

PREDICTION MARKETS VS POLLS – AN EXAMINATION OF ACCURACY FOR THE 2008 AND 2012 ELECTIONS

*Sveinung Arnesen**
Dept. of Comparative Politics
University of Bergen

Ole Bergfjord
Department of Economics
Bergen University College

ABSTRACT

Which technique is more accurate in predicting the outcome of U.S. presidential elections, polls or prediction markets? Several studies on this have been conducted in the past. We use market data and poll numbers, included adjusted version of the poll numbers, to reexamine this question based on the two last American presidential elections, in 2008 and 2012. We find that the market predictions outperformed the polls for these elections, and that adjusting the polls makes them less accurate relative to prediction markets, if anything.

Key Words: Prediction Markets, Election Forecasting, Election Campaign, Polls, Iowa Electronic Markets, US Presidential Election

1 INTRODUCTION

In the course of the last quarter century, prediction markets have received increased attention in academia and elsewhere. Such financial markets rely on traders investing real money into speculations about a future outcome. Prediction markets have been used to forecast everything from movie theater sales to the occurrence of political events. As a matter of fact, it all started with political elections. In 1988, researchers at the University of Iowa set up an experimental market later to be known as the Iowa Electronic Markets (IEM), and invited students at the campus to bet on the outcome of the U.S. presidential election. There were two questions the students should consider. The first was: Who would win the election? This type of market is known as a winner-takes-all market. The second question for the students to consider was: How big shares of the votes would the candidates receive?

* Corresponding author: Sveinung Arnesen, Department of Comparative Politics, University of Bergen, P.O. Box 7802, N-5007 Bergen, Norway. Email: sveinung.arnesen@isp.uib.no

The prices of the vote share contracts translate invariantly into predictions about the candidates' share of the votes on Election Day. If a vote share contract is traded at the price of 0.48, then the prediction of the market is that this candidate will receive 48% of the votes.¹ The last trade represents the current prediction, and as the market is open 24 hours, the predictions are updated continuously.

The vote share predictions are comparable to polls, and their accuracy can be compared directly against the election result. In 1988, the predictions were surprisingly accurate. On election eve, the absolute percentage error of the vote share predictions was 0.2 percentage points, while the polls were off by more than 2.5 percentage points. Since then, the Iowa Electronic Market has apparently continued to outperform the polls. Displaying results from 237 contracts in 49 markets from 13 countries, Berg *et al.* (2008a) find that the average market error for all the predictions was 1.49 per cent when measuring election eve closing prices. The corresponding poll error was 1.91 per cent. Berg *et al.* (2008b) present an analysis of the long-run forecasting ability of markets relative to polls for the US presidential elections from 1988 through 2004. For the final 100 days of the election campaigns, the market predictions were closer to the election results on 74 per cent of the days.

However, not everyone supports the claim that prediction markets forecast US presidential election outcomes more accurately than polls. Erikson and Wlezien have in two separate articles (Erikson and Wlezien, 2008; Erikson and Wlezien, 2012) demonstrated how polls may be adjusted to anticipatable rating developments during the presidential campaigns. Ex post facto, such enhanced polls are more accurate than market predictions. They demonstrated how polls may be forecasting election outcomes just as accurately as the IEM, if the polls are used to make regression model forecasts which take expected poll developments into account.

We complement the analysis and discussion performed by Erikson and Wlezien regarding the relation between prediction markets and polls by applying the adjustment procedure that Erikson and Wlezien suggested *before* the 2008 and 2012 US presidential elections had taken place. We then go on to measure the adjusted polls' accuracy compared against the market predictions for the same elections. Of course, the outcome of those elections are now known to us, but by restricting ourselves to following the procedure suggested before those elections, the two examples will serve truly as out-of-sample observations. By analyzing the predictive power of the markets in 2008 and 2012, and comparing this with various poll-based predictions, we hope to improve the understanding of how market prices in recent elections have predicted the election outcomes.

¹ The IEM operate with share of two-party vote, that is, third party candidates are ignored.

2 PREDICTION MARKETS VS. ENHANCED POLLS

Much of the data are the same as those used by Erikson and Wlezien, and we gratefully acknowledge that they provided much of the data we have used. We were not able to get hand of poll data from the elections before 1988, so our daily intercepts and regression coefficients must be based on data from 1988 onwards.

This methodology is described in detail by Erikson and Wlezien (2008), so we will only outline the framework here.

The incumbent party's two-party vote share is measured (or forecasted), and 50% is subtracted. Hence, what is measured is the incumbent party's support beyond the necessary 50%. If our variables are negative, this indicates a win for the opposition. A projected poll-based vote forecast for year y (V_y) is constructed based on the following equation:


$$V_y = \alpha_T + \beta_T P_{YT}$$

P_{YT} is the raw poll number at time T in year Y . α and β are coefficients based on historical data. One α and one β coefficient is calculated for each campaign day, from 200 to 1 days before the election. If $\alpha_T = 0$ and $\beta_T = 1$, the forecast will be identical to the observed poll results. If $\alpha_T > 0$, this indicates that the incumbent party historically have done better in elections than in polls at time T . This has often been found to be true in modern election research. If $\beta_T < 1$, this means that early leads tend to fade, i.e., that differences between the candidates tend to decrease towards the election. Again, this has often been the case. A two day lag has been used for the polls, as polls on average are estimated to have been conducted two days before publication.

As we only have used data from 1988 and onwards, our data will not be fully identical to those of Wlezien and Erikson. However, more than two thirds of the polls were published between 1988 and 2000. Adding the polls for the presidential races of 2004 and 2008, the missing years constitute around 26% of the data material.

For the new data from the recent elections, we have collected poll results from www.realclearpolitics.com, and adjusted the raw data as described above in order to calculate adjusted poll numbers. For comparison with prediction markets, we have used publicly available market data from the Iowa Electronic Market. In total, we have used data from 919 different polls from 1988 to 2000, 198 polls from 2004, and 217 polls from 2008.

Figure 1: US Presidential election 2008, enhanced polls versus IEM market




Some further description of the figure could be useful. As the "outcome incumbent" line is at around -4 %, this suggests that the incumbent party's share (i.e., John McCain's share) of the vote ended up at around $50 \% - 4 \% = 46 \%$. This also implies that the opposition (i.e., Barack Obama) ended up with around $100 \% - 46 \% = 54 \%$ of the vote. In some sense, this line at around -4 % represents the right answer, and is where perfect polls and market predictions would be situated.

We see from Figure 1 that in 2008, the market predictions from IEM were markedly and consistently closer to the actual outcome for Republican presidential candidate John McCain than the adjusted polls were. Only on very few occasions over the last 200 days before the election did the poll projections outperform the market predictions. Also, it turns out that the enhanced polls were anything but an enhanced version of the raw polls. While the mean absolute percentage error (MAPE) of the enhanced polls for the period was 4.23, the error was less than half the size for the raw polls, at 1.78. Also the raw polls were less accurate than the IEM predictions, however,

which had a MAPE of 1.68.² In short, this means that the election cycle in 2004 unfolded in a different way than what could be expected based on the elections in 1988-2000, so regressions based on these elections did nothing to improve the poll estimates in 2004, rather the opposite. It is worth noting that this not necessarily is unexpected. Based on the limited number of polls, in particular early in the campaigns, many of the α and β coefficients are calculated based on very few observations. A longer timespan would increase the number of observations at each point in the election campaign, but would at the same time introduce polls from a different era, which not necessarily are comparable to current elections.

Figure 2: 2012 US Presidential election, enhanced polls versus IEM market



² The mean absolute percentage error of 1.68 includes some days where the IEM vote share market clearly was out of equilibrium. As is evident from the visual data display in figure 1, the market price bounced early in the observation period up to a level that predicted a vote share that was off by as much as 14 percentage points. These few days had a solid impact on the mean error of the market. If these outlier observations are excluded, the MAPE for the IEM predictions drop to 1.48.

For 2012, the enhanced polls improve compared to the election before, and are somewhat more accurate than the raw polls are. The election campaign thus unfolded more in line with what could be expected based on previous elections. This could be expected, as the numerous and highly relevant data points from 2008 are now included into the regression model. However, measured against the IEM predictions the enhanced poll results are still less accurate. The MAPE of the enhanced polls is 1.31, the raw polls' is 1.42, while the market predictions' MAPE is 1.03. Also, the market predictions are closer to the outcome than the enhanced polls in 59 per cent of the days of the period.

Although mainly beyond the scope of this paper, some further differences between the two years could also briefly be mentioned. In 2008, there appears to be some co-movement between the two estimators. The IEM prediction is generally closer to the final result during the early phase of the campaign, although both the polls and the IEM predictions overestimate John McCain's final vote. Towards the end of the campaign, both polls and IEM predictions move closer to the true results, but still with the IEM predictions closer to the true results. In 2012, the pattern is remarkably different. Early in the campaign, both polls and IEM predictions are fairly stable. The markets were slightly closer to the final result than the polls, but markets overestimated Obama's vote whereas the polls underestimated Obama's vote. From the middle of the campaign, the two estimators move in opposite direction, and away from the final results. The markets overvalue Obama more and more, whereas the polls undervalue him more and more. The market prices moved back towards the true result during the end of the campaign, whereas the polls on average predicted a very close election to the very end of the campaign, i.e., a deviation around 0%, which indicates a clear undervaluation of Obama. These movements, as well as the differences between the two campaigns, could be due to fundamental changes in the nature of the race, how these were perceived, as well as estimator-specific factors and errors or coincidences. Which of these factors were the most important would be an interesting question for further studies, but will not be discussed further here.

Summarized, these comparisons show that prediction markets seemed to be a bit more accurate than polls when it comes to predicting the outcome of the 2008 and 2012 US presidential elections. This is in line with previous studies. Also, the suggested adjustment procedure of the polls to make them better at predicting election outcomes works poorly when applied to the cases of 2008 and 2012. Of course, some of this could be due to the lack of historical data from 1952 until 1984 in our model, yet we believe that the effect of such adjustments is relatively small most of the time.

Regarding the methodology, the choice of a two-day lag period does not seem obvious. It has been argued that this lag (between the last day of a poll and its inclusion in the data set) should correspond approximately to the time period between polling and publication. This means that the timing of a poll

should be roughly equivalent to the time it is published, which also makes comparison to PMs easier. However, the real lag is likely to be very different, decreasing both from election to election as technology has improved, and within election campaigns, as time becomes a more critical factor closer to the Election Day. More generally, we will argue that polls reporting public opinion approximately two days ago not can be compared to market prices reporting the public opinion today. Again, it seems difficult to measure how important this is for comparison.

3 GENERAL COMMENTS AND CONCLUSION

Based on our results above, the analysis shows that the poll adjustment procedure does not work well when tested on the 2008 and 2012 US presidential elections. The adjusted poll projections are not only less accurate than the election market prediction, but also less accurate than the raw polls in one of the cases. We draw from this that poll prediction models that fit historical data well do not necessarily predict future elections accurately. If such adjustments consistently improve the predictive power of polls, why are they not reported or incorporated in the poll models more often?

We would now like to turn to more general thoughts about the applicability of prediction markets in relation to elections, campaigns, and voters, and point out some reasons why prediction market should be relevant for politics and political science.

First, markets and polls have important complementary properties. While the survey respondents *represent* the voters, the market traders *interpret* the voters and anticipate their behavior. Hence, they offer a complementary view to the surveys on the course of election campaigns. The selection of traders, their motivation to reveal information, and their interest in politics are different from those of survey respondents. The little we know about the traders is that they in general score higher than average on socio-economic status measures such as income and level of education. And, to little surprise, they are politically well-informed. Where Downs (1957) argues that staying informed about politics is irrational, this is not true for the market traders. Most likely the traders were interested in political matters before they joined the prediction market, and if they were not, the monetary incentive gave them a good reason to become informed. If they possess unique information, or if they act upon new public information quicker than other traders do, they can make money in the market. Thus, the nature of prediction markets encourages traders to stay well-informed and to react quickly to news.

Under their own volition, the traders have decided to participate in the prediction market. They have a manifest interest in the topic they are contemplating, and they are motivated by monetary incentives or a general desire to test their abilities to predict the election outcome. The information they leave behind when making such predictions can reveal interesting

information about how they perceive the election campaign. For instance, before the 2012 election, traders would speculate about the size of the two-party vote shares Mitt Romney would receive, or how big the chances were for a reelection of incumbent president Barack Obama. As a by-product of their real money trades, the market displayed the traders' consensus on the presidential candidates' chances on Election Day. The current market price indicated the equilibrium where any trader would find it unprofitable both to buy and to sell. This equilibrium becomes the market's current prediction, and is subject to change whenever a trader finds it to be incorrect and is willing to back her opinion with a monetary investment. The traders speculate about the outcome as the election campaign unfolds, and left behind the traders is a trail of trades which shows how the standings of the candidates have evolved in the course of the campaign weeks.

Shaw and Roberts (2000) as well as Kou and Sobel (2004) have previously argued for the benefits of analysing prediction markets in relation to election campaigns. Prediction market prices are well suited to study campaign events, because "movement in the [IEM]... reflects the independent (non-predictable) impact of campaigning and campaign events." . One may perceive the series of polls and the series of market predictions as reflections of two distinct processes during the election campaign: The polls arguably reflect the voters' *enlightening* process where they gradually become aware of which candidate is the best match for their own policy preferences. Within the frames of Gelman and King's characterization of campaigns (Gelman and King, 1993), Shaw and Roberts assert that the movement of the market prices represents an *enlightened* process. That is, in the market some events are already incorporated into the predictions such as the anticipated convention bounces, where American presidential candidates get a popularity boost when their parties' conventions are held and get large media coverage. This distinction between what poll ratings and market predictions represent opens opportunities to study their data in conjunction. These time series have highly frequent observations, and invite scholars to study election campaigns through a new lens. To mention a few: Do the poll respondents react to events that the prediction market traders do not, and vice versa? Do the voters follow a predictable path in the course of the election campaign, which is known to politically well-informed people, or are the traders merely following the tracks of the opinion polls? How do partisan impacts influence the economy?

Data rich and understudied, prediction markets represent an untapped potential for scholars within the political science community who take an interest in election campaigns and electoral behavior. Whereas accurate information has been - and remains - the main criterion when comparing polls and prediction markets, it is also worth pointing out that prediction markets have other advantages compared to polls. The entertainment value of gambling might be hard to quantify, but in an era where it is increasingly hard to get good response rates for polls, it is worth noting that this entertainment

value attract traders, and that any information provided by market prices thus will be free, or at least considerably cheaper than poll information.³ To sum up, the people, the incentives, and the issues are different from one another in polls and prediction markets. What they have in common is that both methods yield data of relevance to the election.

Second, markets have the advantage that several different markets can be constructed, measuring different things. The most obvious example is that separate markets can exist for party/candidate vote share and for the probability of winning. In principle, pollsters could of course ask different questions – for instance “Who will you vote for?” and “What percentage of votes do you think the incumbent will get?” However, such dual polling is rarely done (or at least reported). A vote share poll could of course also be translated into a probability of winning, using for instance the methodology described by Erikson and Wlezien. However, this requires assumptions about distribution, and making those assumptions by definition leads to neglecting important and interesting issues. Assuming a distribution around an average without knowledge about the standard deviation could lead to wrong conclusions, but also ignores uncertainty and polarization. The Iowa Electronic Markets have run both vote share as well as probability markets for the last seven US presidential elections, and given the debate on increased political polarization in that country, it would be interesting to study whether or not that development is accompanied by lower uncertainty about the election outcome. This is because increased polarization is one of the potential causes of lower uncertainty about the outcome, since voters change candidate preferences less frequently when polarization is high.

Third, there are strong economic arguments in favor of markets. In particular, there is the fundamental issue of market equilibrium/efficiency. If polls are believed to provide better information about election outcome than prediction markets, everyone would be free to use the poll results to gain a profit in the prediction markets, and any such effect would be unlikely to persist. In other words: Whereas polls could be used to improve the accuracy of prediction markets, information from prediction markets are unlikely to have a direct effect on polls.⁴ Another benefit of prediction markets is their risk management properties. Although yet perhaps uncommon in practice, such markets provide an opportunity for individuals and firms to hedge

³ Running prediction markets might of course be expensive. However, the mere existence of such markets without subsidies (or, indeed, the need to ban such markets in some countries) is a strong argument for markets being superior to polls from an economic efficiency point of view.

⁴ Of course, prediction markets and media coverage could theoretically affect the polls. However, it is not obvious that any such effect would improve the accuracy of the polls; the effect could just as well be that respondents are prone to tell pollsters they will vote for the seemingly popular party/candidate, without really doing this. If this is the case, the poll accuracy decreases as a result of the extra information.

political risk - risks which in many settings are important, yet difficult to manage. Polls of course provide no such added benefits for stakeholders. It is finally worth mentioning that from an economic perspective, WTA markets look more promising than VS markets (Bergfjord 2013). The easy interpretation of quotes is attractive to market participants, increasing liquidity. Also, such contracts are more attractive to hedgers. If you want to hedge your exposure to the American presidential election, you care about who is elected; whether the vote share is 60-40 or 52-48 is largely irrelevant.

4 REFERENCES

- Berg, J., Forsythe, R., Nelson, F., Rietz, T., 2008a. Results from a Dozen Years of Election Futures Markets Research, *Handbook of Experimental Economic Results*. Elsevier, Amsterdam.
- Berg, J., Nelson, F., Rietz, T., 2008b. Prediction Market Accuracy in the Long Run. *International Journal of Forecasting* 24, 285-300.
- Bergfjord, Ole J. 2013. Prediction Markets and Contract Design. *Journal of Prediction Markets*, 5(2), 1-13.
- Downs, A., 1957. *An Economic Theory of Democracy*. Harper and Row, New York.
- Erikson, R.S., Wlezien, C., 2008. Are Political Markets Really Superior to Polls as Election Predictors? *Public Opinion Quarterly* 72, 1-24.
- Erikson, R.S., Wlezien, C., 2012. Markets vs. polls as election predictors: An historical assessment. *Electoral Studies* 31(3), 532-539.
- Gelman, A., King, G., 1993. Why Are American Presidential Election Campaign Polls so Variable When Votes are so Predictable? *British Journal of Political Science* 23(1), 409-451.
- Kou, S.G., Sobel, M.E., 2004. Forecasting the Vote: A Theoretical Comparison of Election Markets and Public Opinion Polls. *Political Analysis* 12, 277-295.
- Shaw, D.R., and Brian E. Roberts, 2000. Campaign Events, the Media and the Prospect of Victory: The 1992 and 1996 US Presidential Elections. *British Journal of Political Science* 30, 259-289.