficient information, not the sneakiest, fastest, or most sophisticated computations. Prediction markets should incentivize traders to focus on providing information, in whatever form they find convenient, rather than extracting social welfare through uniformed arbitrage.

To address these shortcomings, we advocate for changes in prediction markets design. First, exchanges should move beyond treating every binary outcome as an independent, one-dimensional continuous double auction (CDA). This practice fractures liquidity, limits expressiveness, holds more capital in reserve than necessary, and places undue cognitive burden on traders. Oliven and Rietz [2004] show that, faced with multiple ways to invest in the same prediction, traders don't always choose the cheapest option, defying theories of rationality. Exchanges should begin to enforce logical consistency among related contracts. Linear programming is the right way to generalize a CDA to multiple dimensions [Bossaerts et al. 2002, Fortnow et al. 2004]. When the number of mutually exclusive outcomes is reasonable, in the hundreds or thousands, linear programming is a fast, reliable, and well-understood procedure and there is almost no disadvantage to adopting it. When the number of outcomes is exponential, for example 2^{51} or 2 quadrillion state-by-state election outcomes, the computational complexity of linear programming (or any market mechanism) becomes intractable, though approximation schemes are possible [Dudik et al. 2012, Chen et al. 2008b, Lahaie et al. 2013].

Second, trading wizards should translate human judgments, expressed simply and naturally, into appropriate market orders. The exchange interface should emphasize its primary function—to reward information—and rescue users as much as possible from the swamp of financial or gambling numbers and jargon common today. Wizards allow traders to focus on estimating likely outcomes and ignore the details of particular market mechanics and strategies. For example, buying and selling are logically identical yet almost every exchange makes selling more confusing; ideally, traders should see no difference. Well-designed wizards that incentivize timely information may encourage traders waiting in the shadow order book to reveal themselves, capturing meaningful data currently lost in the market interface. Note that our two suggestions complement each other: combinatorial exchanges using linear programming work best when coupled with a wizard-like interface.

Efficient prediction markets are beneficial for all stakeholders. Investors ultimately benefit from the added ability to match whichever trades provide them the most utility and the added liquidity should allow them to use the markets more efficiently to hedge risk. Exchanges profit from more volume; increased matching of orders supersedes any loses in arbitrage orders. Finally, researchers benefit from more efficient pricing on more questions. The real-time forecasts derived from prediction market data increase efficiency in many domains. Meanwhile, the granular nature of the data is a key ingredient in studying important questions in political science [Snowberg et al. 2007], marketing, public policy [Wolfers et al. 2009], employee psychology [Cowgill et al. 2008] and many other domains [Arrow et al. 2008].

II. Estimation Strategy/Results

Between-Market Arbitrage

We start with the most obvious type of price imbalance; can we buy a single contract in one market and sell it in another market for more money? We looked at matching contracts on the two most liquid exchanges for politics: Betfair and Intrade. Betfair is the world's largest prediction market with FY2012 revenues of £389.7 million and single events that matched over £50 million.[1] Intrade is the most watched and robust political prediction market with over 7.5 million $10 contracts on Obama or Romney to win the 2012 election matched during the 2012 election cycle. First, we verify that the contracts listed on the two exchanges are indeed identical, at least for practical purposes. Second, we catalog the transaction and opportunity costs associated with the contracts and the markets. Third, we follow the lowest buy price and the highest sell price for these contracts in both markets and compare.

The two exchanges listed dozens of related contracts that feature various logical implications and carried similar between-exchange price misalignments; we focus on the most liquid and visible contracts of Obama to win and Romney to win. Some related contracts include "Republican candidate to win", "Republican candidate to win at least 270 Electoral College votes", and similar contracts for the Democratic candidate, and other Electoral College thresholds. For example, if Obama wins at least 270 Electoral College votes: he wins the election, Romney loses the election, and, after the formal nominating conventions, the Democratic Party wins the election, and the Republican Party loses the election.

Whether two contracts are truly identical is a subtle and difficult question, as the contracts often include complex and detailed rules that address rare and varied edge cases (e.g., the death of a candidate); yet, we are confident that all matching contracts noted in this paper were for all practical purposes identical; the chances that their payoffs would differ approached zero and were negligible given the other transaction costs and risks detailed below.

Opportunity cost is the prevailing interest rate over the expected life of the contract on the investment. First, the maximum amount of money that a trader needs to purchase the contracts in both markets to attempt arbitrage is $1 per $1 payout. The cost of buying the less expensive contract is the price $X. The more expensive contract is going to sell for $Y, where $1 > $Y > $X. The cost of covering that sale is $1 – $Y; if the contract hits, the seller needs to cover the difference between the price and $1. Thus, the trader needs to invest $1 - $Y + $X. Since $1 > $Y > $X the cost of covering the two positions can, at most, approach $1 per $1 contract. If the investor does not sell the contracts, but holds everything for the duration, they will expire at the

[1] http://media.investis.com/B/Betfair/PDFs/Annual-Reports/Betfair_Annual_Report_2012.pdf

Conventions or Election Day; one year of time, at most. With very low interest rates, the opportunity cost on the investment was quite low during 2012.

Transaction costs are unique in this setup, because they are unbalanced; Intrade has a single upfront monthly cost while Betfair has a sliding scale of marginal costs. Thus, we compute the final transaction cost in expectation. First, the trader covers all of the fixed fees; for Intrade that includes a rate of $5 per month. Second, the marginal transaction cost is the likelihood that the Betfair contract pays out times the transaction fee. The transaction fee is a sliding scale that starts at 5% and ends at 2%; it decreases with the trader's volume of action. [2] Thus, there is a meaningful distinction between the costs for institutional investors and the small individual investors. For example, if a small investor buys 100 shares of a contract that would pay $1 if it comes true in Betfair at $0.60 per share and sell 100 shares in Intrade for $0.70 per share then her transaction costs, in expectation, is 65% * $40 * 5% = $1.30. Of course, she will either pay $0 if the contract does not pay out or $2.00 if the contract pays out. The highest possible transaction cost is if the contract costs ≈$0 and it hits on Betfair, with a fee of 5% of the total. In summation, the maximum transaction and opportunity cost for a small investor is ≈6% or 5% (transaction) + ≈1% (opportunity). Realistically, the less that can be gained from Betfair the lower the transaction cost; if the contracts are both near $0.50 per $1 payout and we assume a short time period, than the cost is ≈3% or 0.5*5%. An institutional investors' cost is substantially lower.

The currency cost and risk is very divergent for small and institutional investors. A small investor is likely to ride the currency risk from Intrade's U.S. dollars and Betfair's British pounds, while an institutional investor with access to currency futures can hedge the risk at a small cost. Further, the cost of the initial conversion also depends greatly on the access to currency of the traders, from little to nothing for large investors who hold foreign currency already to more costly for small investors.

The counterparty risk is non-negligible and hard to calculate. The traders cover all margin calls at 100% with money in the exchange; by law, the money the exchanges hold is supposed to

[2] They also add additional penalties for highly profitable traders and traders active across many markets, resulting in charges as high as 20% on big wins for frequent and successful traders on Betfair. See http://www.betfair.com/aboutUs/Betfair.Charges.

be held safely, separate from the day-to-day operations of the company. Thus, the counterparty risk during the election was low, because an exchange would not only have to go under, but take the money with it. And, these exchanges are very visible and long-lasting institutions; Betfair is a publically traded company founded in 2000 and the popular press regularly quotes and references Intrade, founded in 1999. Yet, Intrade is a smaller firm than Betfair; it runs more informally and has undergone a number of changes, including shuttering its sports betting operation TradeSports and losing its CEO in a tragic mountain climbing accident. When the United States Commodity Futures Trading Commission (CFTC) went after Intrade shortly after the 2012 election, forcing them to exit the U.S. market entirely, the company was able to quickly refund all U.S.-based investors in full, along with any other investors who wished to cash out at that time. (In contrast, when the U.S. government went after Full Tilt Poker in 2011, that company was unable to pay back investors in full.) After the loss of a substantial portion of their trader base, Intrade uncovered accounting irregularities in their books and, in March of 2013, froze all accounts worldwide; they finally paid all accounts in full in November of 2013. While any 2012 arbitrage investor would have already moved her money out of Intrade, the risk of a 2012 arbitrage investment becoming stuck was obviously non-negligible.[3] Thus, while the absolute percentage of failure is small, a small risk-averse investor may limit their participation as their investment becomes substantial in relation to their wealth. We estimate the currency and counter-party risk to be low for institutional investors.

The exchanges operate on different platforms, with different pricing schemes and other obstacles that make it taxing to investigate and close price misalignments between the markets. Most of this cost is fixed, but since Intrade only sees major liquidity in political markets, traders cannot amortize the cost, for example by writing a computer program that does the search automatically, over many opportunities. Because of its size and diversity, institutional accounts reside almost exclusively on Betfair. Intrade operates in U.S. dollars and caters to U.S. users while Betfair denominates trades in British pounds and explicitly bars access from the United States.[4] There may be few people who regularly maintain accounts in both exchanges.

---

[3] It is important to note that the action for the CFTC after the 2012 election and the subsequent freezing of accounts in March of 2013 was a long-tail event. We first presented this paper in October of 2012 and audiences generally regarded the collapse of Intrade as a non-negligible, but low probability event.

[4] Access for United States-based investors should decrease after the 2012 election due to the CFTC action.

We pulled 10 contracts with high volume during the 2012 Presidential primaries that were listed on both Betfair and Intrade; misalignment of prices occurs in all 10 of these contracts. The contracts are: Gingrich, Romney, and Santorum for the Republican nomination, Obama and Romney for president, and whether Romney would win the Iowa, New Hampshire, South Carolina, Nevada, and Florida primaries.[5] More often than not, it was possible to sell a contract in one exchange for more than it cost to buy the same contract in another exchange. If the markets are perfectly efficient, this should never happen. The left side of Figure 1 shows both the bid and the ask for the Betfair and Intrade contracts on Obama to win the presidency across the general election once per day for all of 2012. With the exception of just 2 days at the beginning of the year, the bid on Betfair (i.e., the price someone is willing to pay) was greater than the ask on Intrade (i.e., the price it would cost to buy). The differences are persistent between the two markets as far back as March of 2011. An interesting phenomenon is that the size of the price misalignment grows towards Election Day, when the markets are more liquid and the opportunity cost of holding the arbitrage shrinks, along with the associated risks. The misalignment peaked at over 23 percentage points at 8:05 PM ET on Election Day.[6]

<Insert Figure 1 Here>

We conclude that there are real and persistent arbitrage opportunities between the two markets that outweigh the transaction and opportunity cost of exploiting them. Figure 1 shows that there are points where the difference is greater than the theoretical maximum cost of 6% and it is regularly greater than the average cost of 3%. The two Obama contracts were over 5 percentage points apart for 64 of last 85 days of the election, a stretch starting just before the conventions on 8/14/2012 through Election Day 11/6/2012.

The stated order book only provides a lower bound to the size of this opportunity; we explore the shadow order book to make a more accurate accounting. Regardless of how the market categorizes it, for easy comparison, we count a share as a contract that pays $1 if it comes true and $0 if it does not come true. At 1:45 PM ET on August 16, 2012, a trader could buy on Intrade 120 shares of Obama to win the presidency for $0.569 per share, then 4,944 shares at

[5] These are the most consistently liquid contracts and the only contracts that span the entire timeframe.
[6] In 2008, Betfair was also more bullish on Obama over McCain by a similar magnitude, but that difference did collapse by mid-October.

$0.570, 210 shares at $0.571, etc. At the same time on Betfair a trader could sell 18,857 shares of Obama to win the presidency at $0.633 per share, 1,858 shares at $0.629 per share, 1,349 shares at $0.625 per share, etc. The minimum amount of money that a trader could gain by closing this price misalignment is to buy shares in Intrade and sell shares in Betfair until the difference matches the transaction and opportunity costs; this assumes the trader has already covered all of fixed costs of being on both sites and finding the price misalignments. But, the order book is only what is sitting visibly to take; the shadow order book includes the market's reaction to new trades. First, on Intrade the bid was $0.562 and the ask $0.569; could the trader get someone to sell shares for a price in the middle? Second, there were 120 shares available at $0.569, what happens if the trader puts in a buy order for 500 or 1,000 shares at that price, are there other traders waiting to take the orders? If the minimum is clearing out the order book, the maximum is an infinite money pit at the marginal difference, or even at more than the marginal difference, by carving into some portion of the bid/ask spreads in both markets.

<Insert Figure 2 Here>

Figure 2 shows the actual order books from Betfair and Intrade referenced above. They demonstrate the complexity of working in both markets. First, Intrade's contracts expire at $10, so each "Qty" corresponds to 10 shares. Second, Betfair trades with odds, where 1/odds equal the cost per $1. The odds of 1.58 equate to 1/1.58 or $0.633 per $1. Third, Betfair trades in British Pounds and it lists the amount of money someone is willing to wager at the current price. Thus, at the exchange rate of 1.5642 American Dollars per Pound someone is will to wager £7,630 or $11,934.85 at $0.633 per $1 return. So, that bet equates to 18,857 shares that would be worth $1 each if Obama wins the election.

On any given day during the late fall, completely closing the stated order book would net an investor between $1,000 and $5,000. In order to test the depth of the shadow order book, we identified several price misalignments where there was arbitrage on the margin; we randomly traded or did not trade in those contracts on any given day to test the cost of buying shares and whether our playing affected the marginal values. First, we identified two sets of related contracts where we could buy the same contract for less than we could sell it in a different exchange and the difference was enough to cover all transaction and opportunity costs. Second, every day for eight days we randomized which market to enter at a random point during the day.

We followed the markets when we did not bid; we captured the order book every 2 minutes during the entire period. Third, if we entered the market we bought matching numbers of shares on each side, to ensure an arbitrage.[7] Fourth, when we bought shares we started by offering at the current bid + $0.001 per share and then moved upwards systematically until we acquired all of the shares. For example if the current bid was $0.230 and current ask was $0.240 we would attempt to buy at $0.231, then $0.232, then $0.232, etc.

The shadow order book exists; we consistently paid less for our contracts then the stated order book indicated and there were more shares available than noted in the order book at the stated prices. As an example, we went into the market shown by Figure 2 and bought 5 shares of Obama to win at $0.566 per share; that was $0.003 per share less than price in the order book. Then, to match this arbitrage opportunity, we sold 5 shares worth of Obama to win in Betfair for $0.637 per share, $0.004 more than the order book's asking price of $0.633 per share. Every time we entered the Intrade market there was evidence of the shadow order book. Six of the eight times we enacted purchases there was space between the bid and ask price in the order book; four of those six times we were able to purchase shares at less than the ask price. Two of the remaining four purchases we bought more shares than were available in the order book at the ask price. Both of those times, after we cleaned out the shares available at the current ask price, the number of shares available recovered within minutes to their original level.

The evidence suggest that the shadow order book multiplies the return from between-market opportunities many times over. We had no measurable influence on the market despite investing $3,686 in arbitrage situations that paid out a return of 6.38% over 3 months.[8] Over the time period that we traded, we examined the starting bid and ask prices and quantities for each day. The difference increased day over day in 6 of the 8 contracts we entered. More telling, after trading for 8 days, and spending $3,686 buying arbitrages in the two contracts, the widest price misalignments to that point came at midnight after we had finished.

Our field experiment shows that utilizing the shadow order book we can estimate a conservative net of $15-20,000 over the course of the last few months of the election; this

[7] For simplicity, we occasionally exchanged buying the mutually exclusive counterpart for selling a contract.
[8] The return is including all transactions costs, but before any opportunity costs.

estimate is at least a magnitude larger than the $1-5,000 that could be gained by closing the arbitrage in the order book during the same time frame. On a daily basis we cleared $15-$25 without making any impact on the markets. We did one large investment of $1,841.80 which had no lasting effect on the price misalignment by the end of the day, for a post-cost return of $135.51. If an investor wanted to invest daily, allowing the market to recover from the direct impact of his investment, he could conservatively net $150 to $200 daily for upwards of the last 100 days. We had no method of investigating what would happen if an investor completely closed the misalignment and whether it might subsequently return.

The investment possibility falls into an awkward mid-size spot; it may be rational for investors to not actively close the price misalignment. Small investors face higher costs and risks in making this investment. With the proper level of risk-aversion they may choose not to invest because of high probability attached to a possible exchange closure and forfeit of capital and currency fluctuations. Institutional investors have lower costs and risks and should be more risk neutral. Yet, they may view this opportunity as too small for their capital investment. The search cost of finding out about the arbitrage opportunity approached zero. This paper was presented at a Nation Bureau of Economics Research meeting in mid-October 2012, several prominent economics bloggers both tweeted and discussed this arbitrage opportunity, and several mainstream media articles highlighted it as well.

Even without investors actively closing the price misalignment, it should close with the dissemination of the price to investors in each market. The price on Betfair and the price on Intrade are valuable data points on the likelihood of the election outcome that should influence traders in each market to move their market in that direction. Yet, Intrade operates with U.S. dollars and advertises towards users in the United States and Betfair operates with British pounds and advertises towards users in Europe. It is possible that persistent informational differences or biases occur between the two geographically distinct populations, similar to the geographic bias observed in sports betting markets [Wong, 2001].[9] In another example, dual-listed companies have a rich history of maintaining different pricing on different exchanges, well beyond any reasonable differences in value [Rosenthal and Young, 1990 and Froot and Dabora 1999]. And,

[9] Although, in the sports wagering example, geographic differences rarely extend beyond the costs.

some researchers suggest local sentiment is a factor. But, there is a unique difference with these stocks versus prediction contracts; the strategy for closing those gaps is not entirely clear as the dual-listed companies could persist in their differences for years and margin calls could eliminate all theoretical gains in buying and selling in the two markets [De Jong et al. 2008]. Yet, in our example, prediction contracts will expire at $0 or $1 in a matter of months, weeks, days, or in the case of Election Day, hours.

This leaves a final concern over manipulation of one market; an investor could decide to maximize something other than return and artificially keep the price up for one of the candidates; that was likely the case in Intrade market in 2012. If the public and press consider the market price valid, manipulating it can be a rational strategy as increased perceived likelihood of victory for a candidate may lead to increased support and engagement for that candidate, thus actually increasing the true likelihood of victory for that candidate [Simon 1954]. Previous research has shown this to be a very difficult task in prediction markets [Rhode and Strumpf, 2008]. To explore this possibility, we examined Intrade data on every sale of “Obama to Win” or “Romney to Win” contracts for the final two weeks of the election, from 5:00 AM ET on Tuesday, October 23 through 9:00 PM ET on Election Day, Tuesday, November 6, 2012. During that time period, one trader spent 35.0% of all money that supported Romney (i.e., either buying Romney to Win or selling Obama to Win). The amount of money is significant enough that we can assume that the trader would have lower costs if s/he chose to enter Betfair rather than Intrade. The trader could have made these purchases at about a 20-25% discount on Betfair (an average of about $0.07 per share on $0.30 per share purchases). We assume if the purchases were made on Betfair, they would have exacted upward pressure on the Romney price there; thus, you can consider this a conservative estimate.

Within-Market Arbitrage

Within-market price misalignment includes two major categories: mutually exclusive contracts (e.g., first, second, and third place in the same contest) and conditional contracts (e.g., victory in the first round and victory in the second round). Within-market price misalignment is when you can buy a contract that is, by definition, at least as valuable as another contract for less money than that second contract. Price misalignment in mutually exclusive contracts occurs when you can sell over $1 worth of contracts for an outcome that can pay a maximum of $1, or you can buy a set of exhaustive contracts for less than $1 that must pay out at $1. Price misalignment in conditional contracts occurs when something that is by definition more likely to happen sells for less money.

Every question in a prediction market has a mutually exclusive list of outcomes. If the question is about winning the presidential election, there is going to be a list of candidates and only one contract can be worth $1 after Election Day and all other contracts will be worth $0. This is the easiest within-market comparison to keep aligned, because the markets generally list all of the contracts for one specific question together. Yet, price misalignment does occur with within these markets.

<Insert Figure 3 Here>

There were two major price misalignments in Intrade's market for the second place position in the New Hampshire primary on January 10, 2012. First, at some points, not only on the day of the primary, as seen in Figure 3, but throughout the previous week, a trader could sell all the contracts on every possible candidate to finish second for more than $1. That includes not just Jon Huntsman and Ron Paul, but Mitt Romney, Rick Santorum, and Newt Gingrich as well. This topped out at a possible sale price of $1.082 and, since only one candidate could capture second place, it was worth $1; the sale of each share bundle guaranteed $0.082 profit. Second, right after the polls closed, Huntsman, who finished third, plunged in both the bid and ask for second, but Paul, who came in second, stayed steady. Thus, for a few minutes a trader could buy all candidates for second place for less than $1; it actually bottomed out at $0.591. Since someone had to finish second, each share bundle purchased for $0.591 was worth $1, for momentary profit of $0.409 per share.

On the margin, mutually exclusive price misalignment happens surprisingly often. Figure 4 is a more common phenomenon where the market is just slightly misaligned. A seller could sell 26 shares in each of candidate in the market for the 2012 president and get $1.003 for something that will cost $1.00 by definition. Intrade makes this very clean by adding the final contract of 2012.PRES.OTHER, making this a fully encompassing market. Betfair does not always include that, leaving the possibly that all contracts for a question could be losers.

<Insert Figure 4 Here>

A more hidden mutually exclusive situation is the set of contracts that cover every possible thing a candidate may do in a situation and we find price misalignment there as well. For example, in the early hours after the polls closed on the Iowa primaries, there were a few minutes were people were willing to buy Romney to finish first for $0.900 and Romney to finish second for $0.110. That was a guaranteed 1 percent return on something that was going to settle that day. On the other side, in the New Hampshire primary, Paul, who finished a convincing second, had points in time where a first, second, or third finish could be purchased for a total of $0.598; this is illustrated in Figure 5. While that is not technically arbitrage, because he could theoretically have not finished in the top three, it was a dramatic price misalignment for a candidate that came in with 22.9 percent of the vote when the first candidate out of the top three had just 9.4 percent.

<Insert Figure 5 Here>

We pulled all of the major candidates for president to win their party's nomination and to win the general election and we did not find points where you could buy them to win the nomination (round 1) for less than you could sell them to win the general election (round 2), but we do find evidence of price misalignment.[10] Barack Obama needed to win the Democratic nomination to run for president in 2012 and as the sitting president there was a very high likelihood of him winning the nomination. If he did not run as the Democratic candidate he would have zero likelihood of winning the election, but the Democratic Party could still win the election; thus, Democratic Party to Win is more valuable than Barack Obama to Win.[11] Yet, there are many times when the price of the Democratic Party winning the election was less than the price of Obama winning the election. For example, on the morning of September 18, 2011 you could buy Democratic Party to Win the Election for $0.491 and sell Barack Obama to Win the Election for $0.501.

Market Inefficiencies

[10] With the exception of Ron Paul, who was the only major party candidate with a non-negligible possibility of running as a third party candidate if he lost his party's nomination.
[11] Intrade provides a little ambiguity in the case of death. The contracts on the Parties would continue, but the individual contracts would likely be suspended and settle at the previous night's price.

Mutually exclusive and conditional contracts both have less salient markets become illiquid when the more salient contracts are moving due to increased information. For mutually exclusive contracts that means markets for first or second place in a contest having robust trading while a third place market develops a large bid/ask spread. Similarly, for conditional contracts, it is means later rounds become illiquid while information is coming in for earlier rounds. Figure 6 shows a very dramatic illustration of this phenomenon in NCAA basketball. The figure charts both the likelihood of Kansas and Ohio State winning their semifinal game in the 2012 NCAA tournament and winning the final game, on Betfair. By definition, if a team loses its semifinal game it cannot win its final game. In the hours before tipoff all contracts are very liquid with tight bid/ask spreads. A little time before tipoff all bids are removed from the contracts for the finals and do not return until after the game is finished. It would be rational to leave ask offers at the value they would be if the team won the game. During the semifinal game there are extremely liquid contracts for the outcome of that game. Several hours after the game is over the market for the final game becomes liquid again for the game's winner, Kansas.

<Insert Figure 6 Here>

This saliency issue is critical in times of lower information flow as well. While neither shuts down, there is an unnatural degree of noise, and occasionally dubious relationship, between first and second round contracts in political events on Intrade. The price of a candidate to win the general election divided by the price of the candidate to win the nomination is the conditional price of the candidate winning the elections should the candidate win the nomination. Figure 7 illustrates this noise by mapping this conditional price for Mitt Romney along with the contract for the Republican Party winning the general election. On a day-to-day basis, the underlying conditional value should be relatively stable, a similar magnitude of volatility to the party's contract, as the likelihood of nomination was stable from day to day, but it bounces around by several points. This is much more extreme on the bottom chart, which shows Newt Gingrich during the period when his likelihood of being the nominee was non-negligible. This conditional value is a very important issue for political scientists and this volatility and noise makes it much harder to track with the precision needed to make strong inferences.

<Insert Figure 7 Here>

A subtle inefficiency occurs with the imbalance in the bid and ask spread; the sum of asks is consistently further from one dollar than the sum of all bids. We see a similar pattern on several exchanges, including Betfair, Intrade, the Iowa Electronic Market, and even virtual exchanges like the Hollywood Stock Exchange [Pennock et al. 2001]. We conjecture this imbalance occurs because people better understand and thus prefer to buy shares rather than to short sell, exacerbated by the often-confusing ways that exchanges implement selling. For example, we sampled the contracts for the winning party in the 2012 election every day at noon from January 1, 2012 through September 30, 2012 and the sum of all bids was further from one dollar 100 times relative to 153 times for the sum of all asks. This imbalance is systematic and especially acute in times of rapid trading. For example, we captured 159 in-game snapshots for the Kansas and Ohio State game shown in Figure 6 and the sum of the asks is further from one 92 times. Like all of the within-market inefficiencies this becomes more extreme in times of high information flow. We show an illustrative example in Figure 8, where the marginal order book for the 2012 Indiana senate race the day after the Republican candidate said some controversial remarks in a debate. The bids are placed up to $0.970 per $1.00, but the asks linger at $1.147.

<Insert Figure 8 Here>

These two inefficiencies help explain the within-market price misalignment we catalog in this paper. First, since users are unable to accurately work in related markets concurrently, they triage contracts that are more secondary, whether they occur later or have less real-time information at that point in time. Much of the price misalignments described above occurs during times of rapid information, like the hours after the polls closed in a primary, as shown in Figures 3 and 5; the misalignments happen because focus is temporary dropped on secondary or tertiary related contracts. Second, traders focus on the bid pricing, more than the ask pricing, because it is harder to think in terms of sales price. Thus, the average of the bid/ask may not be the most efficient price in all contracts and if the ask is moving a little more arbitrarily the midpoint will show inefficient volatility. These two inefficiencies help explain the volatility and occasionally suspect relationships between conditional contracts shown in Figure 7.

This is a satisfying explanation for within-market price misalignment, insofar as it is reasonable that the limits of human computation and multiple screens restrict the ability of people keep too many contracts consistent at any time and the unique nature of Intrade and other prediction market exchanges make it inefficient to pay the fixed costs of overcoming these inefficiencies with customized programs.

During the primary, Betfair consistently lagged behind Intrade, a violation of the efficient market hypothesis thought not strictly an arbitrage. Using the same 10 elections noted earlier, the five biggest primary contests and the five biggest general candidate contracts, we ran a simple set of regression: $P_{at} = \propto + \beta_1 P_{at-1} + \beta_2 P_{at-2} + \beta_3 P_{at-3} + \gamma_1 P_{bt-1} + \gamma_2 P_{bt-2} + \gamma_3 P_{bt-3}$, where P_{at} is the price in market “a” at time “t”. We used a lag of 12 hours and ran this regression for many variations of the lag, but the results are all strikingly the same. Table 1 shows that Intrade’s first lag has a huge correlation with Betfair’s price, but Betfair has a small correlation with Intrade’s price. If we were asked to forecast Intrade’s price in twelve hours, at any given point in this dataset for any given contract, we should say the approximately the current Intrade price; that is efficient. If we were asked to forecast Betfair’s price in twelve hours, at any given point in this dataset for any given contract, we should say the approximately the current Intrade price*0.6 + the current Betfair price*0.4; that is not efficient. There is a similar story if when we regress changes over time by regressing the change in the price in the current time period on the two previous changes in prices: $diff_{at,t-1} = \propto + \beta_1 diff_{at-1,t-2} + \beta_2 diff_{at-2,t-3} + \gamma_1 diff_{at-1,t-2} + \gamma_2 diff_{at-2,t-3}$. Intrade’s difference has a small, but positive and significant correlation with movement in the previous two periods of both markets. This represents the underlying drift in prediction market prices as time runs out on losing candidates and winning candidates run down time (i.e., if nothing happens, the candidates will drift to nearly $1.00 and $0.00 as Election Days approaches). Betfair has massive, significant correlations with difference from the previous time periods in Intrade. A change in Intrade in a previous period is a strong indicator of change in Betfair in the current period. Further, Betfair has a negative correlation with its own previous periods. Any way we slice the data Betfair is trailing Intrade.

<Insert Table 1 Here>

This is not an explanation for the price misalignments that occur between exchanges, because the lagging market could close the gap instantaneously (or within seconds). As we discuss above, key information, in the form of the current price on the other exchange, is readily available to traders in both exchanges.

III. Designing Better Prediction Markets

Industry wide, from the Iowa Electronic Market [Berg et al. 2008] to the Chicago Board Options Exchange, from Las Vegas bookmakers to the Kentucky Derby racetrack, related outcomes are sold as independent instruments with their own order flow and processing. Betfair's Kansas-Ohio State market in Figure 6 is a good example. The two outcomes are mutually exclusive and exhaustive: buying Kansas is equivalent to selling Ohio State. Running separate auctions for both outcomes is redundant. Intrade's slew of presidential election markets and candidate ranking markets, like Figure 5, have more than two outcomes, yet the argument against operating each outcome independently still applies, perhaps even more forcefully. There are five reasons.

1. Splitting up a market can hurt liquidity. In a split market, there are effectively two ways to do everything (e.g., buying the Democratic candidate equals selling the Republican and third-party candidates), so traders may not see the best price for what they want to do, and orders may not fill at the best price available [Oliven and Rietz 2004]. There may even be orders that together constitute an agreeable trade, yet are stuck waiting in separate queues.
2. Splitting the market limits expressiveness. For example, a natural prediction, common at the racetrack, is that a candidate will "place", or finish in first or second place in a race. Expressing this on Intrade requires two transactions, increasing the implied bid-ask spread, and introducing an execution risk that prices will shift in the interim. (Conversely, you can't directly bet on a horse to finish in exactly second place at the racetrack.) A common fix is to open yet another independent market in each popular derivative; however this limits choice and exacerbates the other problems listed here. Bundling is especially useful with interval bets. For example, to predict that a stock will fall within a certain range at a future date requires four options trades, a so-called butterfly spread. When outcomes are disjoint, an interval bet may require dozens of trades to acquire all outcomes in the interval. (For example, if the outcomes are the numbers Electoral

College votes the Democratic candidate might accrue in a Presidential election, the bet "Democrat to win between 200 and 299 electoral votes" requires 100 trades.) Moreover, traders must sum the intervening prices manually to compute a price quote.

3. A split market may slow information propagation. Price changes in one outcome do not directly affect prices of other outcomes; it is left to arbitrageurs to propagate logical implications. As our analysis shows, arbitrageurs are not always instant or efficient.
4. A naïve implementation of a split market may limit traders' leverage, forcing them set aside more money than necessary to complete a set of trades. For example, on the Iowa Electronic Market, short selling one share at $0.99 requires that you have $1 in your account, even though the most you could possibly lose in this transaction is $0.01. The reason is that to short sell on IEM you must first buy the bundle of all outcomes for $1 and then sell off the outcome that you don't want.
5. Simple Internet searches reveal dozens of companies that peddle programs for users to capitalize on potential price misalignments on prediction markets including the simplest mutually exclusive contracts. The price misalignments still exist despite these programs because there is no turnkey solution to the close all of the price misalignments in the constantly evolving set of contracts on a given exchange. Once people invest in these programs, they need to invest further considerable time, effort, thought, and money in detecting and capitalizing on the price misalignments. This cost would be more efficiently spent investigating new contracts and markets. Even when arbitrageurs are effective, they draw rewards away from participants who actually provide information.

The solution is to treat multiple disjoint outcomes holistically rather than separately. The natural generalization of the continuous double auction (CDA) to multiple outcomes is to use linear programming, as several authors have noted. The mechanism has been called combined-value trading [Bossaerts et al. 2002], a pari-mutuel call market [Baron and Lange 2005, Lange and Economides 2005, Peters et al. 2006], and a combinatorial call market [Fortnow et al. 2004].

The underlying principle is straightforward. Let Ω be a set of disjoint exhaustive future outcomes, say all possible numbers of Electoral College votes that the Democratic candidate might receive: that is, all 539 integers between 0 and 538. Following standard industry practice, a betting exchange might list 539 separate CDAs, or partition the outcomes into a few coursegrained ranges. We will instead describe a generalized betting exchange that operates over all 539 outcomes simultaneously and allows traders to price or buy any set of outcomes they want in a single operation. We refer to the operator of the betting exchange as the "auctioneer". Her job is to collect orders from traders, then clear the market by matching acceptable trades together. The auctioneer uses linear programming to compute which subset of orders she can accept and fill. The program optimizes a linear objective function subject to a set of linear constraints or inequalities. We will define the constraints and objective in detail below.

Traders submit their orders to the auctioneer. Each *order* O has three components: the maximum *price* p_O that the trader is willing to pay, the maximum number of shares or *quantity* q_O that the trader is willing to purchase, and the actual event E_O or proposition that the trader wants to bet on. E_O encodes a subset of outcomes $\omega \in \Omega$. For example, E_O={270,271,272,…,330} corresponds to the bet "Democrat to win between 270 and 330 electoral votes".

The auctioneer scans the orders to systematically determine which, if any, she can fulfill. How does she decide? An auctioneer is a neutral third party: her only job is to match willing traders together. She cannot and should not participate in any gamble herself. (In contrast to an auctioneer, a *market maker* does take on risk of his own.) So the auctioneer seeks to find a set of orders such that any trader's gain in any outcome is always fully balanced by some other traders' losses. The auctioneer cannot accept any set of orders that collectively expose her to a loss, even a potential loss in a single outcome.

The auctioneer maintains a decision variable x_O for each order. The linear program will, in the end, assign a real number between 0 and 1 to each variable x_O. The decision variable tells the auctioneer whether to fill the order. If $x_O = 0$, the auctioneer must be reject the order for now. If $x_O = 1$, the auctioneer can accept the order in full and grant the trader all q_O shares they requested. If x_O is something in between, say 0.8, the auctioneer can accept the order only partially, granting the trader the specified fraction of shares, in this case 80% of what they requested.

The auctioneer adds constraints $0 \leq x_O \leq 1$ to the linear program, bounding each x_O between 0 and 1. She also adds one *budget constraint* per outcome ω to encode the requirement that she not take on any risk of loss herself. Each constraint looks like

$$\sum_O x_O q_O (1_{\omega \in E_O} - p_O) \leq 0,$$

where $1_{\omega \in O}$ is the indicator function that equals 1 if the event E_O is true in outcome ω (i.e., E_O contains the outcome ω), and equals 0 otherwise. The left-hand side of the inequality is the total net payment the auctioneer makes to traders in outcome ω. Together, the budget constraints ensure that the auctioneer's net payment is always either zero or negative in every outcome. (A negative net payment means the auctioneer earns a surplus. She can either keep the surplus as profit or return it in full or in part to traders.)

The auctioneer specifies an objective function that she wants to maximize. A natural objective function is volume of trade, or $\sum_O x_O q_O$. Another is fill fraction, or $\sum_O x_O$. A third natural objective worst-case auctioneer profit, or the minimum profit across all outcomes, ensuring that regardless of the outcome the auctioneer does not fare too badly. (To code this, replace the 0 on the right-hand side of each budget constraint equation with *k* and then maximize *k*.) Ultimately, the auctioneer can choose whatever (linear) objective function makes sense depending on her goals.

The auctioneer feeds the objective and constraints into a linear program. The program finds the optimal values for x that maximize the objective subject to the constraints. The auctioneer can implement a *call market* by running the program in batch mode, clearing the auction after waiting to collect a number of orders. Or the auctioneer can implement a continuous auction (analogous to the stock market's CDA) by running the program immediately as new orders arrive.

Traders can bet on any proposition—whether a single outcome, a negation of an outcome, or an arbitrary bundle of outcomes—in a single transaction. Every order goes into the same pool of liquidity.

Note that, if traders are allowed to place all-or-nothing orders (enforcing $x_O \in \{0,1\}$), or more generally allowed to specify any minimum fill constraint, the program becomes an integer program, not a linear program, and the clearing problem becomes NP-hard [Bossaerts et al. 2002, Fortnow et al. 2004]. We don't believe the absence of fill constraints presents a major barrier to trade. We expect many traders willing to risk $50 to win $100 will also be happy to risk $25 to win $50, for example.

For reasonable numbers of disjoint outcomes, say 539, using linear programming is fast, reliable, and well understood. We see almost no disadvantage to using linear programming rather than splitting outcomes into independent markets. Yet industry-wide practice is dominated by independent markets. One exception is the now defunct economic derivatives markets run by Longitude, Goldman Sachs, and Deutsche Bank [Baron and Lange 2005].

When the number of outcomes grows too massive, for example all 2^{51} or 2 quadrillion possible state-by-state election outcomes, an explicit linear program becomes intractable [Chen et al. 2007, Fortnow et al. 2004]. In this case, limiting the expressiveness of bids (i.e., restricting what bundles are allowed) can recover tractable algorithms, though often the limits are severe and impractical [Agrawal et al. 2008, Chen et al. 2007, Chen et al. 2008a, Chen et al. 2008b, Guo and Pennock 2009, Pennock and Xia 2011]. Alternatively, approximation schemes are possible, both stochastic [Chen et al. 2008b] and deterministic [Dudik et al. 2012, Lahaie et al. 2013].

Standard market interfaces create and compound inefficiencies that wizard-style market interfaces can correct. Lowering the barriers to entry in both the market and specific contracts is useful for all stakeholders. Investors have more liquidity and markets have more volume. Researchers are likely to benefit from an increased diversity of the user base. Wizard interfaces gather expectations and convert them into efficient purchases. Tescher and Rothschild [2012], among others, show the advantage of wizard interfaces can be threefold.

1. They lower the barrier of entry by allowing people to provide information without learning the ins and out of trading in a particular market or any market. Markets currently operate with intricate interfaces (see Figure 2) that require high fixed costs for users to learn. Wizards can create simpler environments for traders where they can input their expectations, in ways designed to make to be efficient for both lay traders.

2. Traders can enter multiple contracts with just one expectation, raising the consistency and liquidity of the overall prediction market. First, the purchases would be internally consistent to the users' expectation. Second, the wizard helps the user operate in lower liquidity markets where the trader may not otherwise bother investigating.

3. Although fully rational traders would not alter their behavior based on the market front end, human traders, especially novices, will often provide more information if the process is simple and understandable. Thus wizard interfaces gather more data than standard interfaces, including data from the shadow order book, getting the subjective expectations of those who invest and those who do not invest. In a way, a wizard interface is a cross between a poll and a market. In liquid markets, traders are spending a lot of time and effort to create continuously updated expectations of the outcomes. Yet, they rarely provide this information to the market. A wizard invites the trader to continuously provide their expectations in a setting where they are not necessarily exposing their information to other traders (just the market) or exposing themselves by leaving orders on the order book. This information can be critical to understanding the efficiency of markets and provide better estimations of the outcomes that we design prediction markets to test.

Lowering the barrier to entry can have a cascade effect for the markets. Hanson and Oprea [2007] lays out the argument that manipulators can aid prediction markets by increasing the returns for informed investors. This is the same principle in that new users, possibly less informed users, will increase the return to informed users, thus creating a cascading effect of even more expert users. This effect holds not just for the overall markets, but individual contracts, which will be cheaper to enter because wizards can recommend many trades at once from an individual's expectation.

Investors benefit from the ability to match more trades, in whatever method is most comfortable to them, and the added liquidity should allow them to use the markets more efficiently as a hedge on other investments.

Their pricing schemes indicate that the exchanges are willing to trade arbitrage orders for more matched expectation orders. Prediction market exchanges do not consider the added trades from arbitrage traders a good thing. Intrade operates on a flat monthly fee, so it adds trades with no additional transaction costs. Betfair charges high frequency traders a surcharge, a market reaction to the exchange not viewing them as beneficial to them.

Researchers benefit from more information in the market leading to more efficient prices on more questions. Prediction market prices are not always efficient; this paper shows that in detail. Yet, the current state of information is still extremely useful to researchers in explaining the effect of events on other events, and providing efficiency for decision makers. More traders containing more information can only make the markets more efficient and more useful in both capacities.

IV. Discussion

A meaningful net profit can result in closing price misalignments of the same contracts in different exchanges. Identical contracts on different exchanges can have differences representing between 1 and 5 percent net earnings; these are common and can persist for months, even in the face of high liquidity. Observing the trading of thousands of dollars of contracts in a randomized trial, we demonstrate a significant shadow order book that indicates that the total possible opportunity may be several times the magnitude observable by simply closing the order book. Prices on one exchange have significant correlation with 12-hour lagged prices on a second exchange, but this does not explain the persistence. It is possible that the mid-size arbitrage is too risky for small investors and too small for institutional investors to close actively. Informational difference and/or biases between markets and the size of the opportunities could keep the divide open, be we conclude that there was active manipulation to buoy the perceived likelihood of victory for a chosen candidate.

While we can determine that the shadow order book is large—it seems to account for at least an order of magnitude more volume than the stated order book—we cannot give a concrete measurement of the shadow market without significantly more capital. It will vary considerably based on volume, prices, and timing. Further research will help create a more universal understanding of its size.

Within-market price misalignment, both on mutually exclusive and conditionally related contracts, occurs due to several inefficiencies of traders. High levels of activity and information on some contracts cause confusion on the related secondary contracts; traders unnecessarily withdraw orders or respond too slowly to changes in the primary contract. Further, there is a consistent asymmetry between buying and selling across many exchanges, leaving the average return for selling higher than for buying. Both of these inefficiencies lead to short-term price misalignment and extra noise in the relationship between contracts.

Overall, the markets studied function well considering the sometimes complex and subtle relationships among contracts and that has proven to be true historically [Rhode and Strumpf 2004, Rothschild 2009]; yet, changes in prediction market design can minimize within-market inefficiency, provide more information to researchers, and provide more utility for both the traders and exchanges. We suggest moving the burden of finding and fixing logical contradictions into the exchange, making buying and selling symmetric, and providing trading wizards, thus freeing traders to focus on providing meaningful information in the form they find most natural. This would at least eliminate within-market price misalignments, allowing the market to provide more useful information to both researchers and traders. It would also bring liquidity to contracts that currently lie fallow, providing expanded opportunity for traders to trade and profit for the exchange.

There are two reasons that current exchanges have not already adopted these logical relationships between contracts. First, the market for prediction markets is an oligopoly, dominated by a few major providers. We firmly conclude that logically related contracts and wizards will provide utility for investors, increase the quantity of investors, and make more money for the exchanges. Yet, there is no expectation that oligopolists should be rational once they have consolidated market share. Second, a properly logical market maker is very complex. In a political market it would not only tie mutually exclusive and conditional contracts, but all 51 Electoral College elections. Researchers are very eager to learn about the relationship between states in the same way that they are interested in to know the relationship between economic indicators or product launches. In the case of elections, it is a non-trivial matter keeping over 2^{51} contracts coherent, and while new approximation techniques [Dudik et al. 2012, Lahaie et al. 2013] are possible, they need more time in the field before for-profit exchanges utilize them.

Many research studies have shown that prediction market data can be utilized for accurate forecasting, regardless of its inefficiencies [Rothschild 2009]. Yet, the data has limits. First, when prices diverge between exchanges, the true best estimate is unclear. During the primary season of 2012, as Table 1 shows, Intrade was leading Betfair, so Intrade's price was likely more efficient at that point, yet during the general election the relationship appeared to reverse, with Betfair leading Intrade. In many situations, a simple average performs surprisingly well. Second, within-market inefficiencies are a significant concern in times of high information flow, especially if research includes secondary or tertiary markets related to a primary information source. The mid-point of the bid-ask spread, a good proxy for price in high-liquidity environments, can become meaningless when liquidity dissolves in the face of a rapid influx of information. The problem is exacerbated when computing conditional probabilities that involve two or more securities, as the worst-case bound on each individual price must be respected.

References

AGRAWAL, S., WANG, Z., AND YE, Y. 2008. Parimutuel betting on permutations. In *International Workshop on Internet and Network Economics*. 126–137.

ARROW, K. J., FORSYTHE, R., GORHAM, M., HAHN, R., HANSON, R., LEDYARD, J. O., ... AND ZITZEWITZ, E. 2008. The promise of prediction markets. *Science. 320*(5878), 877.

BARON, K. AND LANGE, L. 2005. *Parimutuel Applications In Finance: New Markets for New Risks*. Palgrave Macmillan, New York, NY, USA.

BERG, J., FORSYTHE, R., NELSON, F.D., AND RIETZ, T.A. 2008. Results from a dozen years of election futures markets research. In *The Handbook of Experimental Economics Results*, Eds. PLOTT, C.R. and SMITH, V.L., Elsevier Science, Amsterdam.

BOSSAERTS, P., FINE, L., AND LEDYARD, J. 2002. Inducing liquidity in thin financial markets through combined-value trading mechanisms. *European Economic Review* 46(9): 1671–1695.

CHEN, Y., FORTNOW, L., LAMBERT, N., PENNOCK, D. M., AND WORTMAN, J. 2008a. Complexity of combinatorial market makers. In *ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 190–199.

CHEN, Y., FORTNOW, L., NIKOLOVA, E., AND PENNOCK, D. M. 2007. Betting on permutations. In *ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 326–335.

CHEN, Y., GOEL, S., AND PENNOCK, D. M. 2008b. Pricing combinatorial markets for tournaments. In *ACM Symposium on Theory of Computing*. ACM, New York, NY, USA, 305–314.

COWGILL, B., WOLFERS, J., & ZITZEWITZ, E. 2008. Using prediction markets to track information flows: Evidence from Google. *Dartmouth College*.

DE JONG, ABE, LEONARD ROSENTHAL, AND MATHIJS A. VAN DIJK. 2009. The risk and return of arbitrage in dual-listed companies. *Review of Finance* 13.3: 495-520.

DUDIK, M., LAHAIE, S., PENNOCK, D. M. 2012. A tractable combinatorial market maker using constraint generation. In *ACM Conference on Electronic Commerce*. ACM, New York, NY, USA.

FROOT, KENNETH A., AND EMIL M. DABORA. 1999. How are stock prices affected by the location of trade? *Journal of Financial Economics* 53.2: 189-216.

FORTNOW, L., KILIAN, J., PENNOCK, D. M., AND WELLMAN, M. P. 2004. Betting Boolean-style: A framework for trading in securities based on logical formulas. *Decision Support Systems* 39(1): 87–104.

FRANCK, E., VERBEEK, E., NUESCH, S., 2013. Inter-market arbitrage in betting. *Economica* 80(318): 300-325.

FREY, S., SANDAS, P., 2008. The impact of hidden liquidity in limit order books. CFS Working Paper No. 2008/48.

GRUCA, T., BERG, J.E., AND CIPRIANO, M. 2007. Public information bias and prediction market accuracy. *Journal of Prediction Markets,* 1(3): 219–231.

GUO, M. AND PENNOCK, D. M. 2009. Combinatorial prediction markets for event hierarchies. In *International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 201–208.

HANSON, R., AND OPREA, R. 2009. A manipulator can aid prediction market accuracy. *Economica* 76.302: 304-314.

LAHAIE, S., DUDIK, M., ROTHSCHILD, D., AND PENNOCK, D. 2013. A combinatorial prediction market for the U.S. elections, In *ACM Conference on Electronic Commerce*. ACM, New York, NY, USA.

LANGE, J. AND ECONOMIDES, N. 2005. A parimutuel market microstructure for contingent claims trading. *European Financial Management* 11(1): 25–49.

OLIVEN, K. AND RIETZ, T.A. 2004. Suckers are born but markets are made: Individual rationality, arbitrage, and market efficiency on an electronic futures market. *Management Science* 50 (3): 336–351.

PENNOCK, D. M. AND XIA, L. 2011. Price updating in combinatorial prediction markets with Bayesian networks. In *Conference on Uncertainty in Artificial Intelligence*, 581–588.

PENNOCK, D.M., LAWRENCE, S., NIELSEN, F.Å., AND GILES, C.L. 2001. Extracting collective probabilistic forecasts from web games, In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 174-183.

PETERS, M., SO, A.M-C., AND YE, Y. 2006. A convex parimutuel formulation for contingent claim markets. In *ACM Conference on Electronic Commerce*. ACM, New York, NY, USA.

RHODE, P. AND STRUMPF, K. 2004. Historical presidential betting markets. *Journal of Economic Perspectives* 18(2): 127-142.

RHODE, P. AND STRUMPF, K. 2008. Manipulating political stock markets: A field experiment and a century of observational data. Working Paper. http://www.unc.edu/~cigar/papers/ManipIHT_June2008(KS).pdf

ROSENTHAL, L., AND YOUNG, C. 1990. The seemingly anomalous price behavior of Royal Dutch/Shell and Unilever NV/PLC. *Journal of Financial Economics* 26.1: 123-141.

ROTHSCHILD, D. 2009. Forecasting elections: Comparing prediction markets, polls, and their biases. *Public Opinion Quarterly* 73(5):895-16.

SIMON, HERBERT A. 1954. Bandwagon and underdog effects and the possibility of election predictions." *Public Opinion Quarterly*. 18(3): 245-53

SNOWBERG E., WOLFERS J AND ZITZEWITZ E. 2007. Partisan impacts on the economy: Evidence from prediction markets and close elections. *The Quarterly Journal of Economics*, *122*(2): 807-829.

STRUMPF, K. 2003. Illegal sports bookmakers. Working Paper. http://www.unc.edu/~cigar/papers/Bookie4b.pdf

TESCHNER, F. AND ROTHSCHILD, D. 2003. Simplifying market access: A new confidence-based interface. *The Journal of Prediction Markets*, forthcoming.

WOLFERS, J., AND ZIZEWITZ, E. 2009. Using markets to inform policy: The case of the Iraq war. *Economica*, *76*(302), 225-250.

WONG, S. 2001. *Sharp Sports Betting*. Pi Yi Press, La Jolla, CA, USA.

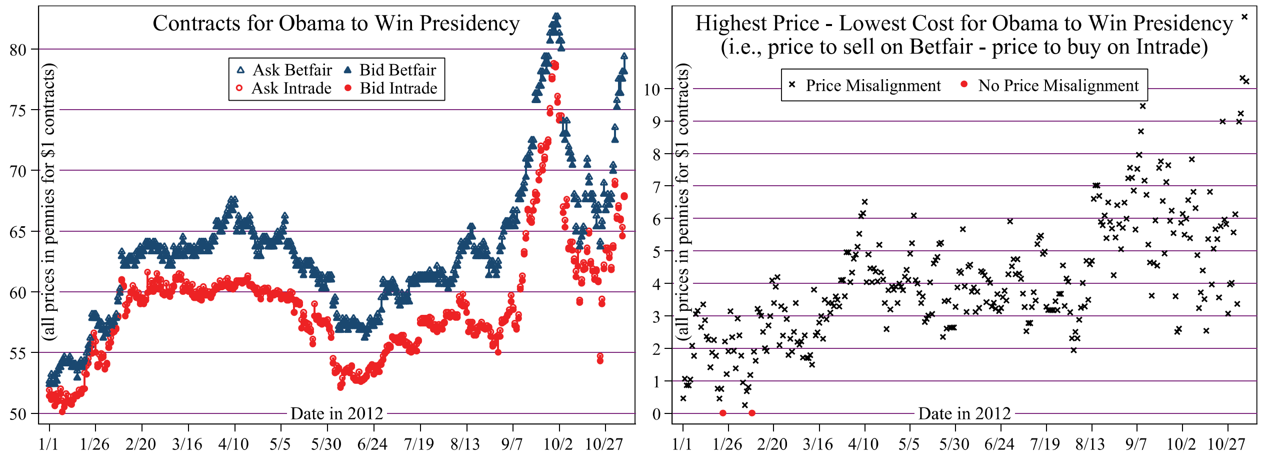

Figure 1: The difference between the prices for Betfair and Intrade (left) lead to the price misalignment (right), which is the highest price to sell minus the lowest cost to buy the same contract in the 2012 presidential election

Intrade

Order Book			
2012.PRES.OBAMA			
Bid		Offer	
Qty	Price	Price	Qty
1	56.2	56.9	12
51	56.1	57.0	494
156	56.0	57.1	21
132	55.9	57.4	6
987	55.8	57.5	523
36	55.7	57.9	45
348	55.6	58.0	104
1050	55.5	58.5	133
600	55.4	59.0	100
500	55.3	59.3	60
500	55.2	59.5	100
625	55.1	59.8	100
1124	55.0	59.9	215
100	54.9	60.0	200
104	54.8	60.2	207

Betfair

1.53	£130		£110,216
1.54	£992		£44,612
1.55	£6,573		£58,305
1.56	£3,499		£85,048
1.57			£122,367
1.58		£7,630	£63,837
1.59		£747	£64,479
1.60		£539	£70,414
1.61		£78	£63,195
1.62		£121	£55,921
1.63		£95	£58,138
1.64		£127	£86,808
1.65		£1,598	£92,735
1.66		£2,673	£29,193
1.67		£182	£21,888
1.68		£228	£30,840
1.69		£50	£41,203
1.70		£129	£25,480
1.71		£114	£16,005
1.72		£50	£17,674
1.73		£2	£54,606
1.74		£1,525	£71,804
1.75		£263	£62,116

Figure 2: The order books for the contract that paid out if Obama won the presidency on Intrade (left) and Betfair (right), at 1:45 PM ET on August 16, 2012

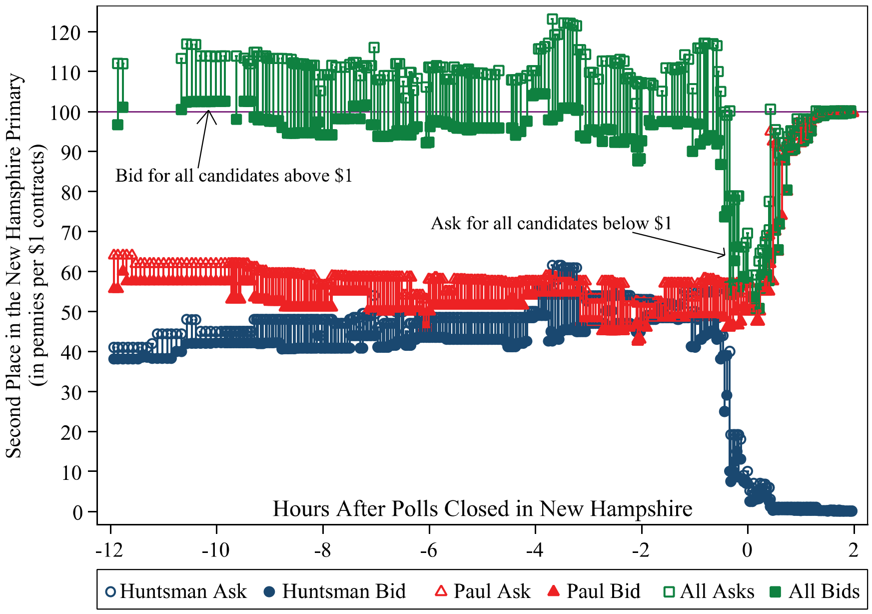

Figure 3: Intrade's bids and asks for Jon Huntsman, Ron Paul, and the sum of all competitors for second place in the New Hampshire primary

–Visited Events–

Add to Google

6:07:42PM GMT

Refresh

	Best to Sell		Best to Buy				
Contract	BQty	Bid	Offer	AQty	Last	Vol	Chge
Trade 2012.PRES.OBAMA	575	56.5	56.9	20	56.6	452.2k	-1.7
Trade 2012.PRES.ROMNEY	26	42.5	42.6	35	42.5	412.5k	+2.0
Trade 2012.PRES.PAUL(RON)	846	0.7	0.8	1188	0.7	421.6k	-0.2
Trade 2012.PRES.CLINTON	3579	0.2	0.3	1074	0.3	316.1k	-0.2
Trade 2012.PRES.HUCKABEE	999	0.2	0.3	1141	0.2	258.1k	+0.0
Trade 2012.PRES.JOHNSON	27.4k	0.1	0.2	14.1k	0.1	207.5k	-0.1
Trade 2012.PRES.SANTORUM	7360	0.1	0.2	3081	0.1	421.2k	-0.2
Trade 2012.PRES.BIDEN	0	-	0.1	9850	0.1	156.5k	0
Trade 2012.PRES.GINGRICH	0	-	0.1	26.2k	0.1	253.7k	+0.0
Trade 2012.PRES.PALIN	0	-	0.1	34.8k	0.1	159.9k	0
Trade 2012.PRES.CHRISTIE	0	-	0.1	54.1k	0.1	80.9k	0
Trade 2012.PRES.DANIELS	0	-	0.1	29.1k	0.1	106.0k	0
Trade 2012.PRES.BLOOMBERG	0	-	0.1	39.2k	0.1	85.5k	+0.0
Trade 2012.PRES.TRUMP	0	-	0.1	48.6k	0.1	74.5k	0
Trade 2012.PRES.HUNTSMAN	0	-	0.1	47.8k	0.1	94.2k	0
Trade 2012.PRES.THUNE	0	-	0.1	51.0k	0.1	57.3k	0
Trade 2012.PRES.CAIN	0	-	0.1	64.4k	0.1	70.5k	0
Trade 2012.PRES.PERRY	0	-	0.1	63.1k	0.1	82.9k	0
Trade 2012.PRES.BACHMANN	0	-	0.1	71.6k	0.1	52.0k	0
Trade 2012.PRES.BARBOUR	0	-	0.1	67.2k	0.1	4574	0
Trade 2012.PRES.PAWLENTY	0	-	0.1	67.4k	0.1	21.3k	0
Trade 2012.PRES.OTHER	0	-	0.1	2345	0.1	160.7k	+0.0
		100.30	102.80				

Figure 4: Intrade's order book for the 2012 presidential election on August 13, 2012.

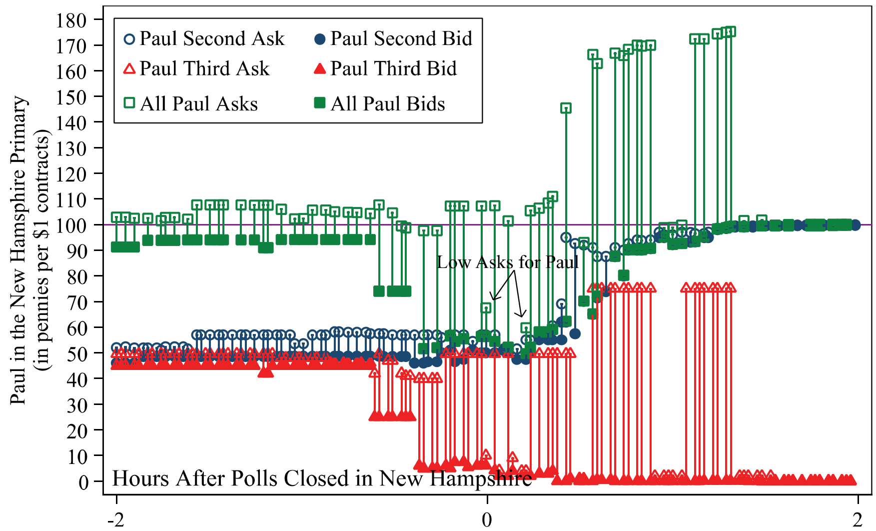

Figure 5: Intrade's bids and asks for Ron Paul for second and third place, separately, and Ron Paul in any position in the top three, combined, in the New Hampshire primary

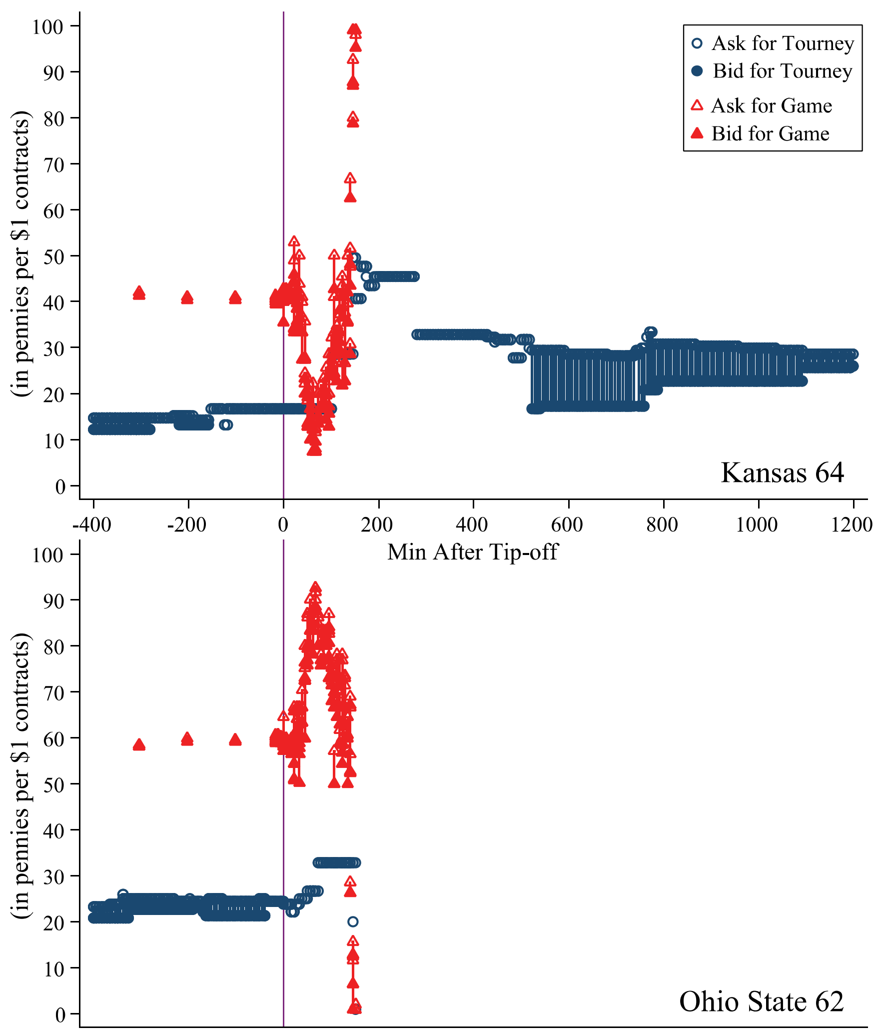

Figure 6: Kansas and Ohio State's contracts for winning their semifinal and final games during the 2012 NCAA championships

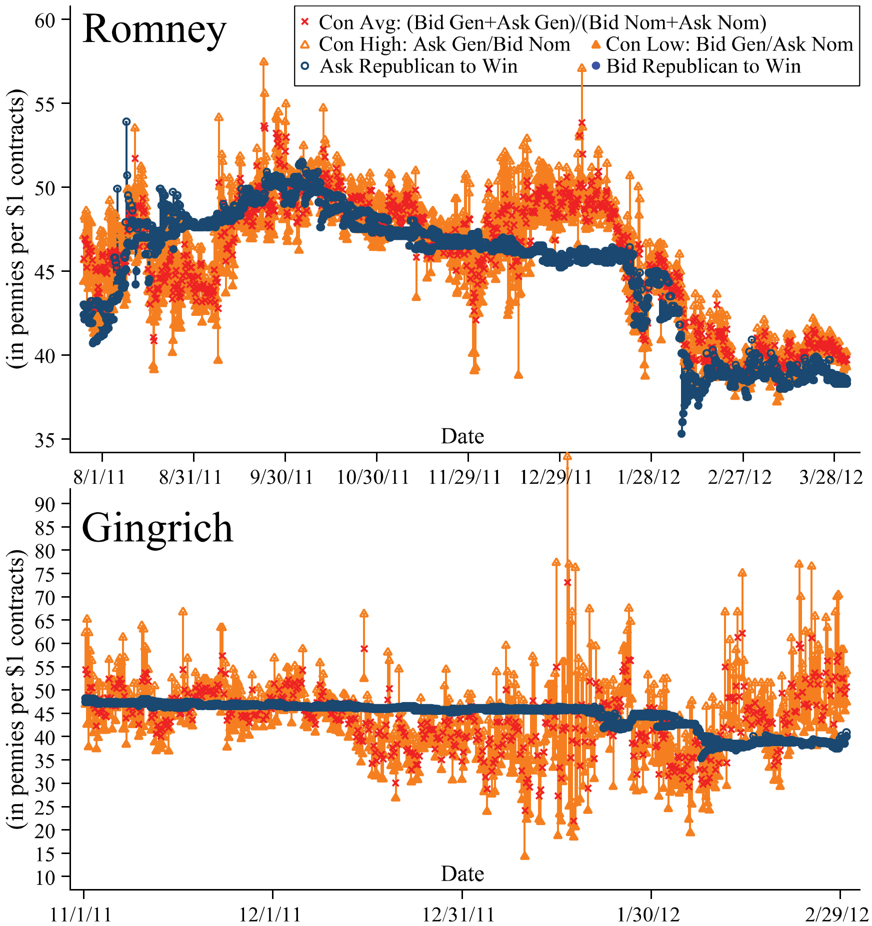

Figure 7: Imputed Intrade prices on Romney and Gingrich winning the presidency conditional on their nomination, along with the bid and ask for the Republican Party capturing the presidency.

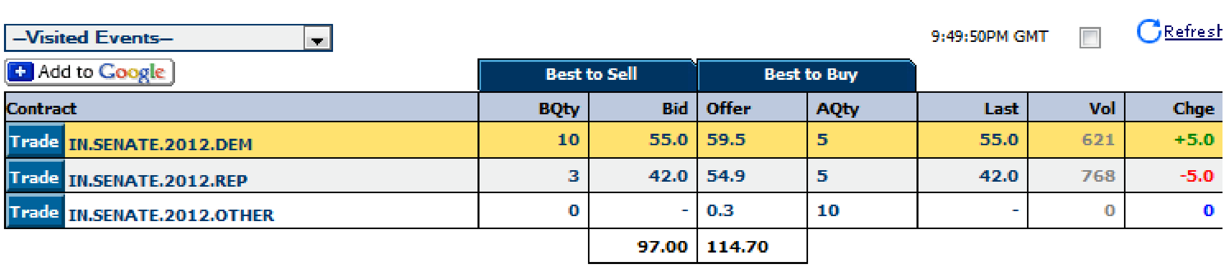

--Visited Events-- 9:49:50PM GMT Refresh

Add to Google

		Best to Sell		Best to Buy				
Contract		BQty	Bid	Offer	AQty	Last	Vol	Chge
Trade	IN.SENATE.2012.DEM	10	55.0	59.5	5	55.0	621	+5.0
Trade	IN.SENATE.2012.REP	3	42.0	54.9	5	42.0	768	-5.0
Trade	IN.SENATE.2012.OTHER	0	-	0.3	10	-	0	0
			97.00	114.70				

Figure 8: Intrade's marginal order book for the 2012 Indiana senate race at 5:49:50 PM ET on October 24, 2012

	Betfair	Intrade
Betfair, 1 lag price	0.422* (0.024)	0.078* (0.018)
Betfair, 2 lag price	0.288* (0.025)	0.061* (0.019)
Betfair, 3 lag price	0.077* (0.023)	-0.101* (0.075)
Intrade, 1 lag price	0.638* (0.033)	1.076* (0.025)
Intrade, 2 lag price	-0.084 (0.047)	0.014 (0.036)
Intrade, 3 lag price	-0.347* (0.035)	-0.129* (0.027)
Betfair, 1 lag difference	-0.494* (0.023)	0.063* (0.017)
Betfair, 2 lag difference	-0.137* (0.023)	0.112* (0.017)
Intrade, 1 lag difference	0.576* (0.033)	0.087* (0.025)
Intrade, 1 lag difference	0.441* (0.035)	0.111* (0.026)

Table 1: Lagged Betfair and Intrade prices and differences on Betfair and Intrade prices and differences

Notes: there are 2,175 observations and the lag is 12 hours. Prices are regressed on prices from 12, 24, and 36 hours ago. Differences are the difference between the current time period and 12 hours ago, regressed on the difference between 12 and 24 hours and 24 and 36 hours.