

Are Political Markets Really Superior to Polls as Election Predictors?*

Public Opinion Quarterly, forthcoming

Robert S. Erikson
Columbia University
rse14@columbia.edu

Christopher Wlezien
Temple University
wlezien@temple.edu

Abstract

Election markets have been praised for their ability to forecast election outcomes, and to forecast better than trial-heat polls. This paper challenges that optimistic assessment of election markets, based on an analysis of Iowa Electronic Market (IEM) data from presidential elections between 1988 and 2004. We argue that it is inappropriate to naively compare market forecasts of an election outcome with exact poll results on the day prices are recorded, that is, market prices reflect forecasts of what will happen on Election Day whereas trial-heat polls register preferences on the day of the poll. We then show that when poll leads are properly discounted, poll-based forecasts outperform vote-share market prices. Moreover, we show that win-projections based on the polls dominate prices from winner-take-all markets. Traders in these markets generally see more uncertainty ahead in the campaign than the polling numbers warrant—in effect, they overestimate the role of election campaigns. Reasons for the performance of the IEM election markets are considered in concluding sections.

* Previous versions were presented at the Annual Meetings of the World Association of Public Opinion Research, Cannes, 2005, the American Political Science Association, Chicago, 2005, the Midwest Political Science Association, Chicago, 2006, the American Association for Public Opinion Research, Montreal, May, 2006, and the Elections, and the Public Opinion, and Parties (EPOP) Group of the Political Studies Association, Bristol, 2007, and the First World Meeting of the Public Choice Society, Amsterdam, 2007. For helpful comments, we thank seminar participants at Dartmouth College, Harvard University, University of Leiden, New York University, Northwestern University, Princeton University, Washington University and Yale University, and also Joseph Bafumi, Neal Beck, Tereza Capelos, Dennis Chong, James Druckman, Sean Gailmard, Andrew Gelman, Jeff Jones, George Krause, Michael Laver, David Lowery, Peter Miller, Rebecca Morton, Jonathan Nagler, Benjamin Page, Thomas Palfrey, Mark Pickup, Colin Provost, Michael Sobel, and the three anonymous referees. The research was supported by National Science Foundation Grants SBR-9731308 and SBR-0112856.

Election markets have recently emerged as an intriguing new tool for predicting elections. These markets—made possible by the Internet—now present the possibility that electoral trends can be discerned well in advance by simply consulting the candidates’ latest market prices. At least, that is a popular belief.

Modern political markets originated with the 1988 launching of the Iowa Electronic Market (IEM) at the University of Iowa. The first IEM market was a “vote-share” market for the 1988 Bush-Dukakis presidential contest, in which Internet traders electronically bought and sold futures contracts based on their forecasts of the candidates’ actual vote percentages. Since 1992, IEM offers both a “vote-share” market and a higher-volume “winner-take-all” presidential market, in which payoffs go to contracts on the popular vote winner.¹ Along the way, IEM has offered occasional markets on non-presidential races. Most recently, commercial political markets on the Internet have entered the field. Most notable is the Tradesport.com winner-take-all market on the 2004 US presidential election. Unlike IEM, which has a \$500 limit on individual investments, Tradesport.com has no limit on the amount invested, making it the thicker and arguably more efficient market.²

Election markets have drawn considerable favor in both the popular and academic press as an alternative to public opinion polls as a method of predicting elections. As James

¹ In the IEM vote-share markets, one share of a candidate pays off in proportion to the candidate’s final vote share. For instance, one unit of a candidate who obtains 44 percent of the vote is worth 44 cents. A portfolio of one unit of each candidate pays exactly one dollar. A candidate’s unit price therefore represents the market’s expectation of the final vote. If a trader buys a candidate at, say, 40 cents per unit, and the candidate wins the 44 percent as in our example above, the profit is 4 cents on the dollar. If our trader buys at 40 and sells at 50, the profit is 10 cents. In the IEM winner-take-all market, one share of a candidate pays off one dollar if the candidate wins and nothing if the candidate loses. A portfolio of one unit of each candidate pays exactly one dollar. A trader who buys one unit of a candidate at, say 40 cents on the dollar, wins either one dollar (a 60 cent profit) or nothing (a 40 cent loss) if the contract is held until market closing following the election. If our trader buys at 40 cents and sells at, say 60, the profit is 20 cents. For further details, consult the IEM website, <http://www.biz.uiowa.edu/iem/>.

² These “modern” election markets were not the first. Before the development of scientific polling, high-volume Wall Street election markets were an important means of gauging election trends. See Rohde and Strumpf (2004).

Surowiecki (2004: pp. 35-36) popularizes the argument in *The Wisdom of Crowds*, IEM traders’ “predictions of what the voters of the country will do are better than the predictions you get when you ask the voters themselves what they are going to do.” Across a wide spectrum of academia, one finds this view repeated—as if it has now entered the domain of common knowledge—that the daily prices in the election markets dominate public opinion polls in terms of forecast accuracy. From economists Justin Wolfers and Eric Zitzewitz (2004: p. 112), we learn that the IEM presidential election market has “outperformed large-scale polling organizations.” Law professor Cass Sunstein (2005: p. 1030) echoes this sentiment: that the IEM markets “have produced extraordinarily accurate judgments...far better than professional polling organizations.” Political scientists have also begun to see election markets as superior to polls. Gregory Caldeira (2004: p. 779) puts this view to print, asserting that IEM prices “are amazingly stable and close to the final outcome, in contrast to polls, which bounce around, by day.”³

The theory of market superiority is seductive: Trial-heat polls are distorted both by their inherent sampling error and their transient reactivity to short-term stimuli that expire before Election Day. They at best capture preferences on the day of the poll, i.e., “if the election were held today.” In theory, disinterested investors in election markets, while certainly incorporating contemporary opinion trends, are capable of discounting short-term shifts in the polls, such as convention bounces, that dissipate by Election Day. Moreover, in theory, election traders hold knowledge about the future that allows them to anticipate subtle electoral forces in advance of their actual impact on public opinion. For instance, a common belief

³ One readily finds the idea of election markets superiority to the polls in popular magazines. Consider the headlines, “Why Vote When You Can Bet?” 2007. *Slate* <<http://specials.slate.com/futures/2008/>> or “Punters v. pollsters: Are betting markets a better guide to election results than the polls?” *The Economist*, April 14 2005 or “The ‘Election Futures’ Market: More Accurate than the Polls? As The U. of Iowa Goes, So Goes the Nation?” *Business Week*, November 11, 1996.

within political science is that electoral shifts from the early polls to Election Day are readily predictable from observable variables. (See Gelman and King, 1992, for the classic statement.) Variables such as presidential approval and economic conditions, which are available information to political market investors, help to predict the presidential vote beyond what trial heat polls show.⁴ The seeming success of election markets at anticipating election outcomes even served as inspiration for expanding the realm of political information markets to predict phenomena outside the electoral realm. The idea of using markets for predicting terrorism and other international political events—while provoking public outrage from politicians—remains the subject of serious discussion in academic circles (Wolfers and Zitzewitz 2004; Meirowitz and Tucker 2004).

Of course, believers in election markets do not draw their enthusiasm solely from theory.⁵ They can and do also cite the available empirical evidence from studies that show that markets do in fact predict better than the polls. This evidentiary trail leads back to the organizers of the Iowa Electronic Market themselves. In a series of papers, Berg, Forsythe, Nelson and Rietz (Berg, Nelson, and Rietz 2003; Berg and Rietz 2006; Berg, Forsythe, Nelson, and Rietz N.d) show that daily prices contain only half the forecast error of the daily polls. Indeed, three days out of four, a poll will be less accurate than the vote-share market price at predicting the election outcome. Someone who played the vote-share market based on the expectation that the division in the latest polls would translate one-to-one into the final vote division would lose decisively in the long run.

⁴ See Campbell (1996; 2000), Holbrook (1996), and Wlezien and Erikson (2004).

⁵ For a formalization of the idea that election markets are superior to the polls, see Kou and Sobel (2004). For a discussion of election markets in the context of information markets more generally, see Wolfers and Zitzewitz (N.d). On behavioral anomalies that hinder the efficiency of markets, see Thaler (1991).

The substance of the IEM authors' test of the market versus the polls is accurate and not in dispute. The market price is superior to a naïve reading of the polls. For instance, if the incumbent leads 60-40 in the polls in May while the market says the incumbent will win with 55 percent, the market price is likely to be closer to the Election Day vote division. But this is not the appropriate test. Whereas market prices reflect forecasts of what will happen on Election Day, trial-heat *polls register preferences on the day of the poll*. It is well-known that vote divisions in polls on any given day in advance of the election do not directly translate into the final vote outcome. Specifically, the hypothetical 60-40 lead in May is likely to fade over the course of the campaign (see Campbell 2000; Wlezien and Erikson 2002). A proper comparison of market prices and preference polls would thus ask: Based on an assessment of the historical record of the polls, what would be the expected November vote division, say, given a 60-40 incumbent lead in May, and does that offer a superior or inferior prediction compared to the May vote-share prices? Moreover, a thorough test of market superiority would also include an evaluation of the higher-volume winner-take-all market where the idea is to pick winners instead of point spreads. We could ask, for instance, what an analysis of polling history would show to be the odds of the incumbent winning in November given a 60-40 lead in May, and whether this prediction based on polls offers greater certainty than the May winner-take-all price.

This paper offers these further tests of the relative accuracy of the IEM presidential markets versus presidential election polls.⁶ Our results put the polls in a much more favorable light than the claims of market enthusiasts. Based on our analysis, an investor with a *modest*

⁶ Although IEM has conducted markets on other elections besides US presidential elections, testing must be limited to presidential markets because they are the only markets where the daily prices can be compared to a density of parallel polling data.

knowledge of how trial-heat polls translate into Election Day outcomes would reap handsome profits from the IEM presidential market. The implication is that where candidate market prices depart from where the polls project that they should be, these deviations contain more noise than signal.

Methodology: An Overview

We apply two tests of the IEM presidential markets versus trial-heat polls as electoral predictors. First, we apply a new test to the vote share market. Since market prices reflect forecasts of what will happen on Election Day and trial-heat polls only register preferences on the day of the poll, it is inappropriate to naively compare them on any given day in advance of an election. Accordingly, we transform raw poll vote divisions into projections of the Election Day outcome and compare these projections to vote-share prices—in effect, putting them on equal footing.⁷ We find that these daily poll projections are superior to IEM prices. In three of the five presidential elections with IEM vote share markets, poll projections are more accurate than market prices. In four of five elections (with one tie), the week’s average poll projection dominates the daily market price.

Secondly, we assess the polls versus market prices in IEM’s thicker winner-take-all markets. For this test, we start by converting our vote projections into probabilities of incumbent party victory, based on the projected vote share outcome and the days to the election. Then, we compare the incumbent win probabilities with the win prices in the IEM market to see which ones are closer to the actual outcomes. Here our test shows the polls systematically dominating market prices in all four elections with IEM winner-take-all

⁷ This is why we also do not discount market prices, which constitute forecasts by definition. There are other reasons as well, most notably that prices in theory already reflect any systematic discounting (Kau and Sobel 2004). On the possibility of discounting market prices, see note 12 ahead.

contests. The implication is that markets are slower to recognize election winners than what can be learned by applying a reasonable understanding of polling history to the interpretation of current polls.

Performing the test of the vote-margin market requires empirical estimates of how—given the number of days before the election—raw poll results translate into expectations of the actual vote division. And the test of the winner-takes-all market requires empirically-based conversions of how expected vote margins translate into probabilities of incumbent victory and defeat. For these tasks, we use a data set of virtually all national presidential trial-heat polls conducted since 1952.

Estimating the projected vote from poll results works as follows. For all days within 200 days of each election, starting with 1952, we record the two-party vote division in the latest trial poll. Where there is more than one poll ending on a specific date, we pool the polls ending on that date. On dates for which we have no polls ending, we use the most recent poll from preceding days. Then for each day before the election (-1 to -200), we regress the actual vote margin on the latest polls. The predictions from these 200 equations provide the vote projections.

In making these vote projections, we use only the historical data that would be available to observers of that election. Thus, to estimate the daily vote projections for each year 1988-2004, the regression equations incorporate only observations through the preceding presidential year. For instance, the 1988 equations incorporate only information from polls 1952-1984 while the 2004 equations are based on polls through 2000.

The generic vote projection equation is:

$$V_y = \alpha_t + \beta_t P_{yt} + e_{yt}, \quad \text{Equation 1}$$

where V_y = the actual incumbent-party percent of the two-party vote in year y minus 50, and P_{yt} = the corresponding trial heat poll division minus 50 in year y on day t of the campaign.

Separate equations are drawn for each t from 1 to 200 days before the election for years 1952- y_{Y-1} where y_{Y-1} is the presidential year preceding the election to be predicted in year Y . For interpretive convenience, the vote division is measured as deviations from a tied 50-50 vote.

Given equation 1 based on electoral history, we compute the projected vote. This poll-based forecast for V_Y , the actual vote in year Y , from the current polls in year Y at date T is:

$$\hat{V}_Y = \alpha_T + \beta_T P_{YT} \quad \text{Equation 2}$$

If α were zero and β were unity, the projection \hat{V} would be identical to the raw two-party vote division of the polls. But, as students of political campaigns know, early leads fade (Campbell 2000; Wlezien and Erikson 2002). The daily β estimates are thus all below 1.0. As we will see, the in-party also sometimes matters, as the α 's are positive early in the campaign, when the incumbent party's poll numbers tend to underestimate Election Day support. (The daily regression results are displayed in Figure 2 ahead.)

The vote projection equations can be used not only to obtain an expectation of the vote but also the variance around that expectation. The estimated variance in the error term, or σ_T^2 , can be used to estimate the daily forecast errors predicting V_Y in year Y from the out-of-sample P_{TY} in year Y and date T :

$$\phi_T = \sqrt{\sigma_T^2 + \text{Var}(\beta_T)(P_{YT} - \bar{P}_T)} \quad \text{Equation 3}$$

Knowing the forecast error we estimate the cumulative normal density Φ_{YT} at zero (50-50 split), that is, the probability of an incumbent victory in year Y based on the polls at time T .

Thus, for each date during the campaign we have two poll-based projections—the projection of the Election Day vote and the projection of the probability that the incumbent party candidate will win. Because our poll data are reported in terms of the beginning and end of the polling period rather than their release dates, we lag the polls’ projections two days when comparing them to market prices.⁸ Market prices are the daily closing prices in the vote-share and winner-take-all polls. The only complication in determining market prices is that where there are separate markets for the two major-party candidates, we ignore any third options (e.g., Perot) in determining the relative market prices. For instance, if the winner-take-all market prices for a day in 1992 are 0.30 for Bush, 0.60 for Clinton, and 0.10 for Perot, we would ignore Perot and treat the net price as $0.30/0.90 = 0.33$ for Bush. The measure thus is directly comparable to our (two-party) measures of the vote and the polls.

Armed with market prices and poll-based projections, our task is to compare the accuracy of each. For each date of five campaigns with an active IEM market, we compare prices with projections in terms of their match with the election outcome. We start with the vote-share market, comparing the accuracy of vote-share prices for the incumbent candidate with the accuracy of the projected vote, using polls lagged two days from the final date of polling. We then turn to the winner-take-all market, comparing the accuracy of winner-take-all prices for the incumbent party candidate with the accuracy of the projected probability of an incumbent party win, based on our analysis of polls ending two days earlier.

The IEM Vote Share Markets

We start with the IEM vote share market. Just as the IEM organizers claim, daily prices in this market more accurately reflect the final election outcomes than do the raw poll

⁸ To be clear, we treat each poll as being released two days after the final date in the field. Thus for any date, the poll results are based on polls that left the field at least two days earlier.

divisions. Figure 1 shows this. For each of five elections, the market prices are persistently closer to the actual outcome than the raw division in the latest polls. The mean absolute error for market prices was a mere 2.41 percentage points compared to 4.46 for the division in the latest polls for the same date. The market's error was the smaller of the two 75 percent of the time.⁹ Table 1, which shows the yearly details, makes clear that the market advantage was decisive over all five elections.

-- Figure 1 and Table 1 about here --

The reason for this seeming victory of markets over the polls is that the markets discounted the size of leads in the polls, especially when the lead was held by the out-party, as in 1988 and 2000. This is exactly how poll results should be interpreted. Poll leads tend to dissipate and early in the campaign, polls inflate the support for the out-party. We will show that when we adjust for these factors, the polls outperform the vote-share market as an election predictor.

Our task is to project the vote based on the incumbent party's share of the two-party vote in the latest poll and the date of the poll. Following equation 1 from above, we do this by regressing the actual incumbent vote on the poll results for each date of the campaign, using data only from previous election years. With 200 dates to cover (from 1 to 200 days before the election), this is 200 x 5 or 1000 equations. We have separate sets of 200 equations for each election because each new election expands the moving wall of prior information. For 1988, the information set is all polls 1952-1984. By 2004, the set expands to all polls 1952-2000.¹⁰

⁹ Note that when weekly averages of polls are used, the advantage is only slightly less dramatic, as market prices win 73 percent of the time, so sampling error makes little difference.

¹⁰ These data were originally collected for a larger project analyzing campaigns and the polls (Wlezien and Erikson 2002). For each date, all polls ending on that date are pooled to form a poll of polls. The polls on each date are averaged, weighting by the number of respondents with a major party preference. Note that we do not attempt to adjust for design or house effects (Traugott 2001).

-- Figure 2 about here --

In practical terms, these different equations for different years are very similar. Figure 2 shows the changing parameters of the regression equation from 200 days to 1 day before the election, as information for projecting the 2004 election. The first two panels present the constant term and the regression coefficient for the incumbent party percent of the two-party “vote” in the latest polls, all for years 1952-2000. Notice that the regression coefficient starts below 0.4 and ends up at about 0.8, with the polls gaining weight over time. The intercept is generally positive, especially during the summer, as a corrective for the poll’s tendency to deflate the incumbent party’s support.¹¹ Late in the campaign, however, the intercept turns slightly but not significantly negative. (To interpret the intercept, recall that both the vote and the poll support for the incumbent party are measured as a deviation from 50 percent.) The third panel shows the root mean squared error, which reflects the error in predicting the vote from the polls. The error declines over the course of the campaign as the polls become increasingly informative about the outcome.

Figure 3 compares the market prices with the poll projections where the equations predicting the vote from date specific polls in past elections are used to project the vote. For instance, Figure 2’s parameters are plugged into the 2004 polls to project the 2004 vote on a daily basis from the daily polls. As with the comparison of the market with the raw polls, the projections are generated using polls lagged two days from their final date in the field.

-- Figure 3 and Table 2 about here --

When we use the daily projected vote from the polls rather than the daily raw poll predictions, the market’s seeming advantage for forecasting elections vanishes. Indeed, in three of the five elections, the advantage goes to the projected polls. In 55 percent of the cases,

¹¹ In effect, the challenger gets a boost in the polls that does not stand the test of time.

the projected poll results are closer to the actual vote than the market price for that date.

Whereas the market price's mean absolute error is 2.41 percentage points, the mean error of the projected vote is a slimmer 2.13. Table 2 presents the details.¹²

We can pursue this matter further. It should be evident that our measure of the projected vote, while more sophisticated than the raw division of the poll results, is cruder than it has to be. This is because it is based solely on the most recent polls with no consideration of polling history. One could remark that *despite* the crudity of the measure, the projected vote generally is superior to the market prices. The dominance of the polls comes into sharper view when we substitute a more sophisticated index—the weekly average of vote projections—that smooths out the occasional wiggles in the polls, much of which is survey error (Wlezien and Erikson 2002).

Figure 4 compares daily market prices to the week's average poll projection of the Election Day vote. On 63 percent of the dates, the market price was farther from the actual vote than was the seven-day average of the projected polls. The mean 7-day average of the projected polls is closer to the outcome than the daily market prices in 4 of the 5 elections, losing only in 2000 by a fraction of a point. Whereas the market is off by 2.4 points on

¹² It might seem that our test is unfair in that we discount the magnitudes of leads in the polls but not the markets. The difference, as we have noted, is that the poll division represents the hypothetical vote if held on the specific date while the market price represents informed opinion on the specific date about the actual vote on Election Day. Still, we can consider how one might discount market prices in a fashion analogous to our discounting of poll leads. We clearly cannot use the same technique of relying solely on the prior history of market prices in each year, as there is no history to discount 1988 prices, only one year in 1992, and so on. However it has been suggested that we can apply out-of-sample techniques. For instance, to discount daily 1988 prices, for each campaign date one uses the four observations from 1992-2004 to generate a regression equation predicting the vote, and then inserts the 1988 market price for each day into each daily equation. We computed these discounted prices for the five years and compared them with the discounted poll margins in terms of closeness to the vote. For four of the five years and 64 percent of the dates, the projected vote margin predicts better than the projected prices. We also could create parity in the projection methodology by switching to new poll projections constructed in using out-of-sample techniques, e.g., by discounting each date's polls in 1988 from the 1992-2004 data rather than the 1952-1984 data. This comparison also favors the polls. The new poll projections win in four of the five years and for 62 percent of the dates. In short, the markets perform worse compared to the discounted polls when the market prices are themselves discounted.

average, the 7-day average of the poll projection is off on average by only 1.7 point, for about a 30 percent improvement over the market. One might argue that a fair test would average market prices as well as poll projections. However, if mean market prices provide an improved signal to investors, investors can take further advantage by exploiting short-term variation around the moving average. Regardless, *using averaged market prices has no effect on the relative dominance of poll projections*. With each variable measured as a seven-day average, the vote projection beats the market price 62 percent of the time.

-- Figure 4 about here --

We can glean some details from the yearly graphs. In 1988, the market correctly ignored the convention-season Dukakis bubble (when he famously led Vice-President Bush in one poll by 17 percent) but generally acknowledged Bush's lead once it became evident in the polls. In 1992, the market overestimated Bush's strength until the very last days of the campaign. In 1996 the market was more bullish on Clinton's vote margin early and late in the campaign but more accurate than the poll projections for the middle dates of the 1996 summer. The 2000 election saw the market's best performance, as it ignored Gore's brief lead in the polls but then briefly and inexplicably surged to Gore just as Gore was fading in the polls. In 2004, while the polls consistently showed a close election with a slight Bush edge, the market tended to overvalue Bush in the spring but undervalue him in the summer. By the fall, both the polls and the market had the race pretty much correct.

In theory, the market's posited advantage over the poll-based projections should be greatest early in the campaign, when there is a longer time interval until the election for trends to develop that the market can anticipate that contemporary polls cannot. Figure 5 shows the relationship between vote-share prices and seven-day vote projections by time in the campaign,

merging observations for different elections. The first panel shows that early in the campaign—more than 90 days before the election—a considerable scatter of vote-share prices around the seven-day average of vote projections. This is as if, early in the campaign, the market anticipates the nature of future shocks to the vote that contemporary polls are not equipped to show. Meanwhile, as the third panel shows, late in the campaign—within 30 days of the election—the share prices begin to fall in line with vote projections, as if the market tracks the polls now that there are few shocks to come. Do the market's departures from the poll projections early in the campaign signal that the market sees future events that the polls cannot see?

-- Figures 5 and 6 about here --

Markets actually perform most poorly relative to projections using polls in the early stages of the campaign. This is clear in Figure 6, which displays the average error of vote-share prices and vote projections over the timeline of the campaign. Early on, market prices respond to information that is—judging by the actual vote—often quite wrong. Only toward the end of the campaign do the markets catch up. That they do catch up makes considerable sense, because toward the campaign's end market actors can efficiently absorb the wealth of poll information without being distracted by other possible electoral forces. As Election Day approaches, there is little new for the market to anticipate, correctly or incorrectly.

What do we learn about the vote-share market? Election markets have been touted as augurs because they are capable of taking into account information beyond the electorate's preferences at the moment as reflected in the polls. In other words, they supposedly incorporate not only current sentiment but also future shifts in sentiment that can be anticipated by knowledgeable observers. We have learned, however, that prices in the IEM vote share

market are no better, and in fact a bit worse, at predicting the vote than are projections based on the day's most recent polls. Where the market holds to a view of the election at odds with the poll projections, it is somewhat more likely to be wrong than right.

As a final comment on the vote share market, let us consider the profit one could make from knowing that informed projections from the polls are better estimates of the final vote spread than the market's prices. Suppose for instance, one set up a robotic trading program to always buy vote-share stock in the candidate whose poll-based prospects appear better than the market price. (Suppose also that this market were thick enough that one's trading actions did not affect market prices.) Let us say that every day the market is open, one buys one unit of the candidate who our poll analysis suggests is underpriced. Alas, the profit rate would be only about 1.4 percent of investment. The meagerness of the profit is due to the conservative nature of the vote-share market. On average, one purchases a unit of a candidate at a price of about 50 cents (a market expectation of winning half the votes). Considering the average edge of poll projections over the markets of 0.7 points, the expected return for a \$0.50 investment is about \$0.507, for a somewhat meager 1.4 percent profit. Although the rewards from our market strategy would accrue slowly, they should be quite steady. Based on a *t*-statistic of over 11 on the weekly average of the projected poll's net advantage over the market price—the difference in absolute errors—the probability (less than 0.0001) of a net loss from our poll-based program trading strategy is infinitesimally low.¹³

Winner-take-all markets

With candidate shares either paying off at full value or no value on Election Day, the

¹³ While a *p*-value of 1 in 10,000 certainly passes the threshold of statistical significance, we should not pretend that the separate daily observations comprise independent events. Our inferences are based on multiple observations over only five elections.

winner-take-all market offers both greater risk and reward than the vote-share market. The trader's challenge is also more difficult in the winner-take-all market: Rather than wagering based on an expectation of the final vote, a trader must also consider the variance around this expectation in order to assess the probabilities of a Democratic and Republican victory. Our vote projections include a measure of the variance around the expectation, as shown in the third panel of Figure 2. Based on the degree of fit between the vote and the projections from each date's regression, one can calculate the forecast error around the point predictions from the poll projections. In turn, one can estimate the probability that the candidate's margin is greater than zero—that is, 50 percent—and thus wins the popular vote, as per equation 3, above (see the earlier discussion).

Figure 7 shows the winner-take-all prices and the daily projections of the probability of an incumbent party victory. The more erratic of the two series is the poll projection, as the bounces in this series reflect the shifting electoral odds that would be inferred by always projecting from the latest poll readings. Even with these bounces, however, projections from the polls are more accurate than the winner-take-all market prices fully 87% of the time.¹⁴ For the 1992 election, a good showing is a score near zero—a low probability that incumbent Bush 41 would win. For the other three elections, a good showing is a score near one—a high probability of an incumbent party popular vote victory. In all four elections (recall there was no winner-take-all market in 1988), the poll-based winner projections dominate the market winner-take all prices.

-- Figure 7 about here --

The general explanation for the poll projections' dominance over the market prices is that the polls were ahead of the market in terms of projecting the winner. In 1992, once

¹⁴ Using weekly averages the poll win-projections do even better, winning over 90 percent of the time.

Clinton led in the polls, the polls projected a more certain win than the market foresaw. In 1996, the polls projected a Clinton victory with a near certainty that the market price did not match until very late in the campaign. In 2000, both the final poll projections and the late market prices failed to see the Gore “win” in the popular vote. During much of the campaign, the poll projection was much more favorable to Gore’s chances than was the market, allowing the polls to score a technical victory in 2000. For much of 2004, the polls persistently translated Bush 43’s typical slight lead as an incumbent into a far greater likelihood of victory than the market projected.

The market’s sluggish response seems due to its assigning a wide variance to the expected vote. Put another way, in addition to the vote-share market being somewhat behind in estimating the point spread based on information of the moment (see the discussion of the vote-share market), the winner-take-all market overestimates the degree to which unexpected events can overtake the market projection of the point spread. In effect, the market greatly overvalues long-shots.

Figure 8 demonstrates this. The top three panels show the daily projected probabilities of incumbent-party victories as a function of the daily projected vote, merging years but sorting by days before the election. These panels reveal the typical S-shaped curves depicting a probability from an expectation. The fuzziness is due to different prediction equations for different years and dates. As one goes from early polls to late across the three top panels, we see that the S-curve becomes steeper, reflecting the fact that polls are more predictive over time.

-- Figure 8 about here --

Meanwhile, the middle panels show the winner-take-all prices as a function of vote-share prices for the comparable time periods. The curves are well-behaved in the sense that for a given time-interval, the market's winner-take-all probability assessment is predictable from its vote-share expectation. Note, however, that the slopes of the market curves are less steep than their comparable poll-based curves. This is exactly what would happen if the market puts a wide variance around its vote-share assessment. The market's degree of uncertainty about its own vote-share expectations is greater than the uncertainty about the poll projections. It is as if the market overestimates the degree to which unanticipated campaign shocks will upend expectations between then and Election Day.

It follows that the slope of the winner-take-all prices on the poll-based projections should be even less steep than the slope of the winner-take-all prices on the vote-share prices. This pattern is borne out in the bottom panel. Especially in the early days of the campaign, when the polls project the verdict to be one-sided, the market is quite uncertain, with variation in the poll projections having little bearing on winner-take-all prices.

Figure 9 presents the one punch-line of this exercise. It displays the winner-take-all prices as a function of our objective probabilities based on the vote projections. For the early panel, based on data 91 days or more before the election, the prices bear virtually no relation to the probabilities that can be projected from the polls. Indeed, the market cautiously puts the odds at about 50-50 no matter how certain is the poll projection of the likely outcome. By the middle dates (31-90 days before the election), the responsiveness of the market to the poll probability projections improves slightly. Within 30 days of the election, the market moves almost all the way to the poll-based probabilities.

-- Figure 9 and Table 3 about here --

How decisive is this victory of the poll projections over the winner-take-all markets? A trader who would have picked the undervalued candidate in the market according to the poll projections would have made the winning choice an astounding 87 percent of the time. As Table 3 shows, the result is a rout in each year. Overall, a trader who had bought one unit of the undervalued candidate (according to the poll projections) every day would have reaped a 15 percent profit on investment.

Discussion: Where Election Markets Go Wrong (and Right)

Why has the Iowa vote-share market failed to do what by theory it is supposed to do—utilize information beyond the snapshot of current sentiment as revealed in the polls—to forecast elections accurately? In theory, market prices incorporate the signal of information about the future course of voter preferences beyond what we would predict from current polls. In practice, the market’s reading of the signal contains considerable noise. That much we know because market prices perform more poorly than our poll projections. But can we find evidence of the signal embedded in the noisy market prices? We ask, in other words, do market prices actually contain information about the future election outcome beyond what is evident in the polls?

To fix ideas, let us consider a different, parallel, kind of vote share market—an “over-under” market regarding whether the incumbent party vote will be greater or less than a specific vote share posted by an election handicapper. In our hypothetical market, the handicapper selects a posted vote share based on our poll projection—the final vote share predicted by the latest polls. The gambler then is offered an even money wager whether the incumbent party vote will be higher or lower than the posted vote share.

We ask, given the setup of this hypothetical game, would the gambler gain an advantage by wagering based on the IEM market price—“over” if the price is higher than the poll projection and “under” if the price is lower than the vote projection? Suppose markets contain no information beyond what the polls contain—as if market fluctuations respond only to polls plus irrelevant variables. Then, using the market as the auger will lead to a winning bet 50 percent of the time. But if markets contain positive information beyond what is in the polls, then betting based on the market leads to a profit. The final possibility is that the market contains negative information—as if market traders make systematic mistakes. Then betting based on the market results in loss.

If the posted over-under vote share is based on the latest polls, betting in the direction of the market price leads to a win 58 percent of the time. Much of the reason, however, is that the market discounts the occasional outlier poll that is riddled with sampling error or capturing a very short term phenomenon. If our handicapper sets the posted over-under vote share based on the seven-day average of the poll projection, the edge from following the market is much less clear: following the market’s deviation from the posted vote share leads to a win only 52 percent of the time.¹⁵

Interestingly, the degree to which the market price deviates in the correct direction from the projected vote varies by year. As Table 4 shows, the market’s negative performances were concentrated in the early years of the IEM market—1988 and 1992. In 1992, for instance, the

¹⁵ It could be the case that even if market prices are limited in predicting the final vote, that they are sensitive to short-term movement in voter preferences. For instance, it could be that the differential between market prices and the projected vote is a good predictor of changing preferences as reflected in the following week’s polls. However, the (seven-day average) polls’ movement over the next seven days is in the direction of the market no more than 50.1 percent of the time. Over the following two-week period, the polls move in the direction of the market only 46.3 percent of the time.

market waited for a Republican trend that never arrived.¹⁶ In 1996, however, the market bet that Clinton's vote would be less than the poll projections, which turned out to be correct following a late "collapse" of the Clinton vote. In the close 2000 and 2004 elections, whenever the poll projection wandered far from 50-50, the market price would typically be (correctly) more in the direction of an even split. Of course these patterns do not negate the greater accuracy of poll projections than market prices. When prices depart from the poll projection in the direction of the actual vote, they often overshoot to the extent that they are farther from the vote outcome than the poll projection. This is shown clearly in Table 4.

-- Table 4 about here --

We can pursue a similar analysis involving a hypothetical parallel winner-take-all election market where the election handicapper offers odds in accord not with the IEM winner-take-all prices but rather with our poll-based projections of the probabilities of an incumbent-party win and an out-party win. The new question is: Can one profit in this market by using the information in the IEM winner-take-all market? If the IEM winner-take-all prices are positively informative, the gambler would profit by always wagering on the candidate given more favorable odds in the IEM market than the handicapper offers.

In this case we already know part of the answer. From the previous section, we know that the winner-take-all projected win probability beats the IEM winner-take-all price 87 percent of the time, so our program gambling based on IEM winner-take-all prices wins money 13 percent of the time.¹⁷ This is not quite as bad as it seems, however, because our IEM-based

¹⁶ The market's systematic error in 1992 might be explained by savvy traders incorporating their knowledge that Ray Fair's economic model was forecasting a sizeable Bush victory. Through 1988, Fair's model had an uncannily accurate track record (Fair 1990). For 1992 Fair predicted Bush would defeat Clinton with 56% of the two-party vote, an error of 11 percentage points. Fair (1996) presents the post-mortem with model revisions.

¹⁷ Since the winner-take-all outcome is an incumbent-party loss or win (0 or 1), the handicapper's posted probability and the IEM price both must be in between ($0 > p < 1$). As a result, when the winner-take-all poll

wagerer would pick mainly long-shots, risking on average only 24 cents for a chance to win a dollar. Still, winning 13 cents for every 24 cents invested means retaining barely 50 percent of one's investment. Clearly in a betting market where poll projections determine the odds, one cannot profit by using the information found in IEM winner-take-all prices. These prices contain information that distorts more than informs regarding the forecast of the election winner beyond what we can tell from the polls.

Conclusion

This paper has tested the claim that the Iowa Electronic Market offers superior predictions of election outcomes than the snapshots from public opinion polls. By our tests, the IEM election markets are not better than trial-heat polls for predicting elections. In fact, by a reasonable as opposed to naïve reading of the polls, the polls dominate the markets as an election forecaster. This is true in the sense that a trader in the market can readily profit by "buying" candidates who, according to informed readings of the polls, are undervalued. Moreover, we find that market prices contain little information of value for forecasting beyond the information already available in the polls.

Where then do the markets go wrong? To begin with, consider the vote-share market. The histories of market prices show that traders tend to hold persistent beliefs about the vote division that contradict the polls and that these persistent beliefs are often wrong. Incorrect beliefs get corrected only in the last days before the election, when the polls are difficult to ignore. The winner-take-all market tracks the vote-share market but compounds its errors by overvaluing long-shot candidates' chances of victory, as if the market expects more campaign surprises than occur in reality. The existence of persistent mistakes in the vote-share market

projection beats the winner-take-all IEM price, it follows that using IEM price to bet in the hypothetical handicapper's betting pool results in a loss.

compounded by the degree of uncertainty about the vote-share estimates makes the winner-take-all market a particularly poor forecasting tool. Based on the experience of the IEM, if the polls show a candidate to hold a decisive lead but the market is unconvinced, bet on the polls.

It should be noted that our daily poll projections are themselves rather crude instruments. Our robotic trading programs are informed by a flat prior, relying solely on the current polls and the days until the election but nothing more. Even when we compare market prices to the weekly average of poll-based forecasts, our instrument is primitive in that the week's polls are not weighted for relative recency. But further perfection of our forecasting model from the polls would only advance our central argument. If we were to apply more rigorous modeling to obtain a properly weighted average of current polls and earlier polls, the victory of poll forecasts over the market forecast presumably would be more secure.¹⁸

One could argue that the results are drawn from a limited number of election years from a toy market with thin volume and limits on trader spending. With time, the IEM record could improve, and there is some suggestion that it has. Full-blown markets like Tradesports.com might in the end achieve an efficiency that so far has eluded the Iowa Electronic Market. Additionally, studies like the present one can suggest improved strategies to traders, which in turn improve the efficiency of election markets.

Since our results are confined to a few runs of the toy Iowa market, some might claim a “so what” reaction. To such claimants, an important reminder is that the allegedly uncanny performance of the Iowa market has been touted as the primary evidence for the supposed

¹⁸ Were we to broaden our approach to include other predictors from the election forecasting literature, such as economic performance and presidential approval, the victory over the IEM markets presumably would be even more pronounced. For important guideposts in this literature, see Rosenstone (1983), Lewis-Beck and Rice (1992), Campbell and Garand (2000), and the special October, 2004 issue of *PS: Political Science and Politics* on forecasting the 2004 US presidential election.

superiority of election markets over the polls as an information source. The Iowa election market's performance has not been so special after all. For now, our results suggest the need for much more caution and less naïve cheerleading about election markets on the part of prediction market advocates.

References

- Berg, Joyce, Forrest Nelson, and Thomas Reitz. 2003. "Accuracy and Forecast Standard Errors for Prediction Markets." Working paper.
- Berg, Joyce E. and Thomas A. Rietz. 2006. "The Iowa Electronic Market: Stylized Facts and Open Issues." In Robert W. Hahn and Paul Tetlock (eds). *Information Markets: A New Way of Making Decisions in the Public and Private Sector*. AEI-Brookings Joint Center, Washington, D.C.
- Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz. N.d. "Results from a Dozen Years of Election Futures Market Research." In Charles Plott and Vernon Smith (eds.). *Handbook of Experimental Results*. Elsevier, forthcoming.
- Caldeira, Gregory A. 2004. "Expert Judgment versus Statistical Models: Explanation versus Prediction." *Perspectives on Politics* 2:777-780.
- Campbell, James E. 1996. "Polls and Votes: The Trial-Heat Presidential Election Forecasting Model, Certainty, and Political Campaigns." *American Politics Quarterly* 24:408-433.
- Campbell, James E. 2000. *The American Campaign*. College Station: Texas A&M University Press.
- Campbell, James E., and James C. Garand (eds.). 2000. *Before the Vote: Forecasting American National Elections*. Thousand Oaks, Calif.: Sage Publications.
- Fair, Ray C. 1990. "The Effect of Economic Events on the Vote for President: 1988 Update." Mimeo.
- Fair, Ray C. 1996. "Econometrics and Presidential Elections," *The Journal of Economic Perspectives*, Summer 1996, 89-102.

Gelman, Andrew and Gary King. 1993. "Why are American Presidential Election Polls so Variable When Votes are so Predictable?" *British Journal of Political Science* 23: 409-451.

Holbrook, Thomas. 1996. *Do Campaigns Matter?* Thousand Oaks, CA; Sage.

Kou, Steven and Michael E. Sobel. 2004. "Forecasting the Vote: An Analytical Comparison of Election Markets and Public Opinion Polls." *Political Analysis* 12:277-295.

Lewis-Beck, Michael and Tom Rice. 1992. *Forecasting Elections*. Washington, D.C.: CQ Press.

Meirowitz, Adam and Joshua A. Tucker. 2004. "Learning from Terrorism Markets." *Perspectives on Politics* 2 (June): 331-335.

Rohde, Paul W. and Koleman S. Strumpf. 2004. "Historic Presidential Betting Markets." *Journal of Economic Perspectives* 18 (Spring):127-142.

Rosenstone, Steven J. 1983. *Forecasting Presidential Elections*. New Haven: Yale University Press.

Sunstein, Cass. 2005. "Group Judgments: Deliberations, Statistical Means, and Information Markets." *New York University Law Review* 80:962-1049.

Surowiecki, James. 2004. *The Wisdom of Crowds*. New York: Doubleday.

Thaler, Richard H. 1991. *Quasi-Rational Economics*. New York: Russell Sage.

Traugott, M.W. 2001. "Assessing Poll Performance in the 2000 Campaign." *Public Opinion Quarterly* 63:389-419.

Wlezien, Christopher and Robert S. Erikson. 2002. "The Timeline of Presidential Election Campaigns." *Journal of Politics* 64:969-993.

Wlezien, Christopher, and Robert S. Erikson. 2004. "The Fundamentals, the Polls, and the Presidential Vote." *PS: Political Science and Politics* 37:747-751.

Wolfers, Justin and Eric Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18 (Spring): 103-126.

Wolfers, Justin and Eric Zitzewitz. N.d. "Prediction Markets in Theory and Practice." *The New Palgrave Dictionary of Economics*, 2nd ed. Larry Blume and Steve Durlauf (eds.). London: Palgrave Macmillan, forthcoming.

Table 1. Vote-Share Market Prices versus Vote-Share in Polls, 1988-2004.

	Mean Absolute Error: Vote-Share Market Price	Mean Absolute Error: Latest Poll Vote-Share	Market Error as Proportion of Poll Error	Proportion of Days Market Price Beats Polls
1988	3.93	5.55	.71	.60
1992	5.68	7.06	.80	.67
1996	1.01	5.32	.19	.92
2000	0.84	3.08	.27	.88
2004	0.89	1.52	.59	.65
Total	2.41	4.46	.47	.75

Market price vote shares and poll vote shares are measured in terms of incumbent party share of the major-party vote. Poll data are from the latest date with poll ending at least two days earlier. Annual *N*s are 159 in 1988 and 198 in all other years.

Table 2. Vote-Share Market Prices versus Poll-Based Projections of Election Day Vote-Share, 1988-2004.

	Mean Absolute Error: Vote-Share Market Price	Mean Absolute Error: Latest Projection of Poll Vote-Share	Market Error as Proportion of Poll Error	Proportion of Days Market Price Beats Polls
1988	3.93	2.86	1.34	.28
1992	5.68	4.21	1.35	.25
1996	1.01	1.76	0.57	.64
2000	0.84	1.16	0.72	.62
2004	0.89	0.79	1.13	.41
Total	2.41	2.13	1.13	.45

Market price vote shares and poll vote shares are measured in terms of incumbent party share of the major-party vote. Poll data are from the latest date with poll ending at least two days earlier. Annual *N*s are 159 in 1988 and 198 in all other years.

Table 3. Winner-Take-All Prices versus Poll-Based Projections of the Outcome, 1992-2004.

	Mean Absolute Error: Vote-Share Market Price	Mean Absolute Error: Latest Poll Vote-Share	Market Error as Proportion of Poll Error	Proportion of Days Market Price Beats Polls
1992 (<i>N</i> =116)	.38	.22	1.73	.15
1996 (<i>N</i> =198)	.29	.08	3.63	.01
2000 (<i>N</i> =188)	.50	.43	1.16	.30
2004 (<i>N</i> =154)	.45	.31	1.45	.09
Total (<i>N</i> =656)	.41	.26	1.58	.13

Market winner-take-all prices and poll-based projections of outcome probability are measured in terms of the probability of an incumbent party win. Because the winner-take-all market pays off the popular-vote leader rather than the Electoral College winner, Gore is coded as the 2000 winner. As an example of scoring absolute error, the average market price for Bush 41 was .62 in 1988, producing a corresponding error in probability of .38 in the upper left cell.

Table 4. Market Error Relative to the Error of Poll-Based Projections of the Vote

	Market Price Deviates in the <i>Correct</i> Direction Relative to Poll Projection		Market Price Deviates in <i>Incorrect</i> Direction Relative to Poll Projection: <i>Poll Projection</i> is Closer to the Vote
	Market Prices Overshoot: <i>Poll Projection</i> is Closer to the Vote	<i>Market Price</i> is Closer to the Vote	
All cases	14%	44%	41%
<i>Days until Election</i>			
0-30	9	64	27
31-60	12	49	39
61-90	15	44	41
91-120	10	64	26
121-150	17	27	57
151-180	16	32	53
181-210	26	17	57
<i>By Year</i>			
1988	5	28	67
1992	16	25	59
1996	13	64	23
2000	18	61	20
2004	17	41	42

Note: Cell entries are percentages, reading across. Poll predictions are daily projections of the vote from the polls. Market price vote shares and poll vote shares are measured in terms of incumbent party share of the major-party vote. Poll data are from the latest date with poll ending at least two days earlier.

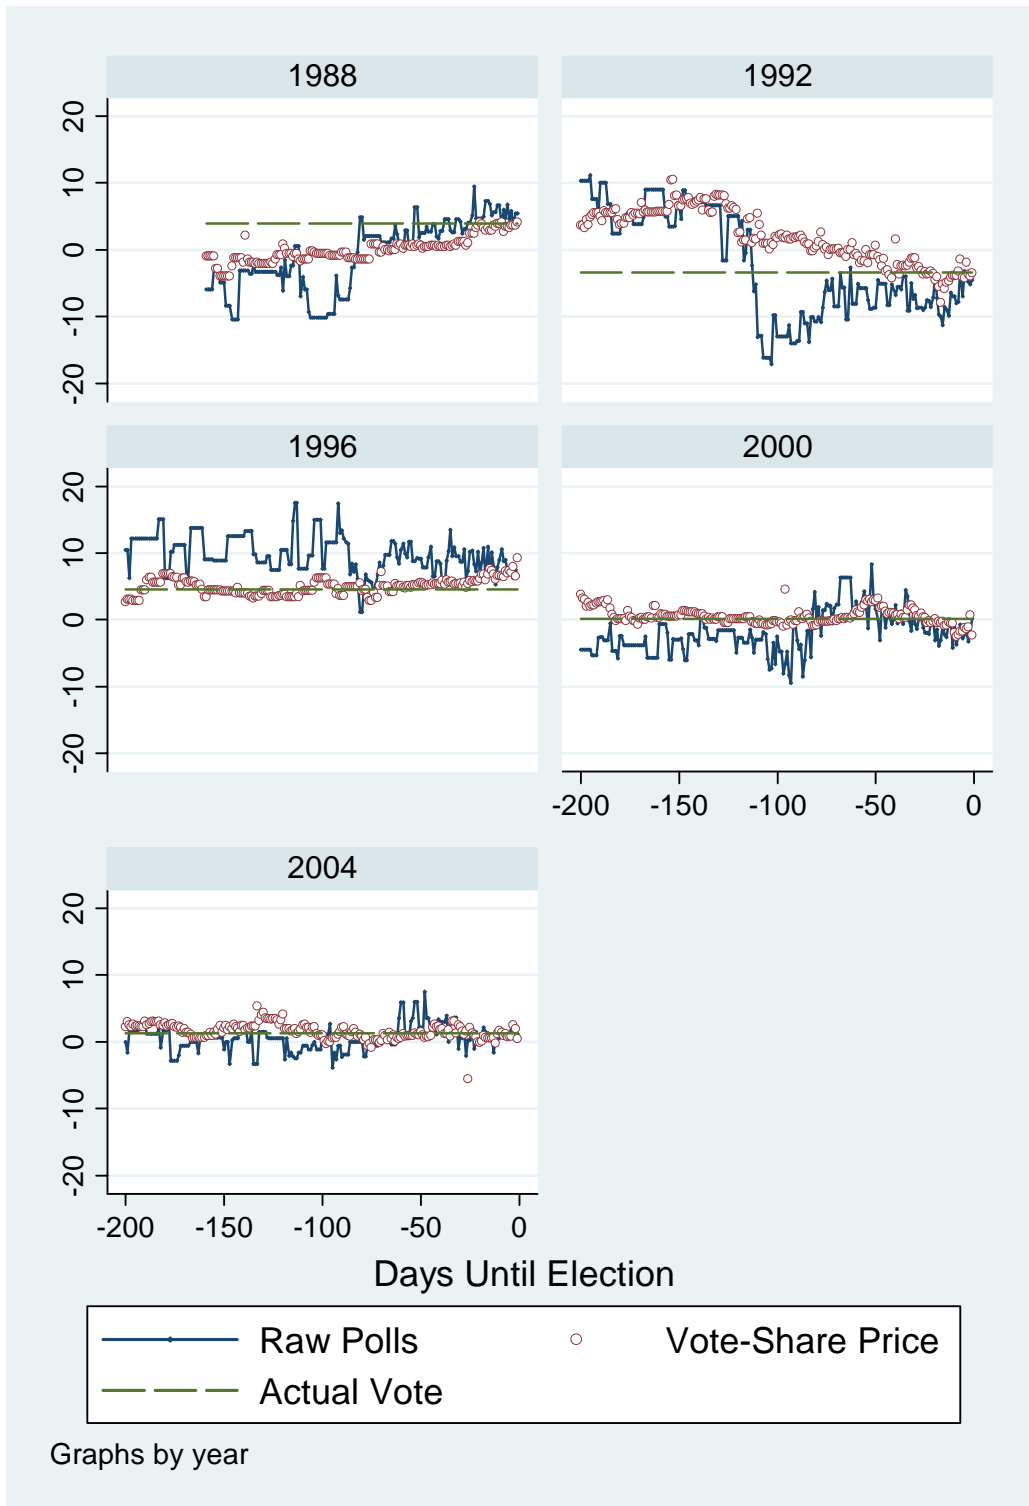


Figure 1. Vote-Share Market Prices and Trial-Heat Poll Results by Days Until the Election, 1988-2004. Poll results are the incumbent-party share of the two-party vote in the poll of polls ending two days earlier or (if no poll ended two days earlier) the latest previous poll date. Vote-share market prices are the incumbent-party share of the prices for the two major-party candidates.

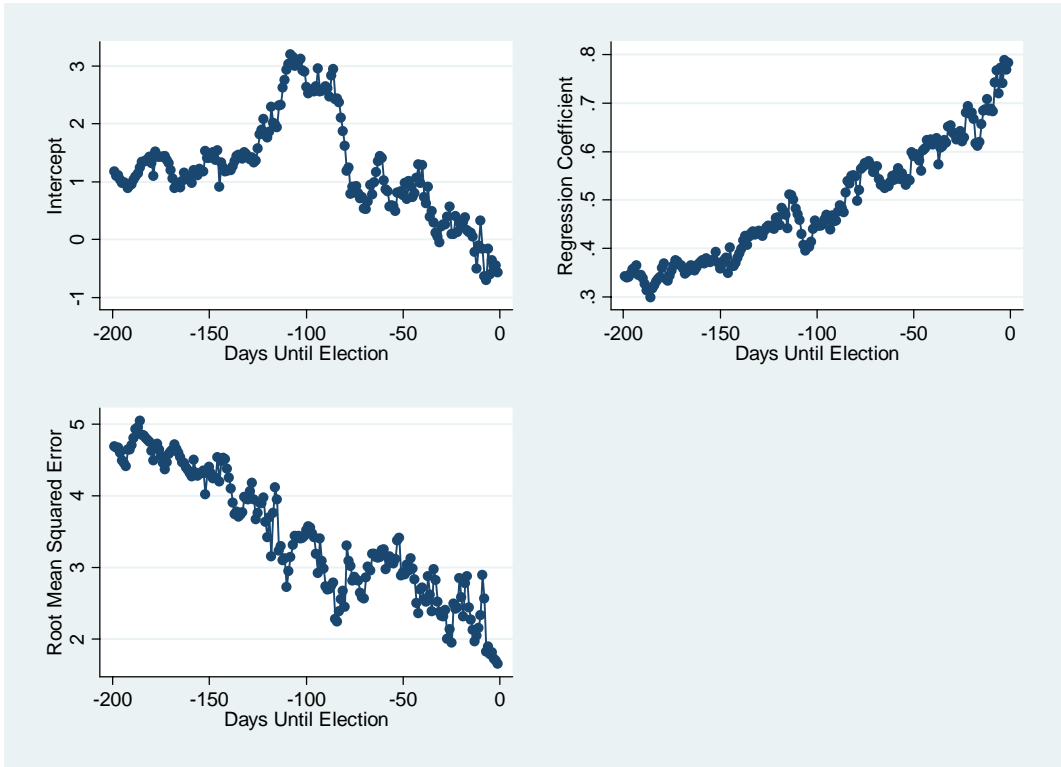


Figure 2. Results of Regression Equations Predicting the Vote from the Polls for Each Date, 1 to 200 Days before the Election. The first panel shows the intercept of the date's equation. The second panel shows the regression coefficient predicting the vote from the polls. The third panel shows the root mean squared error of the prediction. Data are from years 1952-2000.

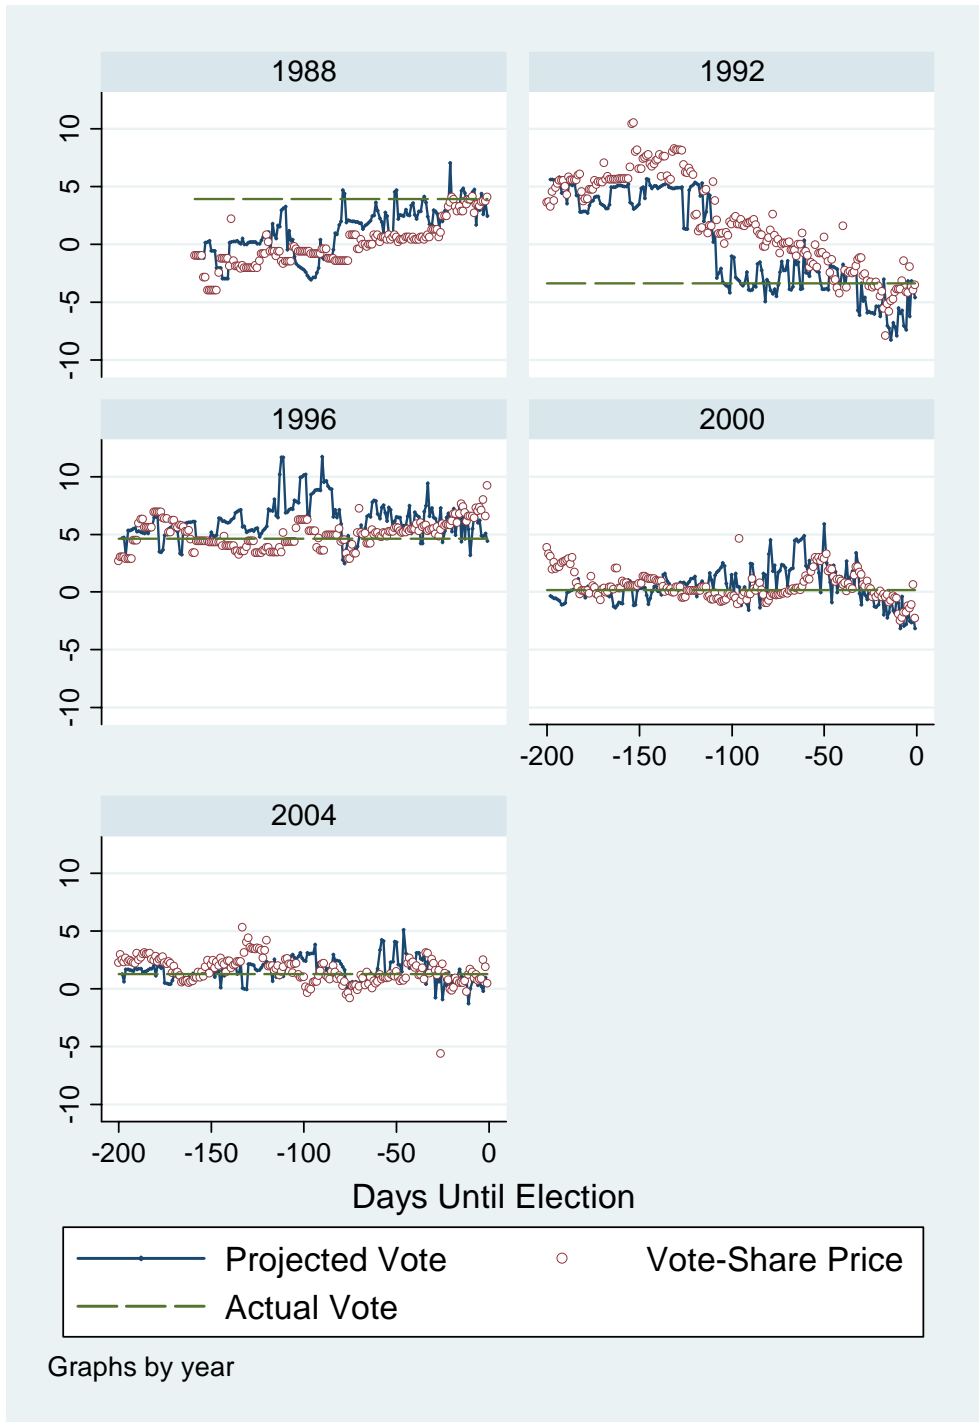


Figure 3. Vote-Share Market Prices and Poll-Based Projections of Election Day Two-Party Vote by Days Until the Election, 1988-2004. Poll-Based Projections are from the poll of polls ending two days earlier or (if no poll ended two days earlier) the latest previous poll date. Projections are based on regression equations predicting the incumbent-party vote from the polls on date t for earlier elections 1952 to the previous presidential election year. Vote-share market prices are the incumbent-party's share of the prices for the two major-party candidates.

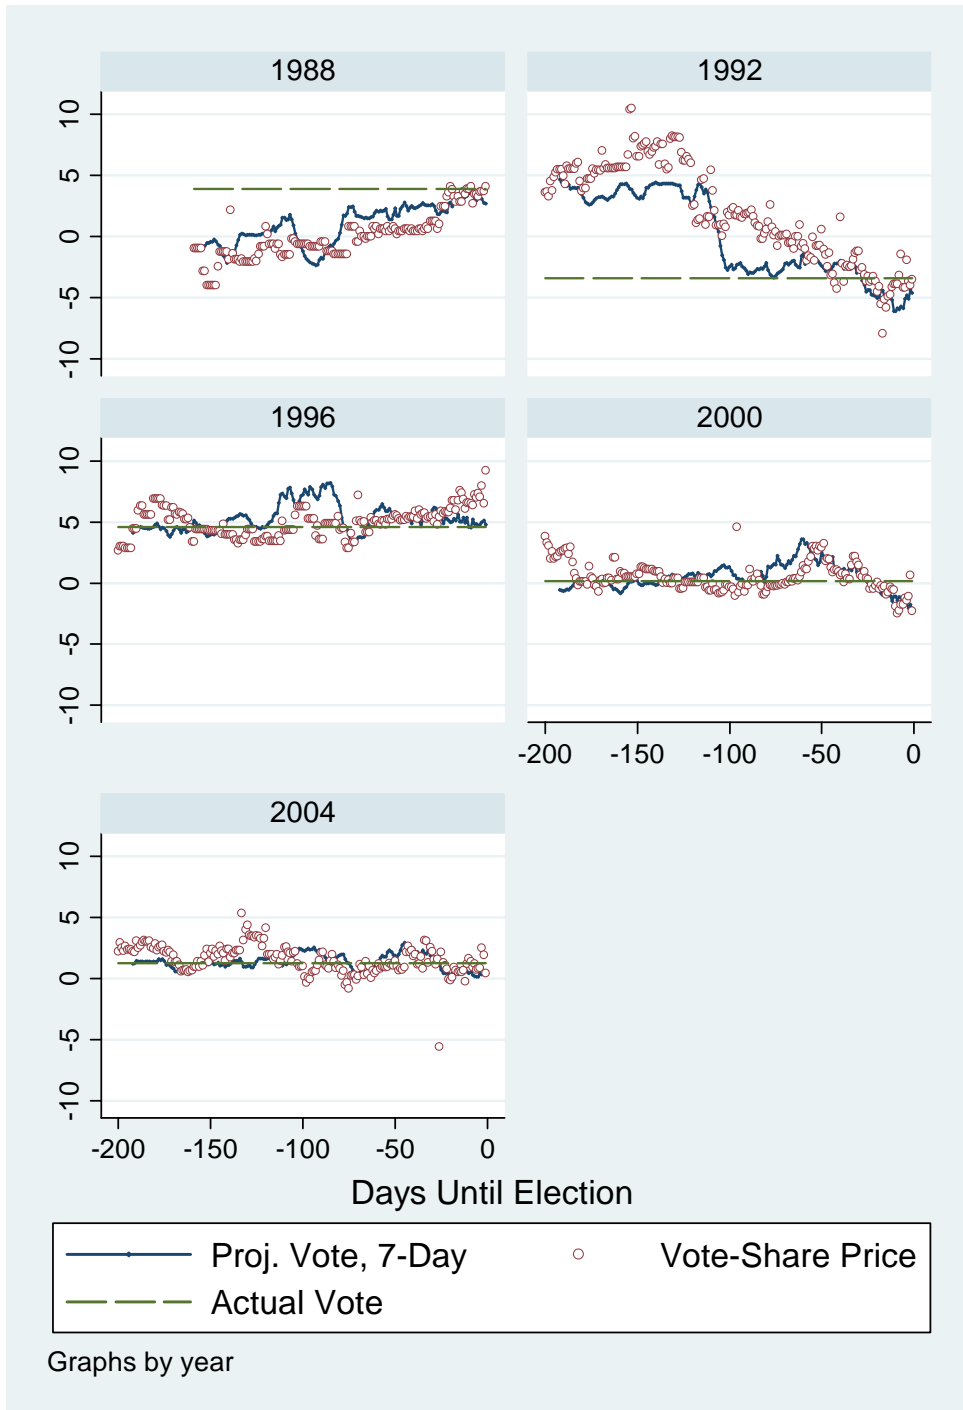


Figure 4. Vote-Share Market Prices and Weekly-Averaged Poll-Based Projections of the Election Day Vote by Days Until the Election, 1988-2004. Poll-Based Projections are the seven-day average of the projections from the latest polls ending two days earlier, shown in Figure 3. Projections are based on regression equations predicting the incumbent-party vote from the polls on date t for earlier elections 1952 to the previous presidential election year. Vote-share market prices are the incumbent-party's share of the prices for the two major-party candidates.

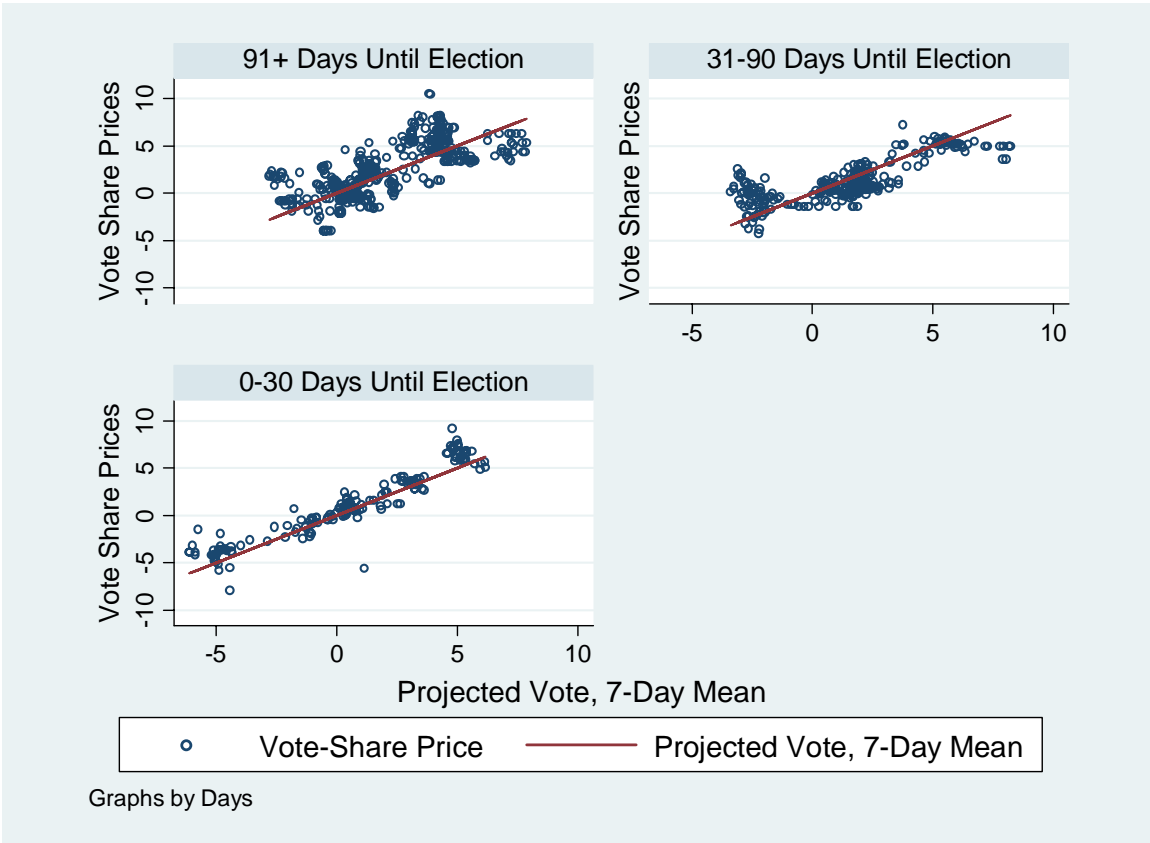


Figure 5. Vote-Share Market Prices by Weekly-Averaged Poll-Based Projections of the Election Day Vote, 1988-2004, Arrayed by Days before the Election. Poll-Based Projections are the seven-day average of the projections from the latest polls ending two days earlier, shown in Figure 3. Projections are based on regression equations predicting the incumbent-party vote from the polls on date t for earlier elections 1952 to the previous presidential election year. Vote-share market prices are the incumbent-party share of the prices for the two major-party candidates only. The diagonal line is a 45 degree line, not a regression line.

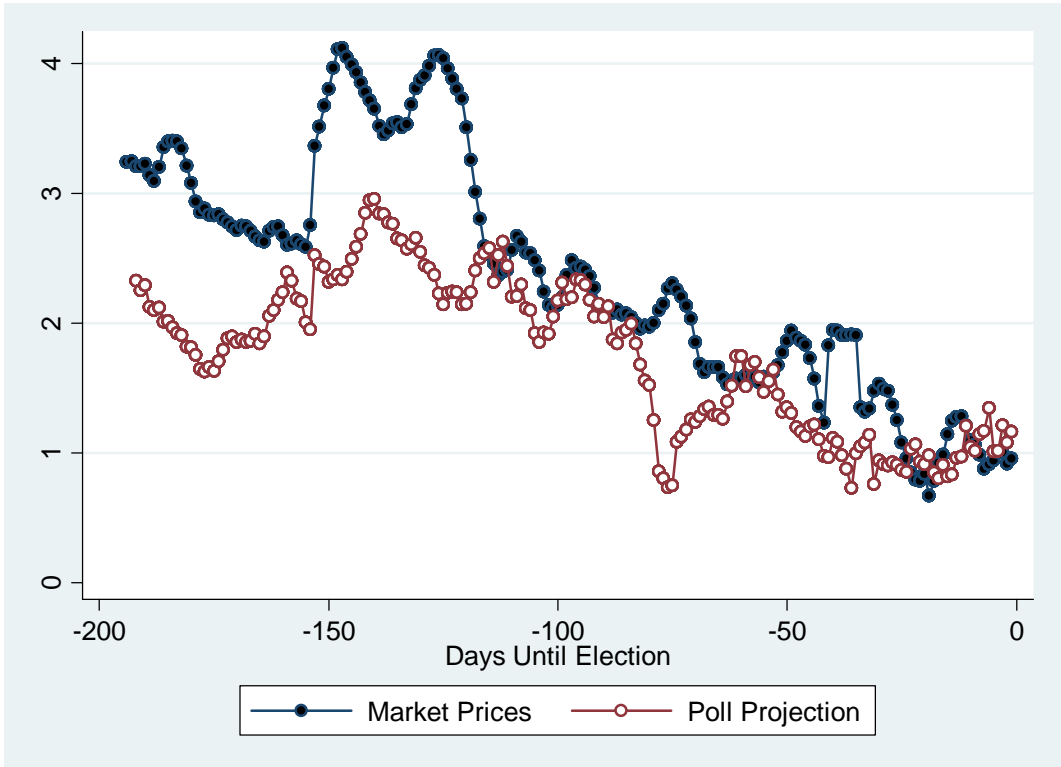


Figure 6. Mean Absolute Error of Vote-Share Market Prices and of Weekly-Averaged Poll-Based Projections of the Election Day Vote, 1988-2004. Poll-Based Projections are the seven-day average of the projections from the latest polls ending two days earlier, shown in Figure 4. Projections are based on regression equations predicting the vote from the polls on date t for earlier elections 1952 to the previous presidential election year. Vote-share market prices for the two-party vote are the relative prices for the two major-party candidates only. Averages incorporate 1988 data beginning with the date 159 days before the election, which was the date the 1988 vote-share market opened. For other years, the first data are reported for the date 192 days before the election.

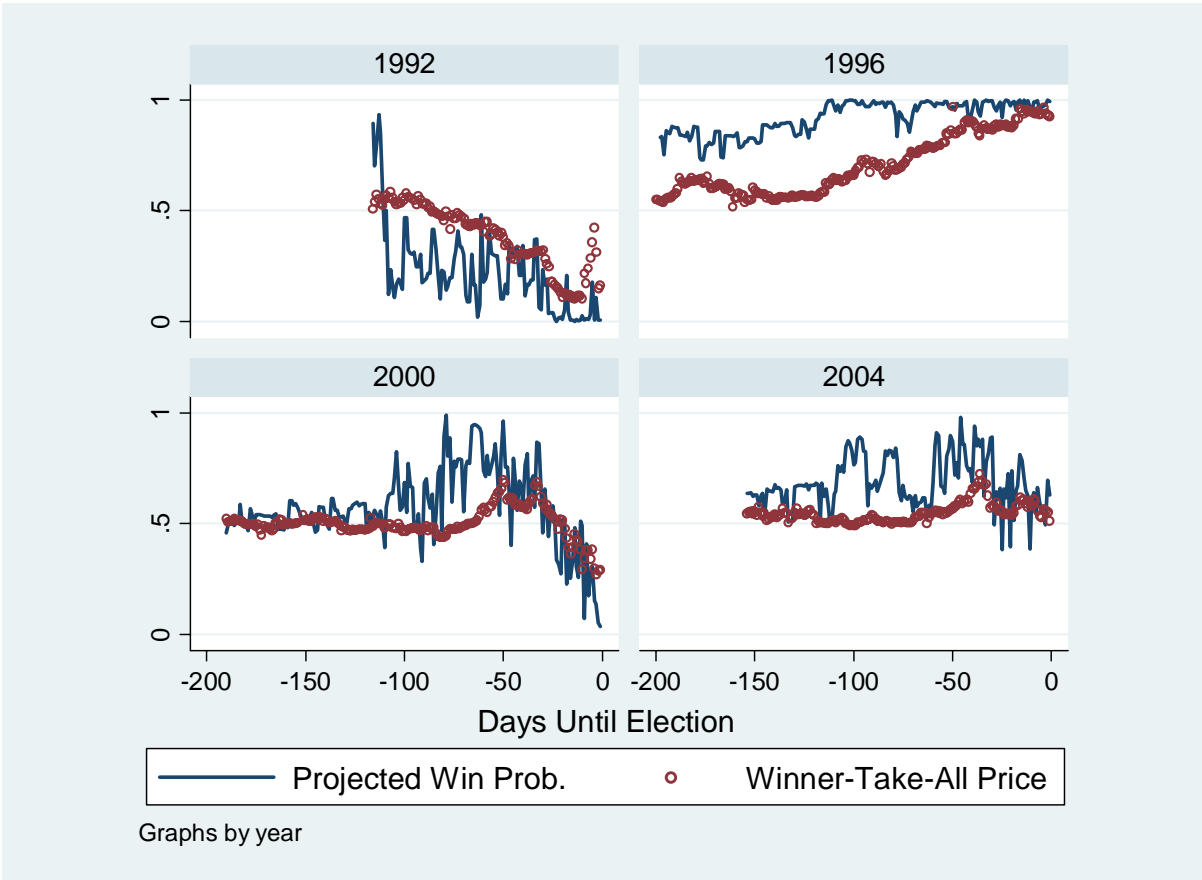


Figure 7. Market Winner-Take-All Prices and Poll-Based Projections of the Popular Vote Winner by Days Until the Election, 1992-2004. Poll-Based Projections are from the poll of polls ending two days earlier or (if no poll ending two days earlier) the latest previous date. Projections are based on regression equations predicting the incumbent-party vote from the polls on date t for earlier elections 1952 to the previous presidential election. The probability of an incumbent-party win is based on the expectation and the variance. Winner-take-all market prices are the incumbent-party share of the prices for the two major party candidates only.

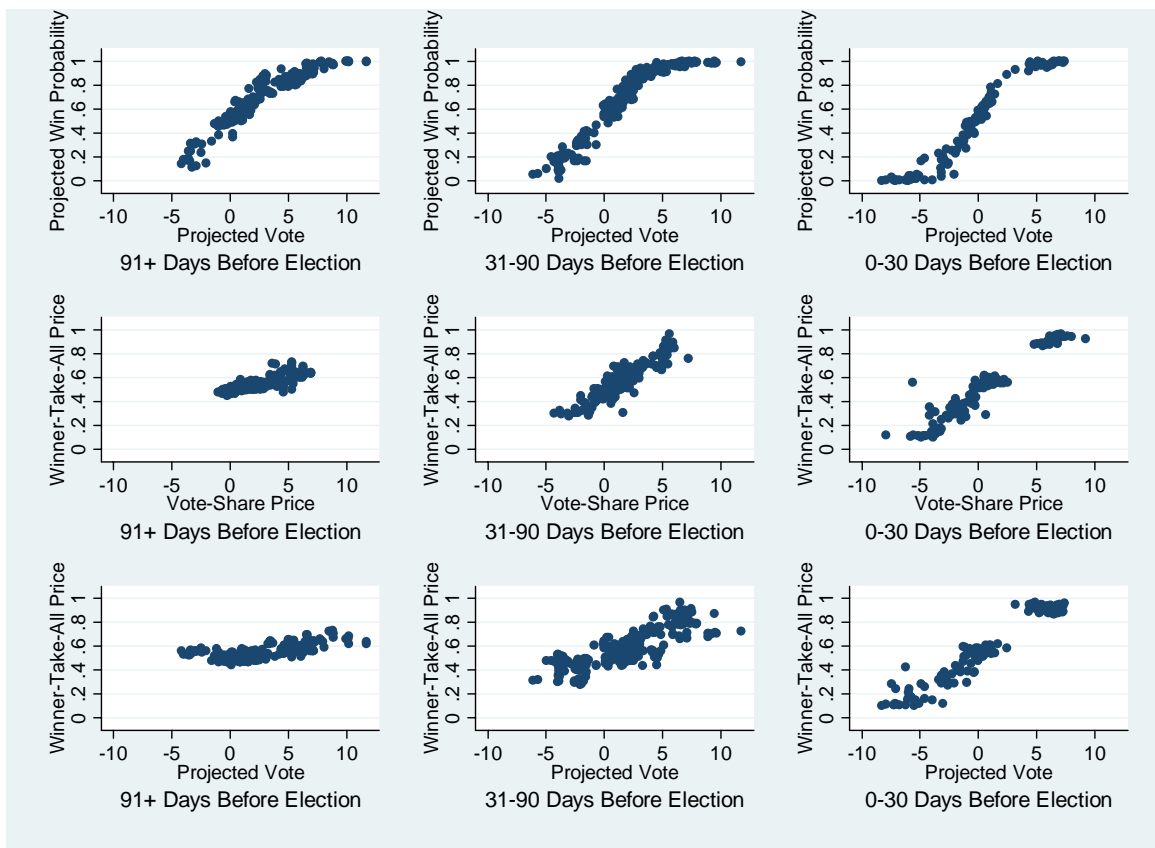


Figure 8. Probability of Incumbent Party Win by Projected Vote (top panels); Winner-Take-All Prices by Vote-Share Prices (middle panels; and Winner-Take-All Share Prices by Projected Vote (bottom panels). Data are for all days when winner-take-all markets were open, 1992-2004. The projected vote and candidate prices represent the value for the incumbent-party.

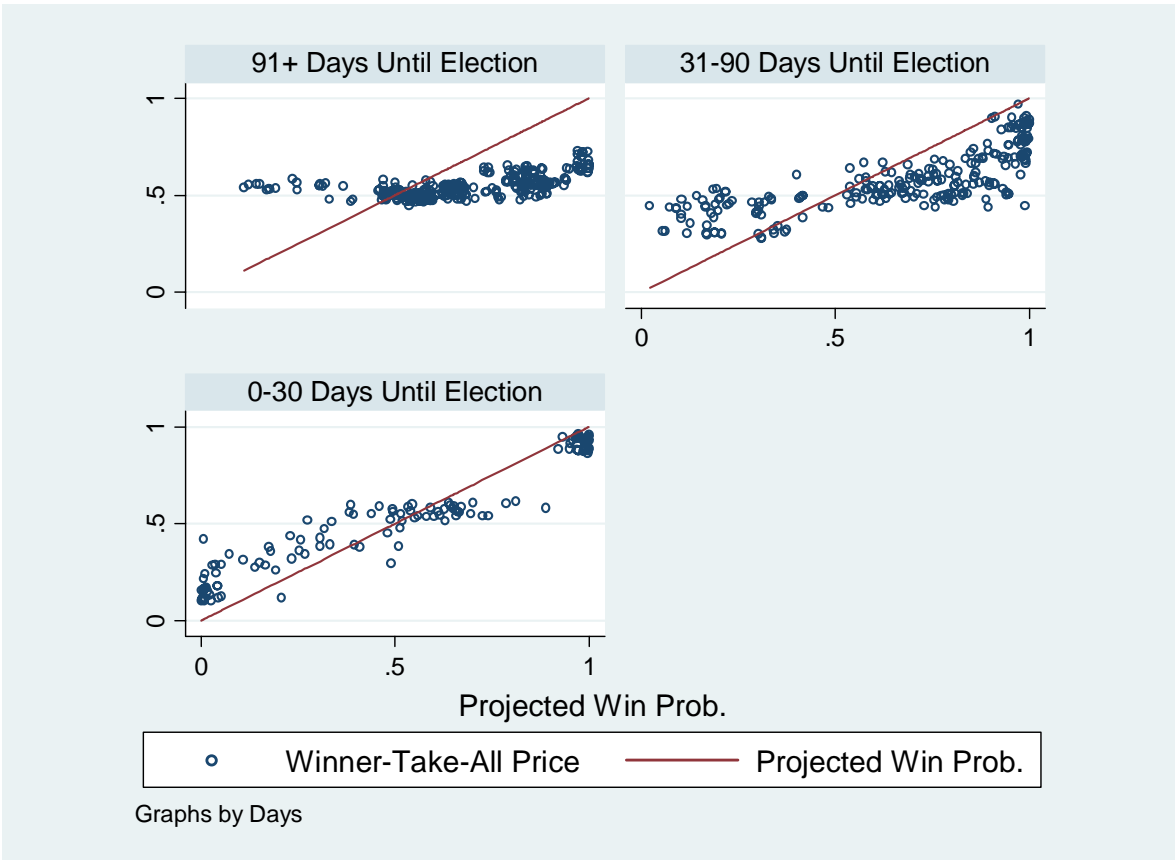


Figure 9. Winner-Take-All Prices by Poll-Based Probabilities of an Incumbent Party Win, by Days before the Election. Data are for all dates with a winner-take-all market price, 1992-2004. Diagonal lines are 45 degree lines, not regression lines.