

Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets*

ANDREW LEIGH

*Research School of Social Sciences, Australian
National University, Canberra, ACT, Australia*

JUSTIN WOLFERS

*University of Pennsylvania, Philadelphia, PA, USA;
Centre for Economic Policy Research, London, UK;
Institute for the Study of Labor, Bonn, North
Rhine-Westphalia, Germany and National Bureau of
Economic Research, Cambridge, MA, USA*

We review the efficacy of three approaches to forecasting elections: econometric models that project outcomes on the basis of the state of the economy; public opinion polls; and election betting (prediction markets). We assess the efficacy of each in light of the 2004 Australian election. This election is particularly interesting both because of innovations in each forecasting technology, and also because the increased majority achieved by the Coalition surprised most pundits. While the evidence for economic voting has historically been weak for Australia, the 2004 election suggests an increasingly important role for these models. The performance of polls was quite uneven, and predictions both across pollsters, and through time, vary too much to be particularly useful. Betting markets provide an interesting contrast, and a slew of data from various betting agencies suggest a more reasonable degree of volatility, and useful forecasting performance both throughout the election cycle and across individual electorates.

I Introduction

* J. Wolfers gratefully acknowledges the support of a Hirtle, Callaghan and Co.-Arthur D. Miltenberger Research Fellowship, and the support of Microsoft Research and the Zell/Lurie Real Estate Center. We are grateful to Lisa Cameron, Mark Crosby, Gerard Daffy (Centrebet), Mark Davies (Betfair), Murray Goot, Simon Jackman, Sol Lebovic (Newspoll), Bryan Palmer, Euan Robertson (ACNielsen) and Mark Worwood (Centrebet) for generously sharing data with us. Thanks to David Bednall, Brett Danaher, Sinclair Davidson, Murray Goot, Simon Jackman, John Quiggin, Betsey Stevenson, Howard Wang and Eric Zitzewitz for useful insights and valuable feedback on an earlier draft. All remaining errors and omissions are our own.

JEL classification: D72, D84

Correspondence: Andrew Leigh, Research School of Social Sciences, Australian National University, Canberra, ACT 0200, Australia. Email: andrew.leigh@anu.edu.au

There has recently been a surge of interest in forecasting election outcomes (Cameron & Crosby, 2000; Fair, 2002a; Wolfers & Leigh, 2002; Leigh, 2004a; Cuzán *et al.*, 2005; Jackman, 2005; Wolfers & Zitzewitz, 2006a). Following the 2004 Australian federal election we now have even more valuable data with which to evaluate competing theories. In light of this new data, we assess the relative efficacy of three approaches: econometric models based on recent economic data; betting markets (also called 'prediction markets'; Wolfers & Zitzewitz, 2004) and opinion polling.

Australia's 2004 poll proved to be especially interesting for several reasons. John Howard was re-elected with an increased majority – an outcome that few pundits expected. In addition, we saw

important innovations in the application of each of the forecasting technologies we examine in this paper. The 2004 election occurred during arguably the best macroeconomic times in a generation, but even so, Howard's vote exceeded the expectations of economic forecasting models. This election cycle saw the establishment of several new polling firms, and the use of new polling technologies. Yet at the same time pollsters faced new challenges, such as declining response rates and rising mobile phone penetration. On net, we find little evidence that polls are becoming more accurate. Finally, election betting became even more widespread, and we now have data from five betting markets run in Australia and in the UK.

To preview our results, we found that economic models provide useful forecasts, particularly when making medium-term predictions. Consistent with international evidence (Berg *et al.*, 2008, forthcoming), betting markets provide extremely accurate predictions, and in this election they not only predicted a Howard victory, but also the outcome in three-quarters of marginal seats. The performance of the polls was quite variable, and one's assessment really depends on which pollster one relies on: Morgan's poor performance in the 2001 election was repeated again, potentially casting doubt on their emphasis on face-to-face polling. We also present a method for converting polling results into an assessment of the likelihood of victory. These results strongly suggest that the 'margin of error' reported by the pollsters substantially overstates the precision of poll-based forecasts. Furthermore, the time-series volatility of the polls (relative to the betting markets) suggests that poll movements are often noise rather than signal.

II The 2004 Election Forecasts

Table 1 shows key forecasts from the 2004 election cycle. We begin by presenting forecasts of the share of the two-party preferred vote won by the Coalition at three time horizons: election-eve, 3 months prior to the election, and a year prior to the election. By construction, forecasts of the vote share of the Australian Labor Party are equal to 100 minus the Coalition prediction.

Panel A shows several forecasts of the two-party preferred vote share. The Coalition government was ultimately re-elected with 52.74 per cent of the vote, an increase of 1.79 per cent on their performance in 2001. The economic models correctly picked the election winner, but surprisingly (given the strength of the economy) still underpredicted the Coalition's performance. That

said, these models did predict a swing towards the Coalition: having predicted the 2001 election almost perfectly, the election-eve economic models suggested that the improved state of the macroeconomy would lead to a swing in the order of 0.3–0.7 per cent towards the government. Importantly, these models were all pointing to a Coalition victory from at least a year prior to the election – a reflection of the stable and robust macroeconomic conditions during the entire election cycle.

The polls paint a more variable picture. By election eve, ACNielsen and Galaxy were predicting a Coalition victory, Newspoll declared a dead heat, and Morgan predicted a Labor victory. Three months earlier, only Galaxy was predicting a Howard victory, and a year prior to the election the polls suggested, on average, a dead heat. The variability of predictions both across polls, and within polls over time, points to the difficulty of using them to form reliable forecasts.

In Panel B of Table 1, we convert each of our measures into an implied prediction of the probability of a Howard victory. Doing so requires us to estimate the probability that each indicator pointing to a Coalition victory is due to chance. For the polls, our numbers are guided by the estimated margin of error suggested by the polling companies. These estimates typically refer to sampling error, and given the binomial nature of the elicited voting intentions (when stated on a two-party preferred basis), the standard error (SE) is calculated simply as $SE = \sqrt{q(1-q)/n}$, where q is the proportion of respondents who say they will vote for the Coalition, and n is the sample size. For example, assuming only classical sampling error, a poll with a sample size of 1000, in which the parties were evenly matched, would have a standard error of 1.6 per cent (and a corresponding 95 per cent confidence interval of ± 3.2 per cent), while a poll with a sample size of 2000 would have a standard error of 1.1 per cent (a 95 per cent confidence interval of ± 2.2 per cent).

Thus $z = (q - 0.5)/SE$ yields a normally distributed test statistic for whether or not a majority of the population intend to vote for the Coalition. The estimates reported in Panel B are the P -value for a test of whether the Coalition will receive more votes than the ALP.¹ In simple terms, if we

¹ In classical statistics, the P -value describes the probability of observing the polling data given the null hypothesis, while the statement that we are interested in making refers to the probability that the null hypothesis is true given the polling data. A Bayesian analysis allows

TABLE 1
Forecasts for the 2004 Australian Election

	Forecast horizon		
	Election eve	3 months prior	1 year prior
Panel A: Predicting coalition two-party preferred vote share (actual: 52.74 per cent)			
Economic Models ^a			
<i>Cameron and Crosby</i>	51.6	51.3	51.1
<i>Jackman and Marks</i>	51.7	51.7	51.5
<i>Jackman</i>	51.3	51.2	51.1
Polls ^b			
<i>ACNielsen</i>	54.0	48.0	52.0
<i>Galaxy</i>	52.0	51.0	— ^d
<i>Morgan</i>	49.0	46.0	50.0
<i>Newspoll</i>	50.0	49.0	49.0
Panel B: Predicted probability of Howard victory, per cent (correct prediction)			
Economic Models ^a			
<i>Cameron and Crosby</i>	70.8	66.7	64.0
<i>Jackman and Marks</i>	70.4	70.3	68.5
<i>Jackman</i>	65.7	65.5	64.0
Polls ^b			
<i>ACNielsen</i>	99.9	6.6	93.4
<i>Galaxy</i>	89.7	73.6	— ^d
<i>Morgan</i>	17.1	0.6	50.0
<i>Newspoll</i>	50.0	24.4	24.4
Prediction Markets ^c			
<i>BetFair</i>	78.7	60.6	66.7
<i>Centrebet</i>	76.9	57.9	73.7

Notes:

^a Economic models: Election eve estimate uses data from the quarter prior to the election; estimate 3 months prior uses data two quarters prior to the election; estimate 12 months prior uses data five quarters prior to the election.

^b Poll dates are ACNielsen 5–7 October 2004, 16–18 July 2004, 19–21 September 2003; Galaxy 1–3 October 2004, 24–25 July 2004 (this was the first poll produced by Galaxy); Morgan 2–3 October 2004, 10–11 July 2004, 11–12 and 18–19 October 2003; Newspoll 5–6 October 2004, 2–4 July 2004; 17–19 October 2003. Poll sample sizes used in Panel B were: ACNielsen 1400 (2000 for election-eve poll), Galaxy 1000, Morgan 1000, Newspoll 1200 (1700 for the election-eve poll).

^c Prediction markets: *Sources:* www.betfair.com; www.centrebet.com.au. Betfair data for 12 months prior is from 17 November 2003, the date on which the Betfair market opened.

^d Galaxy did not conduct a poll a year prior to the federal election.

assume only classical sampling error, we can take any two-party preferred poll result and sample size, and use these two figures to calculate the probability that the Coalition's two-party preferred vote exceeds 50 per cent.

We should also note several assumptions inherent in this conversion. Importantly, our estimates account only for sampling error. Other possible

the latter to be computed from the former. Assuming an uninformative prior yields a posterior for the probability of a coalition victory that in large samples is well approximated by a normal distribution with mean p and variance $p(1-p)/n$, and this provides a more coherent foundation for our probabilistic interpretations of the polls. We thank an anonymous referee for suggesting this point to us.

forms of error may exist. For example, voters may respond to polls strategically rather than truthfully; polling samples may be skewed towards a particular kind of voter; and it may also be incorrect to assume that respondents who say that they 'don't know' which way they plan to vote can be folded into all other respondents. In addition, we do not account for the fact that the parliamentary election system in Australia (like the electoral college system in the USA) means that it is possible to win a majority of the two-party vote but lose the election. This has occurred in five of the 41 Australian federal elections (1954, 1961, 1969, 1990 and 1998). As such, the probabilities implied by the polls should be interpreted as the probability of a party winning the two-party preferred vote, rather than the probability

of that party forming the government. It is possible that polls may represent voting intentions accurately, but that these polls may nevertheless yield poor forecasts of actual voting behaviour in the event that voters change their minds between the poll and the election.

A related approach can be applied to convert the vote-share estimates from the economic models into a probability of winning more than 50 per cent of the two-party preferred vote. In this case the standard error of the forecast is calculated from past forecast errors, using the formula $se(\text{Forecast}) = s\sqrt{1 + x_j(X'X)^{-1}x_j'}$, where s^2 is the mean square error of the prediction, column vector x_j represents the n observations on independent variable x_j , and X is an $n \times J$ data matrix.² These estimates – derived by applying the standard error of the forecast to a normal distribution to generate a probability – should again be interpreted as the probability of the Coalition winning a majority of the popular vote.

Finally, betting markets effectively involve trading a binary option on the re-election of the Howard government. Wolfers and Zitzewitz (2006b) argue that the price of such a security can be directly interpreted as the probability of the event occurring. Unlike the polls and economic models, this price can be interpreted directly as the probability of a Coalition prime ministership.

We believe that presenting election forecasts as probabilities is useful because it focuses attention not only on the point estimate of the likely winner, but also on the uncertainty surrounding that forecast. As such, we believe it important that accounts of polls and other forecasts focus on their probabilistic interpretations. Equally, the validity of this exercise depends on the validity of estimates of the errors associated with these forecasts, an issue we return to below.

We now turn to analysing the new data in somewhat more detail. We begin with economic models, then turn to prediction markets and finally opinion polls.

III Economic Models

The logic of the economic models is simple: voters are more likely to re-elect incumbents who deliver a robust economy. This pattern can be motivated either as voters providing an incentive

² The standard error of the forecast is calculated based on a regression assuming homoscedasticity (in contrast, note that the estimates in Table 2 use robust standard errors).

for politicians to deliver good outcomes, or as voters using available information to discern high-ability incumbents. Fair (1978) assembles evidence showing that the state of the economy was an important factor in US presidential elections. The subsequent Australian literature includes papers by Jackman and Marks (1994), Jackman (1995), and Cameron and Crosby (2000).

The first column in Table 2 shows the economic models, using updated economic data, but with the sample restricted to those elections in the models as originally published.

We updated these economic models to include all data available prior to the 2004 election so as to allow a true 'real-time' forecast; these estimates are reported in the second column. Adding the 1996, 1998 and 2001 elections to the sample substantially undermines the explanatory power of the Jackman and Marks (1994) and Jackman (1995) specifications. Updating Cameron and Crosby yields smaller changes in their estimates, reflecting their longer sample.

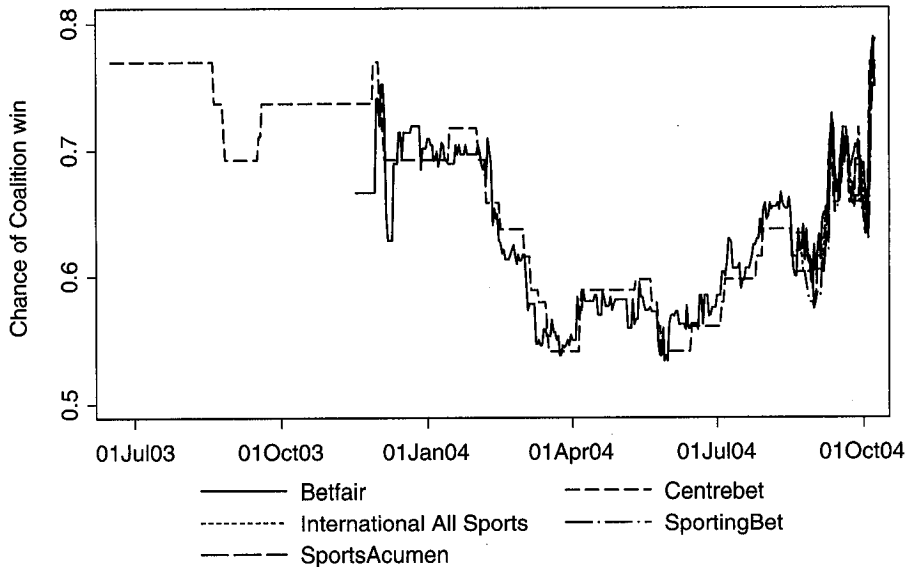
In general, Cameron and Crosby's model fits the Australian data well, but at the cost of parsimony: they include four economic measures, a honeymoon dummy, and five dummy variables for particular elections or wars. Of the four economic coefficients, unemployment and inflation are significant and in the expected direction, while real wage growth has a marginally significant but negative effect. Taken together, these estimates suggest that incumbents are more likely to be re-elected if inflation is low, and the real economy (measured in terms of unemployment) is near capacity.

We use these updated models to generate the predictions shown in Table 1. These different models yield substantially similar predictions, and on election-eve suggest a Coalition vote share of 51.3–51.7 per cent. These strong predictions reflect the robust state of the economy during the 2004 election cycle. Even so, they all underpredicted the performance of the Coalition, suggesting that perhaps these models still understate the importance of robust economic conditions. As such, we further update these models in light of the 2004 election result (column 3), and the economic variables become marginally more important.³

If anything, the real puzzle in these data is how unresponsive Australian voters appear to be to economic conditions, relative to their American

³ This is consistent with the findings of Davidson *et al.* (2004), who conclude that economic conditions are the main factor explaining the outcome of the 2004 election.

FIGURE 1
Comparing Betting Markets Over the 16 Months Before the Election



counterparts. To make a more precise comparison, we estimate each of these econometric models on US presidential elections since 1892. Comparing these results, shown in column four, with earlier columns suggests that US voters are more responsive to economic conditions than Australian voters. This is a particularly surprising result, considering that these specific models were (presumably) estimated to maximise their fit to the Australian data.

IV Prediction Markets

The 2004 election cycle saw robust interest in election betting, and we have obtained data from four large Australian bookmakers: Centrebet, International All Sports, SportingBet and SportsAcumen. We also have data from a British sports betting exchange, Betfair, which operates in a manner akin to a futures exchange, in that participants buy or sell a contract paying \$100 if a candidate wins.

Compared with other betting and prediction markets, total volume was substantial: Betfair saw turnover of \$A705 544 (£307 601),⁴ while Centrebet held \$A1.9 million in bets (up from \$A1.5 million in 2001). The other three Australian bookmakers

⁴ Based on the currency exchange rate on the day of the Australian election.

would not disclose their totals. By way of comparison, the Iowa Electronic Market (which caps bets at \$US500) saw total turnover of \$A438 039 (\$US327 384)⁵ in its winner-take-all market for the 2004 US presidential election. Incidentally, increased competition in the election betting market in 2004 did not appear to have an effect on profit margins. Centrebet's overround (or vigorish) was 8 per cent in both the 2001 and 2004 elections.

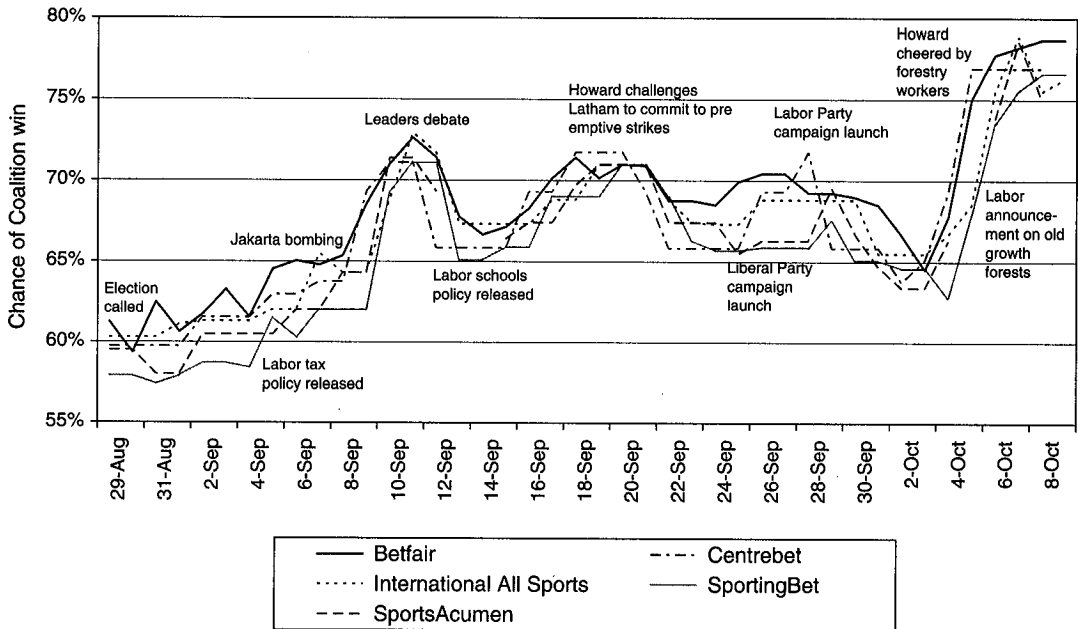
We convert the betting odds to implied winning probabilities, and show them in Figure 1.⁶ Clearly, the Coalition was the favourite from July 2003 (when the first election betting market opened) until polling day.

Did the betting markets predict that Howard would increase his majority? One way of answering this question is to compare the election-eve odds in this election to those just prior to the last poll. On election eve in 2001, the markets assessed the Coalition as having a 60 per cent chance of winning, while on election eve in 2004, the market believed the Coalition had a 77 per

⁵ Based on the currency exchange rate on the day of the US election.

⁶ Where the return from a winning \$1 bet is denoted as the payout: $p(\text{Coalition}) = \text{Coalition Payout}^{-1} / (\text{Coalition Payout}^{-1} + \text{Labor Payout}^{-1})$.

FIGURE 2
Comparing Betting Markets Over the Election Campaign



cent chance of winning. Clearly, the markets thought it more likely that the Coalition would be re-elected in 2004 than it did in 2001. While we believe that this probably reflects a view that the Coalition would increase their majority, alternative interpretations are that markets were more certain about their ability to assess voting trends, or that the market underpredicted the 2001 result.

We can also make similar comparisons on a seat-by-seat basis. Twenty-two electorates were the focus of Centrebet's attention in both the 2001 and 2004 elections, and as at election eve on both elections, the Coalition was the betting favourite in 15 of these seats.⁷ However, the average probability of a Coalition win in these seats was 55 per cent in 2001 and 60 per cent in 2004, which suggests that the market expected Howard to be returned with an increased majority.

Figure 2 shows just the election campaign (the last 6 weeks of the 15-month period depicted in

⁷ These seats were Adelaide, Ballarat, Banks, Bass, Canning, Chisholm, Deakin, Dunkley, Eden-Monaro, Herbert, Hindmarsh, Hinkler, Kingston, La Trobe, McEwen, McMillan, Moreton, Page, Parramatta, Paterson, Richmond and Stirling.

Figure 1). This figure provides an intriguing narrative of the key events during this election campaign.

Four facts are immediately obvious from Figure 2. First, once one takes account of the overround (the betting market equivalent of a bid-ask spread), there were very few arbitrage opportunities during this election cycle.⁸ Second, these markets respond very quickly to campaign news, and different bookmakers respond in similar ways. Third, the betting markets yield predictions that exhibit a plausible degree of volatility; the probability of a Coalition victory fluctuated from around 60 per cent to around 80 per cent. And fourth, they appear to respond to identifiably important political news. Consistent with campaign commentary, the betting markets suggest the importance of arguably exogenous factors to re-election of the Coalition government, with the Jakarta bombing giving the Coalition a lift in the 2004 election cycle just as September 11 did in the 2001 election cycle. These data also suggest that Labor's

⁸ The only arbitrage opportunity we observed was between Centrebet and SportingBet on 5 October 2004, yielding an expected profit margin of 2 per cent.

TABLE 3
Statistical Tests of Pricing Efficiency

	Betfair	Centrebet
Panel A: Dickey-Fuller Test: $\Delta p(\text{Coalition})_t = \beta * p(\text{Coalition}_{t-1})$		
$p(\text{Coalition}_{t-1})$	-0.028	-0.0073
Test statistic (<i>t</i> -stat)	(-1.482)	(-1.328)
(10 per cent critical value)	(-2.571)	(-2.570)
Reject random walk?	No	No
<i>P</i> -value	0.140	0.185
Panel B: KPSS Tests		
Number of daily lags	14	9
Test statistic	0.369***	0.0988
(1 per cent critical value)	(0.216)	(0.216)
Reject trend stationarity?	Yes	No
Panel C: Weak-Form Efficiency		
$\Delta p(\text{Coalition})_t = \beta_1 * \Delta p(\text{Coalition}_{t-1}) + \beta_2 * \Delta p(\text{Coalition}_{t-2}) + \beta_3 * \Delta p(\text{Coalition}_{t-3})$		
$\Delta p(\text{Coalition}_{t-1})$	0.096 (0.096)	0.13 (0.091)
$\Delta p(\text{Coalition}_{t-2})$	0.060 (0.096)	0.086 (0.058)
$\Delta p(\text{Coalition}_{t-3})$	-0.041 (0.080)	-0.14 (0.10)
<i>F</i> -test of joint significance	0.49 (<i>P</i> = 0.68)	1.66 (<i>P</i> = 0.17)
Panel D: Is Polling Information Efficiently Priced?		
$\Delta p(\text{Coalition})_t = \beta_1 * \Delta \text{Polls}_{t-1} + \beta_2 * \Delta \text{Polls}_{t-2} + \beta_3 * \Delta \text{Polls}_{t-3}$		
$\Delta \text{Polls}_{t-1}$	0.088 (0.074)	0.033 (0.039)
$\Delta \text{Polls}_{t-2}$	0.056 (0.058)	0.065 (0.030)**
$\Delta \text{Polls}_{t-3}$	0.035 (0.075)	0.039 (0.026)
<i>F</i> -test of joint significance	0.97 (<i>P</i> = 0.40)	3.14 (<i>P</i> = 0.02)

Notes: ***, ** and * denote statistical significance at the 1 per cent, 5 per cent and 10 per cent levels, respectively (robust standard errors in parentheses).

In KPSS tests, number of daily lags is chosen by Schwert criterion.

Tasmanian forestry announcement provided the Coalition with a substantial boost.


These four facts provide suggestive evidence in favour of the view that these prediction market prices are efficient. Table 3 provides further evidence on this score, presenting several formal statistical tests of market efficiency. We focus on the two prediction markets for which we have the most data: Betfair and Centrebet. Our Betfair data is trade-by-trade, and we convert this into daily prices, taking the last trade prior to 16.00 hours Australian Eastern Standard Time on a given day (on days where no trade occurred, we treat the market as closed). Centrebet prices are based on the bookmaker's price at the end of that day.

In Panels A and B, we test whether the price of a contract paying \$100 if the Coalition wins the election evolves in a manner consistent with the famous random walk hypothesis. Panel A shows that we cannot reject the null hypothesis of a unit root (even at a 10 per cent level of significance), while Panel B tests the alternative null that the price series is stationary. For Betfair data we can

reject this null (at the 1 per cent level); however, for Centrebet data we are unable to reject the null even at the 10 per cent level. In Panel C we test weak-form market efficiency: can one predict today's price change based on the price history? We find little evidence of predictable price changes. Panel D presents a specific test of semistrong form efficiency: can we predict today's price changes on the basis of publicly available polling data? We implement this test by taking the weighted average of the last 7 days of published polls. In the case of Betfair data, the small and statistically insignificant coefficients suggest that the market prices this publicly available information efficiently. For Centrebet, the coefficients are jointly significant – indicating some lag in the incorporation of polling data – but sufficiently small that they are unlikely to yield profit-making opportunities from normal movements in opinion polls.

What about the betting odds in the marginal seats? In the 2001 election favourites won in 43 of the 47 seats where Centrebet had allowed punters to place a bet (Wolfers & Leigh, 2002). This

FIGURE 3
Seat-by-seat Betting



time, Centrebet offered betting in 33 marginal seats.⁹ Based on the odds reported on the morning of polling day, the favourite won in 24 seats, and lost in eight seats, while in one seat (Hindmarsh) there was no favourite, with Centrebet offering the same odds on the two major parties when the polling booths opened. Figure 3 shows the explanatory power of these prices when the betting market opened, and when it closed.

These charts show clearly the forecasting power of the betting markets: on election eve, most seats were clustered in the bottom left and top right quadrants implying that they correctly forecast the winning and losing candidates. The market's prediction of the likely winner also predicts actual vote totals reasonably well. Of those cases where the market-favoured candidate lost, the margin was generally less than a couple of percentage points. Russell Broadbent's victory in the seat of McMillan stands alone as the only case where the underdog won by a substantial margin. However, the predictive performance of the markets 3 weeks prior to the election was substantially poorer. The improvement in fit over the ensuing 3 weeks reflects the markets aggregating possibly new information about voter preferences and the quality of campaigns. We now turn to evaluating the polling data in greater detail.

⁹ Centrebet's total marginal seat betting was approximately \$600 000 in 2001 and \$350 000 in 2004.

V Polls

The most commonly used means of forecasting election outcomes remains opinion polling. In Australia, three major pollsters have provided regular polls through at least the last seven elections: ACNielsen (formerly known as AGB McNair), Roy Morgan and Newspoll. The 2004 election saw the emergence of a fourth major pollster, Galaxy. In addition, the Australian National University, in conjunction with the Bulletin and the Nine Network, conducted an experimental internet election poll. We do not deal with the results of that poll here; a fuller discussion is found in Gibson (2004) and Jackman (2005).

Of the four major Australian pollsters, all except Morgan conduct their polling by telephone.¹⁰ Although the major pollsters do not publish their response rates, we can glean some evidence from the results of market research surveys. Bednall and Shaw (2003) found that the response rate to a short telephone survey (of the type most likely to be conducted by opinion pollsters) averaged 23 per cent. Another factor is that pollsters do not typically attempt to contact those with mobile telephones.

Overall, the predictive power of election-eve polls remains low. Goot (2005) notes that over the

¹⁰ Morgan typically employs face-to-face polling, but conducted a telephone poll 7–8 October 2004. On a two-party preferred basis, the results of this poll were only 0.5 per cent different from the Morgan face-to-face poll conducted on 2–3 October 2004. Both predicted a Labor victory.

seven elections held between 1987 and 2004, the mean absolute error in predicting primary vote shares of the major parties was 1.8 per cent for ACNielsen, 1.3 per cent for Morgan and 1.6 per cent for Newspoll. (The error margin of polls taken months or years before the election is larger still; Wolfers & Leigh, 2002). Since 1993, the major pollsters have produced two-party preferred estimates, and over the five elections from 1993 to 2004, Goot calculates that the mean absolute error was 1.6 per cent for ACNielsen, 2.9 per cent for Morgan and 1.7 per cent for Newspoll. Averaging these figures, the mean absolute error for the major Australian pollsters is 2.1 per cent – the same as the mean absolute prediction error of Gallup, the largest US pollster (Wolfers & Zitzewitz, 2004). Given that voting is not compulsory in the USA, this comparison does not favour the Australian pollsters.

Figure 4 shows the major opinion pollsters over the full election cycle, with a solid line depicting the average result from all polls taken over the previous 7 days, weighted by their sample size (e.g. a poll of 2000 people would receive twice the weight of a poll of 1000 people). The solid line suggests that the Coalition's fortunes peaked in early 2003, and again in late 2004, while Labor performed well in 2002, and during the first half of 2004. The polls concur with the betting markets that the Coalition outperformed Labor during the election campaign. This remains true even if we ignore the Coalition's poll-bounce that followed the bombing of the Australian embassy in Jakarta.

It is also interesting to compare the performance of the pollsters to the betting markets. We begin by comparing the time series properties of each. In order to put polls and markets on the same metric, we convert the polls into an implied probability of a Coalition victory, using the method outlined above (treating the poll average as a cumulative sample). The implied probability of a Coalition victory is shown in Figure 5. To smooth out volatility in individual polls, we analyse the probability implied by the 7-day, sample-weighted poll average.

The most striking aspect of this figure is the extreme volatility of the implied probability of a Coalition victory suggested by the polls. Specifically, if we believe that polls are afflicted only with sampling error, then for 4 months of the last election cycle the chance of the Coalition being returned was less than 1 per cent, while for 9 months of the election cycle, the chance of the Coalition being returned was less than 3 per cent.

It strains credulity to believe that the Coalition's probability of winning could have been this low over a substantial period. It is also difficult to believe what a weighted aggregation of the polls suggests: that over the course of the 6-week election campaign, the Coalition's probability of winning rose from 0.7 per cent to 98.3 per cent. By contrast the betting markets suggest a more reasonable degree of variation in election probabilities (54–77 per cent).

We interpret the excess volatility of the polls as suggesting that published error margins are a substantial underestimate of the true forecast errors.¹¹ As such, we also experiment with our own guess as to the relevant standard error. The dashed line in Figure 5 uses the polling results, but we inflate the standard error substantially, simply assuming that it is 10 per cent. In other words, we assume that the true standard error of the polls is equivalent to a poll of 25 voters that suffered only from sampling error. While this seems a rather extreme assumption, it generates a far more plausible series than when we take the pollsters' estimates of their standard error at face value. The resulting estimate ranges from 26 per cent to 66 per cent, still demonstrating more movement than the betting markets, but now moving within a more credible range. Moreover, this estimate appears to move in lockstep with the betting markets.

Even though this assessment of the polls strikes us as more credible, Figure 5 still suggests that the polls were systematically more pro-Labor than were the betting markets. Jackman (2005) contains a thorough assessment of the bias inherent in polls from each of the major polling organisations.

What does this suggest for those producing and publishing polls in future elections? Given their low predictive power, we propose that pollsters provide more guidance to their clients as to their (in)ability to forecast election outcomes. Our results indicate that – for forecasting purposes – the pollsters' published margins of error should at least be doubled. This is true for each of the four major pollsters covered here. At some point in the 2001–2004 election cycle, ACNielsen, Galaxy and Newspoll each published figures which, if interpreted as a forecast, suggested that the Labor Party had less than a 0.5 per cent chance of winning, while ACNielsen, Morgan and Newspoll each published figures suggesting that the Coalition had less than a 0.5 per cent chance of being

¹¹ For a related discussion of this issue in the US context, see Martin *et al.* (2003).

TABLE 4
Comparing Bookies with Marginal Seat Polling

Seat	Poll prediction (Coalition 2PP) (per cent)	Bookmaker (Coalition probability) (per cent)	Result (Coalition 2PP) (per cent)	How did the pollsters and bookies perform?
Adelaide	51 (Adelaide Advertiser, 14 September)	65	48.67	Both wrong
Bass	54 (Launceston Examiner, 14 September)	30	52.63	Polls right, bookies wrong
Eden-Monaro	54 (Canberra Times, 2 October)	60	52.14	Both right
Herbert	56 (Courier Mail, 8 October)	71	56.20	Both right
Hinkler	61 (Courier Mail, 8 October)	76	54.81	Both right
La Trobe	51 (ACNielsen 24–27 September)	65	55.83	Both right
McMillan	48 (ACNielsen 24–27 September)	30	54.99	Both wrong
Moreton	53 (Courier Mail, 8 October)	68	54.17	Both right
Parramatta	53 (ACNielsen 24–28 September)	43	49.23	Bookies right, polls wrong
Richmond	48 (Gold Coast Weekend Bulletin 14 September)	76	49.81	Polls right, bookies wrong
Wentworth	Coalition 50, Labor 50 (ACNielsen 17–20 September)	Coalition 40, Labor 20 (plus Independents 40)	55.48 Picked coalition equal favourite	Both half-right

Notes:

1. Betting odds for Richmond and Wentworth are for 21 September 2004. In all other cases, betting odds are from the first day on which the poll was conducted.

2. Poll results from 14 September, 2 October and 8 October are taken from the archives of www.pollbludger.com.

returned to power. Consequently the media needs to display substantially greater caution in interpreting changes from one poll to the next. Indeed, even with the published margins of error, a 1 per cent movement from one poll to the next is unlikely to be anything more than noise. But with the margins of error implied by our results, even very large movements are likely to be mere measurement error. Journalists who write about changes in poll movements without discussing the margin of error may well be guilty of misleading their readers.¹²

Comparing the performance of marginal-seat polls and marginal-seat betting also yields some interesting comparisons. Table 4 provides the full set of comparisons that we were able to track down. Excluding the seat of Wentworth (where a

prominent independent candidate makes it difficult to compare pollsters and bookmakers), the pollsters correctly predicted 7 of the 10 seats (mean absolute forecast error = 3.0 per cent; root mean square error = 3.9 per cent), while early odds from Centrebet correctly predicted 6 of the 10 seats. In a regression of the final two-party result on the two sets of predictions, neither pollsters nor bookmakers were (individually or jointly) a statistically significant predictor of the size of the Coalition's eventual majority, although it is difficult to draw strong conclusions from such a small sample.

VI Conclusion

Writing after the 2001 election, we concluded by stating our belief that betting markets and economic models both merited greater prominence in the media and in public discourse. In the case of betting markets, we are pleased to see some evidence that this has occurred. We carried out a citation search across the major Australian newspapers, looking for the name of the most prominent election betting firm, Centrebet, plus the name of one of the major political parties: the number of such stories in the 3 months prior to

¹² For a discussion of this issue during the election campaign, see Leigh (2004b). One could even regard this issue as a one of ethics. The Australian Journalists' Association Code of Ethics begins 'Report and interpret honestly, striving for accuracy, fairness and disclosure of all essential facts. Do not suppress relevant available facts, or give distorting emphasis.'

polling day more than doubled, from 66 in 2001 to 136 in 2004.¹³

By contrast, economic models of voting behaviour still languish in relative obscurity. While a good deal of commentary during each election season is devoted to how the economy will affect the outcome, such discussion is generally conducted outside the more rigorous framework offered by the economic voting models. This appears to be equally true of election commentary proffered by most economists, suggesting that perhaps the complexity of economic voting models is not the only factor impeding their growth.

Overall, most Australian election commentary remains dominated by opinion polls, or even more informal ('finger in the breeze') methods. A survey of 10 experts published on the Sunday before polling day found that three thought Latham would win, while seven thought Howard would win, but with a smaller majority than in 2001. None forecast the true result – a Howard victory with an increased majority. Less poll- and gossip-dominated journalism would be a boon to Australian election commentary, and might even free up space for more substantive discussion of policies.

REFERENCES

- Bednall, D.H.B. and Shaw, M. (2003), 'Changing Response Rates in Australian Market Research', *Australian Journal of Market Research*, **11**, 31–41.
- Berg, J., Forsythe, R., Nelson, F. and Rietz, T. (2008), 'Results from a Dozen Years of Election Futures Markets Research', in Plott, C. and Smith, V. (eds), *Handbook of Experimental Economic Results*, Elsevier, Amsterdam, forthcoming.
- Bureau of the Census (1975), *Historical Statistics of the United States: Colonial Times to 1970*, Bureau of the Census, Washington, DC.
- Bush, G.W. (2005), *Economic Report of the President 2005*, United States Government Printing Office, Washington, DC.
- Cameron, L. and Crosby, M. (2000), 'It's the Economy Stupid: Macroeconomics and Federal Elections in Australia', *The Economic Record*, **76**, 354–64.
- Cuzán, A., Armstrong, S. and Jones, R. (2005), 'Combining Methods to Forecast the 2004 Presidential Election: The Pollyvote', Mimeo, University of Pennsylvania, Philadelphia, PA.
- Davidson, S., Farrell, L., Felvus, C. and Fry, T. (2004), 'Trust Me: It's Still the Economy', Mimeo, Royal Melbourne Institute of Technology, Melbourne, Vic.
- Fair, R. (1978), 'The Effect of Economic Events on Votes for President', *The Review of Economics and Statistics*, **60**, 159–73.
- Fair, R. (2002a), 'The Effect of Economic Events on Votes for President: 2000 Update', Yale University, New Haven, CT. Available at <http://fairmodel.econ.yale.edu/RAYFAIR/PDF/2002DHTM.HTM>.
- Fair, R. (2002b), *Predicting Presidential Elections and Other Things*, Stanford University Press, Stanford, CA.
- Gibson, R. (2004), 'Online Polling in Election Studies: Lessons to be learned – Publicise Early and Publicise Often', Political Science Workshop on the 2004 Election, 19 November, Australian National University, Canberra, ACT.
- Goot, M. (2005), 'The Polls: Liberal, Labor or Too Close to Call?' in Simms, M. and Warhurst, J. (eds), *Mortgage Nation: The 2004 Australian Elections*. API. Network/Edith Cowan University Press, Perth, WA; 55–70.
- Hipel, K. and McLeod, A. (1994), *Time Series Modelling of Water Resources and Environmental Systems*, Elsevier, Amsterdam.
- Jackman, S. (1995), 'Some More of All That: a Reply to Charnock', *Australian Journal of Political Science*, **30**, 347–55.
- Jackman, S. (2005), 'Pooling the Polls over an Election Campaign', *Australian Journal of Political Science*, **40**, 499–517.
- Jackman, S. and Marks, G. (1994), 'Forecasting Australia Elections: 1993, and All That', *Australian Journal of Political Science*, **29**, 277–91.
- Leigh, A. (2004a), 'Does the World Economy Swing National Elections?' Centre for Economic Policy Research Discussion Paper 485, Australian National University, Canberra, ACT.
- Leigh, A. (2004b), 'Bookies Are a Better Bet Than Pollsters', *Sydney Morning Herald*, 1 September (p. 17).
- Martin, E., Traugott, M. and Kennedy, C. (2003), 'A Review and Proposal for a New Measure of Poll Accuracy', Mimeo, University of Michigan, Ann Arbor, MI.
- Reserve Bank of Australia. (1997), *Australian Economic Statistics 1949–50 to 1996–97*, Occasional Paper No. 8 (Tables revised to 1997). RBA, Sydney, NSW.
- Wolfers, J. and Leigh, A. (2002), 'Three Tools for Forecasting Federal Elections: Lessons from 2001', *Australian Journal of Political Science*, **37**, 223–40.
- Wolfers, J. and Zitzewitz, E. (2004), 'Prediction Markets', *Journal of Economic Perspectives*, **18**, 107–26.
- Wolfers, J. and Zitzewitz, E. (2006a), 'Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections', National Bureau of Economic Research, Working Paper No. 12073, Cambridge, MA.
- Wolfers, J. and Zitzewitz, E. (2006b), 'Interpreting Prediction Market Prices as Probabilities', National Bureau of Economic Research, Working Paper No. 12200, Cambridge, MA.