

Markets vs. Polls as Election Predictors: An Historical Assessment

Robert S. Erikson
Columbia University
rse14@columbia.edu

Christopher Wlezien
Temple University
wlezien@temple.edu

Abstract

Prediction markets have drawn considerable attention in recent years as a tool for predicting elections. But how accurate are they? Do they really outperform the polls, as some would have us believe? Do prices in election markets carry information beyond the horserace that is portrayed in the latest polls? And how well do election markets perform when there are no public opinion polls to provide the cues?

This paper compares the accuracy of presidential betting markets in years before and after public opinion polls were introduced. For data, we rely on prices from modern internet-based election markets since 1992 but also from the vigorous election markets that thrived on Wall Street going back from at least as far as 1880 up to 1960. Our results comparing the performance of markets and polls are provocative. First, we find that market prices are far better predictors without polls (1880-1932) than with polls available (1936-2008). Second, we also find that market prices of the pre-poll era predicted presidential elections at least as well as polls have done following the introduction of scientific survey research. Finally, we find that last-minute market prices add nothing to election prediction once we control for trial heat polls during the final week of the campaign. The ordering of these results is not transitive.

We attempt a reconciliation of theory and evidence. In the end we find no reason to challenge the value of contemporary polling. At the same time, it is clear that early election markets before polls were surprisingly good at extracting campaign information without “scientific” polling to guide them.

Earlier versions of this paper were presented at the Annual Meeting of the American Association for Public Opinion Research, Hollywood, Florida, May, 2009, the University of Manchester and the University of Sussex. For helpful comments, we thank Tereza Capelos, Ed Fieldhouse, Robert Ford, and Will Jennings.

In recent years, prediction markets have drawn considerable acclaim as a tool for forecasting future events (e.g., Arrow et al., 2008; Wolfers and Zitzewitz, 2008). Markets on election outcomes are a centerpiece of this discussion. A common claim is that prices in election markets, such as the Iowa Electronic Market (IEM), predict elections better than the polls (Berg and Rietz, 2006; Berg, Forsythe, Nelson, and Rietz, 2008a, 2008b). This idea that election markets are superior to polls has traveled from economics into other academic realms (e.g., Caldiera, 2004; Sunstein, 2005) and popular literature as well (e.g., Surowiecki, 2004).

It is not exactly clear what the empirical test should be for gauging the success of prediction markets, particularly when their forecast accuracy is pitted against that of public opinion polls. Market proponents cite the fact that the IEM “vote share” prices during campaigns correspond more closely to the final vote than do poll margins measured at the same time in the campaign. Erikson and Wlezien (2008), however, show that by discounting early polls for the inflated sizes of poll leads, polls dominate market prices in the sense that historically one could profit in the Iowa market by betting based on the size of the discounted poll results relative to the market prices. One challenge to election markets is that they appear to be subject to an underdog bias similar to that in horseracing (Wolfers and Zitzewitz, 2004, 2008). For instance, throughout most of Bill Clinton’s two victorious presidential campaigns, the Iowa winner-take-all market underestimated his chances compared to what a reasonable interpretation of the polls would suggest (Erikson and Wlezien, 2008).

The best test would be to have both market prices and poll results for a series of elections and then pit them against each other as independent variables to see which works best as an election forecaster in a standard multivariate analysis. Unfortunately, modern electronic markets have been around for only a handful of elections. If, for instance, we want to compare winner-

take-all market prices from the Iowa Electronic Market (IEM) with poll results as predictors of presidential elections in the US, we have only the five presidential elections between 1992 and 2008.

As it turns out, we have much more data set to draw on. Thanks to Rohde and Strumpf (2004), we now have information about a vigorous election markets that thrived on the Wall Street Curb—that is, not on the New York Stock Exchange—going back from at least as far as 1880 up to 1960. Snowberg, Wolfers, and Zitzewitz (2004) have pulled together election-eve prices from these Wall Street Curb markets, contemporary electronic markets, and, for most of the intervening years, London bookmaker odds. This makes it is possible to compare election-eve market prices and election-eve trial-heat polls as predictors of presidential elections for 16 data points—all presidential elections from 1936 through 2008 except for 1964, 1968, and 1972, for which we have no market data. This is more than the usual number of cases that presidential election forecasters typically use for their augury.¹ Curb market prices from earlier years yields 14 “control” cases for 1880-1932—presidential elections with prediction markets but no scientific polls.²

As mentioned, Snowberg, Wolfers and Zitzewitz (2004) report election-eve market prices from Rhode and Strumpf’s Wall Street Curb markets (through 1960) with electronic markets (1992 onward), supplemented by information from London betting markets for 1976-1988. For 1932-2008, we can compare the final pre-election prices with late trial-heat poll preferences as electoral predictors. We also can compare the performance of market prices during the poll era

¹ See, e.g., the forecasting symposium in *PS: Political Science and Politics*, vol. 41, issue 4 (October), 2008.

² For details on the conduct of the early election markets, see Rohde and Strumpf (2003; 2004).

with market prices during the pre-poll era to see whether market prices improve when the feedback of scientific polls is available.

With more than a century and a quarter's worth of the prices of election outcomes on the day before each election, we possess the evidence regarding popular expectations of presidential election outcomes at the moment of the election. This paper compares the accuracy of presidential betting markets in years before and after public opinion polls were introduced. And, for the modern polling era, we compare the predictive power of polls versus markets.

It is commonly understood that before scientific polling was invented, public opinion was difficult to gauge (Geer, 1996). Almost certainly, it would seem, elections of the pre-poll era were conducted under greater uncertainty about the outcome than elections today, but with observers monitoring various indicators for cues (Kernell, 2000; Karol 2007; see also Robinson and Chaddock, 1932). In the absence of scientific polls, election market prices were eagerly studied for evidence of election trends. As Rohde and Strumpf (2004) point out, these early markets had a strong record of accuracy in terms of final prices predicting the winner. Rohde and Strumpf also show that they performed impressively on certain tests of market efficiency. Still, one might expect that with an absence of poll information, these early market prices could not gauge presidential election outcomes as well as polls do today. A second obvious hypothesis is that the accuracy of market prices would improve once scientific polls were established. Based on the conventional wisdom, therefore, one would expect polls to dominate election markets as the election predictor when we only have one but not the other, but election markets should improve once polls are available to provide more reliable information to investors.

What should be our expectations of market prices versus polls when both are available?

We might expect that market prices add information about the forthcoming election beyond what evident from the polls. Enthusiasts for contemporary election markets claim that market prices are superior to the polls for forecasting presidential elections. Tell us the betting line, say market believers, and we tell you the outcome with greater accuracy than the latest polls (Berg and Rietz 2006; Berg, Forsythe, Nelson, and Rietz, 2008a; Wolfers and Zitzewitz, 2004; Page, 2008).

Although these may all be reasonable, if not obvious, expectations, it turns out that *none* holds up when put to the test of data analysis. We find that market prices are far better predictors without polls (1880-1932) than with polls available (1936-2008). We also find that market prices of the pre-poll era predicted presidential elections at least as well as polls have done following the introduction of scientific survey research. Finally, we find that last-minute market prices add nothing to election prediction once we control for trial heat polls during the final week of the campaign.

At first glance, this is quite topsy-turvy. The first of these upsets would have us believe that markets perform better when polls are not available as a guide. The second would have us believe that we can just as easily predict elections with only election markets as with polls. The third would have us believe that election market prices are not informative when polls are available. That is, our preliminary results show that (1) markets without polls beat markets with polls; (2) markets without polls are as good as polls; and (3) polls beat markets when both are present. This ordering is not transitive.

In the sections below we first present the data analysis supporting the odd set of results. Then we attempt a reconciliation of theory and evidence, and consider implications of the

findings and the future of forecasting using markets and polls. In the end we find no reason to challenge the value of contemporary polling. At the same time, early election markets before polls were surprisingly good at extracting campaign information without scientific polling to guide them.

A Technical Note: *Comparing Apples to Oranges*

Before proceeding further, we must consider the challenge that market prices are measured as probabilities of a Democratic (Republican) win while vote margins and poll margins are measured as percents of the vote. Thus, market prices will be related to the vote margin and poll margin in a nonlinear fashion. This is a handicap when comparing the sizes of correlation coefficients. Fortunately, however, it is possible to use the market price to impute a measure of the market's *expected* vote, which should be linearly related to the vote and poll margins.

As the price setter, the marginal trader derives a probability of a Democratic win from an expectation of the vote margin and a normally distributed error that represents uncertainty about the expectation. Figure 1 shows some hypothetical examples. For each of three hypothetical elections (E1, E2, and E3), marginal trader observes an expected vote (x - axis) and, considering the uncertainty of the expectation, converts that to a probability of a Democratic win—the portion of the disturbance term that is greater than 50 percent of the vote. We observe the marginal trader's belief about the probability of a Democratic win (the p -value) but of course not his/her expected vote. But by knowing the p -value, we can back out the estimate of the expected vote, corresponding to the x -axis on Figure 1. For identification, all that is required is the assumption that the variance of the error term is constant over the series of elections examined, e.g., the elections with markets but before scientific polling.

-- Figure 1 about here --

Knowing the prices, we can readily impute the marginal trader's expectation of the vote. This conversion allows both the poll margins and the beliefs represented by the market to be presented in a linear relationship with the actual vote. We impute the value of the expected vote as a distance from the fifty-percent mark where the scale of units is in terms of the standard deviation of the disturbance. For instance, if the p -value equals .975, the imputed expected vote of the price setters is +1.96 standard deviation units. (This is the familiar upper cutoff for the .05 level of statistical significance.)

The one limitation is that the scale of the imputed expected vote is measurable only in standardized units where one unit equals one standard deviation of the expected error. The imputed expected vote is not measureable as a percent of the vote because the scale of the error term is unknown. Fortunately, this presents no loss of generality when estimating correlations between the imputed market expectation of the vote and either the actual vote or the vote margin in the polls.

Election Markets—Then and Now: *Comparing Prediction Markets 1880-1932 vs. 1936-2008*

Recall that we have three different kinds of market data. First, for the period between 1880 and 1960, we have the prices from the real Wall Street Curb markets (Rohde and Strumpf, 2004). Second, for the period from 1992 to the present, we have the prices from online markets, specifically the Iowa Electronic Markets through 2000, Tradesports in 2004, and Intrade (formerly Tradesports) in 2008. Third, for some of the intervening years, where we have neither the old or new markets, we have the London betting odds—specifically, from Snowberg et al. we have these betting odds for 1976-1988 but not 1964-1972. Each provides a winner-take-all price.

These can be interpreted as the market's judgment of the probability of victory, e.g., a price of 34 cents registers a 34% probability of victory.³ For this analysis we rely mainly on the election-eve prices available the day before the election. We convert the prices into two-party probabilities of a Democratic win.⁴ And ultimately we convert them into the election market's imputed expectation of the vote—shorthand as the “imputed vote.”

Figures 2 and 3 show election-eve probabilities from the betting markets as a function of the actual vote. Figure 2 plots these data for the 15 elections between 1880 and 1932, before the advent of polling. Figure 3 plots the data for the 16 subsequent elections for which we have market data—excluding 1964, 1968, and 1972—through 2008. In each figure we overlay an ogive curve to fit the data.⁵ Clearly, the statistical relationship between winner-take-all market prices and vote margins is non-linear. With prices on the vertical axis and the vote margin on the horizontal axis, the relationship approximates the cumulative normal distribution, consistent with our theoretical discussion. Our interest is in whether markets do better in the more recent period where poll results are available.

-- Figures 2 and 3 about here --

The figures show that prices respond most crisply to the signal of the actual vote during the early, pre-poll era. Remarkably, early prices correlate at +.93 with the vote. As has been noted by Rohde and Strumpf (2004), market prices on election eve predicted the correct winner in

³ This seemingly straightforward interpretation of prices is the subject of some dispute—see Manski (2006) and Wolfers and Zitzewitz (2007).

⁴ For measurement details, see Snowberg, Wolfers, and Zitzewitz (2007) and Rohde and Strumpf (2004).

⁵ The curves shown in Figures 1 and 2 (and subsequent figures) are derived from the estimation of the “imputed vote” from the actual vote. The first step is transforming the price into the imputed vote, as described earlier. Second, the imputed vote is regressed on the actual vote to obtain a least squares equation predicting the imputed vote in terms of the actual vote. Third, the predicted imputed vote is de-linearized in terms of the p -value that corresponds to the predicted imputed vote. For instance, if the *predicted* imputed vote is +1.96, then the corresponding value for the curve is .975, or 97.5 percent of the vote.

every one of the 15 pre-poll elections. This is not true for the later period, however. For the era beginning in 1936 when polls became available, the vote-price correlation is a more modest $+0.80$ —see Table 1 for relevant correlations—and the market prices got the wrong popular vote winner in three of the 16 post-poll elections (1948, 1976 and 2000). From these results it is pretty clear that markets have done fairly well forecasting elections on election eve, but especially *before the advent of polling*. This implies that polls may have had a distorting effect on markets. Consider that in 1948, Dewey's late pre-election market price was 89 cents (making him an almost 9 to 1 favorite), pretty much as the polls (and the *Chicago Tribune*) had it.

-- Table 1 about here --

For the pre-poll era (1880-1932), our imputed vote correlates at a striking $+0.95$ with the actual vote outcome. The slight increase from the $+0.93$ correlation between raw market prices and the vote reflects the correction for the nonlinear relationship shown in Figure 2. For the post-poll era (1932-2008), using the imputed vote instead of the raw prices keeps the correlation with the vote at $.80$. Election markets clearly did much better at predicting election outcomes during the era in which they did not have polls for guidance than later, when polls could be a guidepost for setting market prices. *Election markets clearly performed better without polls.*

Markets Then and Polls Now: *Prediction Markets 1880-1932 vs. Polls 1936-2008*

Which leads to better election forecasts? Is it polls or is it election markets that operate in the dark, without opinion polls? For this test we compare the predictive power of early election markets before the advent of scientific polling with that of polls during the modern poll era. We measure election-eve trial-heat poll results as the average of all polls during the final week of the

campaign (or the most recent poll if none are available during the final week). The procedure used follows our earlier work (Wlezien and Erikson, 2002).

Figure 4 demonstrates the strong relationship between the vote and the polls for the 19 post-1932 presidential elections. The correlation is an impressive +.91. This still is slightly less than the pre-poll correlation of +.93 between market prices and the vote and even further below the +.95 correlation between 1880—1932 vote margins and the imputed expected vote. Pre-poll era election markets were at least as good at predicting presidential vote margins as polls have been during the current era of modern public opinion polling.

-- Figure 4 about here --

This completes our second set of counter-intuitive results. First we saw that early markets performed much better than current markets. Now we see evidence that early (pre-polling) markets performed a shade better than the polls when polls came on line.

Polls Now and Markets Now: *Polls 1936-2008 vs. Election Markets 1936-2008*

Next we consider which indicator avoids the booby prize of worst predictor from the following set—contemporary polls or contemporary markets? From our discussion above, we already know that polls win this contest. As indicated above and shown in Table 2, final-week election polls have performed well at predicting the vote (correlation = 0.91). And, as we saw above, market prices (measured either as raw prices or imputed vote) correlate at no more than 0.70 with the vote. Polls clearly are the better predictor.

A second question though is whether market prices contribute useful information about predicting the election that is not apparent from the polls. Ideally we would answer this question

by comparing polls and prices months before the election, when there still are events to affect the outcome. But for the poll era, except for the five most recent elections we only have prices for election eve. The analysis must be limited therefore to comparing prices on the eve of the election with polls during the final week. Do these late market prices contain information not found in the late polls?

-- Table 2 about here --

The answer appears to be no. The multivariate equations shown in Table 2 provide the evidence. The test is not whether the predictive power of market prices exceeds that of the polls. The much more modest test is simply whether the coefficient for market prices is positive in sign when the polls are held constant. In fact, when polls and prices are raced in a multivariate equation predicting the vote, the poll coefficient is positive and significant while the price coefficient is actually slightly negative but not statistically significant. Switching to the imputed vote as the market measure in the second column of Table 2 produces very similar results. From this exercise we see that election eve market prices do not provide information beyond election-eve polls.⁶

This admittedly is a hard test because if there is information about the election, it is likely to be reflected in final polls. Market prices clearly reflect the polls: raw prices correlate with poll margins at +.87 and the imputed vote correlates at +.86, both considerably higher than their correlations (+.70 and +.67) with the vote. Thus, market prices follow the polls when polls are available.

⁶ Because market prices and poll results are highly correlated, there is a high degree of multicollinearity in the data. Rather than being a problem, this only makes the argument more persuasive that market prices add little beyond what the polls show; that is, the polls are highly significant and the coefficient for prices not only is statistically insignificant in the multivariate equation but in fact has the “wrong” sign.

Prediction Markets and the Polls, 1952-2008

Our analysis so far suggests that as tools for election forecasting, early (pre-poll) election markets dominate modern-day polling which dominates modern-day election markets. How can this be? While some of these distinctions are trivial in magnitude, there is little doubt that the quality of election markets declined when scientific polls became available as a cueing device for inferring election outcomes. With poll information available, markets declined in volume and became dependent on polls. Polls provided the leading information for election-eve markets. Under these circumstances, election markets could hardly be expected to perform better than the polls. Rather, later markets reflected the polls plus error.

To understand the relative predictive power of “modern” polls and markets relative to the early pre-poll markets, it is crucial to take into account the election years we include for the analysis of “modern” polls and markets. We begin with 1936, a year when even the heralded Gallup poll considerably underestimated Roosevelt’s vote strength and when there was considerable market uncertainty reflecting the huge difference between the Gallup and the Literary Digest poll predictions. The period also includes the polling disaster of 1948 of “Dewey beats Truman” fame. We know that polling performance changed dramatically, particularly in the wake of the 1948 debacle. Perhaps the markets improved as the polls themselves improved.

The evidence supports this explanation. The correlation between the polls and the vote for the post-1948 period is a near-perfect $+0.97$, appreciably larger than the correlation ($+0.91$) for the full 1936-2008 period. Likewise, the correlation between the imputed vote from market prices and the actual vote is a healthy $+0.94$, strikingly larger than what we get ($+0.80$) including 1936-1948. As the accuracy of polls became clearer, so did the accuracy of markets.

Comparing the relevant correlation coefficients, we now see the following when we measure the modern polling eras as starting in 1952 rather than in 1936. With the new time frame, modern markets predict the vote ($r = .94$) with an accuracy comparable to that of the pre-poll era ($r = .95$). And when defined as starting in 1952, modern polls predict better than pre-poll markets (r 's of .97 vs. .95).

When we set the start of the modern polling era as 1952, polls predict better than the imputed vote from market prices but not by much (r 's of .97 vs. .94). We can ask with our new start period, do election eve market prices add relevant information not available in the late polls? Table 3 compares the polls and the market-imputed vote as predictors of the actual vote once again, this time for the 1952-2008 period. Again we find no evidence that market prices add information beyond the late polls. With poll results held constant, the coefficient for the imputed expected vote from market prices (though positive) is far from statistically significant..

-- Table 3 about here --

The relationship between post-1948 market prices and the vote is shown graphically in Figure 5. (Recall that we do not have market data in 1964, 1968, and 1972.) Prices not only correspond with the election outcome; they demonstrate a much higher level of certainty than we see over the full post-1932 period—compare Figures 3 and 5. Indeed, there is more certainty in the post-1948 period than in the pre-poll period. This can be seen in Figure 6, which plots the ogives for the two periods (from Figures 2 and 5). Although markets prices and the vote were more closely correlated in the pre-poll period, prices then were comparatively cautious. That is, the slope relating the victory margin and prices is relatively flat by comparison with the post-poll era. Conditional on the actual vote margin, participants in pre-poll markets hedged their bets

about the likely winner more than participants in post-1948 markets did. Put differently, once gamblers could confirm with reliable polling the electoral cues that were otherwise available, their betting reflected a greater certainty about the likely winner.⁷

-- Figures 5-6, about here --

Pre-Poll Markets in Advance of the Election

So far, we have examined market prices solely in terms of prices on the eve of the election. These inform regarding Election Day expectations before poll readings. We have not yet examined the accuracy of market prices from earlier in these campaigns. Prices from earlier in these campaigns can inform about the crystallization of expectations during campaigns when there is no feedback from the polls.

In the long version of their paper on election markets, Rohde and Strumpf (2003) display a valuable graph of market prices over the course of the campaigns for the years 1884-1940. (No data can be found for the later years regarding gambling odds in the months before the election—that is, until the electronic winner-take-all market was born in 1992.) Rohde and Strumpf show that the earlier in the campaign, the less certain are prices, conditional on the final outcome. Here we extend this analysis to compare the certainty embedded in early market prices with the certainty from election polls for the modern poll era, 1952-2008. From the early market data we extract prices at 1, 30, 60, 90, and 120 days before each election, for the pre-poll years 1884-

⁷ Figure 6 suggests a hint of a differential partisan bias to the pre-poll and post-poll markets, with the latter favoring Democrats. Consider the neutral point, where the Democratic and Republican candidates each receive 50% of the two-party vote. In such an election, the expected price in the pre-poll market would be about 45 cents and in the post-poll market about 60 cents. The “bias” for the latter years’ markets disappears when we include election markets going back to 1936. (See Figure 3.) This Democratic tilt to the ogive curve is due simply to markets being more certain of the Democratic landslides than the comparable Republican landslides. Recent close elections are predicted as close, without evident bias. (See Figure 5.)

1932. We then “linearize” these prices in the usual way to z-scores representing the imputed scale-free expected vote at 1, 30, 60, 90, and 120 days prior to each election. We then observe the correlations of these “readings” of the expected vote with the actual vote. For each modern election, we observe the correlation between the actual vote and poll preferences (7 day averages) for days 1, 30, 60, 90, and 120 days prior to the election. Table 4 presents the results.

--Table 4 about here--

One sees immediately that the farther back one goes in time before the election, the imputed vote from the markets fades as an election predictor. The same is true for polls in the modern era. Each tendency of course is expected—the earlier in the campaign, the harder it is to tell what will happen on Election Day. The drop-off is not the same for markets and the polls, however. That is, a comparison of correlations indicates that contemporary polls clearly beat the markets during the early stages of the campaign. It took longer for markets to reveal the outcome in the era before scientific polling was available.

None of this should detract from the early market prices as an election predictor. They were almost as good as post-1952 polls and arguably better when the comparison years are 1936-2008. Even at the benchmark date of thirty days before the election, this “contest” between early markets and contemporary polls is roughly even. These findings suggest that the knowledgeable political experts, whose wagering determined the betting odds, knew about as much at the end of the campaign regarding the eventual outcome as do contemporary observers with polls in hand.

But early in the campaign, the polls win the contest. In the modern era, polls provide evidence of who is winning and losing months before the votes are cast. Without polls, it took more time for the consensus to emerge regarding who would win. But this consensus did develop

and evolved so that by Election Day it became extremely accurate. Knowledgeable observers were able to gauge from the election campaign and its reception by the voters how the vote would turn out even without the polls.

Discussion and Conclusion

The impetus for this study is the contemporary interest in how well election markets perform as electoral predictors.. It turns out that the availability of public opinion polls affects the accuracy of election markets, but not as one might think. Once polls entered the picture, betting odds on presidential election outcomes were heavily dependent on what the polls were showing. At least that is the case regarding election eve market prices. With little information to go on that was not already captured by polls, the market prices reflected the polls plus some error. It did not help that when early polls were inaccurate (e.g., 1948) the market followed the polls. If there was information beyond the polls in 1948 to suggest that Truman might win, those who set the prices on election bets did not see it coming.

The big story is that before polling, election markets worked remarkably well. This is a story told before by Rohde and Strumpf (2004). As we show here, they worked so well that we are led to believe that the political cognoscenti of the times could read the political tea leaves about as well as modern day observers can from reading the polls. At the same time, it took time for campaign perceptions to crystallize during the early era. Early on in the presidential campaign, say midsummer, outcomes were less easy to anticipate than they are from the same vantage point in the campaign timeline today with observers armed with trial heat polls. Even on election eve, when early markets foreshadowed the correct winner, prices reflected a high level of uncertainty, much more than in the post-poll era.

The early curb markets were “thick” markets, with a high level of participation and investment (Rohde and Strumpf, 2004). Smart money investors were able to read clues from intuitive indicators and seemingly unreliable straw polls. In an atmosphere with many gamblers willing to bet based on their hearts on their political favorites, others could use their heads to exploit the market for personal gain. The latter group’s actions set the prices in a way that reflected the fundamentals of the election in a way that mimicked the understanding of election outcomes in the current era where the campaign news is dominated by stories about who is ahead in the polling horserace. In these respects, the early election markets approached the conditions for an ideal information market to work (Wolfers and Zitzewitz, 2008).

We can speculate about the lessons learned regarding current election markets. By all accounts betting markets withered with the dawn of polls. Interest in them has renewed with the advent in recent years of electronic election markets such as the pioneering Iowa Electronic Stock Market and, more recently Intrade (formerly Tradesports) and other online election betting sites. Although the performance of these new infant markets may not yet have live up to its promise it can improve (Erikson and Wlezien, 2008), with growing interest and higher volume of trading, the performance of future trading markets may surpass that of its ancestral forbearers from the pre-poll era. The same is not true of polls. In theory, the promise is that election markets, informed by polls plus other information, should perform better than polls alone. Whether this happens remains to be seen.

References

- Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz. 2008a. "Results from a Dozen Years of Election Futures Market Research." In Charles Plott and Vernon Smith (eds.). *Handbook of Experimental Economic Results*. Elsevier: North Holland.
- Berg, Joyce, Forrest Nelson, and Thomas Rietz. 2008b. "Prediction Market Accuracy in the Long Run." *International Journal of Forecasting* 24:285-300.
- 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.
- Caldeira, Gregory A. 2004. "Expert Judgment versus Statistical Models: Explanation versus Prediction." *Perspectives on Politics* 2:777-780.
- Erikson, Robert S. and Christopher Wlezien. 2008. "Are Political Markets Really Superior to Polls as Election Predictions?" *Public Opinion Quarterly* 72:190-215.
- Geer, John G. 1996. *From Tea Leaves to Opinion Polls*. New York: Columbia University Press.
- Karol, David. 2007. "Has Polling Enhanced Representation? Unearthing Evidence from the Literary Digest Issue Polls." *Studies in American Political Development* 21:16-29.
- Kernell, Samuel. 2000. "Life Before Polls: Ohio Politicians Predict the Presidential Vote." *PS: Political Science and Politics* 33: 569-574.
- Manski, Charles. 2006. "Interpreting the Predictions of Prediction Markets." *Economic Letters* 91:425-429.
- Page, Lionel. 2008. "Comparing Prediction Market Prices and Opinion Polls in Political Elections." *Journal of Prediction Markets* 2:91-97.
- Robinson, Claude E. and Robert E. Chaddock. 1932. *Straw Votes: A Study of Political Prediction*; Columbia University Press, 1932.
- Rohde, Paul W. and Koleman S. Strumpf. 2003. "Historic Presidential Betting Markets: Wagering on Presidential Elections." Unpublished paper.
- Rohde, Paul W. and Koleman S. Strumpf. 2004. "Historic Presidential Betting Markets." *Journal of Economic Perspectives* 18 (Spring):127-142.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. 2007. "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections." *Quarterly Journal of Economics* 122:807-829.

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.

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

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

Wolfers, Justin and Eric Zitzewitz. 2007. "Interpreting Prediction Market Prices as Probabilities." Working Paper.

Wolfers, Justin and Eric Zitzewitz. 2008. "Prediction Markets in Theory and Practice." *The New Palgrave Dictionary of Economics*, 2nd ed. Larry Blume and Steve Durlauf (eds.). London: Palgrave Macmillan.




Figure 1. Conversion between the Expected Democratic Vote and the Probability of a Democratic Victory.




Figure 2. Market Prices by Vote for President, 1880-1932. (Curved line represents the best-fit of an ogive representing the cumulative normal distribution.)




Figure 3. Market Prices by Vote for President, 1936-2008. (Curved line represents the best-fit of an ogive representing the cumulative normal distribution.)

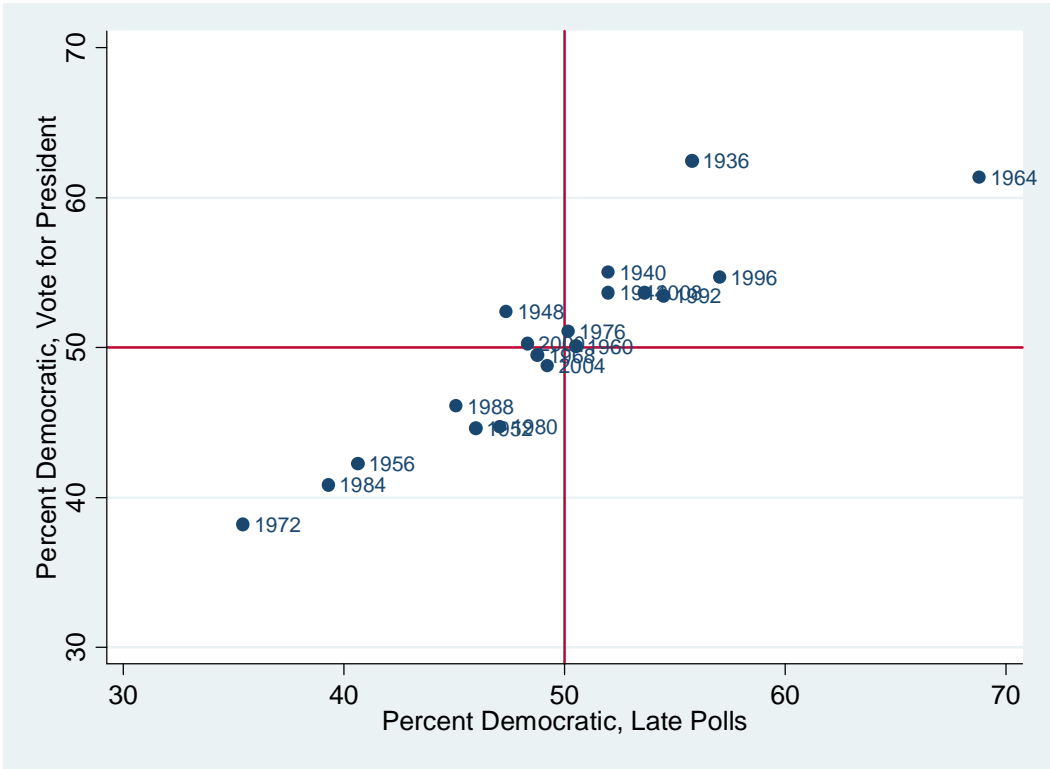


Figure 4. Predicting the Vote from the Polls, 1936-2008.




Figure 5. Market Prices by Vote for President, 1952-2008. (Curved line represents the best-fit of an ogive representing the cumulative normal distribution.)




Figure 6. Market Prices by Vote for President, 1880-1932 and 1952-2008. (Curved lines are best fitting ogives representing the cumulative normal distribution.)

	Correlation with the vote	RMSE predicting vote
Market Prices, 1880-1932 (N=14)	.93	2.88
As Imputed Vote Expectation	.95	2.38
Market Prices, 1936-2008 (N=15)	.80	3.98
As Imputed Vote Expectation	.80	3.92
Polls, 1936-2008 (N=18) ^a	.91	2.72

Market prices are measured in terms of the probability of a Democratic president as seen on election eve. The implied vote from market prices is in units of the “standardized vote” where the price (*p*-value) of a Democratic win is converted to the corresponding value in the normal distribution, calibrated in standard deviation units. Polls represent the Democratic share of the two-party vote in trial heat polls over the final week of the campaign.

^a The correlation is essentially unchanged (.90) when excluding years where market prices are not available—1964, 1968 and 1972.

	Coefficient (std. error)	Coefficient (std. error)
Polls	1.31*** (0.26)	1.38*** (0.24)
Market Price	-0.06 (0.05)	----
Market Price Implied Vote	----	-2.53 ^a (1.35)
Intercept	1.26 (0.66)	1.40 (0.64)
Adjusted R squared	.80	.82
RMSE	2.48	2.36

N=15, ***=significant at .001. .

^a The implied vote from market prices is in units of the “standardized vote” where the price (*p*-value) of a Democratic win is converted to the corresponding value in the normal distribution, calibrated in standard deviation units.

Table 3. Predicting the Vote from Market Prices and the Polls, 1952-2008		
	coefficient (std. error)	coefficient (std. error)
Polls	0.79*** (0.19)	0.81*** (0.19)
Market Price	0.01 (0.03)	----
Market Price Implied Vote	----	0.21 ^a (1.09)
Intercept	-0.41 (0.48)	-0.40 (0.51)
Adjusted R squared	.91	.91
RMSE	1.37	1.37
N=12, ***=significant at .001.		

^aThe implied vote from market prices is in units of the “standardized vote” where the price (p -value) of a Democratic win is converted to the corresponding value in the normal distribution, calibrated in standard deviation units

Table 4. Market Prices and Polls as Predictors of the Vote during the Campaign		
Correlations with %Dem., Pres. 2-Party vote	Democratic Vote Share in the Polls 1952-2008	Imputed Vote from Market Price 1884-1932
1 day	.98	.94
30 days	.91	.91
60 days	.91	.68
90 days	.78	.59
120 days	.79	.63
Number of cases	15	13
<p>Note: Market prices are measured in terms of the probability of a Democratic president as seen at the designated time of the campaign. The implied vote from market prices is in units of the “standardized vote” where the price (<i>p</i>-value) of a Democratic win is converted to the corresponding value in the normal distribution, calibrated in standard deviation units. Polls represent the Democratic share of the two-party vote in trial heat polls over the week ending at the designated time.</p>		