How to Forecast an Election: Polls or Prediction Markets?

Leighton Vaughan Williams
Nottingham Business School*

J. James Reade
University of Reading†

April 16, 2014

Abstract

In this paper we assess polls and prediction markets over a large number of US elections in order to determine which perform better in terms of forecasting outcomes. We consider accuracy, bias and decidedness over different time horizons before an election, and we conclude that prediction markets appear to outperform polls in terms of accuracy, unbiasedness and decidedness. We thus contribute to the growing literature comparing election forecasts of polls and prediction markets.

JEL Classification: C53, D83, D72.
Keywords: Forecasting Models, Information and Knowledge, Elections, Voting Behavior, Prediction Markets.

1 Introduction

There exist many sources of information one could use to forecast the outcome of an election *ex ante*; statistical models, expert opinion, opinion polls, and prediction markets are just four. Any such forecast is dependent on some set of information amassed at a particular point in time prior to the event happening, denoted $\mathcal{I}_t$, and also on the model through which that information is processed, $f_t(\mathcal{I}_t)$. In this paper we consider two of these potential models: *polls*, where information from potential voters is processed by polling companies and released, and *prediction markets*, where agents may also use potentially private information to buy and sell contracts contingent on a particular future event, thus revealing information in the process of doing so. As such, we are providing an additional perspective on the so-called Hayek hypothesis (Hayek, 1945; Smith, 1982) which suggests that markets can work efficiently even when participants have a limited knowledge of the environment or other participants (see also Hurley and McDonough, 1995).

*Nottingham Business School, Nottingham Trent University, Burton Street, Nottingham, United Kingdom, NG1 4BU, Email: leighton.vaughan-williams@ntu.ac.uk, Phone: +44 (0) 115 848 6150.

†Corresponding author. Department of Economics, HumSS Building, Whiteknights Campus, Reading, United Kingdom, RG6 6AA. Email: j.j.reade@reading.ac.uk, Phone: +44 (0) 118 378 5062.

In doing so, we build upon prior literature which identifies different types of prediction market, classified according to type of contract (Snowberg and Zitzewitz, 2005), and which have sought to examine the historical accuracy of election markets (Rhode and Strumpf, 2004) and to compare and/or relate the behavior and performance of these markets to that of opinion polls (e.g. Kou and Sobel, 2004; Leigh and Wolfers, 2006; Berg, Nelson and Rietz, 2008; Rothschild, 2009).

We make use of a vast, novel dataset to conduct a forecast comparison exercise between polls and a number of prediction markets. We assess the two groups based on the same criteria: Accuracy, bias, and precision of forecasts based on past performance. By accuracy we mean how often a forecast correctly predicts the election outcome, by bias whether the expected vote share or outcome probability is equal to the actual vote share or true probability, and by precision the variance of forecast errors. We find that prediction markets are more accurate although not necessarily less biased nor more precise.

In Section 2 we introduce our object of interest, the outcome of an election before we introduce in Section 3 our candidate forecast models and the datasets we have for each forecast model. Section 4 then discusses the methodology we use in assessing these forecast models and Section 5 outlines our results. Section 6 concludes.

2 The Actual Outcome

The outcomes of an election are manifold; more often than not in US elections, there are two candidates (a Republican and a Democrat), and the vote share each receives is one outcome of interest, as well as who actually wins each election.[1]

We think of the two-party vote share for candidate or party i in election j as $V_{i,j,T}$, where T is the date of the election, and we denote forecasts of that vote share made by forecaster f at date t, where $t < T$, as $\widehat{V}_{i,j,f,T|t}$. If we are considering only the two-party vote share, then the alternative outcome of interest is whether or not $V_{i,j,T} > 0.5$, as in this situation party i has won the election in terms of vote share, and hence we might think of a binary variable:

$$W_{i,j,T} = 1_{\left\{V_{i,j,T}=\max_k \{V_{k,j,T}\}_{k=1}^N\right\}}. \tag{1}$$

That is, $W_{i,j,T}$ is 1 if party i wins election j (in terms of vote share), zero otherwise. We define $W_{i,j,T}$ in (1) generally for an N-candidate election, yet often elections in the US involve just two candidates, and in that situation the probability that $V_{i,j,T} > 0.5$, i.e. the vote share on election day is sufficient to win the election popular vote, is what matters. The forecast made at time $t < T$ of whether or not a vote share $V_{i,j,T}$ will be sufficient to win an election we denote as $\widehat{W}_{i,j,f,T|t} = \widehat{\mathsf{P}}_t \left(V_{i,j,T} = \max_k \{V_{k,j,T}\}_{k=1}^N\right)$, or $\widehat{W}_{i,j,f,T|t} = \widehat{\mathsf{P}}_t \left(V_{i,j,T} > 0.5\right)$ in the case of a two-candidate election.

[1] This is particularly so for US Presidential Elections which are determined by the electoral college system and hence anomalies like the 2000 Bush vs. Gore election can happen where the winner was Bush even though Gore gained the larger (popular) vote share.

3 The Candidate Forecast Models

We consider four sources of pre-election forecasts in this paper:

1. Opinion polls as collated by *Real Clear Politics*, a website that collects historical and current polling information surrounding elections.[2]
2. Price data from Iowa Electronic Markets, an online prediction market for various political (and other) events.[3]
3. Price data from Betfair, an online betting company that offers markets which include political events or else which include election outcomes.[4]
4. Price data from Intrade, an online betting company that offers predominantly political markets or perhaps politically related markets.[5]

We consider each to be a forecast model; a mechanism that transforms information available at time t, $\mathcal{I}_t$ into a forecast for either a vote share $\widehat{V}_{i,j,f,T|t}$ or a probability of the election outcome $\widehat{W}_{i,j,f,T|t}$. In the next four subsections we describe each of these sources of data and comment on the mechanisms that generate forecasts from information available at time t.

3.1 Opinion Polls

Opinion polls are conducted by numerous companies in the US surrounding all sorts of elections and political questions (e.g. presidential approval). In the case of elections, polls are forecasts of vote shares conducted at some point $t < T$ by polling company f, hence they are denoted as $\widehat{V}_{i,j,f,T|t}$. Notionally, polls reflect public opinion regarding voting for particular candidates, and assuming the sample upon which they are based is representative, they can be seen as some reflection of voting intentions at time t, something which we denote $V_{i,j,t}$. As such, to treat a poll as a forecast of the eventual election outcome, we assume thus that such voting intentions do not change in the intervening time period. Hence there are at least two sources of error: the first is that the vote share forecast by the poll ($\widehat{V}_{i,j,f,T|t}$) may not be a true reflection of $V_{i,j,t}$; and/or $V_{i,j,t}$ may differ substantially from $V_{i,j,T}$ due, for example, to the learning process that takes place during an election campaign on the part of voters.

Furthermore, political candidates are very keen observers of polls and thus to some extent there may be endogeneity; candidates may respond to poll outcomes when $t < T$, increasing or decreasing effort levels. For example, a particularly disappointing set of polls may lead to a candidate increasing his or her effort in an election, which may thus impact $V_{i,j,T}$ causing it to differ from $V_{i,j,t}$. Furthermore, the success of campaign fundraising efforts may also be affected by poll outcomes (and potentially prediction markets also). As some polling companies are known or suspected to favor one political party or the other, it may also be that there is some strategic behavior on the part of pollsters in the timing and nature of their polls.

[2]See http://www.realclearpolitics.com/ for details.
[3]See http://tippie.uiowa.edu/iem/.
[4]See http://www.betfair.com.
[5]See http://www.intrade.com/.

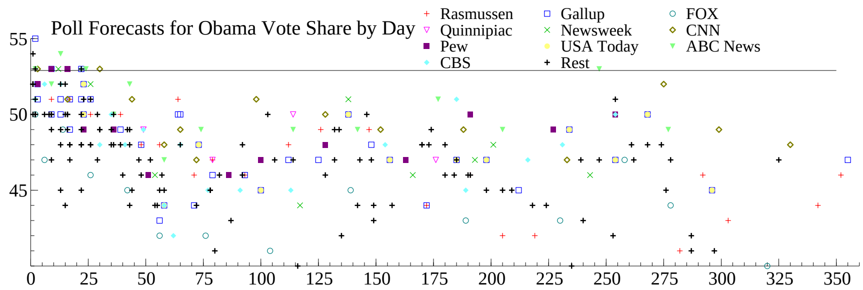

Figure 1: Plot of poll forecasts for Obama vote share in 2008 Presidential Election by day.

Nonetheless, considering the data at our disposal, we take polls to be forecasts of voting intentions *on election day*, T, as expressed at time t, and we analyse the extent to which they are effective forecasts of $V_{i,j,T}$. We take our data from *Real Clear Politics* (RCP) which compiles polling data from thousands of US elections over recent years. Table 5 summarizes the elections from which we have collected data from RCP; overall we have 19,277 observations from 394 different elections ranging from presidential elections in 2004 and 2008 both at the national and state levels, senate, governor and house elections and also Republican presidential candidate selection processes in 2008 and 2012, and the Democratic selection process from 2008.

We collect information on the polling company, the length of time the poll was conducted over, size and type of audience polled (likely voters or registered voters), forecast vote share for each candidate, and we also record the final outcome of each election.[6] There are averages of polls that are constructed by various groups, such as RCP themselves, and also others such as Nate Silver at the blog *FiveThirtyEight*. Although averaging can be a useful tool, particularly if the weights are appropriate (see, e.g., Bates and Granger, 1969; Graefe et al., 2012), it can only outperform the best individual forecast within the pool of forecasts being averaged in the presence of systematic bias (for example if one forecast is known to be positively biased and another negatively biased). Hence given this and our particular loss function in this paper, we focus on individual polls and forecast methods, seeking to understand better any bias that may exist.

Figure 1 plots poll outcomes for Obama's vote share during 2008 for the 2008 Presidential election; his final 52.9% vote share is denoted by the solid black line. The plot should be viewed from right to left, as the horizontal axis is the number of days remaining until the election takes place. Different polling companies are represented by different colors and symbols. Gallup is one of the most frequent pollsters and its polls (blue empty squares) appear to become more accurate as election day nears. All polls throughout the campaign appear from Figure 1 to underpredict Obama's eventual vote share, and even in the

[6] We have data on 446 individual poll producers, however many of the producers of polls are collaborations, such as Reuters and Zogby or Reuters and Ipsos. It is hard to get a precise number of the different forecasting companies involved because RCP often lists them abbreviated, but it appears there are around 200 distinct companies or organizations reflected in our dataset.

final few days the majority of polls announced suggest a vote share lower than what eventually results.

Gelman and King (1993) investigate the observed variability in polls despite the fact that election outcomes are particularly predictable at the outset of campaigns. They find that voters learn over the campaign which contributes to some extent to the variability of polls, meaning that early polls are less reliable relative to those conducted nearer to election date. Additionally their research suggests that poll forecasts should be dominated at all stages by expert opinion, statistical and other types of forecast models that embody some subset of that information. Gelman and King focus on Presidential elections, and the 1988 election in particular, although it is undoubtedly the case that many of their conclusions generalize. Nonetheless they do note that some of the effects they emphasize will likely be different for primary elections, and presumably for Senate, Governor and House races.

3.2 Prediction Markets

Prediction markets are markets in which participants buy and sell contracts in event (including election) outcomes. For example, if the market was the 2008 Presidential Election, the contracts would be for the Democratic candidate to win, or the Republican candidate to win. Prediction markets have attracted a great deal of attention in recent years from academic economists because, as Berg, Nelson and Rietz (2008) note, their primary role is as a forecasting tool rather than a resource allocation mechanism (although to some extent it can be said that they are part of the portfolio allocation problem of participants since there exists some a priori expected rate of return). Nonetheless, provided their design mechanism is effective, the prices produced will reflect expected probabilities of outcomes. Furthermore, in general the markets are short-term (the majority of those we consider in this paper last for considerably less than a year), and once the outcome is realized, the true value of the contract is known. This property enables researchers to consider whether or not prediction markets forecast events well.

Berg, Nelson and Rietz (2008) note the important differences between polls and prediction markets as forecasting devices. The former, at least in principle, are representative samples of the population (or deliberately selected sections of the population), whereas prediction markets are self-selected in that market participants must actively choose to take part. As a result, prediction market participants are anything but representative of the general population; as Berg *et al* point out, "traders are typically young, white, well educated and have high family incomes". Nonetheless, it is clear that this ought to be irrelevant for the accuracy of prediction market forecasts since the payoff structure means that market participants must put aside their own particular preferences over candidates and predict the voting behavior of the electorate at large if they are to make non-negative returns.

Prediction Markets (PMs) have been up and running for over 20 years at this point; with the first market having been established for the 1988 Bush-Dukakis contest (Rothschild, 2009). Berg et al. (2008), in their meta-analysis of the performance of PMs in elections in the USA and in other established democracies, found that, in terms of predicting the final result, "in the majority of (...) cases the market does about as well as the average poll, sometimes worse but often better, even if by a small margin" (p. 747); a finding that builds on a previous papers comparing PMs data to poll data. Erikson and Wlezien (2008) note that electoral markets have gained intellectual traction both in academic circles and in the popular press, with Surowiecki's (2004) *The Wisdom of Crowds* popularizing the idea that aggregated predictions of voting outcomes, which ask individuals to evaluate likely electoral outcomes can be 'better' (p. 35) than polls, which ask voters how they themselves will vote. Indeed, futures markets have been extended to predict non-electoral political phenomena, most controversially the likelihood of terrorist attacks (Wolfers and Zitzewitz, 2004).

Dissenting voices have questioned the alleged superiority of election markets to polls, and some of the most recent published research comparing polls to electoral markets has sought to 'discount' (Rothschild, 2009) or 'de-bias' (Erikson and Wlezien, 2008) poll data, in order to account for observed early poll margin overestimation and anti-incumbency biases in polling data. However, while Erikson and Wlezien found that de-biased poll data outperform national-level electoral market data for US Presidential elections between 1988 and 2004, especially in winner-takes-all predictions, Rothschild found that de-biased market-based data outperforms de-biased poll data in state-level forecasts in the 2008 US Presidential and Senatorial elections. Additionally, Lee and Moretti (2009) use a model of Bayesian learning to suggest that information passes from polls to PMs, while Sjöberg (2009) also challenges the notion of the 'wisdom of crowds' by looking at a range of different groups of forecasters for Swedish elections. Finally, similar comparisons of prediction market forecasts to more traditionally generated forecasts have been carried out in sports betting, looking at prices posted by bookmakers against prediction markets (Spann and Skiera, 2009; Croxson and Reade, 2011; Franck, Verbeek and Nüesch, 2011, see, for example). As such, the relative performance of poll versus electoral market data is still open to debate.

In the next three sections we introduce in turn the three prediction markets (*Iowa Electronic Markets*, *Intrade* and *Betfair*) we will examine for their performance in predicting election outcomes.

3.2.1 Iowa Electronic Markets

Iowa Electronic Markets (IEM) are not-for-profit operated prediction markets generally linked to political elections (but also markets have existed for box office movies and other one-off events), and have been running since 1988. On IEM participants are limited in their exposure in any trade to \$500. Markets have been set up for all major elections for over a decade, and in particular we have collected data on their prediction markets since 2000. Their markets tend to have two forms, either a winner-takes-all (WTA) or a vote-share (VS) format. The former corresponds to forecasting the election outcome and hence would be described as $\widehat{W}_{i,j,f,T|t}$, while the latter corresponds to providing forecasts of the form $\widehat{V}_{i,j,f,T|t}$.

We have 45,590 observations covering 38 elections; those elections are Presidential (2000, 2004 and 2008), Congressional (House and Senate and a joint market, 2000–2010), and a somewhat ad hoc collection of mayoral elections and primary elections (alongside Democratic and Republican Conventions since 2000). See Table 6 for more details. For each market and contract IEM makes available on a daily basis the number of trades (units and dollar volume), the

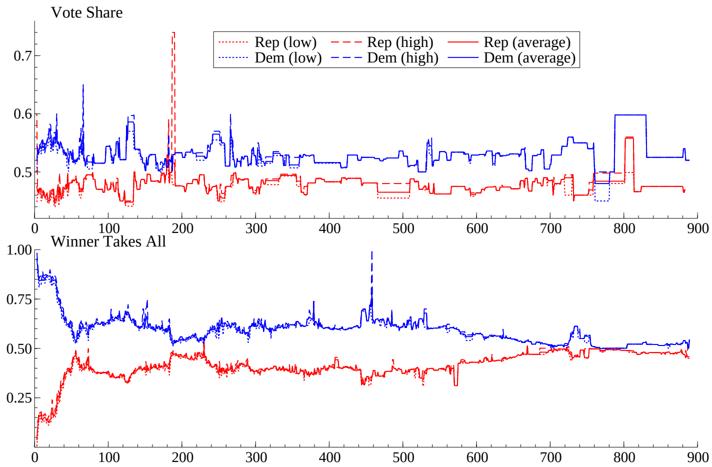

Figure 2: Vote share (top panel) and winner takes all (bottom panel) markets for 2008 US Presidential Election on IEM.

highest and lowest prices traded at, and the average price. An important distinction between IEM markets for House and Senate elections is that the contracts bought and sold are for macro outcomes: Either the Democrats or the Republicans have a majority in the House or Senate as a result of the election. Similarly, the Republican and Democratic Convention markets allow the trading in contracts about the eventual outcome rather than individual primaries. As Table 6 shows, there are a couple of exceptions (e.g. New York Senate), but generally IEM does not provide markets for individual elections outside Presidential elections.

Figure 2 presents prices from IEM for the 2008 Presidential election; on the top panel the VS market prices are plotted (high, low and average), while on the bottom panel the WTA prices are plotted. These two graphs visualize the difference between the vote share type of forecasts that polls constitute, $\widehat{V}_{i,j,f,T|t}$, and the probability of outcome, $\widehat{W}_{i,j,f,T|t}$, that prediction markets usually provide. Viewing from right-to-left, as the election day draws near, although the vote shares forecast don't diverge particularly strongly (top panel), the probability of each outcome does diverge substantially, and in the final days of the election the probability of a Democratic victory is around 85% and above.

Berg, Nelson and Rietz (2008) compare IEM to polls for Presidential elections back to 1988, and find that IEM outperforms the polls in head-to-head comparisons. In relation to their study, we consider a much broader selection of recent elections of all types for both polls and IEM (again see Tables 6 and 5), a strategy that affords us a larger dataset of more recent polls. Berg, Nelson and Rietz (2008) note how the demographic of market participants on IEM has changed over the years since 1988, and hence by considering elections only after 2000 we expect to have data more representative of the current IEM demographic.

Erikson and Wlezien (2009) on the other hand contends that polls are actually more informative than prediction markets making use of novel data on informal prediction markets for presidential elections going back to 1880, and

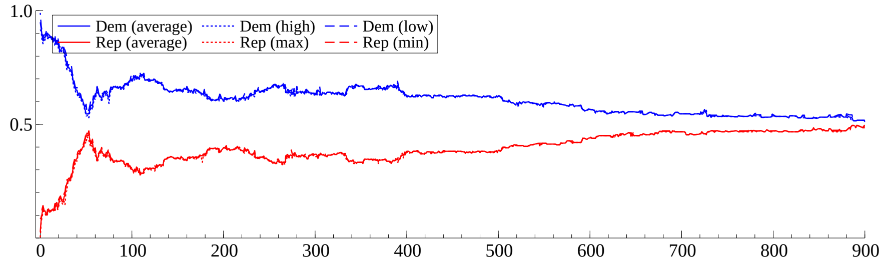

Figure 3: Plot of implied probability of each party winning 2008 Presidential Election from Betfair.

using multivariate methods. Although we cannot match Erikson and Wlezien (2009) for sample length, we have substantially broader depth in that we consider here various types of election other than Presidential elections giving us the sample size we mentioned above. Additionally, we have multiple observations per election whereas Erikson and Wlezien (2009) only use one observation, taken immediately prior to each election. Given King and Gelson's findings regarding the accuracy of polls as election date nears, it seems likely that this is a favorable comparison for polls; we will be able to shed light on this using our dataset. Erikson and Wlezien (2008) do consider polls with longer time horizons until the election when considering Presidential elections between 1988 and 2004, and conduct an empirical bias correction for these polls.

3.2.2 Betfair

Betfair is an online betting company providing markets primarily in sports events but also increasingly in political events such as elections. In the jargon, participants either *back* or *lay* bets on events, equivalent to buying or selling contracts paying out contingent on that event happening, such as a politician to win an election. Betfair operates a limit order book, as it matches participants willing to buy and sell contracts at particular prices. In contrast to IEM, participants are not restricted in their potential exposure on Betfair to any arbitrarily imposed limit, and Betfair is a for-profit company; it seems likely that this would influence the self-selection that takes place for potential market participants. When applied to our context of election outcomes with two parties, Betfair yields observations corresponding to $\widehat{W}_{i,j,f,T|t}$.

As our objective is to consider what publicly available information could be used to best forecast an election, although our data is very rich, only certain aspects of it are relevant. Market participants using Betfair can see what prices are available to buy and sell contracts in an event, and how much money (liquidity) is available at each price (buy or sell).

Figure 3 shows the evolution of the implied probabilities (reciprocal of the market prices for contracts) for each party to win the 2008 Presidential election over the 900 days prior to the election. In Figure 3 the maximum and minimum prices for a given day, as well as the average price, are plotted, but as these are very similar to each other it is almost impossible to distinguish them in the plot. This plot can be compared to the bottom plot in Figure 2 which shows the same price evolution for IEM. As with IEM, the Democrats are always the favorites throughout the 900 days shown, and by a slightly larger margin consistently than IEM, with a similar pattern of divergence in probabilities in the final 50 days of the campaign. These plots suggest that, as with polls highlighted by Gelman and King, also with prediction markets learning takes place and the nearer an election is, the more decided becomes the market on the most likely outcome.

3.2.3 InTrade

Intrade is a prediction market specializing in US political elections.[7] There are no limits on the amount that individuals can trade, as opposed to IEM, and the format is essentially identical: market participants trade contracts whose payout is contingent on some event occurring.[8] As such, when thinking specifically about election outcomes with two parties, our Intrade data provides us with observations corresponding to $\widehat{W}_{i,j,f,T|t}$.

We have data from the 2004 and 2008 US Presidential Election; for both years we have all individual state voting and for 2008 we have a range of additional politically related markets.

For the 2004 elections, we have daily data consisting of the high, low and closing prices, while in 2008 we have data on individual trades carried out on the exchange. The 2008 data provides information on whether contracts were bought or sold, the price at which the trade took place and the quantity, alongside a timestamp of when the trade took place. We have 29,196 observations from the 2004 Presidential election (although all of these relate to individual state markets rather than the overall outcome), and 411,858 from the 2008 elections (although not all of these relate specifically to elections — see Table 8 for a breakdown).

The purpose of this study is to find the best forecast method, and since Intrade reports on its website very visibly the price of the last agreed trade, clearly more information exists for 2008 for us to assess the Intrade predictions, but nonetheless the information from 2004 does provide additional information.

Figure 4 shows the Intrade implied probabilities (prices divided by 100) for the same 900 days prior to the 2008 Presidential election as in Figures 2 and 3. As is perhaps clear, the two parties appear a little closer as measured by Intrade; at two years prior to the election, the two are absolutely even, and even with just 50 days to go before the election, the two implied probabilities overlap for a short period. It is quite likely that this overlap with just 50 days remaining was due to market manipulation; one trader apparently traded so as to raise the price on Intrade for the Republican candidate, John McCain.[9] As discussed by Hanson and Oprea (2009), we do not see this as necessarily

[7] Indeed, the perception has long been that Intrade provides for US elections while Betfair does so for UK elections; see, for example, http://www.midasoracle.org/2007/04/24/betfair-vs-tradesports-intrade/ (last accessed April 17 2012).

[8] Servan-Schreiber et al. (2004) compare Intrade to News Futures, a prediction market based on 'play money', using a 'game' format, to ascertain whether "money matters". They find that money doesn't appear to improve the forecast performance of prediction markets. Our analysis, comparing Intrade and Betfair to IEM will shed some light on this question since IEM restricts the amount of money participants are able to bet.

[9] See http://marginalrevolution.com/marginalrevolution/2008/10/manipulation-of.html for more information on this.

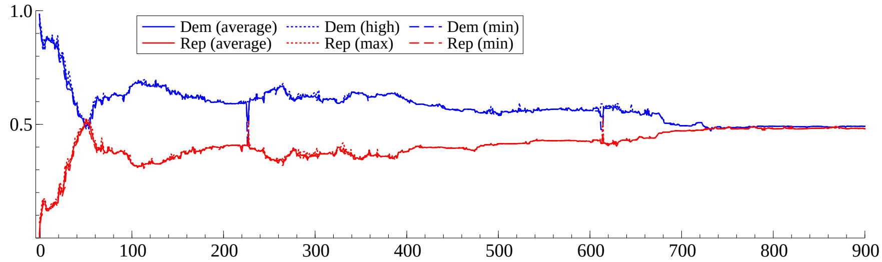

Figure 4: Plot of implied probability of each party winning 2008 Presidential Election from Intrade.

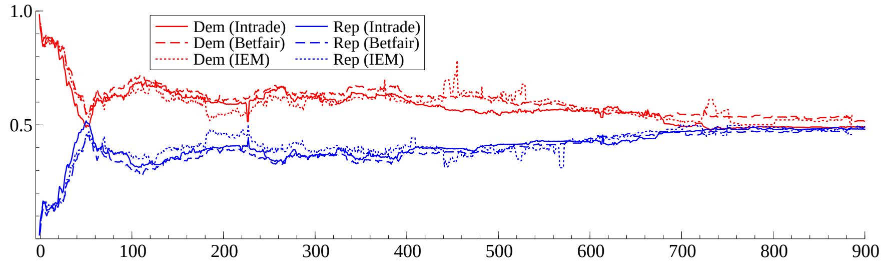

Figure 5: Plot of comparable prices from all three prediction markets from the 2008 Presidential Election.

a problem for our analysis; a manipulator in a liquid market might be viewed as offering other traders a kind of free lunch in correcting that manipulator's attempts to distort.

As with both previous prediction markets, Intrade has also attracted academic interest; Gil and Levitt (2007) investigated market efficiency looking at the 2002 FIFA World Cup, while Hartzmark and Solomon (2008) considered the disposition effect using NFL markets. Snowberg, Wolfers and Zitzewitz (2007) use TradeSports, swallowed up by Intrade in 2008, to infer implications from elections onto the macroeconomy by using the 2004 US Presidential election when unreliable exit polls caused substantial price variation within a single day.

We can now compare all three markets on one plot, in Figure 5, and over a long period of time the co-movement between these series is very clear. A comparison between these three prediction markets is of great interest, not least because the self-selection that takes place in each market will likely be different; Betfair does not allow those based in the US to trade in their markets, while IEM does operate in the US but restricts its participants in their exposure, while Intrade allows those based in the US to trade in its markets but does not restrict the exposure of participants. Hence it is of interest to compare these three markets in their ability to forecast elections; do these differences matter? Furthermore can any of them, as Berg, Nelson and Rietz (2008) assert, improve upon polls?

4 Methodology

We seek metrics to assess each candidate forecast model. Such metrics should be impartial between the different forecast models and hence give us an objective outcome regarding the best forecast model. An immediate obstacle in this pursuit is that we have two types of forecasts; both those of vote share, $\widehat{V}_{i,j,f,T|t}$, and probabilistic forecasts, $\widehat{W}_{i,j,f,T|t}$; with the former the outcome is continuous over the unit interval, whereas for the latter the outcome is a binary variable. Page (2008) considers how to compare such different types of information and concludes that one requires the probability distribution of vote shares, or some approximation or estimation thereof, to make a comparison. In this paper we transform polled outcomes to probabilities using non-parametric estimation. We have hundreds of thousands of polled outcomes and hence we can use these to form a probability distribution, controlling for systematic factors which affect poll accuracy.

As with other attempts to compare forecast methods such as Erikson and Wlezien (2009), if we rely on direct comparison forecast by each forecaster for particular events, we will be severely restricted in our number of elections and hence observations relative to the total datasets we have at our disposal. Instead we assess each forecasting method over all the elections we are able to collect data on *for that method* (see Tables 5–8). Thus we attempt to establish for each forecaster, independent of the others, how well it forecasts election outcomes, before comparing these performances between forecasters. There is considerable overlap between our datasets for each forecaster such that we are considering forecast performance over very similar datasets.[10]

Any forecast assessment is reliant on the loss function assumed; what loss do we suffer if the forecast is wrong in a particular direction? With elections and vote shares, such a loss function is unlikely symmetric since if a forecast is for 51%, then if the outcome is ±2%, thus 49% or 53%, it matters which way — up and the election outcome (in terms of vote share) is unaltered, down and the outcome changes.

Our objective is to pick the winner in a forthcoming election, and hence a very simple metric for forecasts is accuracy: how often does the forecast outcome occur? Hence whether the forecast is for a vote share of 51% or 65% is somewhat irrelevant provided that that event happens. However, it is likely that election outcomes that are nearer to 50% (for a two-party election) will induce lower success rates. Bearing this in mind, and given that often US Presidential elections are very close, we also move to consider forecasts more generally. In this sense, it must be the case that a good forecast is both unbiased, displaying no systematic biases, and precise, and we will outline how we test for this in Sections 4.1 and 4.2.

One of the most important factors determining the accuracy of forecasts is the time horizon, denoted $h = T - t$. As a general rule, forecast accuracy declines with h (Hendry and Clements, 1998), and Gelman and King note this in the context of Presidential election polls. Although we have intra-daily data from Betfair and Intrade, because only a small number of polls are released on

[10]While two or three of our models will have overlapping observations for many elections, the elections for which we have comparable data for all four models is restricted to essentially the 2008 Presidential election (for example, for Intrade we have only state-level 2004 Presidential election data but for Betfair and IEM we have only national level market data).

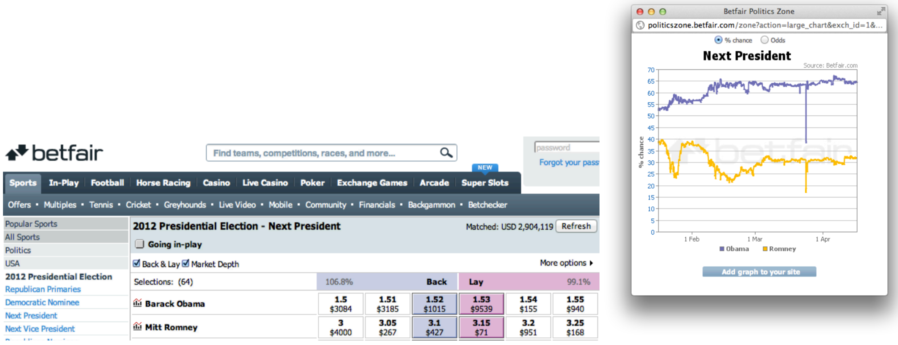

Figure 6: Interface of Betfair (left panel) when attempting to bet on 2012 US Presidential Election outcome, and graph available to Betfair market participant of recent price movements (right panel).

any given day and because IEM data is only available at the daily frequency, we instead consider weekly time intervals when assessing our forecast models, and specifically, weeks until the election takes place. While we could consider information from each day, this would likely lead to rather erratic results due to the lower frequency of polls and also trades as the time interval to an election increases.[11] A weekly frequency also allows us to more easily comprehend general movements and trends between our forecast models over the entire year before an election takes place. Hence from hereon, when a time dimension is mentioned, we are referring to weeks before the election.

Our objective of considering what forecast is best from our candidate forecast models also restricts interest in our datasets to what would be available from our prediction markets for forecasting the election. On Betfair, the immediate interface reveals the three most popular prices for buying or selling each contract (see the left panel of Figure 6), although one can also see a graph of recent price movements, and an example is shown in the right panel of Figure 6 for the 2012 election race. Detail on previous trades matched are not readily available other than their implications on the market price and hence for the purposes of this paper we disregard our data on quantity to focus solely on price.

For IEM, market activity up to the previous day is available on the website even to those who do not participate in the market, as is a graph of all recent price movements, and hence the data we have in our dataset is data that could be used to construct a forecast of the election.

On Intrade, information on the best priced contract and most recently traded contract, both bought and sold, is made available to all visitors to the website. Figure 7 shows the initial interface that greets a willing participant on Intrade if they wish to buy or sell a contract on the 2012 Presidential election outcome; the best available price to buy and sell each contract is provided, as is a graph of recent price movements, if desired. If the participant clicks on one of the candidates, the page on which they buy or sell also reveals the price of the last contract bought or sold, additional potentially important information to be

[11] As an example of this, IEM's 2012 presidential election market has recorded a number of days recently with no trades taking place when there are 28 weeks remaining until the election.

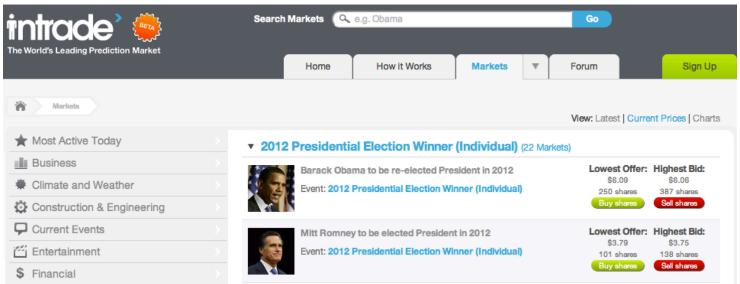

Figure 7: Interface of Intrade when attempting to bet on 2012 US Presidential Election outcome.

used in forecasting. Hence it is important that we make use of our dataset that contains information on every trade that has taken place in that market, although again we can disregard for the purposes of forecasting future elections the quantity data we have.

4.1 Accuracy and Unbiasedness

A very simple measure of accuracy is the percentage of correct forecasts. This is the most direct measure of what minimises our loss function probabilistically; the forecast method that forecasts correctly most often must on average yield the lowest loss. In the case of markets that provide $\widehat{W}_{i,j,f,T|t}$ forecasts, we take a forecast to be predicting a particular outcome if that particular forecast probability is the highest of the candidates in an election. Hence we take, for $i = 1, \ldots, N$ contestants in an election $\widehat{W}^*_{i,j,f,T|t} = \max_k \widehat{\mathsf{P}}_t \left(V_{k,j,T} > 0.5\right)$, the candidate or party with the highest forecast probability of winning, as the forecast outcome at that point. For forecasts of the nature $\widehat{V}_{i,j,f,T|t}$, we take $\widehat{V}^*_{i,j,f,T|t} = \max_k \widehat{V}_{k,j,T|t}$, the maximum vote share, as the favorite and hence predicted outcome. We also denote $V^*_{i,j,T|t} = \max_k V_{k,j,T|t}$ as the candidate with the highest vote share and hence the winner of the popular vote in an election.[12]

We thus calculate, for forecast model f, the percentage of correct forecasts as:

$$\%_f = \frac{\sum_{i=1}^{N_i} \sum_{j=1}^{N_j} \sum_{t=1}^{N_t} 1_{\left\{V^*_{i,j,T} = \widehat{V}^*_{i,j,f,T|t}\right\}}}{N_i N_j N_t}, \tag{2}$$

where N_j is the number of elections considered, N_i the number of candidates and N_t the number of time periods. We compare forecasts along this dimension to assess the *accuracy* of polls.[13]

A related but distinctly different concept to accuracy is that of *unbiasedness*. An *unbiased* forecast can be defined separately for each kind of forecast:

- An unbiased vote share forecast is, on average, equal to the true vote share outcome: $\mathsf{E}\left(\widehat{V}_{i,j,f,T|t}\right) = V_{i,j,T}$.

[12] Which, as noted earlier, need not correspond to the actual winner of the election.

[13] Note that we could be more demanding with our measure of accuracy in (2) and require that forecast models got the final ranking of candidates correct; all we require is that the forecast model correctly identifies the favorite.

- An unbiased probability forecast is, on average, equal to the true probability that that candidate wins the election: $\mathsf{E}\left(\widehat{W}_{i,j,f,T|t}\right) = W_{i,j,T}$.

Hence forecasts that are accurate *can be also biased*, provided the bias is in the correct direction; if polls are consistently upward biased for candidates that eventually win, then despite being biased they will be very accurate in predicting the outcome, whereas polls that are consistently downward biased for candidates that eventually win will be very inaccurate as well as biased.

When we consider vote share forecasts for candidate i in election j, $\widehat{V}_{i,j,f,T|t}$, after an election has happened we observe the true $V_{i,j,T}$ and hence we can evaluate the *forecast error*:

$$\widehat{e}^{V}_{i,j,f,T|t} = V_{i,j,T} - \widehat{V}_{i,j,f,T|t}\,. \tag{3}$$

We can use this forecast error to consider the possibility of biased forecasts. In taking the simple average of (3) we thus learn whether or not a forecast method is unbiased or not. Hence we calculate:

$$MFE_t = \sum_{i=0}^{N_{f,t}} \hat{e}^{V}_{i,j,f,T|t}\,, \tag{4}$$

where $N_{f,t}$ denotes the number of forecasts we have for each forecast method at each time period t.

Although we could substitute W for V in (3) and (4) when we observe forecasts that are probabilities of outcomes, it is likely that because the outcome is binary that a summation such as in (4) would unfairly penalize probabilistic forecasts that are above 50% but not by particularly much. As such we seek an alternative approach, and one method, often referred to as *calibration testing*, is to regress the outcome on the probability as produced by the forecast method:[14]

$$W_{i,j,T} = \alpha_W + \beta_W \widehat{W}_{i,j,f,T|t} + \varepsilon^{W}_{i,j,t}. \tag{5}$$

The assumption we place on $\varepsilon_{i,j,T}$ determines the kind of regression model we employ; although it can be shown that estimating (5) via OLS induces heteroskedasticity, it is most convenient for our analysis to estimate using OLS since that implies an iid (independent and identical distribution) assumption for $\varepsilon^{W}_{i,j,t}$ of $\varepsilon^{W}_{i,j,t} \sim (0, \sigma^2_W)$; we will later make use of this.

In (5), if $\alpha_V = 0$ and $\beta_V = 1$ the forecast method is said to be *unbiased* since then $\mathsf{E}\left(\widehat{W}_{i,j,f,T|t}\right) = W_{i,j,T}$ (because $\mathsf{E}\left(\varepsilon^{W}_{i,j,t}\right) = 0$). Hence the F-test of the null hypothesis $\alpha_V = 0$ and $\beta_V = 1$ is our test of *unbiasedness.* If the constant term $\alpha_V \neq 0$ then the forecast method does exhibit some systematic bias, while if $\beta_V \neq 1$ then if the true probability of an event changes, the forecast method either over- or under-adjusts. This phenomenon is commonly referred to in the betting literature as the *favorite longshot bias* (FLB). The conventional FLB ($\beta_W > 1$) is where bettors relatively over-bet event outcomes with lower implied probabilities of winning (inferred from the odds) and relatively (though not necessarily absolutely) under-bet event outcomes with higher implied probabilities of winning. The reverse FLB ($\beta_W < 1$) occurs where bettors relatively

[14]Note we write $W_{i,j,T}$ here, adding a t to the outcome; the observed outcome does not change through time, we just add this in order that we can run regressions for different forecasts at different time ts before the election occurs.

over-bet and under-bet the converse. Returning to our comparison of bias and accuracy earlier, it ought to be that FLB aids forecast accuracy; a favorite is probabilistically more likely to win and hence if in a market the favorite wins more often than its price (or vote share) implies, that market must predict the correct outcome more often.

We seek a method to assess forecasts that is unitless due to the two different types of forecast in our dataset, and hence we employ the same method outlined in (5) when considering forecasts from polls. Just as $\alpha_V = 0$ and $\beta_V = 1$ implies $\mathsf{E}\left(\widehat{W}_{i,j,f,T|t}\right) = W_{i,j,T}$, and hence that the fitted line through the scatter plot of forecasts against outcomes corresponds to the 45 degree line, we can apply the same methodology to polls; does a polled vote share of, say, 47% imply that on average the resulting outcome is 47%? Hence we run the regression of:

$$V_{i,j,T} = \alpha_V + \beta_V \widehat{V}_{i,j,f,T|t} + \varepsilon^V_{i,j,t}. \tag{6}$$

Equivalently to above, $\alpha_V = 0$ and $\beta_V = 1$ imply that on average polled levels equal actual outcomes and hence the forecast model is unbiased: $\mathsf{E}\left(\widehat{V}_{i,j,f,T|t}\right) = V_{i,j,T}$. The value of this method in comparing our two types of forecast is that for probabilistic forecasts ($\widehat{W}_{i,j,f,T|t}$) we compare to the expected value of, $\mathsf{E}\left(W_{i,j,T}\right)$ rather than the binary variable itself, $W_{i,j,T}$. This reduces a potential distortion when comparing forecast errors from vote shares and probablistic forecasts.

Additionally, if $\alpha_V = 0$ and $\beta_V = 1$ are imposed then $\widehat{e}^V_{i,j,f,T|t} = \widehat{\varepsilon}^V_{i,j,f,T|t}$, our regression model (6) becomes equivalent to the forecast error from (3) earlier for vote shares, and hence we can think about (6) as a generalized forecast error. By running the regression in (6) we learn about the actual relationship between $\mathsf{E}\left(\widehat{V}_{i,j,f,T|t}\right)$ and $V_{i,j,T}$ rather than asserting that the two are equal. Similarly as with (5), if $\beta_V > 1$ we have FLB: the favorite on average gets a higher vote share than the outsider.

Thus in both regression models, (5) and (6), the null hypothesis of $\alpha_g = 0$ and $\beta_g = 1$, $g \in \{V, Y\}$, implies that the forecast method is unbiased — on average it forecasts without error. Although a visual examination of the estimated α and β coefficients will be informative, it is also useful to construct a *direct test of unbiasedness*, and hence we use an F test of the hypothesis that $\alpha_g = 0$ and $\beta_g = 1$ to evaluate the unbiasedness of our forecast methods. Because in both types of forecast the F-test measures departures from unbiasedness (expected values), it should not be influenced by the distinction between $V_{i,j,T}$ being continuous on the unit interval and $W_{i,j,T}$ being binary.

As a final aside on bias, it is worth noting that Erikson and Wlezien and other investigators often de-bias forecasts from polls and prediction markets. We refrain from doing so, preferring instead to compare raw data, noting the biases present in the raw data as we compare forecast models.

4.2 Precision

Having considered unbiasedness, it is now helpful to move on to thinking about precision — how precise are the forecasts we get? A conventional measure of the precision of a forecast is the *mean squared forecast error* (MSFE) — squaring the forecast errors we calculated in (3) and summing:

$$MSFE_g = \sum_{i=0}^{N_{f,t}} \left(\hat{\varepsilon}^g_{T|t} \right)^2, \quad g \in \{V, Y\}, \tag{7}$$

This is an approximation to the variance of the forecast, centred around the outcome, and hence is equivalent to the estimated standard error for our regression model, (6), denoted $\widehat{\sigma}^2_f$ since the formula for that is:

$$\widehat{\sigma}^2_g = \frac{1}{N_{g,t}} \sum_{i=1}^{N_{g,t}} (\widehat{\varepsilon}^g_{i,j,t})^2 = \frac{1}{N_{f,t}} \sum_{i=1}^{N_{g,t}} \left(V_{i,j,T} - \widehat{\alpha}_g - \widehat{\beta}_g \widehat{V}_{i,j,f,T|t} \right)^2 = \sum_{i=0}^{N_g} \left(\hat{e}^g_{T|t} \right)^2 = MSFE_g, \quad g \in \{V, Y\}, \tag{8}$$

provided $\widehat{\alpha}_Y = 0$ and $\widehat{\beta}_Y = 1$. Thus $\widehat{\sigma}^2_g$ is a more general measure of forecast accuracy than MSFE which imposes restrictions on (6).

In essence, $\widehat{\sigma}^2_W$ measures how imprecise a prediction market is at providing probabilistic forecasts, while $\widehat{\sigma}^2_V$ measures how imprecise at providing vote-share forecasts a poll is. However, these two $\widehat{\sigma}^2$ measures consider the precision around the *actual* relationship between forecasts and outcomes, rather than the 45-degree line (which $\alpha_g = 0$ and $\beta_g = 1$ would imply). The equivalent MSFE measures impose the $\alpha_g = 0$ and $\beta_g = 1$ restrictions without testing their appropriateness but nonetheless do provide important information — how dispersed around the 45-degree line are the forecasts.

Hence we assess forecast errors and the F test of $\alpha = 0$ and $\beta = 1$ to assess *forecast unbiasedness* and analyze MSFEs and $\widehat{\sigma}^2$ assess *forecast precision.*

5 Results

We now consider the accuracy, bias and precision of each of our forecasting models. We first consider *accuracy* via the percentage of correct forecasts (2), presenting the results graphically for each market then assessing the markets head-to-head, before considering via regression methods the accuracy and precision of the markets.

5.1 Accuracy

Figure 8 reports the overall percentage of polls that correctly forecast the actual outcome by weeks before the election was due to take place, and the bars represent the number of polls that fall into each category.[15] We then refine by particular types of election.

The overall percentage of forecasts which are correct drawn from polls is 71.0%, increasing to 76.8% if only the Presidential elections of 2004 and 2008 are considered. In Figure 8 we chart the performance of polls as the distance to election, and hence the forecast horizon, increases. There appears to be no particular improvement in poll performance as an election nears, something

[15] Where we point out that the 'winner' in terms of vote share is taken to be the candidate that won the most votes, hence for example in the 2000 Presidential election, Gore is classed as the winner as he won more of the popular vote.

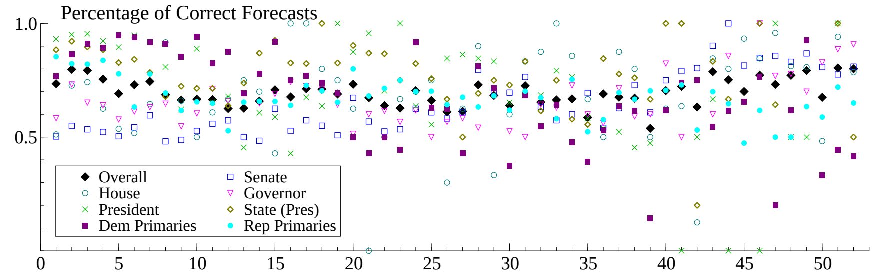

Figure 8: Percentage of polls that correctly predict election outcome by weeks until election.

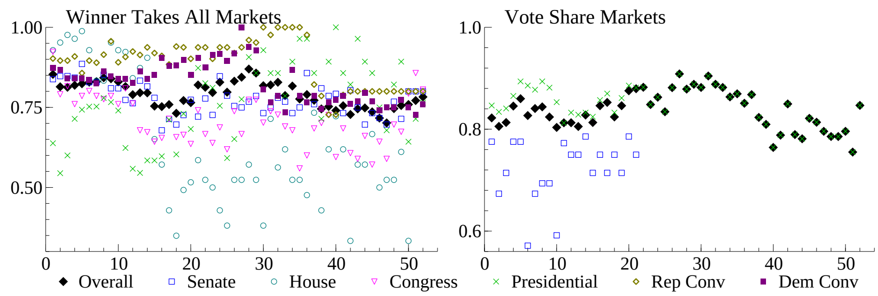

Figure 9: Percentage of IEM prices that correctly predict election outcome by weeks until election.

which contrasts with Gelman and King's findings. Even for presidential elections, for which performance does appear to peak in the 4–5 weeks before an election, performance is actually comparable if not better between weeks 19 and 23 where just under 100 polls record a success percentage of slightly over 90%. The black diamonds in Figure 8 show the overall performance of polls for all elections we consider, and this does improve slightly from a low of just above 60% with 13 weeks remaining to around 80% with two weeks remaining, but this performance is not significantly better than polling performances 30–40 weeks before an election.[16]

Turning to IEM, we split forecasts into vote share (VS) and winner-takes-all (WTA) markets. The percentage of forecasts that were correct is 83.5% (VS) and 79.8% (WTA), changing to 85.2% and 72.3% respectively for presidential elections. Figure 9 presents a graphical breakdown of how these percentages move as the time until election increases with WTA on the left panel, VS on the right. What is perhaps most notable here is that IEM's WTA Presidential elections forecasts seem significantly worse than all of its other election forecasts in the final 20 weeks before the election takes place — before that, its success ratio is considerably higher the IEM average. Intriguingly, the VS markets show Presidential election forecasts better than average, although there is little other than Presidential markets in this category of market (85% of our observations are from the 2000, 2004 and 2008 Presidential elections), but in absolute terms

[16]In Figure 12 we plot standard error bands helping to make such a comparison.

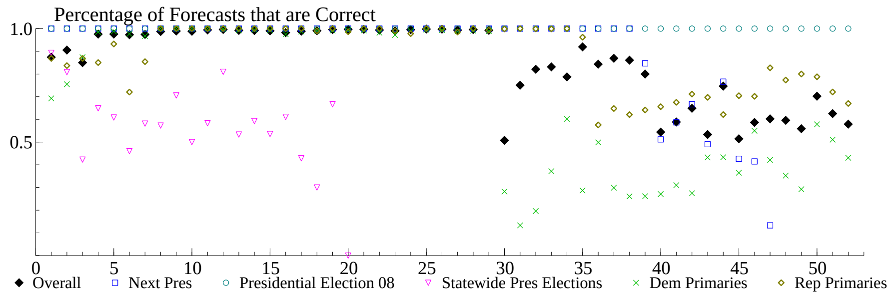

Figure 10: Percentage of Betfair prices that correctly predict election outcome by weeks until election.

the percentage is higher, remaining just shy of 85% up to the final week of the election, as opposed to the WTA percentage of between 55% and 65%.

With Betfair, we display the percentage of correct forecasts in Figure 10 weekly for a year in advance of each election. Betfair's overall percentage of correct forecasts is 85.5% and 99.8% for the 2008 Presidential election (falling to 95.3% for the Next President market), while we present the breakdown by weeks before an election occurs in Figure 10. Aside from this almost perfect record in forecasting the 2008 election, additionally for both Republican and Democratic primaries, Betfair has a success percentage of 87.1% in the final 30 weeks of campaigns. Betfair forecasts of statewide elections for the electoral collage (upside down triangles) improve from essentially zero 20 weeks before election day to 90% in the final week.

Figure 11 shows Intrade percentages in the same format as for the previous three candidate forecast models. Intrade has a percentage of forecasts turning out correct of 84.0%, rising to 88.1% for the 2008 Presidential election. As with Betfair, we see a high level of correct forecasts, particularly for the 2008 Presidential election where again up to around week 38, Intrade prices imply a correct forecast almost every trade. From Table 8 we have a large collection of markets from Intrade related to US Presidential elections in 2008 other than simply the outcome or vote share, and the purple dots in Figure 11 represent these; as can be seen, the prediction record on these more eclectic events (e.g. whether a particular video will be released by the LA Times by a particular date) is dramatically worse than for US elections, as even in the week before the election takes place the percentage of forecasts that are correct is only around 50%. However as these minor markets make up a small fraction of our total observations, their impact on the overall percentage is minimal; it is lower percentages for statewide Presidential elections, Governor markets and other Presidential-related markets that pulls the overall percentage down.

It is perhaps more informative to compare our four candidate forecasts with each other directly on a plot, and Figure 12 does that, plotting the four as different series over the year up to elections. We additionally include 95% significance-level standard error bounds around each market's plot, enabling us to assess whether differences are significant.[17] The plot indicates that the

[17]The distinctly differently sized confidence bands is more a function of sample size rather than any inherent uncertainty in particular models. This is because we only have one obser-

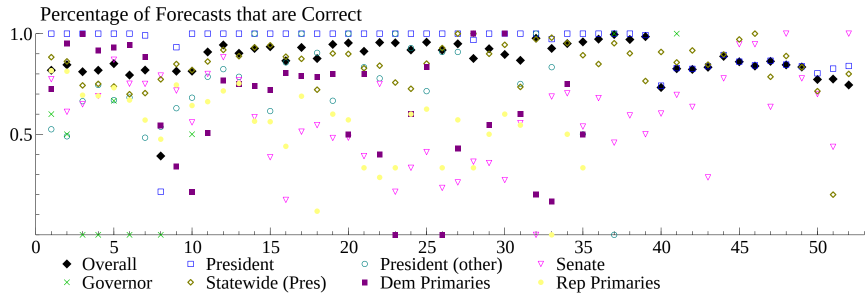

Figure 11: Percentage of Intrade prices that correctly predict election outcome by weeks until election.

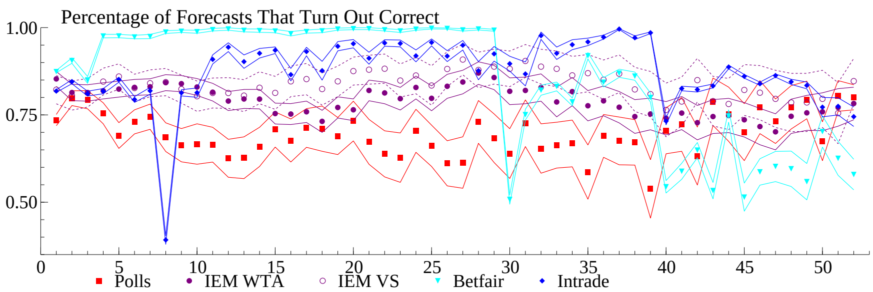

Figure 12: Plot comparing the percentage of correct forecasts for our four different sources of forecast information for US elections.

best forecasting method to choose based on accuracy (percentage of forecasts turning out correct) is Betfair from 30 weeks before an election up to election week. In the three weeks immediately before an election the performance of the four methods becomes much less dispersed, but nonetheless Betfair remains significantly better than polls and Intrade, although not significantly better than IEM WTA. With the exception of three weeks (8, 30 and 40 weeks prior), polls are dominated by prediction markets in providing accurate forecasts in the 40 weeks before an election occurs. In the final 10 weeks before an election, the performance of the IEM (both VS and WTA) and Intrade markets is indistinguishable statistically, and with the exception of forecasts 2 and 3 weeks before an election, significantly superior to polls.

Thus, concluding our discussion of accuracy in terms of the percentage of correct forecasts, we find that prediction markets dominate polls in providing accurate forecasts.

---

vation per day per market for IEM, only a relatively small number of polls per market per week, whereas we have often hundreds and even thousands of trades per day on Intrade and Betfair. We do not reduce our Intrade or Betfair samples down to any kind of daily average in order not to discard any important data.

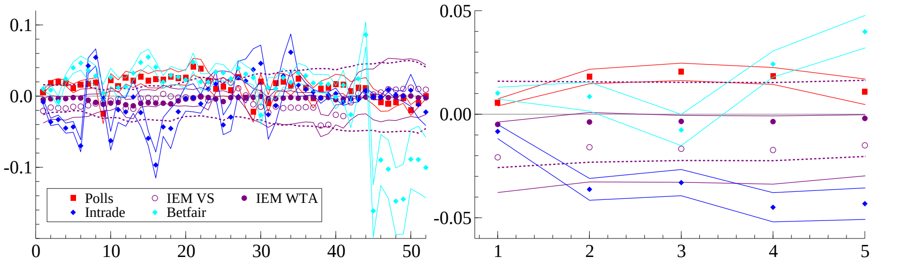

Figure 13: Plot comparing average forecast errors, as calculated in (4), head-to-head between our forecast models, by weeks until election. Left panel is all weeks in the year prior to an election, right panel focuses on final five weeks pre-election.

5.2 Bias and Precision

We next consider *bias*, whether the expected value of a forecast equals the true value, and *precision*, how much variance a forecast model exhibits, graphically before conducting a basic regression analysis help quantify our findings.

Figure 13 plots the average errors with standard error bounds for all forecast models hence giving an idea about the *bias* of forecasts. It is worth noting that the standard error bounds contain information on the *precision* of each forecast since the standard error of the forecast error is equal to the squared root of the MSFE (from (8)) with the restriction $\alpha_g = \beta_g - 1 = 0$ imposed. The left plot shows the entire year before an election, while the right plot zooms in on the final five weeks.

From the left panel what is perhaps most obvious is that over short intervals all forecast models display biases in one direction or another, but over the longer term these biases do appear to cancel each other out. The polls and IEM (both) deliver what appears to be the most consistent performance, with Betfair and Intrade fluctuating markedly around zero. In general polls have a slight upward bias, while IEM has a slight downward bias, moreso in VS than WTA, Betfair upward and Intrade downward. It thus appears that prediction markets provide higher forecast success yet are not necessarily less biased than polls.

Considering the regression model approach outlined in Section 4, Table 1 contains the output from the regression models, while Figure 14 gives a graphical representation, plotting the implied regression lines against a 45-degree line.

The regressions in columns (1) and (2) differ slightly from that of (3)–(5) in that the first two columns are regressions of vote share outcomes ($V_{i,j,T}$) on vote share forecasts ($\widehat{V}_{i,j,f,T|t}$), while in the last three columns contain regressions of actual outcomes of elections ($W_{i,j,T}$) on implied probabilities ($\widehat{W}_{i,j,f,T|t}$). Nonetheless, the principle is the same in both regressions; unbiased forecasts should be reflected in finding that $\alpha_g = 0$ and $\beta_g = 1$, namely that the implied regression line is on the 45 degree line and hence a poll forecasting a vote share on average is correct (columns (1) and (2)), and a contract priced implying a particular probability pays out with that frequency (columns (3)–(5)).

The first row of numbers in each column contains the estimates for α_g, the

	(1) Polls (V)	(2) IEM (V)	(3) IEM (W)	(4) Intrade (W)	(5) Betfair (W)
$\widehat{\alpha}_g$	0.109***	0.135***	0.082***	-0.020***	-0.062***
	(63.974)	(62.113)	(31.031)	(-28.363)	(-47.351)
$\widehat{\beta}_g$	0.802***	0.622***	0.842***	1.044***	1.189***
	(181.418)	(56.241)	(105.674)	(516.245)	(460.119)
p_val	0.000	0.000	0.000	0.000	0.000
F_stat	2699.612	1943.101	482.209	409.130	2849.682
$\widehat{\sigma}_g^2$	0.008	0.042	0.139	0.102	0.113
T	18766	11429	31737	356620	183775

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Regressions for bias and precision for all four firms over all observations.

intercept coefficient, while the second row contains the estimates for β_g, the slope coefficient. Beneath these coefficients is the output of an F-test of $\alpha_g = 0$ and $\beta_g = 1$; the first line is the p-value, the probability of a incorrect rejection of the null hypothesis, and the second row is the F-test statistic itself. In essence, the larger is the F-test statistic, the further away from $\alpha = 0$ and $\beta = 1$ is that particular set of forecasts. Because of the huge sample sizes of our regressions (from the final row), it is expected that p-values will be very small.[18] The largest F-test statistics by some distance are for polls in the first column, and Betfair in the fifth column. For polls this is mainly driven by a departure from unity of the $\widehat{\beta}_g$ coefficient, at 0.802, and the constant coefficient at 0.109, while for Betfair it would appear more a function of sample size since the deviation from $\alpha_g = \beta_g - 1 = 0$ is smaller yet the sample size is ten times as large as for polls. In terms of actual coefficient sizes, the smallest departure from $\alpha_g = \beta_g - 1 = 0$ is for Intrade. Both IEM markets show significant departures also from $\alpha_g = \beta_g - 1 = 0$, with VS dramatically so.

In terms of precision, it is notable that the two vote share regressions have a much smaller $\widehat{\sigma}_g$ than the winner-takes-all regressions; within the vote share models, polls display much more precision than IEM, and within the winner takes all Betfair and Intrade are more precise than IEM, with Intrade appearing most precise. It is here where a correspondence between vote shares and probabilities (WTA) would be most helpful; however as mentioned earlier, to do so requires a number of untestable assumptions to be made regarding the shape of the distribution of vote shares Y_t. Nonetheless a crude comparison we could use here to compare would be to make use of the ratio of $\widehat{\sigma}_g$s for IEM, which is 0.30157, suggesting that polls are more precise than prediction markets in making forecasts.

Finally we note that polls and IEM (both VS and WTA) display evidence of a reverse favourite-longshot bias (FLB), whereas Betfair and Intrade exhibit FLB. The traditional favourite-longshot bias is the observed phenomenon, found in numerous studies dispersed in time and across the world, that 'longshots' (out-

[18](Campos, Hendry and Krolzig, 2003) discuss this problem with inference in large samples, and suggest adjusting significance level to $T^{-0.8}$, where T is sample size.

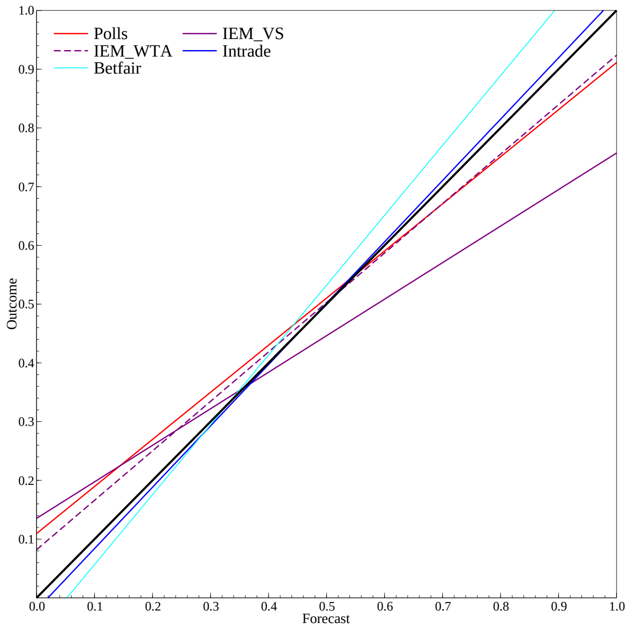

Figure 14: Plot of implied slope for each forecast against 45-degree line.

comes quoted at high odds) tend to win less often than implied in the odds while 'favourites' (outcomes quoted at low odds) tend to win relatively more often than implied in the odds (e.g. Sung and Johnson (e.g. 2010); Snowberg and Wolfers (e.g. 2010). This may help explain the difference in accuracy noted in Figure 12, where Betfair and Intrade appear to dominate the other three forecast models. If favourites win more often than their forecast suggests, as with Betfair and Intrade, then the percentage of forecasts that turn out correct must be higher, and vice versa with polls and IEM.

Thus our regression models from Table 1 lend support to the conclusions drawn from Figure 13: over the longer horizon all forecast models appear to exhibit quite substantial bias, with the exception of Intrade, while prediction markets appear to provide the least precise forecasts.

These results are plotted graphically in Figure 14; we plot each implied regression line against the 45-degree line, as the 45-degree line signifies unbiasedness in each forecasting method. The thick black line is the 45-degree line, and hence the degree of anticipated bias in each of our forecast models can be read off Figure 14. We can read off, for whatever forecast a model provides (horizontal axis), the expected actual outcome. We could make use of these lines to de-bias our forecast models as Rothschild, and Erikson and Wlezien do, since departures from the 45-degree line in Figure 14 reveal the extent of bias in our forecast models. Rothschild uses a recursively estimation version of Figure 14 based on previous elections to de-bias his 2008 election data.

The red line for polls crosses for vote shares in the mid-50s, suggesting that in general the outturn will be less conclusive than polls suggest. The bias in polls appears to be to over-predict the vote share for favorites, and under-predict it for outsiders in a political race; reverse FLB. The IEM lines are

	(1) Polls (V)	(2) IEM (V)	(3) IEM (W)	(4) Intrade (W)	(5) Betfair (W)
$\widehat{\alpha}_g$	0.013***	0.105***	0.067***	-0.009***	-0.059***
	(6.495)	(4.926)	(5.530)	(-3.418)	(-23.634)
$\widehat{\beta}_g$	1.008***	0.855***	0.868***	1.003***	1.130***
	(191.199)	(8.413)	(28.222)	(200.593)	(290.260)
p_val	0.000	0.000	0.000	0.000	0.000
F_stat	186.279	12.618	16.328	12.015	584.645
$\widehat{\sigma}_g^2$	0.002	0.128	0.111	0.127	0.082
T	3007	360	1110	43295	41882

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Regressions for bias and precision for all four firms for observations within one week of an election.

purple and both deviate quite substantially from the 45-degree line and in the same direction as polls — the reverse FLB already noted. Intrade and Betfair both display steeper slopes than the 45-degree line, such that for larger implied probabilities, contracts priced at these levels pay out more often than they ought to. Betfair's slope is steepest out of the two.

Thus to a large extent, this plot helps explain the patterns observed for the accuracy of these models in Figure 12; those models exhibiting FLB are able to pick the favorite accurately, and since the favorite wins with a higher probability, so these models forecast more accurately. Conversely, those models exhibiting reverse FLB are unable to pick the favorite accurately, and by the same reasoning these models must forecast less accurately.

As with *accuracy*, we can refine somewhat our analysis of *bias* and *precision* by looking at forecasts made at various points before an election. Tables 2–4 show the same regressions as Table 1 but for forecasts made within the final week of an election (Table 2), forecasts made between 2 and 10 weeks before and election (Table 3), and forecasts made between 11 and 40 weeks before an election (Table 2).

From Table 2 polls display much less bias in the final week than in the overall regressions in Table 1, and Intrade still displays little bias also. IEM markets still display appear biased (reverse FLB) while Betfair still displays FLB.

As we move to longer horizons in Tables 3 and 4, we observe that both Intrade and Betfair display a more pronounced bias, IEM WTA display less evidence of bias over this slightly longer time interval, polls slightly more bias and IEM VS behavior is essentially unchanged from the final week. Polls also depart further from unbiasedness as the time horizon increases; although at 2–10 weeks the α and β coefficients are close to reflecting an unbiased forecast, by 11–40 weeks they have departed significantly. The prediction markets also display a markedly stronger bias (FLB) and lower precision over these longer periods also. Again referring back to Figure 12, we note that over weeks 2–40, Intrade and Betfair, the two models exhibiting strongest FLB, forecast most accurately.

	(1) Polls (V)	(2) IEM (V)	(3) IEM (W)	(4) Intrade (W)	(5) Betfair (W)
$\widehat{\alpha}_g$	0.057***	0.124***	0.051***	-0.082***	-0.179***
	(24.079)	(17.194)	(11.412)	(-32.560)	(-63.753)
$\widehat{\beta}_g$	0.930***	0.784***	0.943***	1.136***	1.422***
	(157.873)	(21.271)	(74.193)	(218.537)	(282.062)
p_val	0.000	0.000	0.000	0.000	0.000
F_stat	699.494	149.530	70.016	533.289	3602.357
$\widehat{\sigma}_g^2$	0.005	0.136	0.121	0.150	0.076
T	6269	3218	9597	82135	35054

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Regressions for bias and precision for all four firms over observations within 2 and 10 weeks of an election.

	(1) Polls (V)	(2) IEM (V)	(3) IEM (W)	(4) Intrade (W)	(5) Betfair (W)
$\widehat{\alpha}_g$	0.149***	0.124***	0.087***	-0.069***	-0.088***
	(47.760)	(23.438)	(24.190)	(-46.585)	(-44.505)
$\widehat{\beta}_g$	0.721***	0.708***	0.831***	1.289***	1.301***
	(87.155)	(26.797)	(73.972)	(281.890)	(296.885)
p_val	0.000	0.000	0.000	0.000	0.000
F_stat	1418.768	274.713	292.717	2007.052	2509.161
$\widehat{\sigma}_g^2$	0.011	0.140	0.141	0.079	0.113
T	6608.000	6485.000	17317.000	67949.000	81133.000

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Regressions for bias and precision for all four firms over observations within 11 and 40 weeks of an election.

The information presented here suggests that firstly in terms of correct forecasts, prediction markets dominate polls. Consideration of bias and precision shows that all forecast models are shown to be biased in different directions and magnitudes at different times, while levels of precision also vary with polls being the most precise and IEM the least. Nonetheless, it seems evident that models exhibiting FLB are thus more able to identify favorites and hence forecast more accurately.

6 Conclusions

In this paper we have investigated a number of information sources that might be used to form a forecast of an election outcome. We consider the forecasts of opinion polls and three different commonly used prediction markets. We assess these forecast models in terms of accuracy, bias and precision, noting that an accurate forecast can be biased and also imprecise, whilst an unbiased forecast can be inaccurate. We make use of very large datasets recording the forecast performance of these different models over a large number of elections since 2000 in the US. Our analysis suggests that prediction markets tend to provide more accurate forecasts, although poll forecasts appear more precise, and in the final weeks before an election are fairly unbiased. In particular commercial prediction markets display distinct favorite longshot bias, suggesting that they are more able to identify favorites that subsequently win the election, which helps explain why these models forecast more accurately.

References

Bates, J., and C. W. J. Granger. 1969. "The Combination of Forecasts." *Operations Research Quarterly*, 20: 451–468.

Berg, J.E., F.D. Nelson, and T.A. Rietz. 2008. "Prediction market accuracy in the long run." *International Journal of Forecasting*, 24(2): 285–300.

Campos, J., D.F. Hendry, and H.-M. Krolzig. 2003. "Consistent Model Selection by an Automatic Gets Approach." *Oxford Bulletin of Economics and Statistics*, 65(s1): 803–819.

Croxson, K., and J.J. Reade. 2011. "Exchange vs. Dealers: A High-Frequency Analysis of In-Play Betting Prices." Department of Economics, University of Birmingham Discussion Papers 11-19.

Erikson, R.S., and C. Wlezien. 2008. "Are Political Markets Really Superior to Polls As Election Predictors?" *Public Opinion Quarterly*, 72(2): 190–215.

Erikson, R.S., and C. Wlezien. 2009. "Markets vs. Polls as Predictors: An Historical Assessment of US Presidential Elections."

Franck, E., E. Verbeek, and S. Nüesch. 2011. "Sentimental Preferences and the Organizational Regime of Betting Markets." *Southern Economic Journal*, 78(2): 502–518.

Gelman, A., and G. King. 1993. "Why are American Presidential Election Campaign Polls so Variable when Votes are so Predictable?" *British Journal of Political Science*, 23(4): 409–451.

Gil, R., and S. Levitt. 2007. "Testing the Efficiency of Markets in the 2002 World Cup." *The Journal of Prediction Markets*, 1: 255–270.

Graefe, A., J.S. Armstrong, R.J. Jones, and A.G. Cuzan. 2012. "Combining Forecasts: An Application to Elections." APSA Annual Meeting Paper.

Hanson, R., and R. Oprea. 2009. "A Manipulator can Aid Prediction Market Accuracy." *Economica*, 76(302): 304–314.

Hartzmark, S., and D. Solomon. 2008. "Efficiency and the Disposition Effect in NFL Prediction Markets." *Working Paper.*

Hayek, F.A. 1945. "The Use of Knowledge in Society." *The American Economic Review*, 35(4): 510–530.

Hendry, D.F., and M.P. Clements. 1998. *Forecasting Economic Time Series.* Cambridge:Cambridge University Press.

Hurley, W., and L. McDonough. 1995. "A Note on the Hayek Hypothesis and the Favorite-Longshot Bias in Parimutuel Betting." *The American Economic Review*, 85(4): 949–955.

Kou, S., and M.E. Sobel. 2004. "Forecasting the Vote: A Theoretical Comparison of Election Markets and Public Opinion Polls." *Political Analysis*, 12(3): 277–295.

Lee, D.S., and E. Moretti. 2009. "Bayesian Learning and the Pricing of New Information: Evidence from Prediction Markets." *The American Economic Review*, 99(2): 330–336.

Leigh, A., and J. Wolfers. 2006. "Competing Approaches to Forecasting Elections,: Economic Models, Opinion Polling and Prediction Markets." *Economic Record*, 82(258): 325–340.

Page, L. 2008. "Comparing Prediction Market Prices and Opinion Polls in Political Elections." *Journal of Prediction Markets*, 2(1): 91–97.

Rhode, P.W., and K.S. Strumpf. 2004. "Historical Presidential Betting Markets." *Journal of Economic Perspectives*, 18(2): 127–142.

Rothschild, D. 2009. "Forecasting Elections." *Public Opinion Quarterly*, 73(5): 895–916.

Servan-Schreiber, E., J. Wolfers, D.M. Pennock, and B. Galebach. 2004. "Prediction Markets: Does Money Matter?" *Electronic Markets*, 14(3): 243–251.

Sjöberg, Lennart. 2009. "Are All Crowds Equally Wise? A Comparison of Political Election Forecasts by Experts and the Public." *Journal of Forecasting*, 28(1): 1–18.

Smith, V.L. 1982. "Markets as Economizers of Information: Experimental Examination of the "Hayek Hypothesis"." *Economic Inquiry*, 20(2): 165–179.

Snowberg, E., and J. Wolfers. 2010. "Explaining the Favorite-Longshot Bias: Is It Risk-Love or Misperceptions?" *Journal of Political Economy*, 118(4): 723–746.

Snowberg, E., J. Wolfers, and E. Zitzewitz. 2007. "Partisan impacts on the economy: evidence from prediction markets and close elections." *The Quarterly Journal of Economics*, 122(2): 807.

Snowberg, E., Wolfers J., and E. Zitzewitz. 2005. "Information (In)Efficiency in Prediction Markets." In *Information Efficiency in Financial and Betting Markets.* , ed. L. Vaughan Williams, 366–386. Cambridge:Cambridge University Press.

Spann, M., and B. Skiera. 2009. "Sports Forecasting: A Comparison of the Forecast Accuracy of Prediction Markets, Betting Odds and Tipsters." *Journal of Forecasting*, 28(1): 55–72.

Sung, M., and J.E.V. Johnson. 2010. "Revealing Weak-Form Efficiency in a Market for State Contingent Claims: The Importance of Market Ecology, Modelling Procedures and Investment Strategies." *Economica*, 77: 128–147.

Surowiecki, J. 2004. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations.* Little.

Wolfers, J., and E. Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives*, 18(2).

For Online Publication

A Data Information Tables

The Tables on the following pages contain information on the make-up of each of our datasets introduced in Section 3 and analysed in Section 5.

Democratic Primaries	**Governor** 2006	**Governor** 2010	**House 2006**	**House 2010**	**House 2010 (cont)**	**House 2010 (cont)**	**Republican Primaries 2008**	**Senate 2006**	**Senate 2010**	**State 2008**
Alabama	Alabama	Alabama	Arizona 8	Alabama 2	Indiana 8	New York 2	Arizona	Arizona	Alabama	Alabama
Arizona	Alaska	Alaska	Colorado 4	Alabama 5	Indiana 9	New York 4	California	Connecticut	Alaska	Alaska
California	Arizona	Arizona	Colorado 7	Arizona 1	Iowa 1	New York 13	Colorado	Maryland	Arizona	Arizona
Connecticut	Arkansas	Arkansas	Connecticut 2	Arizona 3	Iowa 2	New York 19	Connecticut	Michigan	Arkansas	Arkansas
Delaware	California	California	Connecticut 4	Arizona 5	Iowa 3	New York 20	Delaware	Minnesota	California	California
Democratic	Colorado	Colorado	Connecticut 5	Arizona 7	Kansas 3	New York 22	Florida	Missouri	Colorado	Colorado
Florida	Connecticut	Connecticut	Florida 13	Arizona 8	Kentucky 3	New York 23	Georgia	Montana	Connecticut	Connecticut
Georgia	Florida	Florida	Florida 16	Arkansas 1	Kentucky 6	New York 24	Illinois	New Jersey	Delaware	Delaware
Illinois	Georgia	Georgia	Florida 22	Arkansas 2	Louisiana 2	New York 25	Iowa	Ohio	Florida	Florida
Indiana	Idaho	Hawaii	Georgia 8	Arkansas 4	Louisiana 3	New York 27	Massachusetts	Pennsylvania	Georgia	Georgia
Iowa	Illinois	Idaho	Georgia 12	California 3	Maine 1	New York 29	Michigan	Rhode Island	Hawaii	Hawaii
Kansas	Iowa	Illinois	Illinois 6	California 11	Maine 2	North Carolina 2	Minnesota	Tennessee	Idaho	Idaho
Kentucky	Kansas	Iowa	Illinois 8	California 18	Maryland 1	North Carolina 4	Missouri	Virginia	Illinois	Illinois
Maryland	Maine	Kansas	Indiana 2	California 20	Massachusetts 4	North Carolina 7	Nevada	Washington	Indiana	Indiana
Massachusetts	Maryland	Maine	Indiana 8	California 45	Massachusetts 5	North Carolina 8	New Hampshire	Senate Elections 2008	Iowa	Iowa
Michigan	Massachusetts	Maryland	Indiana 9	California 47	Massachusetts 6	North Carolina 11	New Jersey	Alaska	Kansas	Kansas
Minnesota	Michigan	Massachusetts	Iowa 1	Colorado 3	Massachusetts 10	North Carolina 13	New York	Colorado	Kentucky	Kentucky
Missouri	Minnesota	Michigan	Kentucky 3	Colorado 4	Michigan 1	Ohio 1	Republican Nomination	Georgia	Louisiana	Louisiana
Nevada	Nebraska	Minnesota	Kentucky 4	Colorado 7	Michigan 7	Ohio 6	South Carolina	Kentucky	Maryland	Maine
New Hampshire	Nevada	Nebraska	Minnesota 6	Connecticut 1	Michigan 9	Ohio 10	Tennessee	Louisiana	Missouri	Maryland
New Jersey	New Hampshire	Nevada	Nevada 2	Connecticut 2	Michigan 15	Ohio 12	**Republican Primaries 2012**	Maine	Nevada	Massachusetts
New Hampshire	New Mexico	New Hampshire	Nevada 3	Connecticut 4	Minnesota 1	Ohio 13	Alabama	Minnesota	New York (Dio Vs Gil)	Michigan
New York	New York	New Mexico	New Hampshire 2	Connecticut 5	Minnesota 6	Ohio 15	Arizona	Mississippi Special	New York (Tow Vs Sch)	Minnesota
North Carolina	Ohio	New York	New Jersey 7	Florida 2	Minnesota 7	Ohio 16	Colorado	Nebraska	North Carolina	Mississippi
Ohio	Oklahoma	Ohio	New Mexico 1	Florida 8	Minnesota 8	Ohio 18	Florida	New Hampshire	North Dakota	Missouri
Oregon	Oregon	Oklahoma	New York 20	Florida 12	Mississippi 1	Oklahoma 2	Georgia	New Jersey	Ohio	Montana
Pennsylvania	Pennsylvania	Oregon	New York 24	Florida 22	Mississippi 2	Oregon 1	Hawaii	New Mexico	Oklahoma	Nebraska
South Carolina	Rhode Island	Pennsylvania	New York 26	Florida 24	Mississippi 4	Oregon 4	Illinois	North Carolina	Oregon	Nevada
Tennessee	South Carolina	Rhode Island	New York 29	Florida 25	Missouri 3	Oregon 5	Iowa	Oregon	Pennsylvania	New Hampshire
Texas	Tennessee	South Carolina	North Carolina 11	Georgia 2	Missouri 4	Pennsylvania 3	Louisiana	Virginia	Utah	New Jersey
Virginia	Texas	South Dakota	Ohio 1	Georgia 8	Nebraska 2	Pennsylvania 4	Maine		Vermont	New Mexico
West Virginia	Vermont	Tennessee	Ohio 15	Georgia 12	Nevada 3	Pennsylvania 6	Massachusetts		Washington	New York
Wisconsin	Wisconsin	Texas	Ohio 18	Hawaii 1	New Hampshire 1	Pennsylvania 7	Michigan		West Virginia	North Carolina
	Wyoming	Utah	Pennsylvania 6	Idaho 1	New Hampshire 2	Pennsylvania 8	Minnesota		Wisconsin	North Dakota
		Vermont	Pennsylvania 7	Illinois 8	New Jersey 3	Pennsylvania 10	Mississippi			Ohio
		Wisconsin	Pennsylvania 8	Illinois 10	New Jersey 6	Pennsylvania 11	Nevada			Oklahoma
		Wyoming	Pennsylvania 10	Illinois 11	New Jersey 12	Pennsylvania 12	New Hampshire			Oregon
			Texas 22	Illinois 14	New Mexico 1	Pennsylvania 13	Ohio			Pennsylvania
			Washington 8	Illinois 17	New Mexico 2	Pennsylvania 15	Oklahoma			Rhode Island
			Wisconsin 8	Indiana 2	New Mexico 3	Pennsylvania 17	South Carolina			South Carolina
					New York 1	Rhode Island 1	Tennessee			South Dakota
						South Carolina 5	Vermont			Tennessee
						Tennessee 4	Virginia			Texas
							Washington			Utah
										Vermont
										Virginia
										Washington
										West Virginia
										Wisconsin
										Wyoming

Table 5: Summary of data publicly available from *Real Clear Politics* on polling for US elections.

Election	Start	Finish	Freq.	Percent
Congress 2000	28jan1999	08nov2000	2,592	5.69
Congress 2002	19jul2002	07nov2002	448	0.98
Congress 2004	17jun2004	05nov2004	568	1.25
Congress 2006	01jun2006	12nov2006	664	1.46
Congress 2008	22aug2008	07nov2008	312	0.68
Congress 2010	24nov2009	04nov2010	1,344	2.95
Senate Elections 2004	17jun2004	05nov2004	426	0.93
Senate Elections 2006	01jun2006	10nov2006	492	1.08
Senate Elections 2008	22aug2008	07nov2008	234	0.51
Senate Elections 2010	24nov2009	04nov2010	1,003	2.20
Florida Senate Election 2010 (vote share)	04jun2010	30nov2010	720	1.58
Florida Senate Election (winners takes all)	04jun2010	30nov2010	720	1.58
Minnesota Senate Election 2008 (vote share)	20aug2008	08nov2008	243	0.53
Minnesota Senate Election 2008 (winners takes all)	20aug2008	08nov2008	243	0.53
New York Senate Election 2000	14jun1999	08nov2000	2,725	5.98
House Elections 04	17jun2004	05nov2004	426	0.93
House Elections 06	01jun2006	10nov2006	492	1.08
House Elections 08	22aug2008	07nov2008	234	0.51
House Elections 10	24nov2009	04nov2010	1,004	2.20
Presidential Election 2000 (vote share)	03jan2000	05nov2000	920	2.02
Presidential Election 2000 (winners takes all)	24apr2000	10nov2000	597	1.31
Presidential Election 2008 (vote share)	01jun2006	07nov2008	1,830	4.01
Presidential Election 2008 (winners takes all)	01jun2006	07nov2008	1,830	4.01
Presidential Election 2004 (vote share)	20feb2003	31jul2004	7,026	15.41
Presidential Election 2004 (winners takes all)	26may2004	05nov2004	426	0.93
Democratic Convention 2000	14jun1999	17aug2000	3,219	7.06
Democratic Convention 2004	20feb2003	30jul2004	3,419	7.50
Democratic Convention 2008	24feb2007	28aug2008	2,300	5.04
Republican Convention 2000	14jun1999	03aug2000	2,339	5.13
Republican Convention 2008	24feb2007	10sep2008	3,141	6.89
Reform Convention 2000	03jan2000	12aug2000	1,105	2.42
Iowa Republican Caucus 12	29aug2011	05jan2012	889	1.95
New York City Mayoral Election 2001	03oct2001	09nov2001	108	0.24
Philadelphia Mayoral Election 2007 (vote share)	02apr2007	02jul2007	534	1.17
Philadelphia Mayoral Election 2007 (winners takes all)	02apr2007	02jul2007	529	1.16
Mexican Presidential Election 2000 (vote share)	01may2000	02jul2000	244	0.54
Mexican Presidential Election 2000 (winner takes all)	01may2000	02jul2000	244	0.54
Total			45,590	

Table 6: Summary of data publicly available from *Iowa Electronic Markets* on polling for US elections.

Type	Market	Observations	Type	Market	O
Presidential Election	2004	44,462		Republican Candidate	
	2008	11,831		California Primary	
	Democratic Candidate	39,012		Florida Primary	
	Alabama Primary	123		Iowa Caucus	
	Alaska Caucus	2	Republican Candidacy	Michigan Primary	
	Arizona Primary	132		Nevada Caucus	
	Arkansas Primary	4		New Hampshire Primary	
	California Primary	1,564		New Jersey Primary	
	Colorado Caucus	7		South Carolina Primary	
	Connecticut Primary	78		New York Primary	
	Delaware Primary	13		Total	
	Georgia Primary	63		Additional Runners	
	Idaho Caucus	4		Joe Biden	
	Illinois Primary	60		Michael Bloomberg	
	Indiana Primary	1,860		Mike Huckabee	
	Iowa Caucus	941		Mitt Romney	
	Kansas Caucus	6		Ron Paul	
	Kentucky Primary	88	Next President	Rudy Giuliani	
	Massachusetts Primary	227		Sarah Palin	
	Minnesota Caucus	13		Al Gore	
Democratic Candidacy	Missouri Primary	184		Barack Obama	
	Nevada Caucus	516		Hillary Clinton	
	New Hampshire Primary	1,710		John Edwards	
	New Jersey Primary	250		John McCain	
	New Mexico Caucus	19		Total	
	New York Primary	249		Arkansas	
	North Carolina Primary	551		Indiana	
	North Dakota Caucus	2		New Mexico	
	Ohio Primary	1,041		North Dakota	
	Oklahoma Primary	19		Nevada	
	Oregon Primary	148		Colorado	
	Pennsylvania Primary	1,265		Florida	
	Tennessee Primary	55	Elections 2008	Georgia	
	Texas Primary	2,305		Kentucky	
	Utah Primary	17		Missouri	
	Washington Caucus	69		Montana	
	West Virginia Primary	110		Nebraska	
	Wisconsin Primary	460		Ohio	
	Total	53,167		North Carolina	
				Pennsylania	
				Total	
Grand Total	228,264				

Table 7: Data from Betfair on US Elections

Market	Obs.
Presidential Election - Main	
Winner (Indiv.)	334,286
Winner (Party)	18,007
Rep. Elec. College Votes	1,225
Electoral College Tie	90
Alabama	38
Alaska	209
Arizona	784
Arkansas	248
California	293
Colorado	926
Connecticut	109
Delaware	29
Florida	2,535
New Jersey	405
Nevada	877
Nebraska	63
Montana	1,270
Missouri	2,803
Mississippi	111
Michigan	487
Minnesota	576
Maryland	52
Georgia	1,605
Hawaii	17
Idaho	16
Illinois	57
Indiana	3,956
Iowa	558
Kansas	142
Kentucky	93
Louisiana	194
Maine	152
Massachusetts	26
North Carolina	3,212
North Dakota	1,023
Ohio	1,846
Oklahoma	37
Oregon	247
Pennsylvania	1,404
New Hampshire	796
New Mexico	586
New York	118
Rhode Island	30
South Dakota	120
South Carolina	187
Tennessee	141
Texas	125
Utah	24
Vermont	61
Virginia	2,237
Washington	128
West Virginia	863
Wisconsin	395
Wyoming	22
Total	385,841
Total	411,858

Market	Obs.
Presidential Election - Other	
Bob Barr - Elec. Coll. Votes	14
Bob Barr - Popular Vote	444
Dropouts, April	172
Dropouts, Dec.	210
Dropouts, Feb	519
Dropouts, Jan	349
Dropouts, Jun	72
Dropouts, Jul	1,291
Dropouts, May	186
Dropouts, Mar	568
LA Times Obama PLO video	3
Ralph Nader - Popular Vote	117
Joe Biden to be withdrawn	548
Sarah Palin to be withdrawn	2,599
Michael Bloomberg Independent	1,073
Ron Paul Independent	249
Who benefit most from 1st debate	330
Who benefit most from VP debate	606
Who will run for President?	941
Date of 1st Debate	215
Election Postponed?	48
Obama Touch Mkt	9
McCain Touch Mkt	39
X: Obama Options. F	123
X: Obama Options. M	15
X: Obama Options. T	1
X: Obama Options. W	83
X: Obama Options. W	147
X: Obama Options. W	54
X: Obama Options. W	75
X: Obama Options. W	98
X: Obama Options. W	82
X: McCain Options. Fr	132
X: McCain Options. Mo	1
X: McCain Options. We	21
X: McCain Options. We	17
X: McCain Options. We	7
X: McCain Options. We	19
X: McCain Options. We	12
X: McCain Options. We	34
Total	11523
Other	
Immigration Reform Act 20	24
London Mayoral Election 2008	94
Massachusetts Question 1	8
Media Endorsements	20
New York City Mayoral Term Limits	18
Next Prime Minister of New Zealand	6
Next UK Chancellor	15
Pres. Job Appr. Rating. Dec 3	119
Pres. Job Appr. Rating. Jun 3	38
Pres. Job Appr. Rating. Mar 3	29
Pres. Job Appr. Rating. Sep 3	72
Fairness Doctrine	6
Total	449

Market
House of Representatives
2008 House Control
Dem. Seats in House
Dist. 12 Penn
Dist. 6 Minn
Total
Senate
2008 Senate Control
Dem. Seats in Senate
Alabama
Alaska
Colorado
Georgia
Idaho
Kansas
Kentucky
Louisiana
Maine
Massachusetts
Minnesota
Nebraska
New Hampshire
New Jersey
New Mexico
Mississippi (Class I)
Mississippi (Class II)
North Carolina
Oklahoma
Oregon
South Carolina
South Dakota
Texas
Virginia
West Virginia
Wyoming (Class I)
Total
Governor Elections
Delaware
Kentucky
Louisiana
Missouri
North Carolina
Utah
Vermont
Washington
Total
Party Convention
Brokered Conventions
Clinton Lifeline
Hillary Clinton on Dem.
MI/FL hold new Primar
Most Superdelegates?
Total

Table 8: Summary of data from Intrade for 2008 elections.