Prediction Market Accuracy in the Long Run*

Joyce E. Berg[1], Forrest D. Nelson[2] and Thomas A. Rietz[3]

Working Draft

January 2008

Forthcoming: *The International Journal of Forecasting*

Abstract

"Prediction markets" are designed specifically to forecast events such as elections. Though election prediction markets have been conducted for almost twenty years, to date nearly all evidence on efficiency compares election eve forecasts to final pre-election polls and actual outcomes. Here, we present evidence that prediction markets outperform polls for longer horizons. We gather national polls for the 1988 through 2004 U.S. Presidential elections and ask whether the poll or a contemporaneous Iowa Electronic Markets vote share market prediction is closer to the eventual outcome for the two-major-party vote split. We compare market predictions to 964 polls over the five Presidential elections since 1988. The market is closer to the eventual outcome 74% of the time. Further, the market significantly outperforms the polls in every election when forecasting more than 100 days in advance.

Keywords: Combining forecasts; Evaluating forecasts; Financial markets; Election forecasting; Polls; Comparative Methods; Automatic forecasting; Calibration; Comparative studies; Long-term forecasting; Election market, Political stock market.

* For many helpful comments and suggestions during the course of this research, we thank Bob Forsythe, George Neumann and participants in the Department of Accounting and Department of Finance research seminars at the University of Iowa Tippie College of Business.

[1] Department of Accounting, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242
[2] Department of Economics, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242
[3] Department of Finance, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242

Prediction Market Accuracy in the Long Run

I. Introduction

How does one forecast an election outcome? Authors have suggested (1) naive forecasts (Campbell, 2005, suggests this as a benchmark), (2) polls (e.g., Perry 1979), (3) prediction markets (Forsythe, Nelson, Neumann and Wright, 1993), (4) structural models (e.g., Fair, 1978, or Lewis-Beck and Tien, 2007), (5) time series models (Norpoth, 1996, uses time series elements) and (4) less formal methods such as focus groups, interviews of knowledgeable parties and expert panels (news sources often interview various pundits and experts; on the more formal side Cuzán, Armstrong and Jones, 2005, conducted Delphi Techniques using a panel of experts).

Given a sufficient number of observations under essentially identical conditions, correct specification and sufficient stationarity, parameter estimates from both time series and structural models should converge to their true values, thus eliminating sampling error and leaving inherent randomness as the only error in forecasts of the outcome of the election. However, sufficient data under stationary conditions may be difficult to come by in the political process, and the idiosyncrasies of individual elections may still leave forecasts errors unacceptably high. Given a random sample, accurate responses and a sufficiently static environment, surveys or polls should also accurately predict election outcomes. However, obtaining a truly random sample can be difficult (e.g., the Truman/Dewey race and, now, the prevalence of voters who do not have traditional phone lines) and often the environment can change quickly. Political campaigns are designed to influence how people will vote in an upcoming election. They often react to counter poll results and, if they are effective, essentially invalidate the poll predictions. Expert opinion can be difficult to aggregate in an acceptable manner. The Delphi Technique is designed to overcome many issues with expert opinion, but Cuzán, Armstrong and Jones (2005) and Jones, Armstrong and Cuzán (2007) found no extant studies in the literature of the application of the Delphi method to elections. Ongoing research by these authors into the use of that method suggests promising results but perhaps little gain over simple expert surveys. As with opinion polls, however, the expert surveys and Delphi methods are expensive, and more experience is needed to assess their efficacy.

Here, we extend the research studying whether prediction markets can serve as effective forecasting tools in elections. Prediction markets are designed and conducted for the primary purpose of aggregating information so that market prices forecast future events. These markets differ from typical, naturally occurring markets in their primary role as a forecasting tool instead of a resource allocation mechanism. Beginning in 1988, faculty at the Henry B. Tippie College of Business at the University of Iowa have conducted markets designed to predict election outcomes.[1] These markets, now known as the Iowa Electronic Markets (IEM), have proven accurate in forecasting election vote-shares the evening and week before elections. Here, we show that, well in advance of the elections, these markets dominate polls in forecasting election outcomes.

We report on five markets from the Iowa Electronic Markets designed to predict US Presidential election vote shares and compare them to the obvious alternative: polls. We compare these two techniques specifically because (1) polls and prediction markets are used to forecast the same thing (the vote shares of candidates), (2) in contrast to naive forecasts and typical structural and time series models, they generate a large number of forecasts in each election and (3) unlike expert opinion, they are readily available and can be understood and compared easily.

Prediction markets like the IEM should predict complex phenomena including election outcomes accurately for several reasons. First, the market design forces traders to focus on the specific event of interest, in this case how the entire electorate will vote in the specific election. This requires more than simply building a model based on past elections (because of the large differences across elections) and more than simple consideration of a fictitious election “if it were to be held today" (as polls ask respondents to consider). Second, to voice their opinions, traders must open a position in the market, putting money at stake. Presumably, the more confident they are in their predictions, the more money they will be willing to risk. Third, the market aggregates the diverse information of traders in a dynamic and, hopefully, efficient manner. Finally, the markets provide an incentive to generate, gather and process information across information sources and in a variety of ways. Traders who perform these tasks well prosper. Those who don't may go broke, may drop out of the market and appear less likely to set forecast determining prices (see Oliven and Rietz, 1999).

Existing evidence (e.g., Berg, Forsythe, Nelson and Rietz, 2003, and references cited therein) shows excellent predictive accuracy for election vote share prediction markets in the very short run (i.e., one-day-ahead forecasts using election eve prices). Extending a similar figure from Berg, Forsythe, Nelson and Rietz (2003) to include the 2004 election results, Figure 1 shows this accuracy on election eve, a 1.33 percentage point average absolute error. For the five elections included in that figure, the average absolute error in the market's prediction of the major-party presidential vote share across the 5 days prior to the election was 1.20 percentage points, while opinion polls conducted during that same time had an average error of 1.62 percentage points.

Insert Figure 1 about Here

In this paper, we present an analysis of the long-run forecasting ability of markets relative to polls. Because many of the settings in which prediction markets could be used do not have long histories of results on which to model adjustments to raw data, we compare market prices to raw poll data, adjusting only so that both market prices and poll numbers sum to one.[2] Results show that prediction markets are more accurate long-run forecasting tools than polls across elections and across long periods of time preceding elections (in addition to election-eve). The basis for our statement is a simple one. We compare the market predictions of two-party vote splits to poll predictions, normalizing poll splits to control for third party and undecided votes and comparing to the IEM price on the last day a poll is in the field, so that the market prices and polls are comparable both in measure and in time. We simply ask how often market predictions are closer to the eventual outcome than polls. Aggregating over 964 polls from the five Presidential elections since 1988, the market is closer to the eventual two-party vote split 74% of the time. Further, the market significantly outperforms the polls in every election when forecasting more than 100 days in advance.

In the next section, we briefly describe prediction markets in general and the specific election markets we study. Then, we present our results and end with concluding remarks.

II. Prediction Markets

Since Hayek (1945), economists have recognized that markets have a dual role. They allocate resources and, through the process of price discovery, they aggregate information about the values of those resources. The information aggregation role of some markets seems particularly apparent. For example, corporations cite the value of their publicly traded stock as the consensus judgment of their owners about the value of the corporation's activities and, increasingly, corporations reward managers based on those stock values. Futures and options markets aggregate information about the anticipated future values of stocks and commodities. If it is true that futures prices are the best predictors of actual future spot prices (as the "expectations hypothesis" asserts), then futures prices constitute forecasts.[3] For example, Krueger and Kuttner (1996) discuss how the Federal Funds futures contract can be used to predict future Federal Funds rates and, hence, future Federal Reserve target rates.

In most markets, if prediction uses arise, they do so as a secondary information aggregation role. However, some recent markets have been designed specifically to exploit their information aggregation characteristics for use as dynamic forecasting systems. Examples of such “prediction markets” include numerous markets run under the Iowa Electronic Markets (designed to predict elections, other political events, movie box office receipts, corporate earnings, returns, stock prices, incidences of influenza, hurricane landfalls, etc.; see Forsythe, Nelson, Neumann and Wright, 1992, and Berg, Forsythe, Nelson and Rietz, 2003 for more detailed descriptions), similar markets run in other countries (usually designed to predict election outcomes) and markets cited in Plott (2000) (designed to predict sales at a large corporation). While the majority of such markets use cash payoffs, some similar Internet “games” have been conducted using fictitious currency with prize contests as motivation. These include the Foresight Exchange (http://www.ideosphere.com) with “payoffs” tied to a wide range of social, political and scientific events and issues, the Hollywood Stock Exchange (http://www.hsx.com) with “payoffs” tied to movie box office success and other entertainment events, NewsFutures (http://us.newsfutures.com/) with markets based on politics and other newsworthy events, the (apparently now defunct) Major League Market (http://majorleaguemarket.com) with “payoffs” tied to the performance of teams and athletes, and similar markets with contract “de-listing values” (i.e., liquidating “payoffs”) tied directly to predictable events (in contrast to vague notions of “popularity”). In their survey of prediction markets, Wolfers and Zitzewitz (2004), cite the accuracy of political markets as well as accuracy in markets designed to predict events ranging from printer sales, to macro-economic statistics, to the box office takes of Hollywood movies. In the preface to their book on prediction markets, Hahn and Tetlock (2006) conclude that “the bottom line is that information markets seem to work reasonably well in a wide variety of settings.”

Prediction markets, such as the Iowa Electronic Markets (IEM), represent an important advance in forecasting. The idea is straightforward: trade contingent claims in a market where the claims pay off as a function of something one is interested in forecasting. If structured correctly, the prices should reflect the expected payoffs to the claims and therefore the expected outcome of the event of interest. This relationship can be used for forecasting. For example, the IEM's vote-share markets trade contracts with payoffs that equal $1 times the relative percentages of the vote won by major candidates in a future election. In these markets, prices should converge to the market's expectation of relative vote shares. Though simple in concept, such markets act as complex, dynamic, interactive systems that incorporate information in new ways. Through the actions of traders, prediction markets aggregate information from individuals, polls and other sources of information, weighing all of this information through the price formation process. They compete directly with, and potentially use as information, traditional methods of forecasting such as polls, econometric modeling, panels of experts, and marketing surveys.

Here we ask a simple question: Well in advance of the election, are prediction markets closer to eventual vote-shares than polls? This extends the usual measure of predictive accuracy, which is based on election-eve market forecasts and final polls. To answer this question, we compare (1) the forecasts from IEM prediction markets designed to predict vote shares of candidates in United States Presidential elections since 1988, (2) contemporaneous poll results and (3) the eventual outcomes of the elections. We ask which is closer to the eventual outcome: the poll or the market price at the time the poll was in the field. We find that on average the markets are closer than polls to the eventual election vote-share.

III. The Iowa Electronic Markets Presidential Vote Share Markets

The prediction markets we study in this paper are the Iowa Electronic Markets (IEM) Presidential Vote Share markets conducted since 1988. They are the first and are the longest running set of formal prediction markets known to us. The IEM is a computerized, electronic, real-time exchange where traders buy and sell futures contracts with payoffs based on election outcomes. Traders entering the market are allowed to invest between $5.00 and $500.00. Because real money is used, traders are subject to the real monetary risks and returns that result from their trading behavior.

Contracts in the IEM Vote Share markets are designed to forecast the vote-shares received by candidates. Contracts pay an amount equal to the relative percentage of the popular vote received by a candidate times $1.[4] Table 1 shows the specific contracts for the IEM Presidential Vote Share markets run to date. Appropriate contract specification and vote-share normalization insures that the contract payoffs always sum to $1. Simple no-arbitrage arguments imply that market prices should reflect the traders' consensus forecast of the vote shares taken by each candidate[5]. Thus, vote-share markets provide point predictions about candidate vote shares.

Insert Table 1 about Here

Table 2 shows statistics for the Presidential Vote Share markets for the 1988 through 2004 elections.[6] The number of active traders in the vote share markets ranged from 155 in 1988 to 790 in the 2000 market. Overall volumes ranged from 15,826 contracts worth $8,123 in 1988 to 339,222 contracts worth $46,237 in 2004.

Insert Table 2 about Here

As a prediction system, the IEM differs from expert panels and polls in a number of respects. Instead of being a randomly selected representative sample or a deliberately chosen panel, IEM traders are self-selected. People who are not interested either do not sign up or drop out. Further, the market does not equally weight traders' opinions in the price formation process. Instead, the market price is a metric which, through trading behavior and market dynamics, depends upon the traders' forecasts and the levels of confidence they have in their forecasts as well as an untold number of other factors like aggressiveness, risk aversion, timing, wealth, etc. Unlike polls or expert panels in which participants are asked for their independent opinions, each trader in the market sees the net effect of the beliefs of all other traders, and the time series of changes in those beliefs and can alter his own perceptions accordingly. This makes the market more than a static one-time prediction – it is a dynamic system that can respond instantaneously to the arrival of new information. Unlike polls that ask each respondent how he or she would vote if the election were held today, the market asks traders to forecast how everyone will vote in the actual upcoming election.[7]

As an example of these differences, consider the demographics of IEM traders. A good poll would strive to collect responses from a random, representative sample of voters. This sampling makes a difference: polls of likely voters are more accurate than polls of registered voters (Crespi, 1988). In contrast, IEM traders are self-selected and differ greatly from a representative sample of voters. In 1988, traders included only interested members of the University of Iowa academic community. In the other elections, traders included interested individuals from around the world. For example, in the 2000 vote-share market, 20% of the traders were from Iowa while Iowa only accounted for 1% of the nation's population in 2000. Men constituted 75% of the active traders but only 49% of the overall population (and slightly less of the voting population). IEM traders are typically young, white, well educated and have high family incomes. Thus, IEM predictive accuracy relies heavily on a sample (in practice, a non-representative sample) of interested traders forecasting the behavior of the voting population at large. It does not depend on the traders themselves constituting a representative sample of voters.

IV. Performance Versus Polls

In this paper, we address the question of whether the IEM outperforms polls as a predictive system well in advance of the election outcomes. We use raw polls rather than polls adjusted using mechanical or historical models for three reasons: (1) We want our evidence to be applicable to settings in which there is not a long history of polls, surveys, or other standing forecasting methods, that can be used to build adjustment models. (2) There is no consensus in the election forecasting literature about which poll adjustment mechanism is superior. (3) Raw poll results are the ones reported in the media, are transparent, are readily available and are the results used by the general public as evidence of the likely outcome of the pending election. Indeed the reporting of raw poll results, rather than adjusted results, was the recommendation of the Mosteller Commission, which was formed in 1948 to investigate the failure of opinion polls in the Truman-Dewey election, and that recommendation has been followed religiously by the opinion polling industry since the publication of the commission report (Mosteller, et al. 1949).

The IEM has conducted markets on five US presidential elections. Table 1 summarizes these markets. In 1988, a vote-share market predicted the popular vote shares taken by Bush, Dukakis, Jackson and rest-of-the-field. In 1992, the IEM vote-share market was split between two sub-markets. One sub-market predicted the vote split between the Democrat (Clinton) and the Republican (G.H. Bush). A second sub-market predicted the split between the two major parties and Perot. The 1996 vote-share market predicted the vote split between Clinton as the Democratic nominee and Dole as the Republican nominee. In 2000, the vote-share market forecast the election vote shares for the Democratic, Reform and Republican nominees (Gore, Buchanan and G.W. Bush, respectively). Finally, in 2004, the vote-share market predicted the split between the Democratic and Republican candidates (Kerry and G.W. Bush, respectively).[8] Our analysis focuses on the vote-share markets and the vote splits between just the Democratic and Republican candidates because these are the most directly comparable to polls.[9] We judge the accuracy of these market forecasts by comparing them to the actual election outcomes.

Polls used for comparison with the market include all nation-wide poll reports we were able to find for each of the five elections. For the three elections prior to 2000, polls were collected directly from news reports. For elections in 2000 and 2004, poll results were collected from http://pollingreport.com. Poll reports based on samples of "Likely Voters" were chosen when possible; reports using "Registered Voters" were the second choice; if neither of those were reported then we used reports based on samples of “All Adults.”[10] Many of the polls we found were tracking polls and such polls appear to be reported with greater frequency in more recent elections. Tracking polls use rolling samples with, typically, N/k new subjects added each day to replace the N/k oldest subjects in the sample. These overlapping samples result in a lack of independence from one day to the next. To avoid this dependence, we retain for analysis only every k^{th} report of a tracking poll, working backward from the last report so as to include data as close to the election as possible. Often pollsters will, in the same poll, ask for the favorite candidate from a broad list and again for the favorite from a narrow list, typically just two in the latter question. In such cases we use only the result on the question with the broadest list of candidates. Polls reports with imprecise starting and ending dates, and polls conducted prior to the start of the market were excluded from the analysis. The final sample of polls included 59 polls from 1988, 151 in 1992, 157 in 1996, 229 in 2000 and 368 in 2004.[11]

Figure 2 contains graphs of the margin of victory for the Electoral College winner as predicted by the polls and the market for the five elections.[12] Market predictions are generated from closing prices (the last trade price before midnight each day). Poll outcomes are plotted on the last day that polling took place for that particular poll, which is typically a day earlier than the release of the poll. For both market prices and polls, the outcomes are plotted as the normalized two-party vote margin.

Insert Figure 2 about Here

In order to compare IEM election market prices (where the category, “undecided,” is not applicable) and polls (where “undecided” and even other candidates are possible choices), we normalize using just the Democratic and Republican candidates. For polls this means that the undecided and other candidates are split proportionally as is conventionally done in poll research. For market prices, this means that deviations from a total price of $1.00 (due, for instance, to asynchronous trading) are also spread proportionally to candidates.

Using these conventions, the poll margins in Figure 2 for 1996 are computed as:

$$s^{Poll}_{Clinton-Dole,t} := \frac{r^{Poll}_{Clinton,t} - r^{Poll}_{Dole,t}}{r^{Poll}_{Clinton,t} + r^{Poll}_{Dole,t}}$$

where *s* designates the normalized spread and *r*'s designate the percentages of poll respondents for the indicated candidates at time *t*.[13] Similarly, the IEM market price margins are computed as:

$$s^{VS}_{Clinton-Dole,t} := \frac{p^{VS}_{Clinton,t} - p^{VS}_{Dole,t}}{p^{VS}_{Clinton,t} + p^{VS}_{Dole,t}}$$

where *s* designates the normalized spread and *p*'s designate closing market prices for the indicated candidates at time *t*.[14] In all five graphs, vertical lines indicate significant events (e.g., debates and the start and end of party conventions) and a horizontal line shows the actual election outcome.

Several things are obvious from the five graphs. First, the markets present a very different picture of the elections than the polls. What the polls are measuring as voter sentiment at any particular point in time frequently differs greatly from what the market predicts will actually occur in the election. The market prediction often stays well above or below all contemporaneous polls for extended periods of time. During these periods, the market is typically closer to the final outcome than polls. Second, in each election, we observe the well-known poll phenomenon of "convention bounce" (the tendency for a party to rise in the polls during that party's convention and then fall, see Gelman and King, 1993). These strong effects do not appear in the markets. Third, the market appears to forecast the election outcomes more accurately than polls months in advance.

A fourth observation from the graphs is the striking volatility in polls, both in absolute terms and in comparison to the market. To examine this feature more closely, the graphs of

Figure 2 are repeated in Figure 3 with two differences. First, the plotting symbol used for polls in Figure 3 is a letter indicating the particular polling organization or media outlet. This allows an inspection of the volatility of poll results from a particular organization. Second, the number of polls reported is restricted to avoid graph clutter.[15] The polls chosen for inclusion in Figure 3 are all those conducted by two distinguished polling agencies, Gallup and Harris, and by three broadcast television networks, ABC, CBS and NBC. As is apparent from Figure 3, polls on the same day by different organizations or subsequent polls by the same organizations frequently differ dramatically, generating differences that fall outside the quoted margins of error. There are even cases in which multiple reports by the same polling organization on the same day are noticeably different.[16]

Insert Figure 3 about Here

For a formal comparison of the accuracy of predictions from market prices and polls, we first pair each poll with a set of market prices from the IEM vote-share markets. These market prices are the midnight prices (closing prices) from the last day that the poll was in the field.[17] Note that this choice means traders in the market would not yet have access to the results from that particular poll when they were trading. Next we normalize both the polls and the market prices so that each set of values sums to 1 as we did to create the graphs in Figure 2. Similarly, we normalize the election spread the same way. Then, we compute the average absolute prediction error according to:

$$AAE_t^{Poll} := \frac{\left| s_{Democrat-\text{Republican},t}^{Polls} - v_{Democrat-\text{Republican}}^{Actual} \right|}{2} \text{ and}$$

$$AAE_t^{VSI} := \frac{\left| s_{Democrat-\text{Republican},t}^{VS} - v_{Democrat-\text{Republican}}^{Actual} \right|}{2},$$

where $s^{Polls}_{Democrat-Republican,t}$ is the normalized poll spread as defined above, $s^{VS}_{Democrat-Republican,t}$ is the normalized IEM vote share market spread as defined above and $v^{Actual}_{Democrat-\text{Republican}}$ is the ultimate actual normalized spread for the election.[18] Then, we ask a simple question: which was closer to the actual election outcome, the market or the poll? That is, which has a smaller average absolute error? We use binomial tests to calculate the statistical significance of our results.

Table 3 shows the results of our analysis. For each of several time periods before the elections, Table 3 lists the fractions of times that the markets were strictly closer to the eventual election outcomes than polls (in terms of average absolute prediction error for the two party vote). Ties are broken in favor of polls. We highlight four results from this table.

Insert Table 3 about Here

Result 1: The results from the last five days are similar to prior research on "election eve" forecasts. The market forecasts have a lower absolute prediction error in each election, though not always significantly so. This lack of significance appears to be due to small sample sizes in the last five days - aggregating across elections, the difference becomes significant, with markets closer than polls 68% of the time overall.

Result 2: The markets generally outperform polls over the duration of the markets. In each election except 1988, the market significantly out performed polls overall, coming closer to the eventual outcome from 70% to 87% of the time. Aggregating across all years, the markets were closer to the eventual outcome 74% of the time.

Result 3: Aggregating across all elections, the markets outperformed polls in each time period considered. The advantage of the markets ranged from 68% to 84% depending on the time period. The largest market advantage was in the 66-100 day time range. Interestingly, seven of the ten party conventions occurred during this time period. This accords with the observation from the figures that the markets are less prone to convention bounce than polls.

Result 4: In the longer run, markets perform even better relative to polls. All of the markets, including the smallest market in 1988, significantly outperformed polls that were conducted more than 100 days before the election. If anything, these results suggest that the market improves in relative accuracy the longer the time until the election.

Tables 4 and 5 present robustness checks for these results. The first robustness check is driven by the observation that, since the market aggregates all available information, including recent polls, a fairer comparison may be between the market and some average of recent polls. In Table 4, we present binomial statistics similar to those in Table 3 with one difference: for each poll observation we use a five-poll moving average. Thus, the first poll prediction we use each election year is for the fifth poll and the observation is the average prediction from the first five polls. The next prediction occurs with the release of the sixth poll and the observation is the average of polls two through six, etc. If more than one poll is released in a day, their order is randomized for the purposes of determining the moving average. For each moving average poll prediction, we compute the average absolute error and compare it to the market, generating frequencies of "moving average poll" wins versus market wins, and computing binomial statistics. Again, ties are broken in favor of polls. We put observations into the time period of the last poll in the moving average. Overall, results in Table 4 differ little from results in Table 3.

Insert Table 4 about Here

The second robustness check is driven by the observation that the market gives continuous updates that can always be compared to the most recent poll or, if more than one poll is released on a day, the average of the most recent polls. In Table 5, we present binomial statistics similar to those in Tables 3 and 4 with one difference: each observation is a day. We compare predictions from the closing price in the market each day to the most recent available poll predictions. The first observation is the first day a poll is released. The second is the next day regardless of whether a new poll or multiple polls were released or not. If a new poll is released, that poll is compared to the market. If multiple polls are released, the average prediction from the polls is compared to the market. If no poll was released, results from the prior most recent poll (or average of polls if multiple polls were released on the most recent release day) are compared to the market. For each day, we compute the average absolute error from the most recent poll(s) and compare it to the market, generating frequencies of "most recent poll" wins versus market wins, and computing binomial statistics. Again, ties are broken in favor of polls. Overall, results in Table 5 differ little from results in Tables 3 and 4.

Insert Table 5 about Here

One final robustness check involved the selection of poll types. Crespi (1988) argues that polls using only likely voters are more accurate than those using registered voters or all adults. In selecting polls for the analysis reported in Table 3 above, we included polls of likely voters when that breakdown was included in the poll report. If it was not, we used polls of registered voters if that breakdown was available. If tabulations by neither likely voters nor registered voters were available, we included reports for all adults. If the registered voter and all adult polls are indeed less accurate than those of likely voters, that may have biased the results of Table 3 in favor of the markets. To assess this possibility, we repeated the analysis of Table 3 using only those polls of likely voters. In spite of the reduction in the number of polls from 948 to 536, the results were essentially identical and thus are not reported here. Not one of the "% market wins" entries in the "All years" column of Table 3 fell by more than 2 percent, for example, and more of them increased than fell.

The results above suggest that predictions from markets dominate those from polls about 75% of the time, whether the prediction is made on election eve or several months in advance of the election. To assess the size of the advantage in addition to its frequency, we computed the average absolute error for both polls and markets on each day a poll was released. The mean error for polls across all 964 polls in the sample was 3.37 percentage points, while the corresponding mean error for market predictions was 1.82 percentage points.[19] That advantage persisted for both long term and short term forecasts. Using only those dates more than 100 days prior to the election, the poll error averaged 4.49 percentage points and the market error averaged 2.65 percentage points. Polls conducted within 5 days of the election had an average error of 1.62 percentage points, while the corresponding market prediction error average was 1.11 percentage points.[20]

V. Concluding Remarks

Previous research has shown the absolute and relative accuracy of prediction markets at very short horizons (1 day to 1 week). The evidence we present in this paper shows that the markets are also accurate months in advance and do a markedly better job than polls at these longer horizons. In making our comparisons, we compare unadjusted market prices to unadjusted polls, demonstrating that market prices aggregate data better than simple surveys where results are interpreted using sampling theory. Thus our evidence not only speaks to predicting U.S. Presidential election outcomes, but also offers insight into the likely predictive accuracy of markets in settings where there is not a long history of similar events or a clear model for adjusting survey results.

Obviously, given the success and increasing use of prediction markets, they may attract the interest of political campaigns. How might campaigns put prediction markets to use? Several uses come to mind. We discuss three here.

First, might campaigns want to influence prediction markets in hopes of influencing the future vote? We view this as a scenario that may be attempted, but is unlikely to be successful. It depends on several things that create difficulties. First, the campaign would have to be able to affect market prices. Berg and Rietz (2006) document attempts by campaigns to influence prices, but also cite evidence that suggests this is difficult at best. Individual accounts are limited to $500 and, as a result, are small relative to the market as a whole. Berg and Rietz (2006) show that known deliberate attempts to manipulate prices have little discernable transient effect and no apparent long term effect. One feature of the IEM markets in particular makes manipulation difficult: the unit portfolio method of issuing contracts. A unit portfolio is a set of one of each contract in a market. For example the unit portfolio in the 2004 vote share market consisted of one share each of the Democratic and Republican contracts. Traders create contracts by purchasing a unit portfolio from the exchange and then trading the individual components. With unit portfolios, there are always contracts for all candidates. The payoff for each contract in a unit portfolio will be $1 minus the payoffs of the other contracts in the market and, as a result, the price of each contract should be $1 minus the prices of the other contracts in the market. This means it is not enough to drive one candidate's price up, one must also drive down the prices of all other candidates to effectively manipulate the market. Moreover, beyond believing it can influence prices, the campaign would also have to believe that voters change their votes in response to market prices and do so in a predictable way. We are not aware of any evidence that voters respond to market prices and the direction of any potential response is certainly debatable.[21] Finally, attempts to manipulate prices create profit opportunities for other traders, and exploitation of these opportunities makes manipulation efforts difficult to sustain given account limits -- individual accounts are limited to $500 and, as a result, each trader is small relative to the market as a whole.

Second, might campaigns use prediction markets to assess the campaigns themselves? We view this as quite reasonable. Campaign tactics are designed to change voters' minds and influence their actions. Market participants see the tactics that a campaign is using and will react if they think the tactics influence voters. This is observable by campaigns and can be used as an immediate, low-cost means of assessing the effectiveness of their tactics. Further, campaigns might want to use small, closed markets to test tactics. For example, instead of showing a potential campaign commercial to a focus group, they might show it to a group of market traders. On a small, closed market, these traders could then trade “conditional” contracts that predict the vote share taken by the candidate if the commercial is used versus if the commercial is not used.[22]

Third, might the electorate or parties as a whole use prediction markets to select candidates or policy positions? We view this as potentially quite useful. Berg and Rietz (2003) show how appropriately designed conditional contracts can be used to forecast the relative viabilities of potential candidates. For example, in advance of the primary process, IEM markets indicated that Dole was a relatively weak candidate to run against Clinton in 1996. Markets like these could help primary voters and parties select the strongest candidates. Hanson (1999, 2007) goes so far as to suggest using conditional prediction markets to determine policy. A campaign may want to use this idea to propose policies that are the most likely to achieve particular platform goals.[23]

In this paper, we document that prediction markets are viable election forecasting tools, both in the short run and in the longer run. They outperform the natural alternative, polls, in both cases. Because they react dynamically to information, they can also be used as evaluation tools to assess the impact of decisions such as policy positions, candidate viability, campaign strategies, etc. Current research suggests that these results generalize to other forecasting settings where information is widely dispersed and must be aggregated.

Author Biographies

Joyce E. Berg is an Associate Professor of Accounting and Pioneer Hi-Bred Research Fellow in the Tippie College of Business. She is a director of the Iowa Electronic Markets, using the markets in her research on financial markets, individual decision -making and prediction. Her work has been published in economics, accounting, and electronic markets. She has served as section head for both experimental methods and accounting on the executive committee of the Economic Science Association, the professional organization for experimental economists.

Forrest D. Nelson is a Professor of Economics and Tippie Research Fellow in the Tippie College of Business. He is a director of the Iowa Electronic Markets and one of three originators of the Iowa Political Stock Market.

Thomas A. Rietz is an Associate Professor of Finance and Hershberger Faculty Research Fellow in the Tippie College of Business. He is a director of the Iowa Electronic Markets and uses it in his research on financial markets, individual decisions and political science. His publications span economics, finance and political science. He served three terms as the section head for finance on the executive committee of the Economic Science Association, the professional organization for experimental economists.

References

Berg, J.E., Forsythe, R. Nelson, F.D. & Rietz, T.A. (2003). Results from a Dozen Years of Election Futures Markets Research. Forthcoming in C.R. Plott and V.L. Smith (Eds.), *Handbook of Experimental Economic Results*, Amsterdam: North Holland.

Berg, J.E., & Rietz, T.A. (2003). Prediction Markets as Decision Support Systems, Information *Systems Frontiers*, 5, 79-93.

Berg, J.E., & Rietz, T.A. (2006). The Iowa Electronic Market: Stylized Facts and Open Issues" In *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, R.W. Hahn & P.C. Tetlock (Eds.), Washington DC: AEI-Brookings Joint Center, pp. 142-169.

Campbell, J.E. (2000). Polls and Votes–The Trial-Heat Presidential Election Forecasting Model, Certainty, and Political Campaigns. In J.E. Campbell & J.C. Garand (Eds.), *Before the Vote: Forecasting American National Elections*, Thousand Oaks, CA: Sage Publications, Inc. pp.17-46.

Campbell J.E. (2005). Introduction—Assessments of the 2004 Presidential Vote Forecasts. *PS-Political Science & Politics*, 38, 23.

Crespi, I. (1988) *Pre-election Polling: Sources of Accuracy and Error*, New York: Russell Sage Foundation.

Cuzán, A.G., Armstrong, J.S., & Jones, R. (2005). Combining Methods To Forecast The 2004 Presidential Election. Working paper prepared for presentation at the Southern Political Science Association, New Orleans, January 6, 2005.

Erikson, R.S. & Wlezien, C. (2007). Are Political Markets Really Superior to Polls as Election Predictors? *Public Opinion Quarterly*, forthcoming.

Fair, R.C. (1978). The Effect of Economic Events on Votes for President. *The Review of Economics and Statistics*, 60, 159-173.

Forsythe, R., Nelson, F.D., Neumann, G.R. & Wright J. (1992). Anatomy of an Experimental Political Stock Market, *American Economic Review*, 82, 1142-1161.

Forsythe, R., Rietz T.A. & Ross, T.W. (1999). Wishes, Expectations and Actions: Price Formation in Election Stock Markets, *Journal of Economic Behavior and Organization*, 39, 83-110.

Gelman, A., & King, G. (1993). Why are American Presidential-Election Campaign Polls so Variable When Votes are so Predictable? *British Journal of Political Science*, 23, 409-451.

Hayek, F.A. (1945). The Use of Knowledge in Society, *The American Economic Review*, 35, 519-530.

Hahn, R.W. & Tetlock, P.C. (2006). Introduction to Information Markets. In *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, R.W. Hahn & P.C. Tetlock (Eds.), Washington DC: AEI-Brookings Joint Center, pp. 1-12.

Hanson R.D. (1999) Decision Markets, *IEEE Intelligent Systems*, May/June, 16–19.

Hanson R.D. (2007) Shall We Vote on Values, but Bet on Beliefs? Forthcoming *Journal of Political* Economy.

Hicks, J.R., (1946). *Value and Capital*, 2nd ed., London: Oxford University Press.

Jones, R.J., Armstrong, J.S., & Cuzán, A.G. (2007). Forecasting elections using expert surveys: an application to U.S. Presidential elections. University Library of Munich, Germany MPRA Working Paper No. 5301, prepared for presentation at the annual meeting of the American Political Science Association, Chicago, August 30, 2007

Keynes, J.M. (1930). *Treatise on Money*, 2nd ed., London: MacMillan.

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

Krueger, J.T., and Kuttner, K.N. (1996). The Fed Funds Futures Rate as a Predictor of Federal Reserve Policy," *Journal of Futures Markets*, 16, 865-879.

Lewis-Beck M.S., & Tien, C. (2007). Jobs and the Job of President: Forecast for 2004, *PS: Political Science and Politics*, 37, No.4, 753-758.

Marsh, C. (1984) Do Polls Affect What People Think? In C.F. Turner & E. Martin (Eds.), *Surveying Subjective Phenomenon, Vol. 2*, New York: Russell Sage, pp.565-591.

Mosteller, F., Hyman, H., McCarthy, P.J., Marks, E.S., Truman, D.B., Doob, L.W., MacRae, Jr., D., Stephan, F.F., Stouffer, S.A. & Wilks, S.S. (1949). *The Pre-Election Polls of 1948: Report to the Committee on Analysis of Pre- Election Polls and Forecasts*, Social Science Research Council.

Norpoth, H. (1996). Of Time and Candidates—A Forecast for 1996, *American Politics Quarterly*, 24, 443-467.

Oliven, K., & 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, 336-351.

Panagakis, N. (1997). Incumbent Races: A National Perspective, *Public Perspective*, January, 21-23.

Perry, P. (1979). Certain Problems with Election Survey Methodology, *Public Opinion Quarterly*, 43, 312-325.

Plott, C.R. (2000). Markets as Information Gathering Tools, *Southern Economic Journal*, 67, 1-15.

Simon, H.A. (1954). Bandwagon and Underdog Effects and the Possibility of Election Predictions, *The Public Opinion Quarterly*, 18, 245-253.

Wlezien, C. & Erikson, R.S. (2000). Temporal Horizons and Presidential Election Forecasts. In J.E. Campbell & J.C. Garand (Eds.) Before *the Vote: Forecasting American National Elections*, Thousand Oaks, CA: Sage Publications, Inc., pp, 103-118.

Wolfers J. & Zitzewitz E. (2004) Prediction Markets, *Journal of Economic Perspectives*, 18, 107-126

Footnotes

[1] These markets are the longest running prediction markets known to us.

[2] There are many methods of adjusting raw poll data to arrive at “revised” predictions. For examples, see Crespi (1988), Panagakis (1997), Campbell (2000), and Erikson and Wlezien (2007).

[3]Debate over the ability of futures markets to forecast future prices extends back to Keynes (1930) and Hicks (1946). Many of the arguments result from the secondary nature of information aggregation in these markets. The early “normal backwardization” versus “contago effect” arguments were based on relative power of speculators and hedgers. Today, the idea that “risk neutral” probabilities used to price futures and options differ from the “true” underlying probabilities results from relative levels of hedging demand in the markets. While the IEM markets discussed below could be subject to price deviations due to hedging activities, the narrow scope of the IEM markets, the small size of investments and analysis of individual traders (e.g., Forsythe, Nelson, Neumann and Wright, 1992, and Forsythe, Rietz and Ross, 1999) all lead us to conclude that hedging activities do not affect IEM prices significantly.

[4] In 1988, the contracts paid the vote share times $2.50.

[5] This can be shown in a variety of ways. See Malinvaud (1974) for the general equilibrium proof. One can also price the contracts as assets using the Capital Asset Pricing Model (CAPM) and/or Arbitrage Pricing Theory (APT) model found in the Finance literature. In each, $P_t = E(P_{t+h})/(1+k)^h$, where k (the required expected return) is the sum of the risk free rate and compensation for aggregate risk factors, t is the current date and $t+h$ represents any given future date, up to and including the election date. Since the risk free rate is zero and there are no aggregate risk factors (while individual contract payoffs depend on vote shares received by the candidates, the overall sum of payoffs in an IEM prediction market does not), the expected return on any given asset is zero. That is, k=0, which implies that $P_t = E(P_{t+h})$. Alternatively,

given that the expected market portfolio return is constrained to be zero by design, any factor risk premiums must be zero. Again, this makes for a zero expected return on any given asset. As a result, $P_t = E(P_{t+h})/(1+k)^h = E(P_{t+h})$. Even though traders cannot make the appropriate risk free hedges here (because they cannot trade the underlying fundamental asset), one might be tempted to use the modern portfolio theory futures pricing relationship: $F_{t+h} = E(P_{t+h})x(1+r_f)^t/(1+k)^t$, where F_{t+h} is the time t futures price for delivery at date $t+h$, $E(P_{t+h})$ is the expected future spot price of the underlying fundamental, r_f is the risk free rate and k is the required expected return determined by the risk of the futures position. Again, both the risk free rate and the required expected return are constrained to be zero. This gives: $F_{t+h} = E(P_{t+h})$.

[6] Beginning in 1992 and onward, traders could participate in both vote-share and winner-takes-all markets. Here, we discuss the vote share markets only.

[7] Interestingly, polls did not always ask "if the election were held today…?" Crespi (1988) reports that prior to 1940 pollsters asked who the participant would vote for in the election. Poll designers believed the move to "if the election were held today" would result in more willingness to express preferences. Gelman and King (1993) document that this wording change makes a significant difference in the number of respondents choosing "undecided," but does not affect the relative proportion of respondents choosing Democrat and Republican.

[8] In 2004, the IEM ran dual races between Bush (as the Republican Nominee) and a set of possible democratic nominees. We only report the results for the Bush/Kerry race here.

[9] In 1992, Ross Perot represented a major third party candidate. The IEM predicted his vote share very accurately (e.g., the election eve absolute error in predicting his vote share was 0.3 percentage points). Were we to include this in the analysis, the relative accuracy of the IEM would increase. Omitting it gives an advantage to polls We omit it so that 1992 is directly comparable to other election years where a two-way vote splits are evaluated.

[10] Since Crespi (1988) shows that polls of likely voters are more accurate than polls of registered voters, this selection criterion is advantageous for polls.

[11] All poll data that we collected is available from the authors upon request. This data will also be available at the IEM website (www.biz.uiowa.edu/iem) but is not yet posted.

[12] Note that the IEM payoffs are based on the popular vote, not the Electoral College vote. Polls also predict the popular vote.

[13] Notice that this is equivalent to taking the difference in normalized vote shares, where normalization consists of dividing each share by the sum of the shares for the two candidates.

[14] Closing prices are the last trade price before midnight each day. If no trade occurs in a day, the previous day's closing prices are carried over.

[15] Particularly with recent elections, the number of polls is so great that a graph of all polls results in an indistinguishable cloud near Election Day.

[16] This might result, for example, from reporting responses to two different questions, one in which the respondent is prompted with a list of candidates and another in which no prompts are provided.

[17] We do this to compare what would be predicted from market prices to what would be predicted from polls. The poll prediction could not be made until after all poll results were collected. So, we use a comparable market price – the market price from the last day the poll was in the field. IEM prices are recorded at midnight each day, so our prices are the midnight prices.

[18] This is equivalent to adding up (across the parties) the absolute difference between the actual and predicted normalized vote shares and dividing by two.

[19] Note that this is a matched sample average with polls as the unit of analysis. IEM prices between poll releases are not included. If more than one poll is released in a day, the IEM

prices from that day are included twice in the average. This corresponds to the way binomial statistics are computed in Table 3. Different averaging techniques leave the results unaffected.

[20] The attentive reader will notice that this market figure differs slightly from the figure quoted in the introduction. The prior figure is a 5 day average of the market closing prices. The figure here is the average of the matched sample, that is, one market observation for every poll released during the last five days. See footnote 19.

[21] Voter response to polls has been debated since Simon (1954) outlined bandwagon effects (where voters are more likely to vote for candidates with strong numbers) and underdog effects (where voters are more likely to vote for candidates with weak numbers). Yet, there remains no consensus about the direction of any reaction to polls (see, for example, Marsh, 1984), much less markets. For instance, it is as easy to imagine voters will make an extra effort to support a candidate whose price is low as they will to jump on a front-runner's bandwagon. Of course, it is also easy to imagine other markets where a decision maker's response is predictable (for example, a market where the price will determine a policy and the decision maker follows a known, market based rule). This may make manipulation more viable.

[22] Such conditional contracts would be easy to design. Suppose a campaign wanted to assess the effectiveness of "Commercial A" on "Candidate X." The following four contracts could be issued in a unit portfolio: VSX|A has a payoff equal to Candidate X's vote share IF commercial A is run and 0 otherwise. VSY|A has a payoff equal to the other candidate's vote share IF commercial A is run and 0 otherwise. VSX|NA has a payoff equal to Candidate X's vote share IF commercial A is NOT run and 0 otherwise. VSY|NA has a payoff equal to the other candidates' vote share IF commercial A is NOT run and 0 otherwise. The four contracts span the state space and will have a total liquidation value of $1. The difference in price between VSX|A and VSX|NA tells the campaign the predicted impact of running the commercial.

[23] Again, conditional contracts would be easy to design. Suppose a campaign wanted to determine whether “Policy A” or “Policy B” was more likely to achieve platform “Goal X.” For simplicity, suppose that either Policy A or Policy B will be implemented (but not both, nor neither). The following four contracts could be issued in a unit portfolio: WTAX|A has a payoff equal to $1 if Goal X is achieved after implementing Policy A and $0 otherwise. WTANX|A has a payoff equal to $1 if Goal X is NOT achieved after implementing Policy A and $0 otherwise. WTAX|B has a payoff equal to $1 if Goal X is achieved after implementing Policy B and $0 otherwise. WTANX|B has a payoff equal to $1 if Goal X is NOT achieved after implementing Policy B and $0 otherwise. The four contracts span the state space and in total will liquidate at $1. The difference in price between WTAX|A and WTAX|B tells the campaign the predicted efficacy of Policy A versus Policy B. Of course, if neither, both, or other policies may be implemented, the contracts will need to be adjusted accordingly.

Figures

Figure 1: Election Eve Forecast Vote Shares and Actual Outcomes

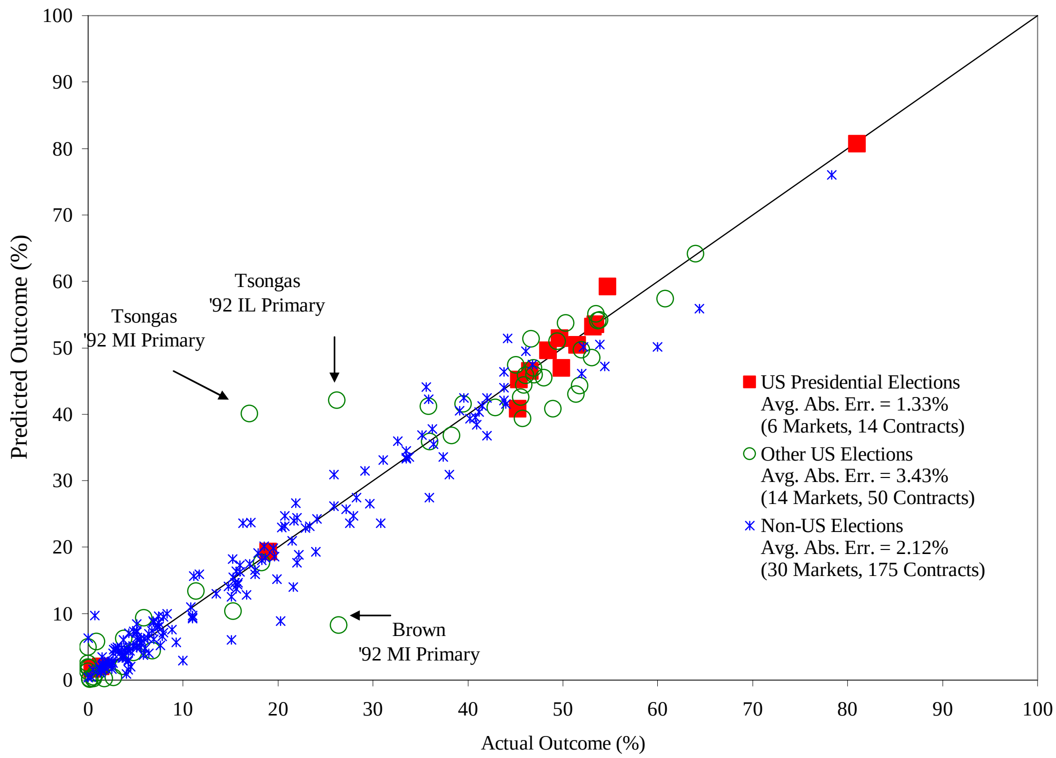

Figure 2: Implied Vote Share Margins for Market and Polls

Vertical axis is vote margin; horizontal axis is date; the margin implied by the market is the solid moving line; poll margins are represented by small circles; horizontal line at election outcome, vertical lines at beginning and end of conventions, debate days and election days

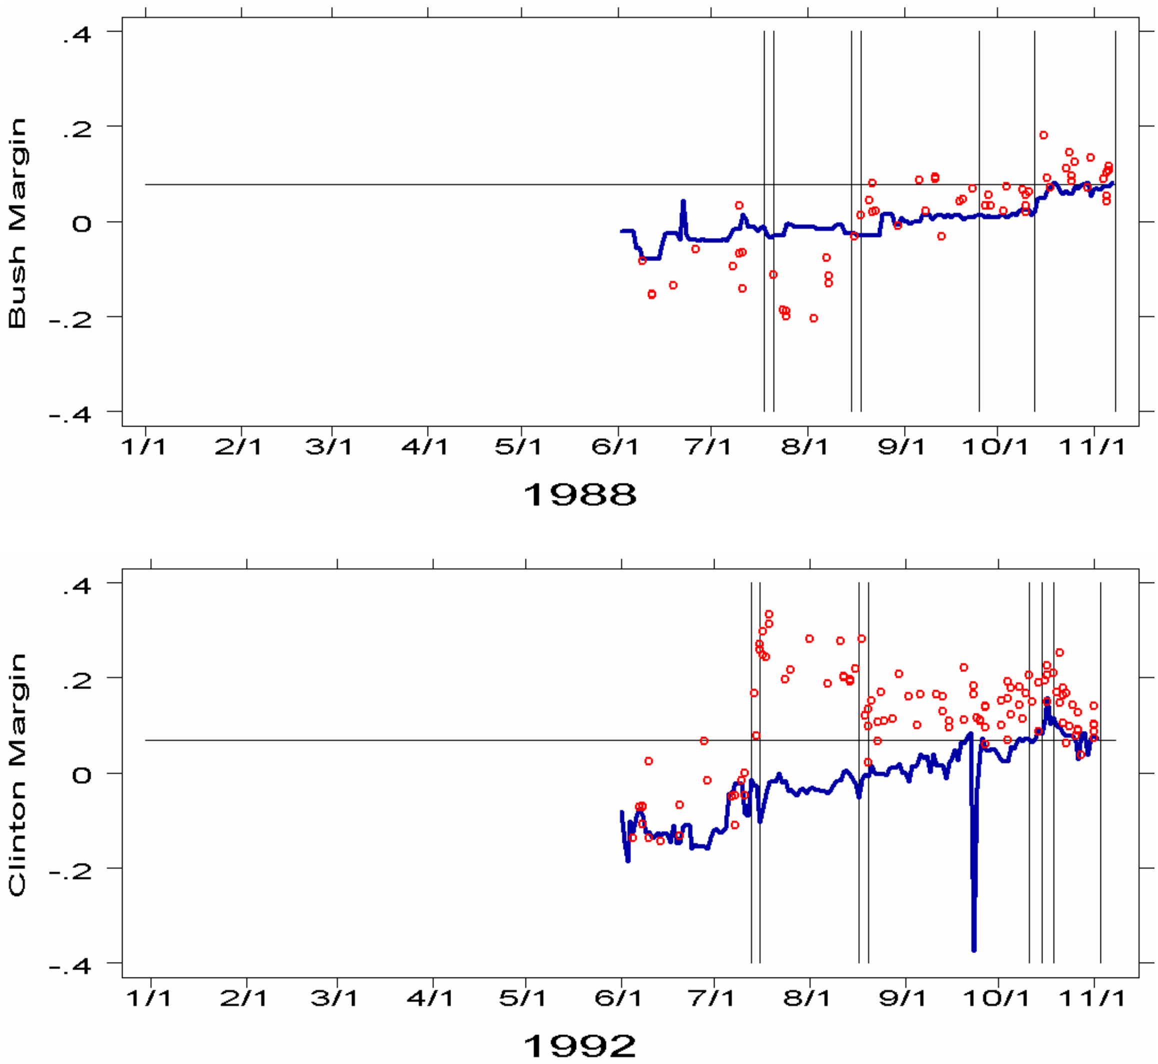

Figure 2 (continued)

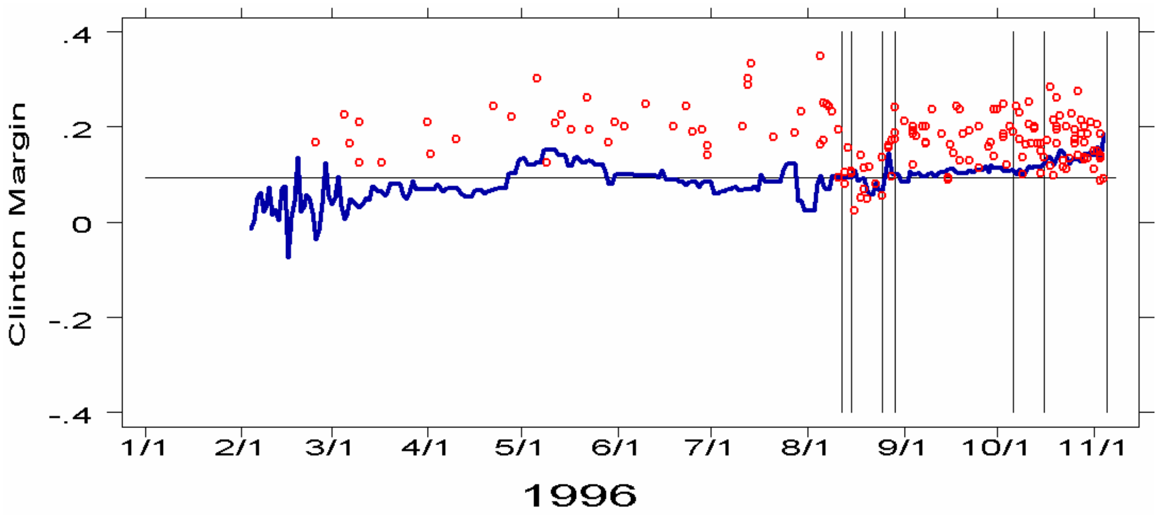

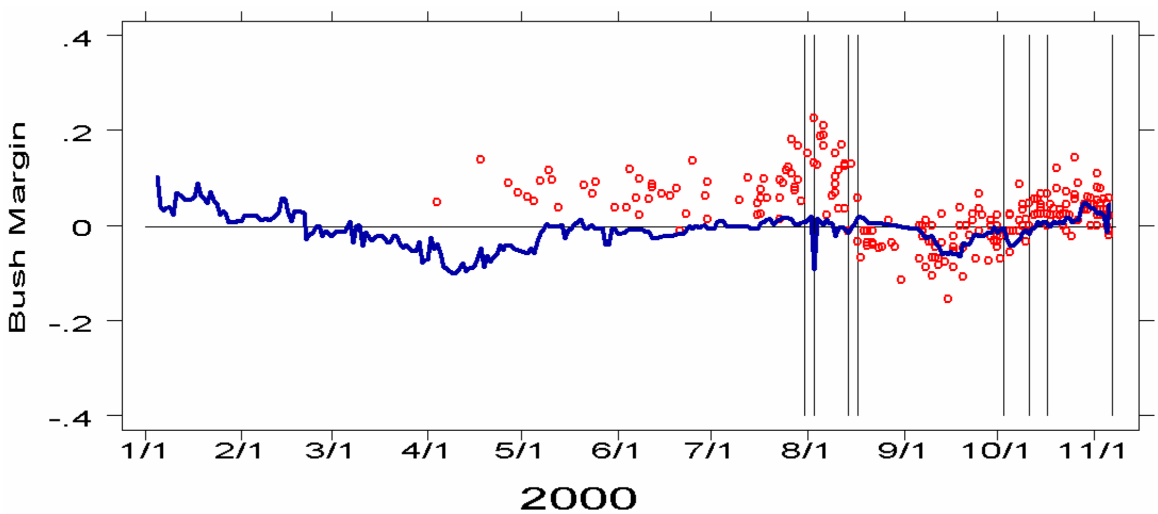

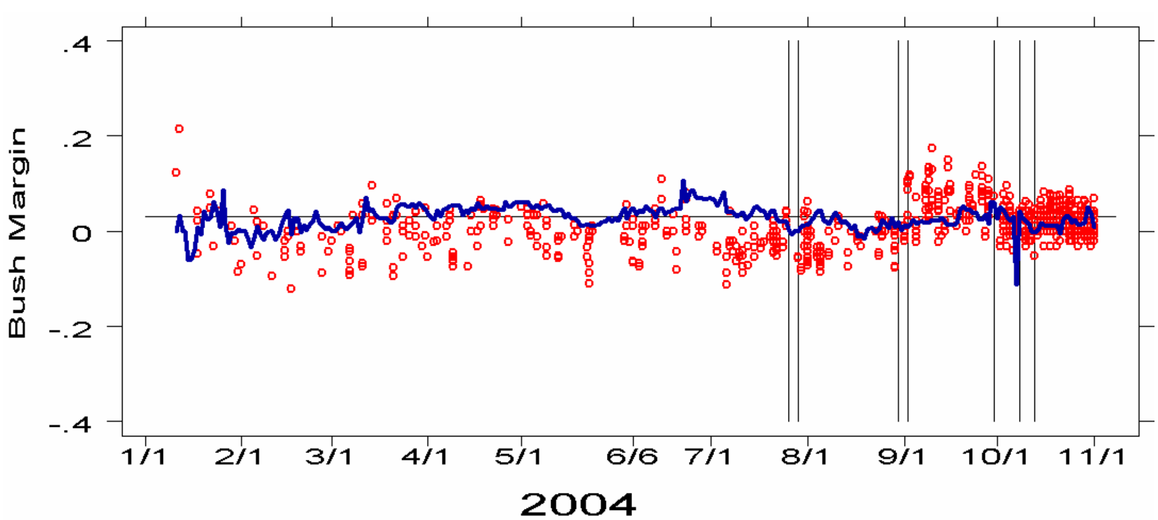

Figure 3: Implied Vote Share Margins for Market and Selected Polls

Vertical axis is vote margin; horizontal axis is date; the margin implied by the market is the solid moving line; poll margins are represented by letters (A=ABC, C=CBS, G=Gallup, H=Harris, N=NBC); horizontal line at election outcome, vertical lines at beginning and end of conventions, debate days and election day.

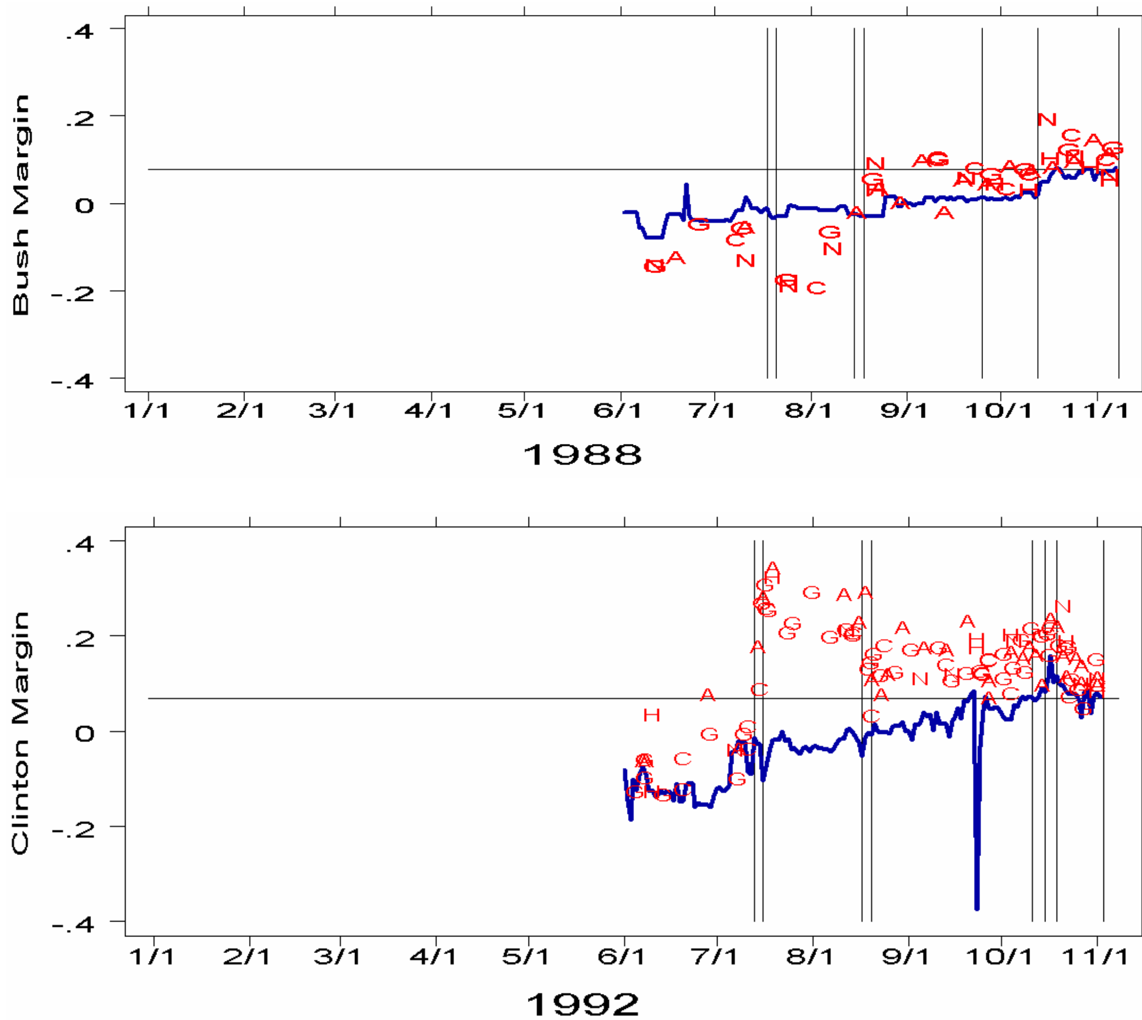

Figure 3: (continued)

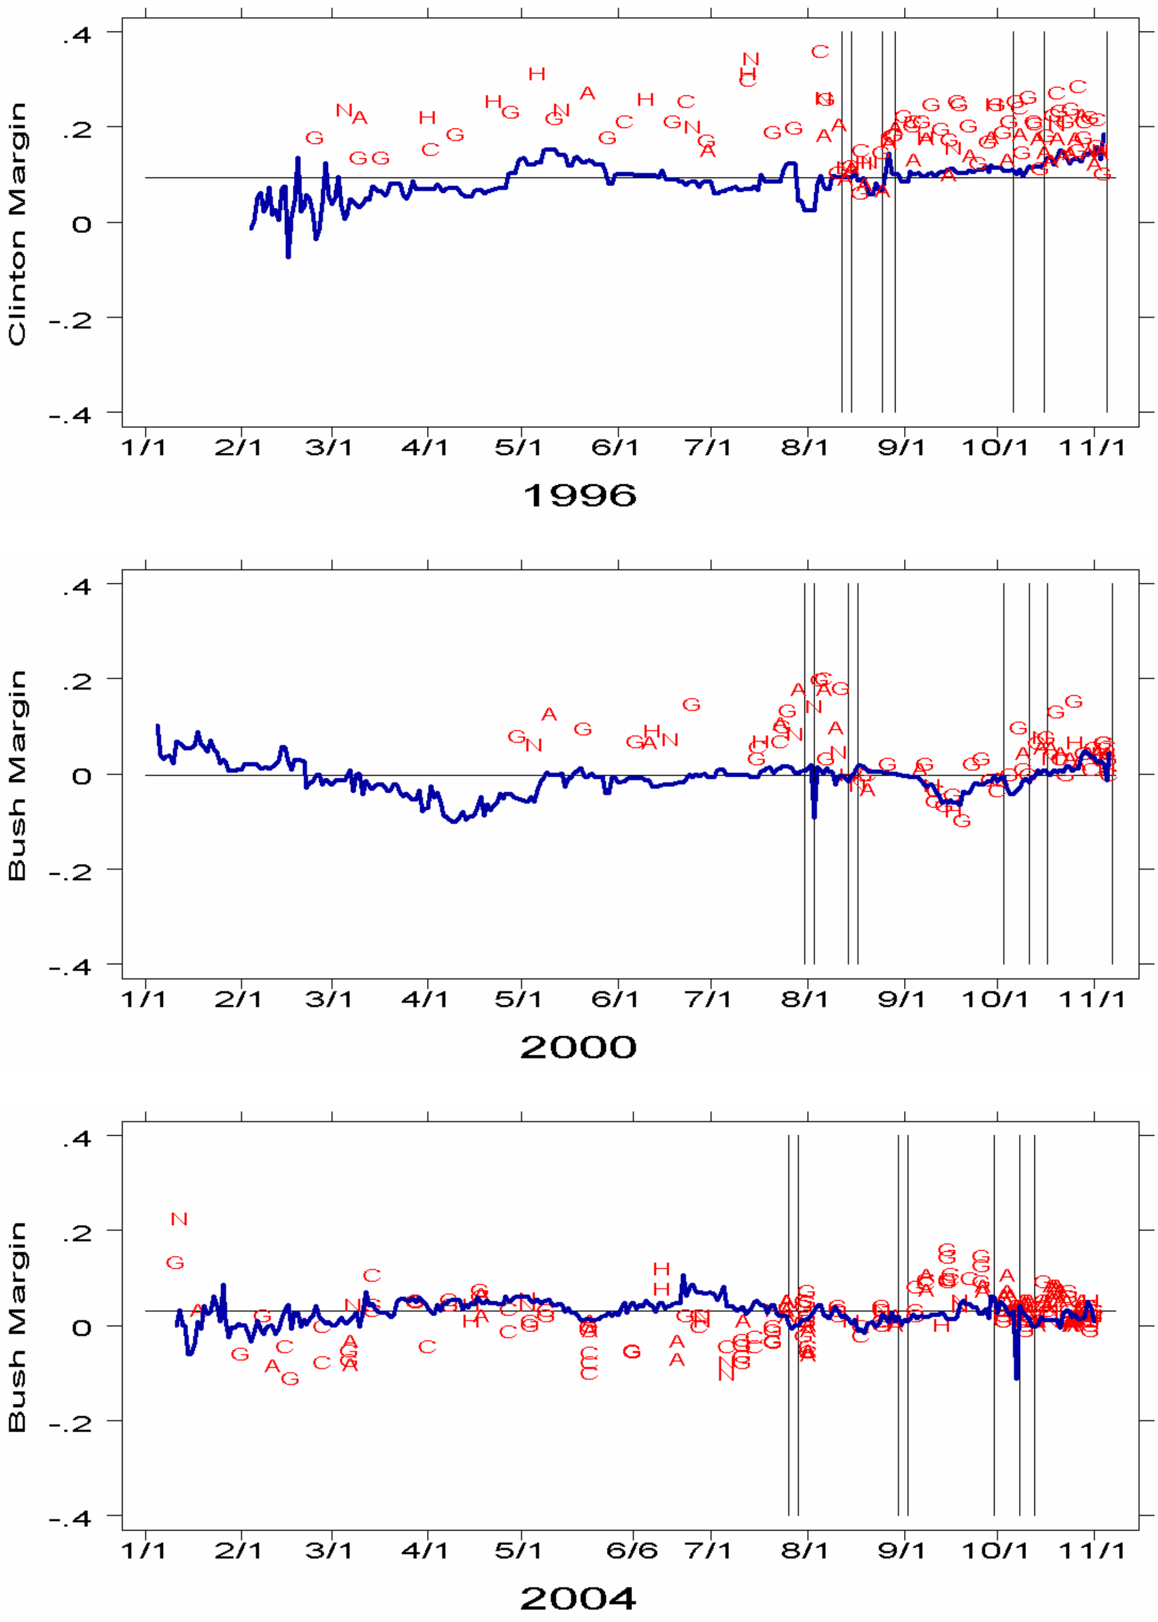

Tables

Table 1: Contracts Traded in IEM
Presidential Vote-Share Election Markets

Year	Final Contracts Traded[1]	Number of Weeks Open Prior to Election
1988	Bush Dukakis Jackson Rest-of-field	23
1992	Democrat Republican ---------- Perot Democrat and Republican	43
1996	Democrat Republican	39
2000	Democrat Reform Republican	45
2004	Democrat Republican	37

[1]In 2004, contracts traded in a number of potential democratic nominees versus Bush. Here, we report only on the Bush/Kerry race.

Table 2: Summary of IEM Presidential Vote-Share Election Market Activity

	1988	1992[1]	1996[1]	2000[1]	2004[1]
Market Statistics					
Opening Date	6/1/1988	1/10/1992	2/4/1996	1/5/2000	2/21/2003
Election Date	11/8/1988	11/3/1992	11/5/1996	11/7/2000	11/2/2004
Weeks Open	23	43	39	44	45
Trader Investments	$4,976	$79,356	$200,000[2]	$210,633.00	$355,281.00
Overall Market Activity					
No. of Active Traders	155	592	264	790	777
Contract Volume	15,826	78,007	23,093	46,820	339,222
Dollar Volume	$8,123	$21,445	$3,628	$13,694	$46,237
Activity in Last Week					
No. of Active Traders	54	114	41	104	124
Contract Volume	962	1,389	592	4,192	4,947
Dollar Volume	$1,924	$569	$312	$609	$2,476

[1]Traders may have been active in multiple markets. Investments were fungible across markets.
[2]Estimated.

Table 3: Binomial Tests for Relative Accuracy of the
Market And Contemporaneous Poll Predictions for Specific Time Ranges
Each Poll is an Observation

Poll predictions come from major polls taken during the election and are the normalized two-party vote shares. The market predictions are the normalized two-party vote share market prices on the last day each poll was in the field collecting data. The binomial variable takes the value 1 if the market prediction is (strictly) closer the actual election outcome and 0 otherwise. Each p-value is the exact binomial probability of a number of 1s that large or larger, given that number of trails and a hypothesized probability of 0.50.) The number of observations is the number of polls in the sample period. If multiple polls are released on the same day, the same market price is compared to each poll.

Days included in sample	Item	1988	1992	1996	2000	2004	All years
All (from the beginning of the market	Number of polls	59	151	157	229	368	964
	poll “wins”	25	43	21	56	110	255
	market “wins”	34	108	136	173	258	709
	% market wins	**58%**	**72%**	**87%**	**76%**	**70%**	**74%**
	p-value (1sided)	0.149	0.000	0.000	0.000	0.000	0.000
More than 100 Days	Number of polls	14	69	33	49	195	360
	poll “wins”	1	20	3	2	66	92
	market “wins”	13	49	30	47	129	268
	% market wins	**93%**	**71%**	**91%**	**96%**	**66%**	**74%**
	p-value (1sided)	0.001	0.000	0.000	0.000	0.000	0.000
66-100 Days	Number of polls	11	20	33	39	28	131
	poll “wins”	5	8	3	2	3	21
	market “wins”	6	12	30	37	25	110
	% market wins	**55%**	**60%**	**91%**	**95%**	**89%**	**84**%
	p-value (1sided)	0.500	0.252	0.000	0.000	0.000	0.000
32-65 Days	Number of polls	13	22	33	57	48	173
	poll “wins”	12	8	2	26	8	56
	market “wins”	1	14	31	31	40	117
	% market wins	**8%**	**64%**	**94%**	**54%**	**83%**	**68%**
	p-value (1sided)	1.000	0.143	0.000	0.298	0.000	0.000
6-31 Days	Number of polls	15	34	47	59	67	222
	poll “wins”	7	6	9	18	21	61
	market “wins”	8	28	38	41	46	161
	% market wins	**53%**	**82%**	**81%**	**69%**	**69%**	**73%**
	p-value (1sided)	0.500	0.000	0.000	0.002	0.002	0.000
Last 5 Days	Number of polls	6	6	11	25	30	78
	poll “wins”	0	1	4	8	12	25
	market “wins”	6	5	7	17	18	53
	% market wins	**100%**	**83%**	**64%**	**68%**	**60%**	**68%**
	p-value (1sided)	0.016	0.109	0.274	0.054	0.181	0.001

Table 4: Binomial Tests for Relative Accuracy of the
Market and Five-Poll Moving Average Poll Predictions
for Specific Time Ranges
Each Poll Release is an Observation

Poll predictions come from major polls taken during the election and are the normalized two-party vote shares. The market predictions are the normalized two-party vote share market prices on the last day each poll was in the field collecting data. The binomial variable takes the value 1 if the market prediction is (strictly) closer the actual election outcome than a five-poll moving average of polls and 0 otherwise. Each p-value is the exact binomial probability of a number of 1s that large or larger, given that number of trails and a hypothesized probability of 0.50.) The number of observations is the number of polls in the sample period, beginning with the fifth poll overall. If multiple polls are released on the same day, the same market price is compared to each poll and the poll order is randomized to determine the moving average.

Days included in sample	Item	1988	1992	1996	2000	2004	All years
All (from the beginning of the market	Number of Obs.	55	147	153	225	368	948
	MA poll "wins"	26	39	9	53	144	271
	market "wins"	29	108	144	172	224	677
	% market wins	52%	73%	94%	76%	61%	71%
	p-value (1sided)	0.394	0.000	0.000	0.000	0.000	0.000
More than 100 Days	Number of Obs.	10	65	29	45	195	344
	MA poll "wins"	0	22	0	0	80	102
	market "wins"	10	43	29	45	115	242
	% market wins	100%	66%	100%	100%	59%	70%
	p-value (1sided)	0.001	0.006	0.000	0.000	0.007	0.000
66-100 Days	Number of Obs.	11	20	33	39	28	131
	MA poll "wins"	4	7	8	1	0	20
	market "wins"	7	13	25	38	28	111
	% market wins	64%	65%	76%	97%	100%	85%
	p-value (1sided)	0.274	0.132	0.002	0.000	0.000	0.000
32-65 Days	Number of Obs.	13	22	33	57	48	173
	MA poll "wins"	13	5	0	31	7	56
	market "wins"	0	17	33	26	41	117
	% market wins	0.00%	77%	100%	46%	85%	68%
	p-value (1sided)	1.000	0.008	0.000	0.786	0.000	0.000
6-31 Days	Number of Obs.	15	34	47	59	67	222
	MA poll "wins"	8	4	0	16	39	67
	market "wins"	7	30	47	43	28	155
	% market wins	47%	88%	100%	73%	42%	70%
	p-value (1sided)	0.696	0.000	0.000	0.000	0.929	0.000
Last 5 Days	Number of Obs.	6	6	11	25	30	78
	MA poll "wins"	1	1	1	5	18	26
	market "wins"	5	5	10	20	12	52
	% market wins	83%	83%	91%	80%	40%	67%
	p-value (1sided)	0.109	0.109	0.006	0.002	0.900	0.002

Table 5 Binomial Tests for Relative Accuracy of the Market and Most Recent Poll Predictions for Specific Time Ranges Each Day is an Observation

Poll predictions come from major polls taken during the election and are the normalized two-party vote shares. The market predictions are the normalized two-party vote share market prices on each day. The binomial variable takes the value 1 if the market prediction is (strictly) closer the actual election outcome than the most recent poll and 0 otherwise. Each p-value is the exact binomial probability of a number of 1s that large or larger, given that number of trails and a hypothesized probability of 0.50.) The number of observations is the number of days in the sample period. If no poll is released on a particular day, the poll results from the most recent day are used. If multiple polls are released on the same day, they are averaged before comparing to the market.

Days included in sample	Item	1988	1992	1996	2000	2004	All years
All (from the beginning of the market	Number of Obs.	152	228	254	214	616	1,464
	MA poll “wins”	50	61	18	37	194	360
	market “wins”	102	167	236	177	422	1,104
	% market wins	67%	73%	93%	83%	69%	75%
	p-value (1sided)	0.000	0.000	0.000	0.000	0.000	0.000
More than 100 Days	Number of Obs.	52	128	154	114	516	964
	MA poll “wins”	0	45	8	16	165	234
	market “wins”	52	83	146	98	351	730
	% market wins	100%	65%	95%	86%	68%	76%
	p-value (1sided)	0.000	0.000	0.000	0.000	0.000	0.000
66-100 Days	Number of Obs.	35	35	35	35	35	175
	MA poll “wins”	12	8	6	3	9	38
	market “wins”	23	27	29	32	26	137
	% market wins	66%	77%	83%	91%	74%	78%
	p-value (1sided)	0.045	0.001	0.000	0.000	0.003	0.000
32-65 Days	Number of Obs.	35	34	34	34	34	170
	MA poll “wins”	26	7	1	11	3	48
	market “wins”	8	27	33	23	31	122
	% market wins	24%	79%	97%	68%	91%	72%
	p-value (1sided)	1.000	0.000	0.000	0.029	0.000	0.000
6-31 Days	Number of Obs.	26	26	26	26	26	130
	MA poll “wins”	12	1	1	6	13	33
	market “wins”	14	25	25	20	13	97
	% market wins	54%	96%	96%	77%	50%	75%
	p-value (1sided)	0.423	0.000	0.000	0.005	0.577	0.000
Last 5 Days	Number of Obs.	5	5	5	5	5	25
	MA poll “wins”	0	0	2	1	4	7
	market “wins”	5	5	3	4	1	18
	% market wins	100%	100%	60%	80%	20%	72%
	p-value (1sided)	0.031	0.031	0.500	0.188	0.968	0.022