


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at SciVerse ScienceDirect

Electoral Studies

journal homepage: www.elsevier.com/locate/electstud

Markets vs. polls as election predictors: An historical assessment

Robert S. Erikson^a, Christopher Wlezien^{b,*}^aDepartment of Political Science, Columbia University, New York, NY 10027, USA^bDepartment of Political Science, Temple University, Philadelphia, PA 19122-6089, USA

ARTICLE INFO

Article history:

Received 16 August 2011

Received in revised form 20 April 2012

Accepted 23 April 2012

Keywords:

Prediction markets

Polls

Presidential elections

Forecasting

Horserace

ABSTRACT

Prediction markets have drawn considerable attention in recent years as a tool for forecasting elections. But how accurate are they? Do they outperform the polls, as some scholars argue? Do prices in election markets carry information beyond the horserace in the latest polls? This paper assesses the accuracy of US presidential election betting markets in years before and after opinion polling was introduced. Our results are provocative. First, we find that market prices are far better predictors in the period without polls than when polls were available. Second, we find that market prices of the pre-poll era predicted elections almost on par with polls following the introduction of scientific polling. Finally, when we have both market prices and polls, prices add nothing to election prediction beyond polls. To be sure, early election markets were (surprisingly) good at extracting campaign information without scientific polling to guide them. For more recent markets, candidate prices largely follow the polls.

© 2012 Elsevier Ltd. All rights reserved.

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). Election markets—betting markets on election outcomes—are a centerpiece of this discussion. Indeed, information-market proponents typically promote the purported success of election markets as a basis for advocating prediction markets in other realms. A common claim is that prices in election markets, such as the Iowa Electronic Market (IEM), predict elections better than “trial heat” polls that register vote preferences prior to Election Day (in addition to Arrow et al., 2008; Wolfers and Zitzewitz, 2008, see Berg and Rietz, 2006; Berg et al., 2008a,b). Rhode and Strumpf (2004) have shown that even the voting markets in the era before scientific polling had a strong record of accuracy.

Why is it widely believed that election market prices forecast better than polls? The rationale is that market price setters can anticipate electoral factors that are not

incorporated in the polls.¹ This idea that election markets are superior to the polls has traveled from economics into other academic realms (e.g., Caldeira, 2004; Sunstein, 2005) and popular literature as well (e.g., Surowiecki, 2004). As mentioned, the evidence most commonly cited as the basis for election market superiority comes from the Iowa “vote share” market. In a vote share market, one share of a candidate pays off on Election Day in direct proportion to the candidate’s final vote share. For instance, one share in Barack Obama at the end of the 2008 campaign paid off at 54 cents. A trader who bought Obama at, say, 50 cents

¹ Consider the earliest stages of an election campaign. At this point in time, poll results bear little resemblance to the final Election Day numbers (Campbell, 2000; Wlezien and Erikson, 2002; Erikson and Wlezien, 2012). But it is known that campaigns tend to steer the vote in the direction of the “fundamentals” of the election that are knowable in advance (see, e.g., Gelman and King, 1993; Bartels and Zaller, 2001; Vavreck, 2009). Election market actors have an obvious incentive to take into account this information and so there is reason to expect early market prices to reflect Election Day fundamentals that are not yet reflected in the polls.

* Corresponding author.

E-mail address: Wlezien@temple.edu (C. Wlezien).

and held the stock to Election Day, would obtain 54 cents for a four cent profit, an eight percent return.²

What market proponents have shown is that IEM vote share prices during campaigns correspond more closely to the final vote than do trial heat poll margins measured at the same time in the campaign (Berg et al., 2008a). In this narrow sense, it appears that markets prices are better than polls at predicting the exact Election Day vote margin at different points in the election cycle. The problem is that market prices and trial heat polls are not directly comparable. Although they are both measured as percentages, market prices and vote margins are measures of different things. Whereas market prices reflect forecasts of the expected vote on Election Day, the polls register vote intentions when the polls were taken. As is well known to those familiar with campaigns, vote margins in polls do not translate directly into the Election Day vote (Campbell, 2000; Wlezien and Erikson, 2002; Erikson and Wlezien, 2012). Specifically, leads tend to shrink over time. A proper comparison would take this into account, and translate the polls into forecasts based on their historical relationship with the vote. When converted in this way, the polls beat market prices as predictors more often than not (Erikson and Wlezien, 2008). That is, historically one could profit in the Iowa market by betting based on the size of the discounted poll results relative to the market prices.

Actually, few election markets are conducted as vote share markets. Most, both today and in the past, are winner-take-all markets (WTA), where investors wager on the winner, given the candidates' odds of victory.³ WTA markets are subject to an underdog bias similar to what we see in horserace betting (Wolfers and Zitzewitz, 2004, 2008). For instance, throughout most of Bill Clinton's two victorious presidential campaigns, the Iowa market overestimated his opponents' chances of catching up, compared to what a reasonable interpretation of the polls would suggest (Erikson and Wlezien, 2008).⁴

While these results are intriguing, the best test of whether markets prices and polls are the stronger predictor would be to pit them against each other as independent variables in a standard multivariate analysis forecasting the outcome. Unfortunately, this test requires a number of elections in which both poll data and market data are available, and 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 2008.⁵

As it turns out, however, we have much more national level data to draw on. Thanks to the efforts of Rhode and Strumpf (2004), we now have information about 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 et al. (2007) assembled 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 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 number of cases that presidential election forecasters typically use for their augury.⁶ Curb market prices from earlier years yield 14 “control” cases for 1880–1932—presidential elections with prediction markets but no scientific polls.⁷

With market prices from the end of each of 30 presidential elections, we are able to compare the performance of markets and the polls. For the modern polling era, we can directly test whether market prices or poll results really are the best predictors. We also can compare the accuracy of presidential betting market prices in the two historical periods, that is, before and after public opinion polls were introduced. This will inform us regarding whether markets improved (or worsened) with the availability of scientific polls.

A limitation of our analysis is that the market price data are all from late in the election campaign, when there is little time for future events to affect the election. Still, we can compare the accuracy of markets from an earlier era without scientific polls with polls from the current era when polls are trusted to forecast the outcome. And we can see whether election markets provide useful forecasting information during the current era of polling over and above what is in the polls from the closing days of the campaigns.

1. Prediction markets and the polls

It is commonly understood that before scientific polling was invented, political observers had a difficult time gauging public opinion on the issues of the day (Geer, 1996). It would seem to follow that elections of the pre-poll era were conducted under greater uncertainty about the outcome than elections today. Given their limiting conditions, observers would monitor various indicators for cues (Kernell, 2000; Karol, 2007; see also Robinson and

² One need not buy and hold, as a share can be sold at any time there is a buyer willing to pay the asking price.

³ With WTA markets, one share of a candidate pays off one dollar if the candidate wins and nothing if the candidate loses. For instance, a trader who bought one share of Obama for 50 cents and held it until Election Day received one dollar, for a 50 cent profit, fully 100% of the investment. In contrast with vote share presidential markets, there is little research on winner-take-all markets, but see Erikson and Wlezien (2008).

⁴ Holbrook (2010) shows that citizen forecasts of who will win (from National Election Study surveys) also performed well by comparison with electronic markets in recent years, at least at the very end of the election cycle.

⁵ In response to this limitation, Rothschild (2009) turned to analysis of polls and markets in a set of state-level elections from 2008, and found strong support for informed markets using poll projections and “debiased” market prices.

⁶ See, e.g., the forecasting symposium in *PS: Political Science and Politics*, vol. 41, issue 4 (October), 2008.

⁷ For details on the conduct of the early election markets, see Rhode and Strumpf (2004).

Chaddock, 1932). In the absence of scientific polls, election market prices were eagerly studied for evidence of election trends. As Rhode and Strumpf (2004) were the first to show, these early markets had a strong record of accuracy in terms of final prices predicting the winner. Rhode 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.

Based on the conventional wisdom, therefore, one would expect polls to dominate election markets as the election predictor when only one but not the other is present. When polling is introduced, election markets become more efficient because market players can incorporate the new type of information from the polls. If nothing else, they should reduce uncertainty about public preferences (see, e.g., Geer, 1996). Given that market prices are informed by both polls and additional information beyond the polls, it is widely believed that markets are superior to the polls for forecasting elections. For instance, market participants can take into account other information that is not beyond expressed preferences of the day. They have the incentive—the motivation—to do so. 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 et al., 2008a; Wolfers and Zitzewitz, 2004; Page, 2008).

Although these all may be reasonable expectations, it turns out that *none* holds up when put to the test of applying them to US history when predicting presidential elections on the eve of Election Day. Our analysis shows that market prices were 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 with an accuracy that is competitive with that of modern polling. 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 polls add only marginally to predict elections beyond the information election markets hold in the absence of polls. The third would have us believe that election market prices are not informative when polls are available. In the sections that follow 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.

2. 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 (Rhode 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. As the final available market prices, they can be interpreted as the market's judgment on the eve of the election regarding the probability of victory, e.g., a price of 34 cents registers a 34% probability of victory.⁸ We convert the prices into two-party probabilities of a Democratic win.⁹

Since winner-take-all prices are related to the actual vote in a decidedly non-linear fashion, it is helpful to transform the price data in some way. The obvious choice is to convert the prices as probabilities into the log of the odds. Accordingly, our dependent variable of interest becomes the log of the market odds (in shorthand the log-odds) that the Democrats will win the election. The log of the odds of a Democratic victory according to the market is related in a linear way to the actual vote.

Figs. 1 and 2 show the election-eve log-odds of a Democratic victory from the betting markets as a function of the actual vote. Fig. 1 plots these data for the 14 elections between 1880 and 1932, before the advent of polling. Fig. 2 plots the data for the 16 subsequent elections for which we have market data—excluding 1964, 1968, and 1972—through 2008. The comparison of the two figures is revealing.

As inferred by the tightness of the correlation between the log-odds and the vote, prices responded most crisply to the signal of the actual vote during the early, pre-poll era. Remarkably, for the 1880–1932 period, log-odd prices correlate at +.96 with the vote. As Rhode and Strumpf (2004) note, market prices on election eve predicted the correct winner in all but one (1916) of the 14 pre-poll elections. This degree of accuracy 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 +.65, and the market prices got the wrong popular vote winner in three of the 16 post-poll elections (1948, 1976 and 2000).¹⁰ (All correlations reported in the text are presented in one convenient table in the appendix.)

The most precise way to compare the patterns shown in Figs. 1 and 2 is to compare market prices in the different eras as a function of the signal of the actual vote. Specifically, for a given vote margin in the election, was the outcome more certain in the earlier or the later era, as seen by betting market investors?. For this test we turn from correlation to regression, and for each era compute the regression equation predicting the log-odds from the vote.

⁸ This seemingly straightforward interpretation of prices is the subject of some dispute—see Manski (2006).

⁹ For measurement details, see Snowberg et al. (2007) and Rhode and Strumpf (2004). All prices we use are from the twilight of the campaign. It is not always clear whether they are literally from the final day or from the final weeks. With a bit of license, the candidate prices of this paper are referred to as being on “election eve.”

¹⁰ The 2000 IEM market was to select the popular vote winner. Since the final pre-election price favored Bush over Gore, this election is classified as a mistake.

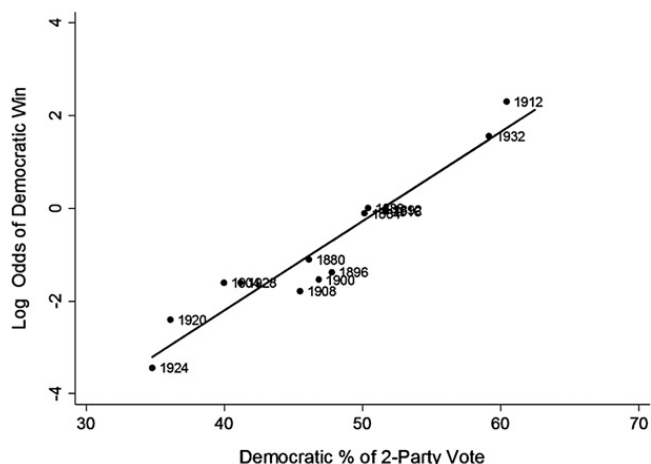


Fig. 1. Election market prices as a function of the actual vote, 1880–1932. Prices are measured as the log of the odds of a Democratic victory. Vote is measured as the Democratic percent of the two-party vote. The line is based on a regression relating prices and the vote.

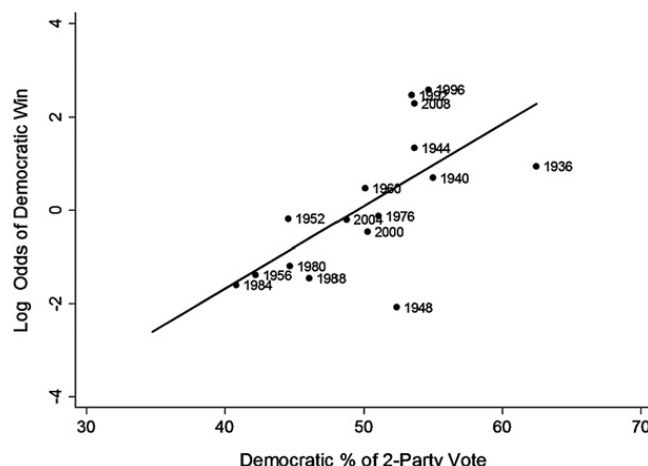


Fig. 2. Election market prices as a function of the actual vote, 1936–2008. Prices are measured as the election-eve log of the odds of a Democratic victory. Vote is measured as the Democratic percent of the two-party vote. The line is based on a regression relating prices and the vote.

For the pre-poll era, the slope is 0.19; for the post-poll era, it is a similar 0.17. Neither the pre-poll nor post-poll equation suggests bias, in that each regression line shows that when the vote was 50-50, the expected log-odds is close to the neutral zero value.¹¹ Not only are the two equations in the different periods strikingly similar. The slight edge in terms of steepness of the regression slope goes to the pre-poll markets.

What do these results imply about election-eve election markets, before and after the advent of polling? First, the correlation and regression-based evidence indicates that the knowledgeable price-setters were able to gauge degrees of relative Democratic vs. Republican strength with amazing clarity—greater, so it appears, than price-setters in the thinner markets of the polling era. Polls may have had a distorting effect on markets. When the polls have been accurate, so too have been market prices. But when polls have gone wrong, so have market prices. 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.

Second, for any given presidential vote outcome, on average the odds were at least as certain during the pre-poll era as later. We know this from the similar regression lines for the two eras. Before polls, landslides were seen as landslides and close elections were seen as close, just as in the modern era. And elections in between were seen with about the same degree of certainty as during the polling era. Compare the market prices for the fabled 1896 election and the 1988 election of recent history. Each was won by the Republican (McKinley and Bush respectively), with McKinley’s victory over Bryan a bit closer than Bush’s over Dukakis. But, the 1896 Wall Street curb market saw McKinley’s victory with greater certainty than did the

London punters regarding Bush’s win. If market prices reflect informed opinion of the moment, it would seem that scientific polls are not necessary for knowledgeable observers to foresee an election outcome.

It will certainly occur to the reader that, while Figs. 1 and 2 display the actual vote as the independent variable and the market prediction as the dependent variable, we could just as easily do it the other way around. Setting the vote as the independent variable is the appropriate way for causal inference, since the purpose is to see how the vote margin works as a signal to account for betting behavior. We can reverse the dependent and independent variables, with no intrinsic gain or loss of information. However, the purpose becomes strictly descriptive for forecasting purposes rather than causal inference (prices do not cause the vote). This reversal of dependent and independent variables allows a comparison of earlier and the later betting markets in terms of their ability to predict the vote division. The test then becomes: which equation has the smaller prediction error on average, as measured by the root mean squared error (RMSE)? The results are from the first two regression equations in Table 1. The early election markets win this test easily, with a RMSE of only 2.30 compared to 4.36 for election markets of the polling era. Clearly, markets with polls were substantially worse than markets without polls.

3. Prediction markets 1880–1932 vs. polls 1936–2008

The advent of scientific public opinion polling supposedly led to more accurate readings of public sentiment in advance of presidential elections. But, have polls performed any better than the early election markets, whose surprising accuracy was the subject of our discussion? If the comparison is with polls over the full period of polling, 1936–2008, the surprising answer is no. For our test, 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). In these polls, respondents were asked about how

¹¹ If the Democratic presidential vote is measured as the deviation from 50 percent, the intercept provides the estimate of bias. With this transformation of the vote, neither intercept is statistically significant at the .05 level.

Table 1

Predicting the presidential vote from election market prices (log of odds of democratic win) during the pre-polling era and from trial-heat polls during the polling era.

	1880–1932	1936–2008	
	(1)	(2)	(3)
	Coefficient (std. error)		
Constant	51.09 (0.70)	49.93 (1.09)	10.53 (4.42)
Market price of dem. win (log odds)	4.75 (0.42)	2.43 (0.75)	–
% Dem. in polls	–	–	0.80 (0.88)
Adjusted R-squared	.91	.39	.82
RMSE	2.30	4.36	2.72
N	14	16	19

Standard errors in parentheses. Market prices are for the last days of the campaign. Poll results are for the final week. Poll results and the vote are measured as the Democratic percent of the two-party vote. The key result is that the RMSEs are comparable in size. Pre-poll markets predicted the vote with almost as little error as do opinion polls. The RMSE for polls is lower (2.57) if the three elections without market data are excluded.

they would vote “if the election were held today” (with minor differences in question wording) or, less frequently, who they “would like to see win.” The exact measure is the Democratic candidate’s share of the two-party “vote” in the polls, ignoring third-party candidates.

Fig. 3 demonstrates the strong relationship between the vote and the polls for the 19 post-1932 presidential elections. The correlation is an impressive +.91 (+.90 excluding the three elections without market data). This still is slightly less than the pre-poll correlation of +.96 between the log-odds of market prices and the vote. But this comparison does not mean that the earlier markets were actually superior to today’s polls; the correlation comparison can be misleading when the two “samples” have different variances. A fairer comparison is to model the vote in the different eras as the dependent variable and predict it separately from the log odds of market prices (pre-polls) and the later poll results. The test is the comparison of the root mean squared error (RMSE) predicting the vote from these variables.

The third equation in Table 1 predicts the vote from the polls for the period 1936–2008. Here we can see that the RMSE of 2.72 is greater than that the 2.30 value we saw for the pre-poll election markets. Whereas the pre-poll market equation projects a confidence interval of about

4½ points (+/– two RMSEs), our estimate of the confidence interval based on actual poll performance is a slightly larger 5½ points. Though the difference is not large, the results imply that pre-poll market prices predicted presidential vote margins with at least the accuracy of contemporary polls. There is no denying that early election markets were surprisingly good at extracting campaign information without scientific polling to guide them. Somehow, the market price-setters were able to weigh the political evidence with a forecasting acumen similar to if not greater than what we learn today from the polls.¹²

4. Prediction markets vs. polls 1936–2008

Having compared prediction markets before and after polls, we next want to compare markets and polls directly. Since polling began in the 1930s, are polls or election markets the better forecaster on election eve? From our discussion of the various RMSEs, we already have the answer. Since the RMSE for polls (column 3 of Table 1) is smaller than the RMSE for the election markets for the same period (column 2 of Table 1), we already possess the information to answer that question. Polls are the better predictor.

There is an important follow-up question, however. That is, do market prices at least 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




Fig. 3. Predicting the vote from the polls, 1936–2008.

¹² One potential challenge to the acumen of early market price-setters is that they simply followed the election returns from the previous election. More than one might think, the current presidential vote historically tracks the previous election’s vote (Bartels, 1998). We can test whether the log of market prices tracked the previous election with a regression equation predicting the log-odds from the signal of the previous election result independent of the signal of the current election about to be held. The lagged vote does not add anything to predictions of markets prices during either the pre-poll era or the poll era. In each case the coefficient is incorrectly negatively-signed but not significant.

Table 2

Predicting the presidential vote from election market prices (log of odds of democratic win) and trial-heat polls during the polling era, 1936–2008.

	Coefficient (std. error)
Constant	17.99 (11.35)
Market price of dem. win (log odds)	−1.54(0.77)
% Dem. in polls	1.39 (0.23)
Adjusted R-squared	.82
RMSE	2.33
N	16

Standard errors in parentheses. Market prices are for the last days of the campaign. Poll results are for the final week. Poll results and the vote are measured as the Democratic percent of the two-party vote.

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?

The answer is no. The multivariate equation in Table 2 provides 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 and significantly different from zero when the polls are held constant. As can be seen, when polls and prices are raced in a multivariate equation predicting the vote, the poll coefficient is positive and statistically significant ($p < .001$). The price coefficient is actually negative, although not to a statistically significant degree. From this exercise we see that election eve market prices have not provided information about the election outcome beyond what is in election-eve polls.¹³

This admittedly is a hard test because, if the polls are accurate, virtually all information about the election will be reflected in final polls. And market prices clearly reflect the polls: election market log odds correlate with poll margins at +.86, as shown in Fig. 4. This correlation is considerably higher than prices' +.65 correlation with the vote. It thus appears that market prices follow the polls when the latter are available. This is true whether the polls get it right, as they usually do, or wrong. The 1948 upset reelection of Truman over Dewey is the classic example of where the polls went wrong. Yet the betting markets at the time did not contradict the flawed polls, and made Dewey the prohibitive favorite instead. Thus, over the era of polling, election-eve market prices have been no more informative about the election outcome beyond what was known from the polls. To the extent market prices have departed from the polls, they added error rather than further accuracy.

5. Prediction markets vs. the polls, 1952–2008

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

¹³ An equivalent demonstration is to flip the vote and the market price, making the price the dependent variable. When the log-odds of market prices are predicted from both the polls and the actual vote (as a signal of what is not in the polls), the coefficient for the vote is negative. This follows from the discussion in the text, and makes clear that market prices failed to incorporate information from the actual outcome that was not already present in the poll results.




Fig. 4. Predicting the market prices from the polls, 1936–2008.

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 interpretation. First, note that the correlation between the trial-heat polls and the vote for the post-1948 period is a near-perfect +.97, appreciably larger than the correlation (+.91) for the full 1936–2008 period, and slightly edges the +.96 correlation between market prices and the vote during the pre-poll era. Second, as seen from the first equation of Table 3, the RMSE predicting the vote from polls drops to 1.39, meaning that post-1948 polls make sharper percentage point predictions than did pre-poll markets (RMSE = 2.30).

With the new time frame, the log-odds of market prices predict the vote ($r = +.88$) with an accuracy far better than for the full 1936–2008 period ($r = +.65$). As shown in Table 3, Eq. (2), the RMSE predicting the vote from the log-odds of prices, declines to 2.30, precisely the same as the 2.30 RMSE for the pre-poll era. This improvement in market prices when starting the clock in 1952 can be traced to the markets’ reliance on the greatly improved polls.¹⁴ As polls improved, market prices did too.

We can also ask, with our new start period, do election eve market prices add relevant information not available in the late polls? The third equation of 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. As for the full 1936–2008 period, we find no evidence that market prices add information beyond the late polls. The polls are a highly significant predictor. With poll results

¹⁴ The polls-prices correlation for the 1952–2008 period is slightly stronger: +.91 instead of +.86 for the 1936–2008 period.

Table 3

Predicting the vote from market prices and the polls, 1952–2008.

	(1)	(2)	(3)
	Coefficient (std. error)	Coefficient (std. error)	Coefficient (std. error)
Percentage in polls	0.74 (0.05)	–	0.81 (0.18)
Log-odds of market prices	–	2.62 (0.45)	0.11 (0.64)
Intercept	12.62 (2.32)	4.81 (0.66)	9.35 (8.17)
Adjusted <i>R</i> -squared	.95	.75	.91
RMSE	1.39	2.30	1.37
<i>N</i>	15	12	12

Standard errors in parentheses. Market prices are for the last days of the campaign. Poll results are for the final week. Poll results and the vote are measured as the Democratic percent of the two-party vote. The three elections (1964, 1968 and 1972) without market data are excluded.

held constant, however, the coefficient for market prices, though positive, is far from statistically significant.¹⁵

6. Discussion and conclusion

The impetus for this study is the contemporary interest in how well election prediction markets perform as electoral predictors. Back in the era before scientific public opinion polls, election markets worked remarkably well. This is a story told before and told well by Rhode and Strumpf (2004). As we show here, early markets 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. Before polling, US presidential elections were not held under a cloud of uncertainty about the outcome. Somehow at least the “smart money” knew.

The curb election markets of the pre-poll era were “thick” markets, with a high level of participation and investment (Rhode and Strumpf, 2004). In an atmosphere with many gamblers using their hearts to bet on their political favorites, others could use their heads to exploit the market for personal gain.¹⁶ The group’s actions set the prices in a way that reflected the fundamentals of the election and 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).

When polls supplanted betting markets as the major mechanism for reading the electoral tea leaves, the quality of election forecasting actually sank, at least for the short run. In the early days of polling, through at least 1948, the polls were less accurate than the earlier markets. It turns out that the availability of public opinion polls affects the accuracy of election markets, but not as one might initially think. Once polls entered the picture, betting markets thinned, possibly due to a decline in their salience. And remaining presidential betting markets (offshore for

a time) became heavily dependent on what the polls were showing. Even when the 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, bettors in the election markets did not see it coming.

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 and other online election betting sites. Although the new infant markets may not yet have lived up to their promise, they can improve. With growing interest and a higher volume of trading, the performance of future trading markets may surpass that of its ancestral forbearers from the pre-poll era. In theory, the promise is that election markets, informed by polls plus other information, will perform better than polls alone.

We close with questions for further research. If election markets were so accurate (at least on election eve), could we say the same thing about public opinion on topical issues? Perhaps we could answer this question if betting markets could be found on referendum elections from the pre-poll era. Ideally we could broaden the analysis to cover election markets prices early in the campaign cycle. This data is harder to come by except for all but the most recent elections. Could it be that pre-polling election markets compared in accuracy to today’s polls when measured in, say, July of the election year? What would the data indicate about the volatility of campaign dynamics of the pre-polling era?

We do know the following: Before polling, election outcomes were not the major surprises we might think they were. Late in the campaign, betting market prices reflected the final outcome with remarkable accuracy—competitive with the election odds we would assign today based on careful readings of late polls.¹⁷ Once polls became available, betting markets on election eve simply follow the polls. This comes as no surprise. If polls are credible, there is little information for market investors to

¹⁵ The result is robust to excluding each election with replacement.

¹⁶ Smart money investors might have relied on the equivalent of modern-day models of election outcomes, focusing on the state of the economy and evaluations of the president’s public standing. For reviews of election prediction models in US presidential elections, see Lewis-Beck and Rice (1992), Campbell and Garand (2000).

¹⁷ Polls provide point predictions plus a margin of error. In this sense they provide a forecast of certainty which is usually interpreted by the media as whether the point prediction is within “the margin of error.” Winner-take-all election markets provide an exact degree of certainty about the election in the report of market prices or odds on victory.

incorporate regarding shocks to the electoral preferences from the time of the final polls to Election Day.

Acknowledgements

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 Texas at Dallas, University of Manchester, the University of Mannheim and the University of Sussex. For helpful comments, we thank Harold Clarke, Larry Bartels, Tereza Capelos, Ed Fieldhouse, Robert Ford, Rachel Gibson, Thomas Gschwend, Michael Hagen, Will Jennings, Marianne Stewart, Thomas Konig, and Justin Wolfers.

Appendix

Correlations among variables			
	Pre-polls (1890–1932)	Full polling era (1936–2008)	Later polling era (1952–2008)
Vote-market	.96	.65	.88
Market-polls	–	.86	.90
Vote-polls, all cases ^a	–	.91	.97
Vote-polls, only if market data	–	.90	.96
Number of elections	14	15	12

Vote = Democratic percent of Two-Party Vote. Market = Log of the Odds of a Democratic Win based on Market Prices. Polls = Democratic percent of Two-Party Vote in Final Week's Polls.

^a Includes 1964, 1968 and 1972, years for which market prices are not available but polls are.

References

- Arrow, K.J., et al., 2008. The promise of prediction markets. *Science* 320, 877–878.
- Bartels, L., 1998. Electoral continuity and change. *Electoral Studies* 17, 301–326.
- Bartels, L., Zaller, J., 2001. Presidential vote models: a recount. *PS: Political Science and Politics* 34, 9–20.
- Berg, J., Forsythe, R., Nelson, F., Rietz, T., 2008a. Results from a dozen years of election futures market research. In: Plott, C., Smith, V. (Eds.), *Handbook of Experimental Economic Results*. Elsevier, North Holland.
- Berg, J., Nelson, F., Rietz, T., 2008b. Prediction market accuracy in the long run. *International Journal of Forecasting* 24, 285–300.
- Berg, J.E., Rietz, T.A., 2006. The Iowa electronic market: stylized facts and open issues. In: Hahn, R.W., Tetlock, P. (Eds.), *Information Markets: A New Way of Making Decisions in the Public and Private Sector*. AEI-Brookings Joint Center, Washington, DC.
- Caldeira, G.A., 2004. Expert judgment versus statistical models: explanation versus prediction. *Perspectives on Politics* 2, 777–780.
- Campbell, J.E., 2000. *The American Campaign*. Texas A&M University Press, College Station, Texas.
- Campbell, J.E., Garand, J.C., 2000. *Before the Vote: Forecasting American National Elections*. Sage Publications, Thousand Oaks, Calif.
- Erikson, R.S., Wlezien, C., 2008. Are political markets really superior to polls as election predictions? *Public Opinion Quarterly* 72, 190–215.
- Erikson, R.S., Wlezien, C., 2012. *The Timeline of Presidential Elections: How Campaigns Do (and Do Not) Matter*. University of Chicago Press, Chicago.
- Geer, J.G., 1996. *From Tea Leaves to Opinion Polls*. Columbia University Press, New York.
- Gelman, A., King, G., 1993. Why are American Presidential Election polls so variable when votes are so predictable? *British Journal of Political Science* 23 (July), 409–451.
- Holbrook, T., 2010. Forecasting US Presidential Elections. In: Leighley, J. (Ed.), *Oxford Handbook of American Politics*. Oxford University Press, Oxford.
- Karol, D., 2007. Has polling enhanced representation? Unearthing evidence from the literary digest issue polls. *Studies in American Political Development* 21, 16–29.
- Kernell, S., 2000. Life before polls: Ohio politicians predict the presidential vote. *PS: Political Science and Politics* 33, 569–574.
- Lewis-Beck, M.S., Rice, T.W., 1992. *Forecasting Elections*. Congressional Quarterly Press, Washington, DC.
- Manski, C., 2006. Interpreting the predictions of prediction markets. *Economic Letters* 91, 425–429.
- Page, L., 2008. Comparing prediction market prices and opinion polls in political elections. *Journal of Prediction Markets* 2, 91–97.
- Robinson, C.E., Chaddock, R.E., 1932. *Straw Votes: A Study of Political Prediction*. Columbia University Press.
- Rhode, P.W., Strumpf, K.S., 2004. Historic presidential betting markets. *Journal of Economic Perspectives* 18 (Spring), 127–142.
- Rothschild, D., 2009. Forecasting elections: comparing prediction markets, polls and their biases. *Public Opinion Quarterly* 73, 895–916.
- Snowberg, E., Wolfers, J., Zitzewitz, E., 2007. Partisan impacts on the economy: evidence from prediction markets and close elections. *Quarterly Journal of Economics* 122, 807–829.
- Sunstein, C., 2005. Group judgments: deliberations, statistical means, and information markets. *New York University Law Review* 80, 962–1049.
- Surowiecki, J., 2004. *The Wisdom of Crowds*. Doubleday, New York.
- Vavreck, L., 2009. *The Message Matters: The Economy and Presidential Campaigns*. Princeton University Press, Princeton.
- Wlezien, C., Erikson, R.S., 2002. The timeline of presidential election campaigns. *Journal of Politics* 64, 969–993.
- Wolfers, J., Zitzewitz, E., 2004. Prediction markets. *Journal of Economic Perspectives* 18 (Spring), 103–126.
- Wolfers, J., Zitzewitz, E., 2008. Prediction markets in theory and practice. In: Blume, L., Durlauf, S. (Eds.), *The New Palgrave Dictionary of Economics*, second ed. Palgrave Macmillan, London.