Forecasting Elections: Do Prediction Markets Tells Us Anything More than the Polls?

Davis, Brent

Australian National University

July 2015

Forecasting Elections:

Do Prediction Markets Tells Us

Anything More than the Polls?

Dr Brent Davis[1]

School of Politics and International Relations

Australian National University

Keywords: voting behaviour/ choice; election forecasting

JEL: D72; C53; C54

[1] The views expressed in this article are solely those of the author, and do not necessarily reflect those of any other organisation with which he may be associated.

Abstract

Election forecasting is an expanding domain within political science, moving from the outer edges (as a novelty pursued by a few 'quants') toward the mainstream of the discipline. Amongst the most high profile of election forecasting techniques are prediction markets and vote-intention polls. While the weight of scholarly opinion appears to favour prediction markets over polls for election forecasting, there remain challengers and critics.

This article joins with the challengers and the critics, looking at whether this 'horse race' competition between election forecasting approaches is valid. Using data from the 2013 Australian federal election, we conclude such comparisons-of-forecasts are misplaced in the Australian context, as prediction markets and vote-intention polls appear to be independent of each other given information from one appears to have no impact on the other.

Introduction

Election forecasting is an expanding domain within political science, moving from the outer edges (as a novelty pursued by a few political 'quants') toward the mainstream of the discipline (and of increasing interest to politicians, strategists and the allied media). The emergence of election forecasting competitions, looking at which forecaster and/or technique delivers the 'best' or the 'better' forecast, underscores the still open nature of the range of alternate approaches. Amongst the most high profile of election forecasting techniques are prediction markets and vote-intention polls. While the weight of scholarly opinion appears to favour prediction markets over polls for election forecasting, there remain challengers and critics. This article joins with the challengers and the critics, looking at whether this 'horse race' competition between election forecasting approaches is valid, and whether one or other method adds any informational value to the other. Using data from the 2013 Australian federal election, we conclude such comparisons-of-forecasts are misplaced in the Australian context, as prediction markets and vote-intention polls are addressing different questions, and they appear to be independent of each other given information from one appears to have no impact on the other.

Prediction Markets

Prediction markets – also known as 'information markets' or 'events futures markets' – have emerged as potential sources of market information, intelligence and as platforms for forecasting the occurrence and/or outcome of an event. Such markets are analogous to the widely used futures markets in commodity and financial markets (Berg et al, 2003.

Prediction markets have emerged in economics, in the form of the "Economic Derivatives" markets founded by Deutsche Bank and Goldman Sachs, covering the outcome of certain prominent macro-economic indicators; in entertainment, with the somewhat misnomered Hollywood Stock Exchange, which focused on film revenues and awards; various closed (intra-firm) markets for prescribed commercial events or outcomes (for example, Hewlett-Packard and sales of information technology products; and, Siemens on computer software) and, potentially most interestingly, the 'terrorism futures' market proposed, but quickly scuttled following political pressure, by the United States Department of Defence (Wolfers and Zitzewitz, 2004).

However, prediction markets appear to have gained greatest traction in the political space, in particular regarding the outcome of elections, whether in presidential or congressional (most notably, the United States) or in Westminster systems (such as Australia). The platforms range from the Iowa Electronic Markets pioneered (and operated) by the University of Iowa in the United States (covering, inter alia, United States presidential elections) to the markets made by professional betting agencies such as Sportsbet; Betfair and others for Australian Federal elections (covering, inter alia, which party will win the election, the number of seats the major players will win, and even individual wins won/lost; see, for example, Leigh and Wolfers, 2006).

Prediction markets tend to be designed around a number of approaches, with market positions (ostensibly bets by the participants) paying off depending on either the outcome of a specific event (for example, Party A wins the election), the realisation of a continuous variable (Party B wins between 70 and 80 seats in the lower house in the election) or with some combination of these two (Party A wins the election, with a majority of between 15 and 20 seats in the lower house). In general terms, if the outcome bet upon occurs, then the 'bet pays off' at the odds on which it is placed (for example, a winning bet of $1 at 3-to-1 returns $4, being the original $1 plus the $3 dividend). For quantitative political scientists, the patterns of the bets placed, and the movements in the odds provide useful information on distribution of probabilities at different points in time, and across time, of the event(s) at hand occurring.

Prediction markets come in a number of market designs, the most common of which are continuous double auction markets (both with and without market-makers), where each seller/buyer of a position is matched by a counterparty buyer/seller of the same position. For example, someone buying a market position of say 3:1 of Party A winning the election is countervailing another person holding the opposing view (that is, they think the odds are 3:1 that Party A will not win the election). Movements in these odds (and their implied probabilities) signal changing opinions within the relevant market (and even between markets, where spread-betting can be undertaken). The odds ratio is essentially the mean price of buyers and sellers, while the spread of the bids/offers is indicative of the variance of beliefs of market participants.

Other forms of prediction markets include pari-mutuel systems, where all the money bet by the participants is placed in a common pool which is then divided amongst the winners. Pari-mutuels have particular application where it is not appropriate for counter-parties to bet against each other, but rather when the objective is to elicit the best forecast from a stable of participants (eg professional forecasters predicting, say, the inflation rate, the unemployment rate or some other defined economic indicator). By comparison, market scoring-style prediction markets, which appear to have found favour in the business world as a means of extracting information from internal markets (essentially their own employees), focus on rewarding a single person for the accuracy of their forecast. For example, a large firms creates an internal market designed to predict says the volume of sales of their product twelve months hence (measured as the number of units sold). While the most accurate forecasts 'wins the pool', other participants are allowed to purchase the right for the reward when they think they have a better forecast;

the size of the winnings-pool often being subsidised by a market-maker (usually the employing corporation).

Prediction markets are seen to have a number of important advantages both per se, and relative to opinion polling. These advantages include: they tend to respond more effectively to new information (Wolfers and Zitzewitz 2006; Rothschild, 2009; Snowberg et al 2012), given prediction markets are 'open' and functioning continuously (making them particularly useful for event impact analysis), compared to polls which are only taken occasionally and at certain points in time; they are consistent with the efficient markets hypothesis (that is, markets consolidate, and appropriately weight, all available information: Wolfers and Ztizewitz, 2006); they are resilient/less vulnerable to manipulation (Rhode and Stumpf, 2005); and, they are better forecasters of election outcomes than alternative mechanisms, most notably opinion polls (Leigh and Wolfers, 2006; Wolfers and Zitzewitz, 2006; Berg, Nelson and Rietz, 2007; Rothschild, 2009).

However, prediction markets also have a number of disadvantages, which can introduce bias and/ or noise into their operation and signal-sending. Amongst these disadvantages are: the tendency for some market participants to take positions based on their exogenous (to the market) existing biases or preference, for example relating to sporting identification (with a given team: Strumpf, 2004) or party identification (with a given political party: Forsythe, Nelson, Neumann and Wright, 1992; Forsythe, Reitz and Ross, 1999); the potential to attract out-right gamblers whose driving motivation is to take 'risky positions' in the market, and/or contrarians who gain some intangible utility from 'leaning against the wind'; and, their vulnerability to 'favourite' or 'long-shot' bias, where market participants tend to over-value very low probability events ('extreme longshots': Thaler and Ziemba, 1988; Erikson and Wlezien, 2008), which means prediction markets tend to perform poorly at forecasting such events, and/or a more likely in 'thinner' markets (that is, where there is a lower volume of transactions (Wolfers and Zitzewitz, 2004). Participants in political prediction markets tend to far from representative of the overall population. One United States study (Berg, Nelson and Rietz, 2007) found participants in political prediction markets where overwhelmingly young, white, educated males with high family incomes, and potentially not even eligible to vote; far from representative of the socio-demographic profile of the American electorate or population.

Prediction Markets and Polls

Prediction markets are becoming increasingly prominent in election analysis and forecasting. A steady literature has emerged examining the relative performance of prediction markets and voting intention polls as predictors of election outcomes; something of a 'horse-race' of prediction markets compared to political polls.

Studies have tended to focus on a number of issues regarding the relative performance of prediction markets and polls: which is more efficient in aggregating and messaging information on the dynamics of voter behaviour and intentions; and, which is better at forecasting the outcome of the event (more often than not a given election). While the weight of scholarly findings tends to show prediction markets are more efficient information collators and deliver better forecasts of election outcomes than are vote-intention polls (Kou and Sobel, 2004; Leigh and Wolfers, 2006; Wolfers and Zitzewitz, 2006), they are not without their detractors (Erikson and Wlezien, 2008).

One study (Leigh and Wolfers, 2006) of the relative performance of prediction markets and polls in forecasting the 2004 Australian federal election found political betting markets far and away out-performed vote intention surveys in forecasting the outcome of that election (both as to the outcome – a Liberal-National Party coalition win; and, the outcomes in three-quarters of marginal seats). The authors were also strongly critical of the value of polls for election forecasting, arguing the margin of error reported by polling houses substantially overstated the precision of poll-based forecasts, and the relative volatility of the polls compared to the prediction markets over time suggests movements in the polls are more about noise than signal.

However, other studies (Erikson and Wlezien 2008) challenge the evolving consensus favouring prediction markets over polls, pointing out analyses of their relative performances in forecasting election outcomes are not comparing like-with-like. When the polling data is rebased (as projections of the election outcome as a whole) and thus made comparable to prediction markets, the former tend to outperform the latter in their forecasting, and tend to be faster at absorbing new information which would likely impact the outcome of the election. The superiority of poll-projections over prediction markets expands further when smoothing methods (such as moving averages) are applied to the two data series. Indeed, they go so far to describe prediction markets as providing *"…information that distorts more than informs regarding the forecast of the election winner beyond what we can tell from the polls"* (Erikson and Wlezien, 2008: 212)

However, an important question which has attracted only limited scholarly attention is whether prediction markets provide additional information above and beyond that which can be extracted from the polling data. One study (Gurkaynak and Wolfers, 2005) comparing the results of the "Economic Derivatives" market to surveys of market economists of certain economic outcomes found the market-based predictions encompassed those of the survey-based forecasts, implying the former provided additional information beyond that available in the latter. However, another study focusing on economic indicators (Snowberg et al 2012) found a survey of business expectations did not encompass either initial unemployment claims or retail sales, suggesting the business confidence, and the labour force and retail series were complementary, and did not overlap in terms of the economic information provided. In essence, each series had informational value for market participants.

Another important issue in comparing the relative performance of prediction markets and polls in forecasting elections concerns the temporal horizons. Interestingly, and far from trivial, studies have found different relative performances at different times in the election campaign cycle, with betting markets performing more strongly than polls closer to the election date, but being "*substantially poorer*" (Leigh and Wolfers, 2006: 18) even three weeks before voters were required to cast their ballots. Erikson and Wlezien (2008) find prediction markets only catch-up to projection polls in the final stages of the election campaign. Similarly, other studies have found prediction markets do not always outperform polls: Berg et al 2003 put the 'win rate' for prediction markets at about 60 per cent for the elections studied using somewhat less interesting election eve data, and Berg et al 2007 put it at 68 per cent for 'week before election day' data. Looked at another way polls outperformed prediction markets between 30 to 40 per cent of the time.

Another important issue assessing the relative performance of election prediction markets and vote-intention polling is the fulcrum issue of whether the comparisons are valid, and in particular whether we are 'comparing like-with-like'. In most election prediction markets, participants are taking a position (making a bet) on what they think will be the outcome of the election; the collective decision of other people (the electorate as a whole). By contrast, vote-intention polls are generally expression of the political preferences of the individual respondent and their expected action on voting day (casting their ballot for Candidate A or Party B); they are expressing an opinion as to who they would like to win the election, not necessarily who they think will win the election. (While Newspoll does publish the results of polls asking respondents 'who do they think will win the coming election', these are only undertaken occasionally – just twice during the 2013 Australian federal election - and are not suitable for time series analysis. Nevertheless, voter expectations of the likely winner of the election broadly align with those of players in prediction markets – and the election outcome.) In short, prediction markets tend toward the macro-political (the behaviour of the electorate as a whole), while polls lean toward the micro-political (the behaviour of the individual voter). However, scholars (Berg, Nelson and Rietz, 2007: 4) have bridged this gap by arguing "*… forecasting is a natural use of poll results and polls are the best alternative forecasts available.*"

Which Leads Which

In contrast to the ample scholarly literature on the relative performance of prediction markets and vote intention polls as forecasting tools (and the not necessarily settled nature of the debate), seemingly scant attention has been paid to the questions of 'which leads which', and does 'one add value to or replicate the information contained in the other'. A priori it would not be unreasonable to expect published vote intention polls, which attract wide media reporting and discussion during the latter part of the electoral cycle and especially during more highly contested election campaigns, to lead prediction markets. In short, prediction market participants use published polls (or forecasts thereof) to drive their position taking. Scholarly studies appear to deliver mixed results on this question: Forsythe, Nelson, Neumann and Wright, 1992, finding prices in prediction markets do not follow polls, but rather tend to lead them; while Kou and

Sobel, 2004, find prediction markets do outperform polls alone, but only when the markets incorporate information from the polls.

The data series used in this study come from two sources: the prediction data from Betfair, a commercial betting house, and are based on bets placed by market participants as to which of the two major parties – the Australian Labor Party, or the Liberal National Party (Coalition) will form government after the 2013 Federal Election; and, vote intention survey from Newspoll, a leading market research organisation, which polls eligible voters on their vote intention 'if an election were held this Saturday?'. The Betfair data were available on a daily basis commencing 1 January 2013 in betting odds form, which were converted by the author to probabilities of each of the major parties winning (summing to 100 per cent; that is, taking out the 'bookies margin').

By comparison, the Newspoll data were generally available on a fortnightly basis commencing 12 January 2013, although became weekly during the election campaign period (that is, from 3 August up to 7 September 2013). The Newspoll data used were the two-party preferred estimates of the vote support for the ALP and for the LNP (that is, after an distribution of preferences from the supporters of other parties and independents, which averaged a not-insubstantial 19 per cent of primary vote support over the election campaign period). The final data set involved matching estimated probabilities from Betfair betting data to the same dates covered by Newspoll data (where Newspoll surveys were taken over three days, the middle date was used for referencing), producing a time series of 20 observations commencing 12 January and ending 7 September (election day) 2013.

Before outlining and reporting the data analytics, it is worth reviewing the political context of the 2013 Australian Federal Election, most notably the 'unusual' political dynamics within the ALP Federal Government. The 2013 political year commenced with what can only really be regarded as an outlier event – then Prime Minister Julia Gillard announcing at the end of January that the 2013 Federal election would be held on 14 September. The announcement effectively launched a more than 7 month election campaign, in contrast to the conventional Australian practice of the incumbent Prime Minister keeping the election date a closely guarded secret to gain tactical political advantage over their opponents, followed by a formal election campaign of around 5 weeks duration.

The second notable feature of the 2013 political year was the instability within the leadership of the ALP Government, in particular the ongoing rivalry, and barely concealed hostility, between former ALP Prime Minister Kevin Rudd and his followers, and incumbent ALP Prime Minister Julia Gillard and her supporters within the parliamentary party (and within the ALP machine). This was reflected in the 'leadership challenge which never happened' in March 2013 (when Rudd did not take up a chance to challenge Gillard when she gave him the opportunity to do so), and then 'the leadership challenge which did happen' in late June 2013 when Rudd responded to an opportunity to challenge Gillard, and defeated her in an acrimonious party room ballot. As such, the ALP Government spent much of 2013 engaged in internecine warfare, which was inevitably reflected in the polls and the prediction markets, as can be seen in Graph 1 following.

Graph 1: LNP Support in Polls and Prediction Markets

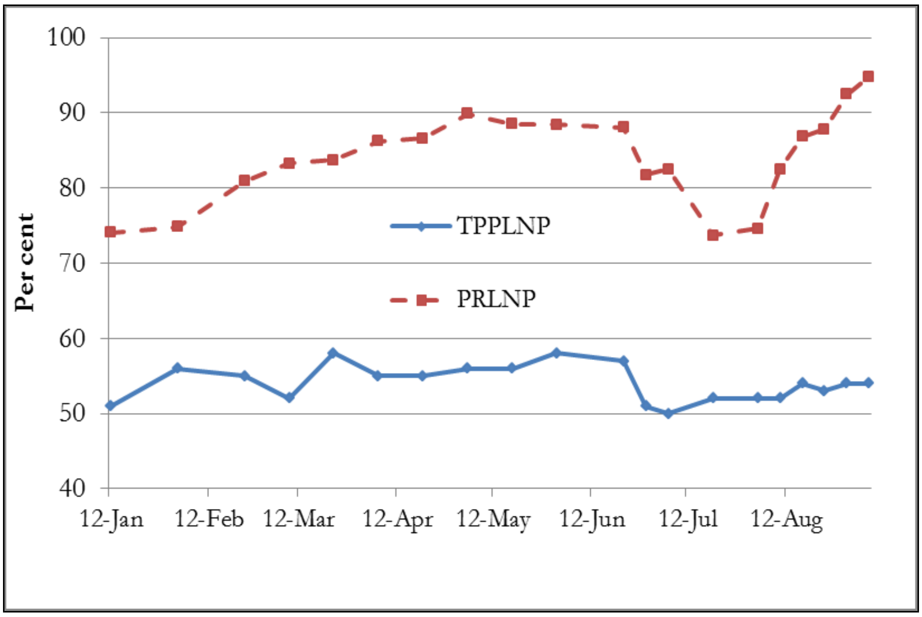

The two party preferred vote for the Liberal National Party (TPPLNP) remained fairly stable for much of the January to June period, if anything drifting upward slightly over the period, averaging around 55.4 per cent (figures for the ALP are just the converse – TPPALP = 100 – TPPLNP). However, with the return of Kevin Rudd to the leadership of the ALP, and thus the Prime Ministership, the TPPLNP dropped from about 58 per cent just before his return to 51 per cent just after, before edging back up to around 54 per cent in the final stages of the 2013 federal election. By comparison, the prediction market saw the probability of an LNP win rise from around 74 per cent in mid-January to 88 per cent just before Rudd's return to the Prime Ministership, after which it dropped back to about 74 per cent before rising to 92 per cent the week before the election and 95 per cent on election day (again, the figures for the ALP are just the converse – probability for the ALP = 100 – probability for the LNP; PRLNP). A pairwise correlation shows the reasonably degree of coincidence between the two series (r = 0.412), which is on the edge of conventional statistical significance (r = 0.071).

While the two series tend to move together, the underlying issue remains: do they essentially report the same information or does one contain additional information beyond that available in the other. This leads to a null hypothesis of H0 = non-additionality, that is the two series essentially contain the same information; and, an alternative hypothesis of HA = additionality, that is each of the series' contain some element of information not available in the other. To examine these hypotheses we will follow, with some extension, the approach taken by Alymer and Gill (2003) in a similar study into the additionality of business confidence surveys to official commercial and economic data releases. In their three step schema: the first step involved examining the pair-wise correlations between the indicators of interest; the second step adding one of the variables to an auto-regression model of the other variable; and, the third step involved performing Granger causality tests between the variables of interest. In this study, our modification is to broaden the first step to include cross-correlations (to identify lagged correlations), and concordance analysis (to see if the two series track each other across cycles).

As reported earlier, a pairwise correlation of TPPLNP and PRLNP shows a reasonable degree of coincidence (r = 0.412), which is on the edge of conventional statistical significance (r = 0.071), suggesting some degree of commonality in the information contained in the two series. A cross-correlation (that is, measuring correlations across different lag-time periods) shows TPPLNP tends to lead PRLNP by one period (r = 0.455; p = 0.05), rather than the other way around (r = 0.239; p = 0.324), suggesting information contained in the polls may feed into prediction markets, but not the other way around. A concordance analysis (rc = 0.012; p = 0.104), however, would indicate to no practical (or statistically significant) agreement in the pattern of the simultaneous movement between the two series. Taken together, these simple, first look tests would suggest at best a weak-to-modest association between the polls and the prediction markets.

The next step involves estimating a set of simple auto-regressive models regressing: TPPLNP on its own lagged values; TPPLNP on its owned lagged values, and on present and lagged values of PRLNP; PRLNP on its own lagged values; and, PRLNP on its own lagged values, and on present and lagged values of TPPLNP. If the present and lagged values of the non-dependent variable in each model do not achieve statistical significance at conventional levels, then this would add weight to the null ('non-additionality') hypothesis; if they are statistically significant, then this would favour the alternative ('additionality') hypothesis. Table 1 following reports the results of these auto-regressive models, with the top panel reporting models where TPPLNP is the dependent variable, and the bottom panel where PRLNP is the dependent variable. (EDS: TABLE 1 ABOUT HERE)

Looking first at Model 1 (TPPLNP regressed on its own lagged values) shows only TPPLNP lagged one period comes close to having any meaningful impact on TPPLNP (b = 0.44; p = 0.09), and overall the model has poor explanatory power (R Sq = 0.19; Adj R Sq = 0.08). Adding present and lagged values of PRLNP (Model 2) shows neither of the newly added variables are statistically significant, with no real movement in the explanatory power of the model. Turning to Model 3 (PRLNP regressed on its own lagged values), it would appear PRLNP lagged one period plays an important deterministic role (b = 0.978; p = 0.001), with the model having a reasonable degree of explanatory power (Adj R Sq = 0.523). Adding current and one period lagged values of TPPLNP (Model 4) again shows neither of the new variables are statistically significant, leading only to a reduction in model goodness of fit. Both sets of results would appear to favour the null ('non-additionality') hypothesis – that is, neither adds anything to the explanatory power of the other.

The third step is to apply Granger-causality tests to the TPPLNP and PRLNP variables to look for additionality. The null hypothesis of Granger-causality is current and past values of one variable contain no additional information for estimating another variable beyond that found in lagged values of the latter variable. In the current case, current and lagged values of TPPLNP/PRLNP contain no additional information in explaining PRLNP/TPPLNP than would be obtained from lagged values of PRLNP/TPPLNP. Rejection of this null hypothesis would signal additionality. The results of Granger-causality analysis can be found in Table 2. (ED: TABLE 2 ABOUT HERE). These analysis considered four possible situations: TPPLNP Granger-causes PRLNP; PRLNP Granger-Causes TPPLNP; there is instantaneous feedback between TPPLNP and PRLNP; and, the total correlation between TPPLP and PRLNP; across lags ranging from 1 to 4 periods. A review of Table 2 indicates we cannot reject the null hypothesis that TPPLNP does not Granger-cause PRLNP, or PRLNP does not Granger-cause TPPLNP at any of the four lags considered, with the single exception of TPPLNP Granger-causing PRLNP at lag 4 ($\chi 2 = 12.36$; $p = 0.02$). At best, it can be said PRLNP has no additionality upon TPPLNP, while TPPLNP has some delayed additionality upon PRLNP.

Summary and Conclusion

There is a growing volume of scholarly literature examining the performance of prediction (also known as betting) markets, both per se and relative to various forms of vote-intention polls, in forecasting election outcomes. Prediction markets are seen to have a number of advantages over vote intention polls, such as they are better able to aggregate all relevant information, and have demonstrated superior forecasting ability. However, prediction markets are also seen to have a number of disadvantages, such as market-positions taken may reflect the underlying partisan biases (or gambling/ risk-seeking inclinations) of players.

This study was initially motivated by a concern that forecasting competitions between prediction markets and vote intention polls may not be comparing like-with-like. As observed, most election prediction markets involve participants taking a position on what they think will be the outcome of the election – that is, how the electorate as a whole will vote: will Party A or B win the plurality; a macro-political perspective. By contrast, vote intention polls ask respondents of their vote preference, usually in the form of 'who will you vote for on election day' or some variant thereof; vote for Party A, or B or C or so on; a micro-political perspective. The latter asks the individual about their personal preference; the former asks participants collectively their expectations of the election outcome.

This study did not attempt to add directly to this literature, but sought to answer a different question. That is, whether the two variables of interest (TPPLNP and PRLNP) contained any additional information not available in the other series: did information on PRLNP help explain TPPLNP; did information in TPPLNP help explain PRLNP. The answer appears, by and large, to be in the negative – information in one does not help improve our knowledge of the other. This situation may reflect one of two possibilities. The first is the two series are largely duplicates of each other, although their respective poor performances in the Granger-causality testing in particular would suggest this is not the case. The second is the two series are largely independent of each other, which would be consistent with the outcomes of the auto-regressive modelling and the Granger-causality testing. At the same time, intuition would likely support the interpretation of independence of the two series, given they essentially measure two different things – TPPLNP, the vote intention of individuals polled (ie micro-political); PRLNP, the expectation of political tragics as to the outcome of the election (ie macro-political). While time series' of each metric can be usefully taken up for forecasting future values of itself, given their seemingly independent nature, comparisons of such forecasts are like comparing the dynamics and the outcomes of football matches with an archery competition.

Table 1: Auto-regressive models for TPPLNP and PRLNP

Dep Var	tpplnp	Model1				tpplnp	Model2		
	b	t	p	beta		b	t	p	beta
Const	31.79	2.16	0.05	…		28.09	1.88	0.08	…
tpplnp(-1)	0.44	1.79	0.09	0.45		0.33	1.28	0.22	0.33
tpplnp(-2)	-0.03	-0.11	0.91	-0.03		-0.11	-0.42	0.68	-0.12
Prlnp	…	…	…	…		0.15	1.01	0.33	0.35
prlnp(-1)	…	…	…	…		0.01	0.08	0.94	0.03
R Sq	0.190					0.300			
Adj R Sq	0.082					0.085			
RMSE	2.25					2.25			
AIC	82.99					84.36			
SIC	85.66					88.81			

Dep Var	prlnp	Model3				prlnp	Model4		
	b	t	p	beta		b	t	p	beta
Const	28.95	1.91	0.08	…		6.56	0.25	0.804	…
prlnp(-1)	0.978	3.89	0.001	0.962		0.871	3.16	0.007	0.856
prlnp(-2)	-0.313	-1.27	0.223	-0.314		-0.286	-1.13	0.278	-0.288
Tpplnp	…	…	…	…		0.444	1.00	0.336	0.195
tpplnp(-1)	…	…	…	…		0.095	0.21	0.833	0.042
R Sq	0.579					0.619			
Adj R Sq	0.523					0.502			
RMSE	3.679					3.76			
AIC	100.70					102.91			
SIC	103.37					107.36			

Table 2: Granger-Causality between TPPLNP and PRLNP

	Granger Causality			Granger Causality			Instant Feedback			Total	Total
	TPPLNP	**TPPLNP**		**PRLNP**	**PRLNP**		**TPPLNP**	**TPPLNP**		**TPPLNP**	**TPPLNP**
	>>	**>>**		**>>**	**>>**		**<<>>**	**<<>>**		**with**	**with**
	PRNLP	**PRNLP**		**TPPLNP**	**TPPLNP**		**PRLNP**	**PRLNP**		**PRLNP**	**PRLNP**
	$\chi 2$	**P**		$\chi 2$	**p**		$\chi 2$	**p**		$\chi 2$	**p**
lag 1	0.57	0.45		0.16	0.69		0.78	0.38		1.51	0.68
lag 2	0.45	0.80		1.20	0.55		1.15	0.28		2.79	0.73
lag 3	2.08	0.56		2.28	0.52		1.21	0.27		5.57	0.59
lag 4	12.36	**0.02**		1.27	0.87		0.12	0.72		13.76	0.13

Bibliography

Aylmer C and Gill T (2003) "Business Surveys and Economic Activity", Research Discussion Paper 2003-01, Reserve Bank of Australia, Sydney

Berg J, Forsythe R, Nelson F and Rietz T (2003) "Results from a Dozen Years of Election Futures Market Research", unpublished monograph, University of Iowa

Erikson R S and Wlezien C (2008) "Are Political Markets Really Superior to Polls as Election Predictors?" 72 Public Opinion Quarterly 190 -215

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

Forsythe R, Reitz T and Ross T (1999) "Wishes, Expectations and Actions: Price Formation in Election Stock Markets" 39 Journal of Economic Behaviour and Organisation 83 - 110

Gurkaynak R S and Wolfers J (2005) "Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts Uncertainty and Risk", Discussion Paper 1899, Institute for the Study of Labor (IZA), Bonn

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

Leigh A and Wolfers J (2006) "Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets", Discussion Paper 1972, Institute for the Study of Labor (IZA), Bonn

Rhode P W and Stumpf K S (2005) "Manipulation in Political Stock Markets: A Field Experiment and A Century of Observational Data", Unpublished Monograph, Department of Economics, University of Arizona.

Snowberg E, Wolfers J and Zitzewitz E (2012) "Prediction Markets for Economic Forecasting", Discussion Paper 6720, Institute for the Study of Labor (IZA), Bonn

Tziralis G and Tatsiopoulos I (2007) "Prediction Markets: An Extended Literature Review" 1 Journal of Prediction Markets 75 - 91

Wolfers J and Leigh A (2002) "Three Tools for Forecasting Federal Elections: Lessons from 2001" 37 Australian Journal of Political Science 223 - 240

Wolfers J and Zitzewitz E (2004) "Prediction Markets", Discussion Paper 03-25, Stanford Institute for Economic Policy Research, Stanford University

Wolfers J and Zitzewitz E (2006) "Prediction Markets in Theory and Practice", Discussion Paper 1991, Institute for the Study of Labor (IZA), Bonn