



The Australian National University
Centre for Economic Policy Research
DISCUSSION PAPER

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

Andrew Leigh* and Justin Wolfers**

DISCUSSION PAPER NO. 502
November 2005

ISSN: 1442-8636
ISBN: 0 7315 3572 3

*Andrew Leigh, Social Policy Evaluation Analysis Research Centre, Research School of Social Sciences, Australian National University, Canberra, ACT 0200 Australia
Email: Andrew.Leigh@anu.edu.au. <http://andrewleigh.com>

** Justin Wolfers, Wharton, University of Pennsylvania, CEPR, IZA and NBER
Email: jwolfers@wharton.upenn.edu www.nber.org/~jwolfers

Acknowledgements

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, Brandon Verblow and Eric Zitzewitz for useful insights and valuable feedback on an earlier draft. All remaining errors and omissions are our own.

ABSTRACT

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 suggests a more reasonable degree of volatility, and useful forecasting performance both throughout the election cycle and across individual electorates.

JEL Codes: D72, D84

Keywords: Voting, elections, prediction markets, opinion polling, macroeconomic voting

I. Introduction

There has recently been a surge of interest in forecasting election outcomes (Cameron and Crosby 2000; Fair 2002; Wolfers and Leigh 2002; Leigh 2004a; Jackman 2005; Wolfers and Zitzewitz 2005b; Cuzán, Armstrong and Jones 2005). Following the 2004 Australian federal election we now have even more valuable data with which to evaluate competing theories. In light of this new evidence, we assess the relative efficacy of three approaches: econometric models based on recent economic data; betting markets (also called “prediction markets”; Wolfers and 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. And we saw important innovations in the application of each of the forecasting technologies we examine. The 2004 election occurred during arguably the best macroeconomic times in a generation, and yet even so, Howard’s vote exceeded the expectations of economic forecasting models. This election cycle also saw the introduction of several new polling firms and polling technologies, but rising challenges for pollsters too, 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 popular, and we now have data from five betting markets run in Australia and the UK.

To preview our results, we find that economic models provide useful forecasts, particularly when making medium-term predictions. Consistent with international evidence (Berg, Forsythe, Rietz and Nelson 2001), 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 you rely on; Morgan’s recent 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. Further, 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 presents key forecasts from the 2004 election cycle. We begin by presenting forecasts for the share of the two-party preferred vote won by the Coalition, providing election-eve forecasts, those taken three months prior to the election, and those taken a year prior. By construction, forecasts of the vote share of the Australian Labor Party are equal to 100 minus the Coalition prediction.

Table 1: Forecasts for the 2004 Australian Election

	Forecast Horizon		
	Election Eve	3 months prior	1 year prior
Panel A: Predicting Coalition 2-Party Preferred Vote Share (Actual: 52.74%)			
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	-
<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, % (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	-
<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 2 quarters prior to the election; estimate 12 months prior uses data 5 quarters prior to the election.
- (b) Poll dates are ACNielsen 5-7 Oct 2004, 16-18 July 2004, 19-21 Sept 2003; Galaxy 1-3 Oct 2004, 24-25 July 2004 (this was the first poll produced by Galaxy); Morgan 2-3 Oct 2004, 10-11 July 2004, 11-12 & 18-19 Oct 2003; Newspoll 5-6 Oct 2004, 2-4 July 2004; 17-19 Oct 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 Nov 2003, the date on which the Betfair market opened.

Panel A shows several forecasts of the two-party preferred vote share. The Coalition government was ultimately re-elected with 52.74% of the vote, an increase of 1.79% on their performance in 2001. The economic models correctly picked the election winner, but surprisingly (given the strength of the economy) still under-predicted the Coalition's performance. That said, these models did predict a swing toward 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% toward 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 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 $se = \sqrt{\frac{q(1-q)}{n}}$, 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 1 000, in which the parties were evenly matched, would have a standard error of 1.6% (and a corresponding 95% confidence interval of $\pm 3.2\%$), while a poll with a sample size of 2 000 would have a standard error of 1.1% (a 95% confidence interval of $\pm 2.2\%$).

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

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 the latter to be computed from the former. Assuming an uninformative prior yields a posterior for the probability of a coalition victory that in

simple terms, if we 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%.

We should also note several assumptions inherent in this conversion. Importantly, our estimates account only for sampling error. Other possible 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 US), means that it is possible to win a majority of the two-party vote but lose the election. This has occurred in 5 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 winning the two-party preferred vote, rather than the probability of picking the correct Prime Minister. Finally, it is possible that polls may represent voting intentions accurately, but these may 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 voteshare estimates from the economic models into a probability of winning more than 50% of the two-party preferred vote. In this case the standard error of the forecast is calculated from past forecast errors, using the formula $se(Forecast) = s\sqrt{1 + x_j(X'X)^{-1}x_j'}$, where s^2 is the mean square error of the prediction.² Again, these estimates should be interpreted as the probability of the Coalition winning a majority of the popular vote.

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.

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

Finally, betting markets effectively involve trading a binary option on the re-election of the Howard government. Wolfers and Zitzewitz (2005a) 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 for politicians to deliver good outcomes, or as voters using available information to discern high-ability incumbents. Fair (1978) assembled 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.

Table 2: Economic Models

Dependent Variable: Incumbent Party's Vote Share				
	(1) Replication	(2) Extend Sample to before 2004 election	(3) Extend Sample to include 2004 election	(4) Estimate on US elections 1892-2004
Jackman and Marks (1994)				
Unemployment (<i>Δ</i> over election cycle)	-0.726** [0.298]	-0.359 [0.364]	-0.383 [0.343]	-0.589** [0.248]
Inflation	-0.199* [0.105]	-0.147 [0.112]	-0.151 [0.108]	-0.634** [0.285]
Constant	52.523*** [0.982]	51.540*** [0.954]	51.612*** [0.883]	54.494*** [1.326]
N	18	21	22	28
R²	0.32	0.12	0.14	0.33
Jackman (1995)				
Unemployment (<i>Δ</i> over election cycle)	-0.536** [0.237]	-0.25 [0.346]	-0.292 [0.328]	-0.579** [0.228]
Inflation	-0.304** [0.118]	-0.208 [0.122]	-0.211* [0.120]	-0.670** [0.255]
Honeymoon	2.804 [1.764]	1.905 [1.779]	1.806 [1.759]	4.930** [1.854]
Constant	52.468*** [1.067]	51.416*** [1.026]	51.529*** [0.946]	52.483*** [1.396]
N	18	21	22	28
R²	0.43	0.18	0.19	0.48
Cameron and Crosby (2000)				
Unemployment (<i>Level</i>)	-0.289* [0.165]	-0.306* [0.160]	-0.307* [0.159]	-0.124 [0.280]
Inflation	-0.428*** [0.102]	-0.374*** [0.115]	-0.376*** [0.114]	-0.513* [0.286]
Real GDP growth	-0.176 [0.164]	-0.207 [0.161]	-0.209 [0.161]	0.441* [0.254]
Real wage growth	-0.277 [0.164]	-0.275* [0.158]	-0.276* [0.156]	0.714* [0.410]
Honeymoon	5.227*** [1.642]	4.313** [1.749]	4.281** [1.737]	3.093 [1.885]
Constant	54.962*** [1.240]	54.980*** [1.237]	55.041*** [1.215]	51.394*** [3.428]
N	37	39	40	29
R²	0.8	0.77	0.77	0.49

Notes:

- (a) ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively [robust standard errors in brackets].
- (b) Models used are the authors' preferred specifications: Jackman and Marks (1994) model 4; Jackman (1995) model 4 (both of which use elections from 1951 onwards); and Cameron and Crosby (2000) model 1.2 (using elections from 1903 onwards).
- (c) Cameron and Crosby's specification also includes indicator variables for 1906, 1931, 1975, plus separate dummies for the two world wars. Replicated specification does not precisely match that presented in Wolfers and Leigh (2002), due to the release by the Australian Bureau of Statistics of revised GDP figures.

- (d) When estimating the Cameron and Crosby specification on US presidential elections, all variables are based on annual data in the election year, or the change from the year before the election to the election year. As Cameron and Crosby do for Australia, we include separate dummies for World War I (1916 election) and World War II (1940 and 1944 elections). Since unemployment figures are first available in 1890, the 1892 election is included in the Cameron and Crosby model only.

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 undermined the explanatory power of the Jackman and Marks (1994) and Jackman (1995) specifications. Updating Cameron and Crosby yielded 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 used these updated models to generate the predictions shown in Table 1. These different models yielded substantially similar predictions, suggesting on election-eve a Coalition vote share of 51.3%–51.7%. These strong predictions reflected the robust state of the economy during the 2004 election cycle. Even so, they all *under-predicted* the performance of the Coalition, suggesting that perhaps these models still under-state the importance of robust economic conditions. As such, we further updated these models in light of the 2004 election result (column three), 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 counterparts. To make a more

³ This is consistent with the findings of Davidson, Farrell, Felvus and Fry (2004), who conclude that economic conditions are the main factor explaining the outcome of the 2004 election.

precise comparison, we estimated each of these econometric models on US Presidential elections since 1892. These results are shown in column four. Comparing these results 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 maximize 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 \$705 544 (£307 601),⁴ while Centrebet held \$1.9 million in bets (up from \$1.5 million in 2001). The other three Australian bookmakers would not disclose their totals. By way of comparison, the Iowa Electronic Market (which caps bets at \$500) saw total turnover of \$438 039 (\$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% 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.

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

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:
 $p(\text{Coalition}) = \text{CoalitionPayout}^1 / (\text{CoalitionPayout}^1 + \text{LaborPayout}^1)$.

Figure 1: Comparing Betting Markets Over the 16 Months Before the Election



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% chance of winning, while on election-eve in 2004, the market believed the Coalition had a 77% 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 under-predicted 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 favourite in 15 of these seats.⁷ However, the average probability of a Coalition win in these seats was 55% in 2001 and

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.

60% in 2004. This suggests that the market expected Howard to be returned with an increased majority.

Figure 2 shows just the election campaign (the last six weeks of the 15-month period depicted in Figure 1). This figure provides an intriguing narrative of the key events during this election campaign.

Figure 2: Comparing Betting Markets Over the 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% to around 80%. 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 in the 2004 election cycle providing a boost the Coalition just as September 11 did in the

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

2001 election cycle. These data also suggest that the Labor's 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 4pm 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.

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 (t-stat)	(-1.482)	(-1.328)
[10% critical value]	[-2.571]	[-2.570]
Reject Random Walk?	No	No
[p-value]	0.140	0.185
Panel B: KPSS Tests		
Number of Daily lags	14	9
Test statistic	0.369***	0.0988
(1% 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)
F-test of joint significance	0.49 (p=0.68)	1.66 (p=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)
F-test of joint significance	0.97 (p=0.40)	3.14 (p=0.02)


Note: ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively. (robust standard errors in brackets). In KPSS tests, number of daily lags is chosen by Schwert criterion.

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% 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% level); however for Centrebet data we are unable to reject it even at the 10% 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 semi-strong 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 seven 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 and Leigh, 2002). This 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.

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

Figure 3: Seat-by-Seat Betting



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. 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. The predictive performance of the markets three weeks prior to the election was substantially poorer. The improvement in fit over the ensuing three weeks reflects the markets aggregating possibly new information about voter preferences and the quality of campaigns.

These cross-sectional results provide us with an interesting point of comparison with the polls. Specifically, we ask: what sorts of sample sizes would we need to obtain to yield a set of predictions this accurate?¹⁰ We begin by converting the election-eve betting odds back into implied two-party vote shares, essentially reversing the process used to generate Panel B of Table 1. That is:

¹⁰ We thank Simon Jackman for suggesting this analysis to us. For a related exercise using data from the Iowa Electronic Markets, see Berg, Nelson and Rietz (2003).

$$p(\text{Coalition}_i) = \Phi \left(\frac{q_i - 0.5}{\sqrt{\frac{q_i(1 - q_i)}{n}}} \right), \text{ where } p(\text{Coalition}_i) \text{ represents the probability of a}$$

Coalition win in electorate i (as inferred from the Centrebet odds on election-eve), q_i represents the proportion who voted for the Coalition in that electorate, Φ represents the cumulative normal distribution, and n is the implied sample size. Thus for any given sample size, n , we can convert betting odds back into predictions of the two-party vote share. Our estimates simply seek the value of n that maximizes the fit of this implied vote share forecast with the actual vote share. That is, we re-arrange the above equation so that we can use OLS to estimate: $q_i - 0.5 = \beta \Phi^{-1}(p(\text{Coalition}_i))$, where β is the implied standard error of the betting market in forecasting the two-party vote share. Our estimates yielded $\beta=0.062$, with a 95% confidence interval between 0.045 and 0.079. This estimate of the standard error allows us to infer the “equivalent sample size” of the betting market, by the transformation $n = \frac{q_i(1 - q_i)}{\beta^2}$. If one could conduct a poll that yielded *only sampling error*, one could obtain a similar degree of accuracy to the Centrebet odds by polling only 60-100 voters.¹¹ At first blush, this suggests that prediction markets in marginal seats provide very little useful information aggregation. However as we will argue below, even polls with much larger samples rarely yield forecasts even this precise, as other sources of uncertainty increase the ratio of noise-to-signal in poll-based forecasts. We now turn to evaluating the polling data in greater detail.

¹¹ Without a constant, the implied sample size is around 60-70. Including a constant in the regression (to correct for either bias in the market measure, or a correlated shock across electorates) yields a smaller implied standard error, and hence an implied sample size around 90-100.

V. Polls

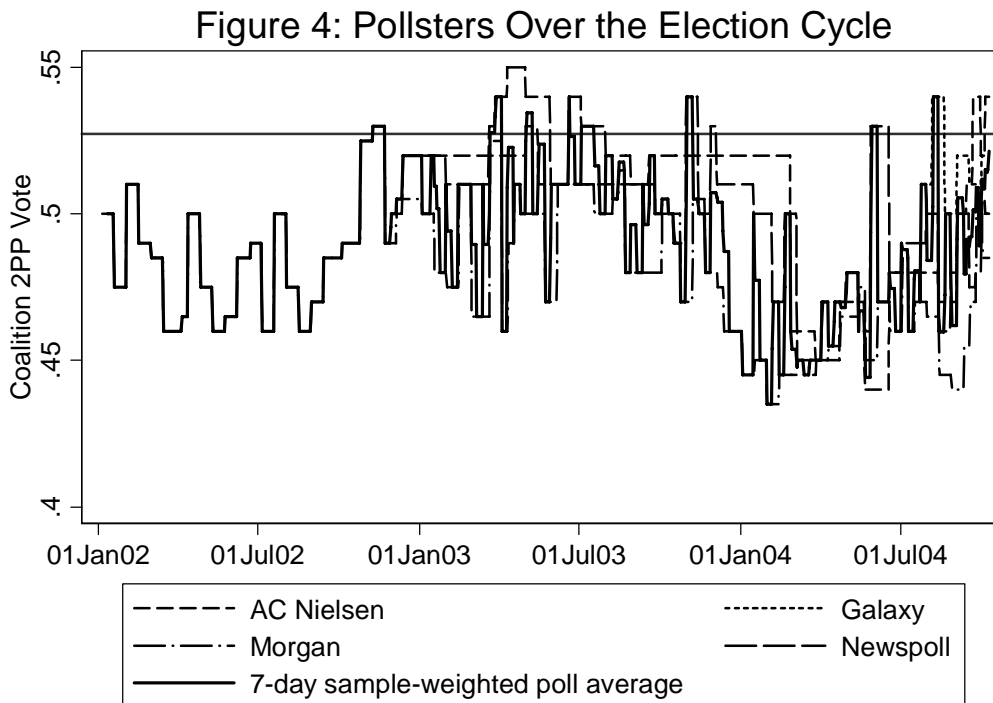
The most commonly used means of forecasting election outcomes remains opinion polling. In Australia, three major pollsters have provided regular polls through at least the last seven elections: ACNielsen (formerly known as AGB McNair), Roy Morgan and Newspoll. The 2004 election saw the emergence of a 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.¹² 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%. 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 seven elections held between 1987 and 2004, the mean absolute error in predicting primary vote shares of the major parties was 1.8% for ACNielsen, 1.3% for Morgan and 1.6% for Newspoll. (The error margin of polls taken months or years before the election is larger still: Wolfers and Leigh 2002). Since 1993, the major pollsters have produced two-party preferred estimates, and over the five elections from 1993-2004, Goot calculates that the mean absolute error was 1.6% for ACNielsen, 2.9% for Morgan and 1.7% for Newspoll. Averaging these figures, the mean absolute error for the major Australian pollsters is 2.1% – the same as the mean absolute prediction error of Gallup, the largest US pollster (Wolfers and Zitzewitz 2004). Given that voting is not compulsory in the US, this comparison does not favour the Australian pollsters.


¹² 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% different from the Morgan face-to-face poll conducted on 2-3 October 2004. Both predicted a Labor victory.

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 seven-days, weighted by their sample size (eg. a poll of 2 000 people would receive twice the weight of a poll of 1 000 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 seven-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 four months of the last election cycle the chance of the Coalition being returned was less than 1%, while for nine months of the election cycle, the chance of the Coalition being returned was less than 3%. 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 six-week election campaign, the Coalition’s probability of winning rose from 0.7% to 98.3%. By contrast the betting markets suggest a more reasonable degree of variation in election probabilities (54%–77%).

We interpret the excess volatility of the polls as suggesting that published error margins are a substantial under-estimate of the true forecast errors.¹³ As such, we also experiment

¹³ For a related discussion of this issue in the US context, see Martin, Traugott and Kennedy (2003).

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%. 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% to 66%, 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 organizations.

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 2001-04 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% chance of winning, while ACNielsen, Morgan and Newspoll each published figures suggesting that the Coalition had less than a 0.5% chance of being 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% 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.¹⁴

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 7/10 seats (mean absolute forecast error=3.0%; root mean square error=3.9%), while early odds from Centrebet correctly predicted 6/10 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.

¹⁴ 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."

Table 4: Comparing Bookies with Marginal Seat Polling

Seat	Poll Prediction (Coalition 2PP)	Bookmaker (Coalition Probability)	Result (Coalition 2PP)	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 Indep 40%)	55.48%	Both half-right: picked coalition equal favourite.

Note:

- (a) 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.
- (b) Poll results from 14 September, 2 October and 8 October are taken from the archives of www.pollbludger.com.

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 three months prior to polling day more than doubled, from 66 in 2001 to 136 in 2004.¹⁵

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

¹⁵ These figures are the combined results of two searches. First, we used Lexis-Nexis to search five newspapers: the *Advertiser*, *Australian*, *Canberra Times*, *Courier Mail* and *Mercury*. Next, we used the Fairfax online search engine to search all Fairfax publications. In each case, we searched for stories containing the word “Centrebet” plus “Labor or ALP or Coalition”. The two date ranges were 10 August 2001-10 November 2001, and 9 July 2004-9 October 2004.

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 ten 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.

**Appendix Table 1: Economic Model Data
Jackman & Marks (1994) and Jackman (1995)**

Year	Incumbent vote share	Δ Unemployment (over the cycle)	Inflation	Honeymoon election?
1951	50.7	-0.6	16.7	1
1954	49.3	0.9	0.9	0
1955	54.2	-0.5	3.5	0
1958	54.1	0.5	1.6	0
1961	49.5	1.6	0.7	0
1963	52.6	-1.9	0.7	0
1966	56.9	0.1	2.6	0
1969	49.8	0.2	3.0	0
1972	47.3	0.8	4.7	0
1974	51.7	-0.7	14.6	1
1975	44.3	3.2	14.4	0
1977	54.6	0.7	9.3	1
1980	50.4	-0.2	9.2	0
1983	46.8	3.8	11.4	0
1984	51.8	-1.0	2.6	1
1987	50.8	-0.7	8.2	0
1990	49.9	-1.7	8.6	0
1993	51.4	4.6	1.2	0
1996	46.4	-2.7	3.8	0
1998	49.0	-0.4	1.6	1
2001	51.03	-0.8	3.1	0
2004	52.74	-1.8	2.6	0

Note:

- (a) Data in the above table are constructed by the authors based on the most recently available economic releases, and following as closely as possible the variable definitions in Jackman and Marks (1994) and Jackman (1995). Due to data revisions, the economic variables do not correspond precisely with those shown in Jackman and Marks (1994, Table 1).
- (b) Incumbent vote share is based on two-party preferred vote, with Labor defined as the incumbent in 1975. Source: Australian Electoral Commission website (www.aec.gov.au)
- (c) Change in unemployment is based on the unemployment rate in the election quarter. Sources: 1951-61 from Jackman and Marks; 1963-77 from Reserve Bank (1997); 1978-2004 from Reserve Bank of Australia economic statistics, Table G07, Labour Force, available at www.rba.gov.au.
- (d) Inflation is the change in the CPI over four quarters ending in the election quarter. Sources: 1951-61 from Jackman and Marks; 1963-77 from Reserve Bank (1997); 1978-2004 from Reserve Bank of Australia economic statistics, Table G01, Measures of Consumer Price Inflation.
- (e) "Honeymoon" is a dummy variable denoting the first election faced by an incumbent government.

**Appendix Table 2: Economic Model Data
Cameron and Crosby (2000)**

Year	Incumbent vote share	Unemployment	Inflation	Δ Real GDP	Δ Real wages	Honeymoon election?
1903	60.76	9.2	0	7.6	-0.1	1
1906	71.44	5.5	-0.5	12.2	1.4	0
1910	52.37	3.1	1.8	7.1	1.5	0
1913	49.89	4.3	2.8	6.1	0.7	0
1914	48.13	5.9	4.3	-6.9	-3.1	1
1917	45.05	3.2	-1.8	-2.5	9.8	1
1919	56.45	3	14.4	-4	-1.4	1
1922	54.64	5.3	1.7	3.7	-4.8	0
1925	54.59	5.2	-0.8	-1.2	-1.4	0
1928	53.71	6.4	-0.5	-1.7	2.5	0
1929	47.97	8.8	2.5	0.8	-3.4	0
1931	41.01	19.3	-8.8	-0.7	1	1
1934	54.11	15.1	1.3	2.8	-0.7	1
1937	53.53	7.2	2.4	5.8	3.7	0
1940	51.17	6.4	4.6	6.7	-1.4	0
1943	61.42	0.9	2.5	2.6	4.3	1
1946	54.73	3	1.8	-3.5	-0.2	0
1949	49	1.8	8.8	7.5	-0.8	1
1951	50.7	1.2	16.7	4.6	6.2	1
1954	49.3	1.8	0.9	6.2	0.7	0
1955	54.2	1.5	3.4	5.2	-1	0
1958	54.1	2.1	1.6	6.4	-1.4	0
1961	49.5	3.6	0.7	-0.31	1.8	0
1963	52.6	1.7	0.7	7.41	5.5	0
1966	56.9	1.8	2.6	4.82	2.5	0
1969	49.8	2	3	5.06	4.2	0
1972	47.3	2.8	4.7	2.32	4.84	0
1974	51.7	2.1	14.6	0.74	4.22	1
1975	44.3	5.3	14.4	1.5	-0.3	0
1977	54.6	5.9	9.3	0.41	0.85	1
1980	50.4	6	9.2	2.65	3.38	0
1983	46.8	9.6	11.4	-2.68	0.45	0
1984	51.8	8.6	2.6	4.91	5.73	1
1987	50.8	7.8	8.2	5.26	-4.53	0
1990	49.9	6	8.6	3.86	-1.34	0
1993	51.4	10.7	1.2	4.23	3.65	0
1996	46.4	8.2	3.7	5.07	1.49	0
1998	49.0	7.5	1.6	5.92	2.4	1
2001	51.03	6.96	3.12	4.32	0.99	0
2004	52.74	5.19	2.59	3.05	0.4	0

Notes:

- (a) Data in the above table are constructed by the authors based on the most recently available economic releases, and following as closely as possible the variable definitions in Cameron and Crosby (2000).

- (b) Incumbent vote share from 1949-2004 is based on two-party preferred vote, with Labor defined as the incumbent in 1975. Source: Australian Electoral Commission. From 1903-46, Cameron and Crosby calculate two-party voteshare by classifying parties into left-wing and right-wing.
- (c) Unemployment is the rate in the election quarter. Sources: 1903-58 from Cameron and Crosby (2000); 1963-77 from Reserve Bank (1997); 1978-2004 from Reserve Bank of Australia economic statistics, Table G07, Labour Force.
- (d) Inflation is the year-ended percentage change in the election quarter. Sources: 1903-48 from Cameron and Crosby (2000); 1949-77 from Reserve Bank (1997); 1978-2004 from Reserve Bank of Australia economic statistics, Table G01, Measures of Consumer Price Inflation.
- (e) Real GDP growth is the year-ended percentage change to the election quarter. Sources: 1903-58 from Cameron and Crosby (2000); 1959-2004 from Reserve Bank of Australia economic statistics, Table G10, Seasonally Adjusted GDP (chain volume measure).
- (f) Real wage growth is the year-ended percentage change in average wages to the election quarter. Sources: 1903-70 from Cameron and Crosby (2000); 1971-2004 from Reserve Bank of Australia economic statistics, Table G06, Labour Costs.
- (g) “Honeymoon” is a dummy variable denoting the first election faced by an incumbent government.
- (h) The Cameron and Crosby dataset also includes three dummy variables for the 1906, 1931 and 1975 elections, plus separate dummies for World War I (1914 and 1917 elections) and World War II (1940 and 1943 elections).

**Appendix Table 3: Economic Model Data
US Presidential Elections 1892-2004**

Year	Incumbent vote share	Unemploy- ment	Δ Unemploy- ment (over the cycle)	Inflation	Δ Real GDP	Δ Real wages	Honeymoon election?
1892	48.3	3	-	0	9.6	0.4	1
1896	47.8	14.4	11.4	0	-2.1	0.2	1
1900	53.2	5	-9.4	1	2.6	1.8	1
1904	60.0	5.4	0.4	0.9	-1.3	-0.2	0
1908	54.5	8	2.6	-1.8	-8.2	0.5	0
1912	54.7	4.6	-3.4	2.6	5.0	0.3	0
1916	51.7	5.1	0.5	7.9	7.5	3.2	1
1920	36.1	5.2	0.1	15.6	-6.5	1.5	0
1924	58.2	5	-0.2	0	0.9	0.8	1
1928	58.8	4.2	-0.8	-1.7	1.0	1.4	0
1932	40.8	23.6	19.4	-9.9	-13.0	-2.2	0
1936	62.5	16.9	-6.7	1.5	12.9	1.7	1
1940	55.0	14.6	-2.3	0.7	8.5	3.1	0
1944	53.8	1.2	-13.4	1.7	8.2	7.0	0
1948	52.4	3.8	2.6	8.1	4.3	4.9	0
1952	44.6	3	-0.8	1.9	4.0	1.0	0
1956	57.8	4.1	1.1	1.5	2.0	3.6	1
1960	49.9	5.5	1.4	1.7	2.5	2.4	0
1964	61.3	5.2	-0.3	1.3	5.8	6.6	1
1968	49.6	3.6	-1.6	4.2	4.8	1.3	0
1972	61.8	5.6	2	3.2	5.4	4.3	1
1976	48.9	7.7	2.1	5.8	5.6	1.5	0
1980	44.7	7.1	-0.6	13.5	-0.2	-5.9	1
1984	59.2	7.5	0.4	4.3	7.3	0.6	1
1988	53.9	5.5	-2	4.1	4.2	-0.9	0
1992	46.5	7.5	2	3	3.1	-0.2	0
1996	54.7	5.4	-2.1	3	3.6	0.4	1
2000	50.3	4	-1.4	3.4	3.8	0.4	0
2004	51.6	5.5	1.5	1.5	5.1	-0.3	1

Notes:

- (a) Incumbent vote share from Fair (2002).
- (b) Unemployment rates are annual averages. From 1890-1970, figures are from Hipel and Mcleod (1994); and 1948-2004 from Federal Reserve Economic Data (<http://research.stlouisfed.org/fred2/>).
- (c) Inflation is the annual inflation rate, from Robert Sahr, "Inflation Conversion Factors for Dollars 1665 to Estimated 2015", available at http://oregonstate.edu/dept/pol_sci/fac/sahr/sahr.htm. Sahr uses data from John J. McCusker and colleagues for the period 1665 to 1912, and CPI-U data from the US Bureau of Labor Statistics for 1913-2004.
- (d) Real GDP growth is the annual change, constructed for 1892-2000 by averaging quarterly GDP from Fair, Appendix Table 2; 2004 from Federal Reserve Economic Data.
- (e) Real wage growth is the annual change in real weekly earnings. 1892-1900 from Bureau of the Census (1975, Series D735-738, nonfarm employees only), 1904-60 from Bureau of the Census (1975, Series D722-727, real earnings when employed), 1964 from Bureau of the Census (1975, Series D722-727, full-time employees only), 1968-2004 from Bush (2005, Appendix Table B-47, private sector only).
- (f) "Honeymoon" is a dummy variable denoting the first election faced by the party of an incumbent president.

- (g) The Cameron and Crosby model also includes separate dummies for World War I (1916 election) and World War II (1940 and 1944 elections).

References

- Bednall, D.H.B. & Shaw, M. (2003), 'Changing response rates in Australian market research', Australian Journal of Market Research, 11(1): 31-41
- Berg, J. Forsythe, R., Nelson, F. and Rietz, T. (2001), 'Results from a Dozen Years of election Futures Markets Research', Handbook of Experimental Economic Results. Plott, C. and Smith, V. (eds). Amsterdam: Elsevier.
- Bureau of the Census (1975), Historical Statistics of the United States: Colonial Times to 1970. Washington, DC: Bureau of the Census
- Bush, G.W. (2005) Economic Report of the President 2005. Washington, DC: United States Government Printing Office
- Cameron, L. and Crosby, M. (2000). 'It's the Economy Stupid: Macroeconomics and Federal Elections in Australia', The Economic Record. Vol 76, No 235: 354-364
- Cuzán, A., Armstrong, S., and Jones, R. (2005). 'Combining Methods to Forecast the 2004 Presidential Election: The Pollyvote', *mimeo*, University of Pennsylvania.
- Davidson, S., Farrell, L., Felvus, C. and Fry, T. (2004). 'Trust me: It's still the economy', *mimeo*, Royal Melbourne Institute of Technology
- Fair, R. (1978) 'The Effect of Economic Events on Votes for President', The Review of Economics and Statistics 60(2): 159-173
- Fair, R. (2002). 'The Effect of Economic Events on Votes for President: 2000 Update', Available at <http://fairmodel.econ.yale.edu/RAYFAIR/PDF/2002DHTM.HTM>
- Fair, R. (2002). Predicting Presidential Elections and Other Things, Stanford University Press: Stanford, CA.
- Fair, R. (2004). 'A Vote Equation and the 2004 Election', November 22. Available at <http://fairmodel.econ.yale.edu/vote2004/vot1104a.pdf>

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

Goot, M. (2005). 'The Polls: Liberal, Labor or Too Close to Call?' in M Simms and J Warhurst (eds) Mortgage Nation: the 2004 Australian Elections. API. Network/Edith Cowan University Press: Perth

Hipel, K. and McLeod, A. (1994). Time Series Modelling of Water Resources and Environmental Systems. Elsevier: Amsterdam

Jackman, S and Marks, G. (1994). 'Forecasting Australia Elections: 1993, and All That', Australian Journal of Political Science 29(2): 277-291

Jackman, S. (1995). 'Some More of All That: A Reply to Charnock', Australian Journal of Political Science 30: 347-355

Jackman, S. (2005). 'Pooling the Polls over an Election Campaign', Australian Journal of Political Science 40(4): 499-517

Leigh, A. (2004a). 'Does the World Economy Swing National Elections?', Centre for Economic Policy Research Discussion Paper 485, Australian National University, Canberra

Leigh, A. (2004b). 'Bookies are a better bet than pollsters', Sydney Morning Herald, 1 September 2004

Leigh, A. (2005). 'Economic Voting and Electoral Behavior: How do Individual, Local and National Factors Affect the Partisan Choice?', Economics and Politics, 17(2): 265-296

Martin, E., Traugott, M., Kennedy, C. (2003). 'A Review and Proposal for a New Measure of Poll Accuracy', *mimeo*, University of Michigan

Reserve Bank of Australia. (1997). Australian Economic Statistics 1949–50 to 1996–97, Occasional Paper No. 8 (Tables revised to 1997). RBA: Sydney

Wolfers, J and Leigh, A. (2002). ‘Three Tools for Forecasting Federal Elections: Lessons from 2001’, Australian Journal of Political Science 37(2): 223-40

Wolfers, J. and Zitzewitz, E. (2004) ‘Prediction Markets’, Journal of Economic Perspectives 18(2): 107-126.

Wolfers, J and Zitzewitz, E. (2005a). ‘Interpreting Prediction Market Prices as Probabilities’, *mimeo*, University of Pennsylvania

Wolfers, J. and Zitzewitz, E. (2005b) ‘Correlation versus Causation in Prediction Markets and the 2004 Elections’, *mimeo*, University of Pennsylvania