

Electoral Forecasting and Public Opinion Tracking in Latin America: An Application to Chile

Kenneth Bunker ^{*}
The London School of Economics

Stefan Bauchowitz [†]
The London School of Economics

August 10, 2015

Abstract

The purpose of this article is to explore electoral forecasting and public opinion tracking in Latin America. *Approach.* We review different approaches used to estimate the true value of public opinion, and assess the range of their application. *Application.* We focus on election night forecasting and campaign variation tracking in Latin America. *Model.* We propose a two-stage model based on poll aggregation and Bayesian inference. *Data.* We present data from two presidential elections in Chile. *Results.* We test the model and show that it provides the most accurate election night forecasting point estimate and the most comprehensive campaign variation tracking method. *Discussion.* We discuss the advantages and limitations of our model, and suggest a route for future research.

KEYWORDS. Bayesian Inference; Chile; Electoral Forecasting; Latin America; Public Opinion Tracking.

^{*}Corresponding author. E-mail: k.a.bunker@lse.ac.uk. Tel: +44 (0) 7979327708. Address: Room CON.H503, Government Department, The London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

[†]E-mail: s.bauchowitz@lse.ac.uk. Address: The London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

1 Introduction

Electoral forecasting dates back to at least 1948, when Lois Bean argued that data from individual states could be used to predict the outcome of national elections in the United States (1948). This approach was complemented in the mid-1970s, when Ray Fair used aggregate economic measures along with political variables to explain electoral returns (1978). Michael Lewis-Beck and Tom Rice produced the first book-length overview in the 1990s, suggesting that forecasting should be treated as a predictive instrument and be evaluated accordingly (1992). Randall Jones extended the scope to trial heats (polls), exit polls, expert judgments, electoral cycles, and the nomination process (2002).

Campbell and Garand suggest that electoral forecasting has three different objectives, related to academia, politicians and the media (2000). They argued that the first objective is to provide academics with evidence, which helps advance scientific knowledge. They can use the data to accept or reject electoral theories. The second objective is to allow politicians to evaluate the effect of campaign events. They can use the input to change or uphold their strategies. The third objective is to provide the media with reference points, which can help them interpret the electoral process and provide voters with information they can use to evaluate the candidates in the race.

Electoral forecasting is synonymous to election night forecasting. It is concerned with the final result of the election. Its objective is to estimate the final state of voting intentions. It attempts to provide an answer to the hypothetical question: “who will win the election?” Traditionally, this has been related to academia. Major symposiums organized by leading political science journals (*Electoral Studies* in 2011, by *The British Journal of Politics and International Relations* in 2005, and by *PS: Political Science & Politics* in 2012) reflect the interest of scholars in predicting electoral results. Forecasting elections provides academics with valuable data and evidence, which may ultimately help them develop theories of electoral behavior.

Electoral forecasting is often – but not always – accompanied by public opinion tracking. In contrast to electoral forecasting, public opinion tracking is concerned with the partial result of the election. Its objective is to estimate the current state of voting intentions. It attempts to provide an answer to the hypothetical question: “who would win the election if it were held today?” Public opinion tracking is usually related to politicians and the media. The recent success of *FiveThirtyEight* or the *Princeton Election Consortium* in the United States reflects the interest of both candidates and news outlets in interpreting campaign variations. Tracking public opinion can provide both politicians and the media with information proxies and short-cuts.

Electoral forecasting and public opinion tracking are intrinsically related to each other. Their ultimate objective is to estimate the true value of

public opinion. They differ in that they each focus on the value of the measure at different points in time. On the one hand, forecasters attempt to estimate vote intention for each candidate on election night, by comparing the forecast to the election result; the veracity of their estimates can only be observed on election night. On the other hand, trackers aim to estimate vote intention for each candidate throughout the campaign; the veracity of their estimates cannot be observed during the campaign, although estimates of public opinion closer to the election can arguably be considered a forecast in their own right.

Electoral forecasting and public opinion tracking are at an incipient state in Latin America. While, scholars, politicians and voters alike turn to polls, expert opinions, quantitative models – and in some cases electronic markets – they rarely use their full potential to track the final and partial state of public opinion during electoral cycles. The contribution of this article is to develop and test a model for electoral forecasting and public opinion tracking sensitive to the particular characteristics of Latin America. As most countries approach three decades of uninterrupted democracy, we believe that there is finally leeway to develop a model that is accurate on both counts. We propose a model that increases the accuracy of election night forecasts and provides a more comprehensive account of campaign variations.

The remainder of this article is organized as follows. In Section 2 we review different approaches used for electoral forecasting and public opinion tracking, and assess the range of their application. In Section 3 we argue that given the structural and circumstantial state of electoral systems and party systems in Latin America, an approach based on public opinion polls is the most appropriate. In Section 4 we propose a two-stage model based on poll aggregation and Bayesian inference. In Section 5 we present data from two presidential elections in Chile. In Section 6 we test the model. In Section 7 we discuss advantages and limitations, and suggest a route for future research.

2 Approach and Geographical Scope

2.1 Approach

Andreas Graefe et. al. identify four approaches that aim to estimate the true value of public opinion during electoral cycles (2014). We follow their insightful summary. The first approach is based on public opinion; it usually relies on information stemming from trial heat polls to make inferences on voting behavior. The second approach is based on expert opinion; it generally uses the judgment from political insiders and experienced observers as reference points to predict the current state of an electoral campaign and the final outcome of an election. The third approach is based on quantitative models; it commonly combines economic indicators and political variables

to explain electoral results. The fourth approach is based on electronic markets; it normally incorporates information from electoral markets to make projections.

The approach based on polls is perhaps the most common (Graefe et al., 2014). It normally relies on data from vote intention polls, where individuals are asked who they would vote for in the upcoming election. Questions are usually open-ended (unlimited number of options) or close-ended (limited number of options). Individual polls can be used on their own or combined into aggregate data. James Campbell and Ken Wink analyzed the accuracy of polls for 11 presidential elections in the US from 1948 to 1988 (1990). They find that errors are often large, increasing as the distance from the election grows. Christopher Wlezien shows that errors in polls are often caused by sampling problems, non-responses, inaccurate measurements and faulty processing (2003).

One method used to increase the accuracy of estimations has been to average polls conducted by different organizations near the same time. Richard Gott and Wesley Colley used this method to correctly predict the outcome of the 2004 and 2008 US elections (2008). This method is known as aggregating. A more advanced form of aggregating is called projecting, in which the historical record of the polls used are taken into account when forecasting electoral results or tracking variations throughout the campaign. Indeed, by using the error of each poll in the previous electoral cycle, a formula can be derived to estimate the expected error the poll will have in the current electoral cycle. Robert Erickson and Christopher Wlezien find that poll projections are much more accurate than individual polls for election night forecasting (2008).

The approach based on experts is the traditional manner of forecasting elections and tracking public opinion (Graefe et al., 2014). Samuel Kernell shows that before the emergence of polls in the 1930s, the opinion of experts was the predominant way of assessing electoral cycles (2000). Experts' influence extends to the present day. They normally know how to read and interpret polls, and sometimes have insider information that enhances their intuition. Some experts take public polls at face value and simply propagate their forecasts through the media. Other experts may completely disregard polls and use insider information to state their predictions. Either way, they do not have to meet the same standard as other approaches with the same objective and disclose their particular methods.

Expert predictions are often a product of different sources. They may rely on polls or on insider information. Experts may rely on both or on none. They may consider only some of the polls or some of the insider information to make their election night forecast and interpret campaign variations. Experts are often considered reliable proxies since they have experience in reading elections. But also because their careers often depend on their accuracy. Experts who get elections right stand out, while experts

who get elections wrong fade away. Their forecasts therefore tend to be more conservative – they do not normally incur in high risks. Since they can change their opinion at their discretion, their readings generally converge towards the average prediction.

The approach based on quantitative models is considered the standard manner approaching elections scientifically (Graefe et al., 2014). For more than three decades scholars have been developing and testing these types of models. Alan Abramowitz (2012), James Campbell (2012), Robert Erickson and Christopher Wlezien (2012) and Michael Lewis-Beck and Charles Tien (2012) all use quantitative models that combine variables that represent economic indicators and political variables. Most of these models include between two and five determinants, and use indicators such as GDP, inflation and unemployment along with variables such as presidential approval and incumbency status.

One way that has been singled out as a manner of increasing accuracy further is to combine individual models. Larry Bartels and John Zaller used various combinations of structural variables representative of the economy to develop and test 48 models (2001). They showed that by combining models together they were able to reduce error of election night forecasts made by individual polls by around 15%. Robert Erickson, Joseph Bafumi and Bret Wilson built on this, and showed that accuracy could be further increased when public opinion variables were added into the combined models (2001). In their study they added presidential approval ratings to each of the 48 models, and showed that accuracy of election night forecasts increased by around 30%.

A fourth approach to forecasting elections is based on betting markets (Graefe et al., 2014). Almost as old as expert opinions, traditional betting markets out-date both public opinion polls and quantitative models. In the time before scientific models, experts opinions were more valuable when they were backed by monetary endorsement. Paul Rhode and Koleman Strumpf look at historical betting markets in US presidential elections between 1884 and 1940 and conclude that they did a “remarkable job” in predicting electoral outcomes (2004, 138). Indeed, by including a monetary penalty betting markets were able to consistently exclude betters with larger margins of error associated to their predictions. On average, only confident betters are expected to participate, increasing the accuracy of the instrument.

A modern version of traditional betting markets are electronic betting markets. Electronic betting markets – such as the now defunct *Intrade* – constitute a reliable proxy to estimate the result of an election, though Robert Erickson and Christopher Wlezien show that they do not outperform polls (Erickson and Wlezien, 2008). Electronic betting markets are, however, appropriate instruments for tracking electoral campaign variations. Indeed, they produce daily estimates. While individual polls only provide snapshots of the moment, markets provide a continuous real-time account of electoral

preferences. At any rate, electronic betting markets are an important source of information for people who do not trust polls, do not believe experts and do not understand quantitative models.

2.2 Geographical Scope

The four approaches outlined above are often used to predict and track elections in parliamentary regimes. They have been used in the United Kingdom, since at least 1987 (Mughan, 1987). David Sanders was one of the first to attempt to predict electoral outcomes in the UK, when he looked at prospects for the hitherto forthcoming general election of 1995 (Sanders, 1995). Most – but not all – electoral forecasting in the UK is based on economic models (Lewis-Beck, Nadeau and Bélanger, 2004). Matthew Lebo and Helmut Norpoth presented an account based on a pendulum model to capture intra party variations (Lebo and Norpoth, 2007). Richard Nadeau, Michael Lewis-Beck and Eric Bélanger presented a two-step model to increase parsimony (Nadeau, Lewis-Beck and Bélanger, 2009). The extent of data and methods has recently widened (Fisher et al., 2011; Gibson and Lewis-Beck, 2011). Most notably, *Election Forecast*, which took a scientific approach to pooling polls to estimate election night results in the United Kingdom general election of 2015.

Electoral forecasting has recently extended to other European parliamentary regimes. One case is Spain, in which Juan José Dolado explored the long-range dependence of polls (Dolado, Gonzalo and Mayoral, 2003). More recently, Pedro Magalhães, Sandra Aguiar-Conraria and Michael Lewis-Beck developed a model from core political economy theory (Magalhães, Aguiar-Conraria and Lewis-Beck, 2012). In Italy, Paolo Bellucci explored and described election cycles and electoral forecasting in 1994–2008 (Bellucci, 2010). In Austria, Julian Aichholzer correctly forecast elections using what he described as the grand coalition model (Aichholzer and Willmann, 2014). In Germany, Helmut Norpoth and Thomas Gschwend picked the incumbent vote share to the decimal in the 2002 election and were just three-tenths of a percent off in the 2005 election using the chancellor model (Norpoth and Gschwend, 2010).

Electoral forecasting expands beyond European borders. Simon Jackman developed a model based on averaging polls for the 2004 election cycle in Australia (Jackman, 2004). Richard Nadeau shows that electoral forecasting dates back to at least 1993 in Canada (1993). Mark Pickup used campaign trial heats to produce electoral forecasts for the 2004 and 2006 Canadian elections (2007). More recently, Eric Bélanger and Jean François Godbout tested a vote function model, using variables related to the economy and presidential approval, on all federal elections between 1953 and 2008 (Bélanger and Godbout, 2010). The website *ThreeHundredEight* has tracked campaign variations and produced accurate estimates for most Canadian

elections since 2008.

Electoral forecasting has also frequently been used in presidential regimes, with most evidence stemming from US elections. Some of the major contributions date back to the early 1980s (Lewis-Beck and Rice, 1984; Rosenstone, 1983). Since 1996, the most stable approaches have correctly forecast the election winner with an average error of less than 3% (Jones and Cuzán, 2008). These forecasts have often been quantitative models based on economic measures (Alesina, Londregan and Rosenthal, 1993; Nadeau and Lewis-Beck, 2001; Lewis-Beck and Stegmaier, 2000). Since the 2004, electoral forecasting has also relied on public opinion polls (Campbell, 2008). In 2012, Drew Linzer combined polls to produce an election night forecast and an account of campaign variations (2013).

France is the leading case in European presidential regimes, with electoral forecasts dating back to 1991 (Lewis-Beck, 1991). While most elections since then have had some sort of forecasting, in 1997 (Jérôme, Jérôme and Lewis-Beck, 1999) and 2002 (Durand, Blais and Larochelle, 2004) the polls failed dramatically. While some academics questioned the quality of the data, others turned to explanations based on political instability (Jerôme and Jerôme, 2000). This turn in events led to the massification of quantitative models. In 2002, vote functions were explored (Dubois and Fauvelle-Aymar, 2004). In 2005, literature surrounding models grew further, when economic conditions were found to be robust predictors (Auberger and Dubois, 2005). Richard Nadeau, Michael Lewis-Beck and Eric Bélanger provide a multi-equation solution (2010).

3 Application

Latin American countries have received significantly less attention from electoral forecasters and campaign trackers. One reason may be because of their young democratic record (Mainwaring, 1998). A second reason may be because of the quality of available data (Lewis-Beck and Rice, 1992). Electoral forecasting cannot take place in systems without elections or in environments with incomplete information. Accordingly, most approaches to electoral forecasting and public opinion tracking in the region have used rudimentary resources. Inaccurate assumptions about the data have often been made in order to generate pseudo-scientific inferences. For example, some experts interpret individual polls as indicators of the current state of public opinion many days after they have been published, or as a forecast for election night many days before it takes place.

The shortcut mentioned in the example above represents a common methodological shortcoming of the public opinion approach. First, because a poll is technically a “snapshot” of the moment in which it is fielded. As such, it only represents the moment in which it is taken. It can only illustrate

the distribution of preferences for each of the candidates for the days the interviews were conducted. The rule of thumb is that as a poll ages, it loses its predictive power. In the example, experts also have fault. If they read polls irresponsibly, they can inadvertently sway the result of an election. An influential expert, in a country sensitive to the bandwagon or underdog effects, could unintentionally bias voters towards a given candidate. Both sources of bias are more common in countries with poorly regulated public opinion research and media.

In the case of quantitative models, most of the research is undertaken ex-post and can thus only be used as a contextual element of electoral events, or as independent variables. These studies – which have grown in quality and quantity in recent years – can be more useful as control variables in electoral forecasts and campaign tracking. Models with few variables tapped at few point in time cannot serve to track campaign trends. Electronic markets do not have this problem. Since they present continuous data points, they usually serve as real time trackers of public opinion. The shortcoming of electronic markets is that they are rare in the Latin America. We are only aware of one such experience in the region, *Bolsa Electoral* in Chile.¹

Any one of the four approaches outlined above could be potentially be developed and tested in Latin America to forecast elections. Indeed, as more elections take place and more data becomes available, they should. However, at this point in time the approach based on public opinion stands out. First, and in terms of data availability, because most countries in the region have pollsters and polls. Only some countries are stable enough to use quantitative models and most countries do not have electronic markets. Second, and in terms of data quality, because most pollsters and polls report their methods. Only some regional experts and electronic markets (if available) do the same. Third, and in terms of data parsimony, because polls are the smallest unit of analysis in an aggregate forecast. Both quantitative models and electronic markets are products of aggregate indicators.

4 Two-Stage Model

In this section we propose a two-stage model based on the public opinion approach to forecast elections and track public opinion in Latin America. We argue that the approach fits in well, given the structural and circumstantial conditions of the region. We believe that data reported by pollsters is the best possible source of information to estimate the true value of public opinion during an electoral cycle. In what follows we propose a model

¹Bolsa Electoral is a project of the *Facultad de Economía y Negocios* of the *Universidad de Chile*. It's main purpose has been to predict the outcome of elections in Chile, such as the share of the popular vote for each presidential candidate, by observing the transaction of fictitious electoral shares.

based on poll aggregation and Bayesian inference capable of both forecasting elections and tracking campaign variations in the region. We borrow from advances made in two-party systems with large banks of poll data. However, we design our model to fit the particular characteristics of the multi-party systems and quantity of information in the region.

The two-stage model we propose was developed and tested in Chile, a stable multi-party democracy with regular public opinion polls across numerous electoral cycles. Chile is an interesting case because polls have often been seen as having predictive power. Yet, evidence from the 1999 and 2005 presidential elections shows the contrary; polls are significantly inaccurate when used as predictors. This originally led us to believe that they were also biased instruments to measure campaign variations. We tested our model live with data from the 2009 and 2013 presidential elections on the website *TresQuintos*.² For both elections we produced accurate estimates of the election night result, along with a comprehensive account of campaign variations.

4.1 Stage 1

4.1.1 Poll Aggregation

In the first stage we propose a model based on poll aggregation. Indeed, combining polls is likely to reduce error. Graefe et al. note that by using different forecasts (in this case polls), both selection bias and omitted variable bias can be reduced (2014). Indeed, when polls are normally distributed, combining them into an aggregate measure will output more accurate estimates. This is true in any electoral cycle in which some polls predict a large outcome for a given candidate and other polls predict a low outcome for that same candidate. In these situations it is said that bracketing occurs, and when this takes place aggregate measures will always perform better than a random poll (Larrick and Soll, 2006). A random poll may outperform the aggregate. But a random poll is often irregular over time, making it hard to predict how it will fare beforehand (Graefe et al., 2014; Pasek, 2015).

While simple averages may perform well in some situations, the scarcity of data intrinsic of Latin America warrants a different approach. When some information about the accuracy of polls is available, Scott Armstrong suggests weighing each observation (2001). Since we find that pollsters in the region are often of uneven quality (some may be politically biased or have significant methodological flaws) and the number of polls they output can be small (sometimes only publishing one poll in an entire electoral cycle)

²TresQuintos is a project of Kenneth Bunker and Stefan Bauchowitz. Its main purpose has been to predict the outcome of elections in Chile, such as the share of the popular vote for each presidential candidate, by aggregating polls. Note that the model presented here and used on the website is a work in progress. Hence there are slight variations between the estimates presented here and those on TresQuintos.

we opt for this weighted approach. In constructing the weight, we take into account each poll’s margin of error, the accuracy of its pollster, and the time at which it was fielded. In what follows we review each of these variables.

4.1.2 Margin of Error

The margin of error is a variable that expresses the proportion of error expected to occur in any given poll. As the margin of error increases, the likelihood the given poll is reporting results representative of the population decreases. In contrast, as the margin of error decreases, the likelihood the given poll is reporting results indicative of the “true” value of public opinion increases.

We apply a standard formula, derived for polls that draw simple random samples from a large population. The logic behind the formula is that the numerator represents the confidence level, while the denominator represents the sample size. If we hold the numerator constant at a given confidence level, as the denominator increases, the margin of error decreases. We propose the following specification:

$$\text{moe}_i \approx \frac{k}{\sqrt{N_i}} \quad (1)$$

where i is a poll, N is its sample size, and k reflects the confidence level.³

4.1.3 Accuracy of a Pollster

There are different methods to assess the accuracy of a pollster. It ultimately depends of the type of accuracy that is sought after (Mosteller, 1948). For example, a pollster can be considered accurate if its polls consistently (1) correctly predict the winner of elections, or if its polls consistently (2) correctly predict the margin of victory of elections (Martin, Traugott and Kennedy, 2005).

We apply a method that accounts for the particular features of multi-party systems with multiple candidates (see Wright and Russell, 2013). We assess the accuracy of a pollster according to the overall error of their last available poll for the immediately previous election. That is, we assess each pollsters’ mean squared error. We propose the following specification:

$$\text{accuracy}_p = \frac{1}{C} \sum_{c=1}^C (\hat{y}_{p,c} - y_c)^2 \quad (2)$$

³We use 1.29 as the numerator for a confidence level of 99%, we use 0.98 as the numerator for a confidence level of 95%, and we use 0.82 as the numerator for a confidence level of 90%.

where p is the pollster, y is the observed share of votes for candidate c in the immediately previous election, \hat{y} is the estimated share of votes for candidate c in the immediately previous election, and C is the total number of candidates in the immediately previous election.⁴

4.1.4 Time of the Poll

The time of the poll is a variable that aims to estimate a poll's predictive utility taking into account its age. It works under the assumption that the predictions in a poll published *today* will be less useful after 30 days, and even less than that after 60 days and so forth. This implies that older information becomes less useful as new information becomes available.

We apply a decay function to account for this issue. In our model, an average poll will lose about half its weight in thirty days, but the exact rate of decay of each poll is itself a function of its moe and accuracy. The age is given by the number of days that have elapsed since the poll was fielded. We propose the following specification:

$$\text{decay}_i = \left(\frac{1}{\text{moe}_i} \times \text{accuracy}_p \right) \quad (3)$$

$$\text{time}_i = \left(\frac{\text{decay}_i}{\frac{1}{n} \sum_{i=i}^n \text{decay} \times 0.5} \right)^{\frac{T_i}{30}} \quad (4)$$

where i is a poll, p is a pollster, and T reflects the number of days that have elapsed since the poll was fielded.⁵

4.1.5 Specification of the Model

Once the weights are constructed, polls can be aggregated. Polls with larger margins of error that are fielded by less accurate pollsters weigh less in the overall than polls with smaller margins of error that are fielded by more accurate pollsters. If the two polls are fielded on the same day, the former poll will decay faster than the latter poll.

An example may help clarify this. The first poll in an electoral cycle weighs 100% in the overall. If a second (more accurate) poll is published on the same day, the first poll will weigh less than the second poll in the overall. If a third (less accurate) poll is published that same day, the first poll will

⁴We assign the average pollster weight as the default for pollsters for which no previous information is available.

⁵We use the median date of fieldwork as the date of the poll.

weigh more than the third, but less than the second, in the overall. Weights are recalculated each time a poll is published and assigned to a vector:

$$\mathbf{W} = \frac{\text{time}_i}{\sum_{i=1}^n \text{time}_i} \quad (5)$$

where i is the poll, p is the pollster, and n is the number of polls fielded to date.

Then, the estimated share of votes for each candidate is multiplied by its respective weight. This process is described by the following specification:

$$\hat{z}_{c,t} = \sum_{i=1}^n (\hat{y}_{i,c} \times \mathbf{W}) \quad (6)$$

where \hat{z} is the predicted share of votes for candidate c at time t , (\hat{y}) is the estimated share of votes for that candidate, (\mathbf{W}) is the respective poll's weight, and n is the total number of polls.

4.2 Stage 2

4.2.1 Bayesian Inference

In the second stage we propose a model based on Bayesian inference. We believe it is well suited to estimate the true value of public opinion. Bayesian inference is normally used to update a previously estimated probability given new information (Gelman et al., 2004). As a general principle, we can state Bayes' theorem as follows:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (7)$$

where $P(A|B)$ is the probability of A given that B is true, $P(B|A)$ is the probability of B given that A is true, and $P(A)$ and $P(B)$ are the probabilities of A and B independent of each other.

We use this theorem to update our predicted share of votes for a candidate (say, B) every time a poll publishes an estimated share of votes for that candidate (say, A). It has been frequently applied in electoral studies, including research that has used public opinion polls to estimate electoral returns (Fernandez-i Marin, 2011; Lock and Gelman, 2010; Jackman, 2004; Linzer, 2013; Strauss, 2007). As such, it has been found to increase the overall accuracy of estimates (see Brooks and Gelman, 1998; Gelman and Rubin, 1992) and (Jackman, 2000, 2009).

4.2.2 Bayesian model

We update the predictions derived from Equation 5 with the predictions derived from Equation 6. We use a dynamic linear model to track latent support between each prediction (see Linzer, 2013). While there are no major changes at the point estimate level, there are significant variations at the margin of error level. We borrow from models developed for elections in two-party systems, but adjust it to fit the particular characteristics of elections in multi-party systems (see Gelman et al., 2004). Instead of using a binomial distribution (that uses a single parameter to estimate the margin between one candidate and another), we use a multinomial distribution (that uses many parameters to estimate the relative difference in preferences among all of the candidates).

We use the adjusted proportion of the vote intention reported for each candidate in each poll as the parameter of interest (see Stoltenberg, 2013). We assign the number of respondents that vote for a certain candidate c in a given poll i to a vector \mathbf{y}_i . For example, if a 1000 individuals are interviewed for poll i , and candidate c obtains .4, we infer that $y_{i,c} = 400$. For the multinomial distribution this means that the event of voting for candidate c is observed 400 times; in each individual poll, the counts themselves are considered a multinomial distribution, with parameter π . Accordingly, we propose the following specification:

$$\mathbf{y}_p \sim \text{Multin}(\pi_1, \dots, \pi_c, N_i) \quad (8)$$

where π is the probability of a respondent stating his intention to vote for candidate c , and N_i is the sample size of the poll.

The resulting posterior distribution for π_1, \dots, π_c takes the form of a Dirichlet distribution. This ensures that all parameters π add up to 1. We propose the following specification:

$$\pi \sim \text{Dirichlet}(\alpha_1 + y_1, \dots, \alpha_c + y_c) \quad (9)$$

where α represents the parameter for a voting proportion.

We use a non-informative prior. This is also known as the “let the data speak for itself” approach (Gelman et al., 2004, 51). The prior distribution is thus based on a “vague” estimation of the possible range of support for each candidate. We use the size of the endorsing party as a proxy. We assign a prior of 30 to 60 percent for candidates endorsed by major parties, and priors of 0 to 30 percent for the candidates endorsed by minor parties.

4.2.3 Markov chain Monte Carlo

We fit the Bayesian model using a Markov chain Monte Carlo (MCMC) sampling procedure to provide posterior estimates. This enables us to produce logically valid, probabilistic inferences about the vote proportion for each candidate. The Markov property is that the probability distribution over the next observation depends only upon the current observation. A Markov chain is the string of these numbers.

Because we already know the distribution of interest (the posterior distribution of the parameters), we use MCMC to find a transition kernel to move from one state to the next. A Markov chain uses an equilibrium distribution. The state of the chain after a large number of iterations is used as a sample of the desired distribution. The quality of the sample improves as the number of iterations increases.

As the number of iterations approach infinity, the sample converges on the posterior density of the distribution of interest (Jackman, 2000). In practice, this requires including uncertainty into the pollsters' estimates, such that we are not concerned with individual probabilities and vote shares but with probability distributions. By doing so, we account for the pollsters' estimates of public opinion.

This procedure accounts for the gaps between polls. We estimate daily unobserved vote shares' using a "random walk" (see Jackman, 2004). For the transition model, the posterior distribution at time t reflects probabilities, given by the voting probabilities at time $t - 1$, which are then fed back into the multinomial model as proportions of a pseudo sample of the size of the previous observation.

4.3 Computing

Data are analyzed using the R statistical package (R Core Team, 2012). We generate Markov chains with the Gibbs sampler JAGS (Plummer, 2012). JAGS takes samples from the individual conditional distributions of each parameter of interest. We estimate the sampling distribution for each day in the electoral cycle.⁶

A caveat is warranted. Our model is not designed to produce an election night forecast, since we cannot anticipate the events that take place between the final poll in an electoral cycle and election night. It is designed to track public opinion, since we produce an estimate of vote intention for each candidate during the entire electoral cycle.

⁶We iterate the process 100,000 times, with a 10,000 burn-in.

5 Data

We use data from the 2009 and 2013 elections in Chile. In the first round of the 2009 presidential election (15 December) four candidates ran for office: Sebastián Piñera, Eduardo Frei, Marco Enríquez-Ominami, and Jorge Arrate (see Table 1). Since none of the candidates secured the absolute majority of votes needed to take the presidency outright, a run-off between the two most-voted candidates – Piñera and Frei – was held on 17 January, 2010. Piñera was elected to the presidency with 51.61% of the vote (see Toro and Luna, 2011).

Table 1: Results of the 2009 Chilean Presidential Election

Candidates	Votes	Percent	Votes	Percent
Piñera	3,074,164	44.06	3,591,182	51.61
Frei	2,065,061	29.60	3,367,790	48.39
Enríquez-Ominami	1,405,124	20.14		
Arrate	433,195	6.21		
Total	6,977,544	100.00	6,958,972	100.00

Source. Servel.

In the first round of the 2013 presidential election (17 November) nine candidates ran for office: Michelle Bachelet, Evelyn Matthei, Marco Enríquez-Ominami, Franco Parisi, Marcel Claude, Alfredo Sfeir, Roxana Miranda, Ricardo Israel and Tomás Jocelyn Holt (see Table 2). Since none of the candidates secured the absolute majority needed to take the presidency outright, a run-off between the two most-voted candidates – Bachelet and Matthei – was held on 15 December, 2013. Bachelet was elected to the presidency with 62.16% of the vote (see Bunker, 2014).

Table 2: Results of the 2013 Chilean Presidential Election

Candidates	Votes	Percent	Votes	Percent
Bachelet	3,075,839	46.70	3,470,379	62.16
Matthei	1,648,481	25.03	2,111,891	37.83
Enríquez-Ominami	723,542	10.98		
Parisi	666,015	10.11		
Claude	185,072	2.81		
Sfeir	154,648	2.34		
Miranda	81,873	1.24		
Israel	37,744	0.57		
Jocelyn-Holt	12,594	0.19		
Total	6,585,808	100.00	5,582,270	100.00

Source. Servel.

We consider 11 major national pollsters (see Table 3). In 2009 we consider polls from: Centro de Estudios Públicos (CEP), Centro de Estudios de la Realidad Contemporánea (CERC), El Mercurio/Opina, Giro País/Subjetiva, Imaginación, IPSOS, La Tercera/Feedback, UDD/La Segunda and UDP/ICSO. Three of the 10 pollsters conducted face-to-face interviews (CEP, CERC and UDP/ICSO) in all of their polls. The other seven pollsters conducted telephone interviews. We considered the question: “If the election were held next Sunday, which of the following candidates would you vote for?” We coded answers: “Sebastián Piñera”, “Eduardo Frei”, “Marco Enríquez-Ominami”, “Jorge Arrate”, and “Other/None”.

In 2013 we consider polls from: CEP, Conecta, El Mercurio/Opina, ICHÉM/UA, IPSOS, UDD/La Segunda and UDP/ICSO. Three of the seven pollsters conducted face-to-face interviews (CEP, ICHÉM/UA and UDP/ICSO) in all of their polls. The other four pollsters conducted telephone interviews. We considered the question: “If the election were held next Sunday, which of the following candidates would you vote for?” We coded answers: “Michelle Bachelet”, “Evelyn Matthei”, “Marco Enríquez-Ominami”, “Franco Parisi”, “Marcel Claude”, “Alfredo Sfeir”, “Roxana Miranda”, “Ricardo Israel”, “Tomás Jocelyn-Holt”, and “Other/None”.

Table 3: Pollsters and Polls in Chilean Presidential Elections

Pollster	2009	MSS	2013	MSS
CERC	2	1200		
Giro País/Subjetiva	1	810		
Imaginación	3	1092		
La Tercera/Feedback	1	801		
CEP	2	1505	1	1437
El Mercurio/Opina	3	1200	1	561
Ipsos	1	1522	4	690
UDD/La Segunda	2	1197	2	969
UDP/ICSO	1	1302	1	756
Conecta			1	601
Ichem/UA			1	1708
Total	16		11	

Source. Authors.

Legend. MSS: Mean Sample Size.

For both elections we ignore “Other/None” answers, and use “likely-voter models” when available.⁷ The estimates for each of the candidates are thus rescaled to reflect the proportion of valid answers only. We register the total number N as the sum of the preferences for all of the registered candidates. While this inevitably reduces the size of the samples, it makes theoretical sense. Undecided voters are, on average, normally distributed, and likely-voter models generally help distinguish signals from noise.

For both elections we use data that covers the final 100 days of the campaign. This is the date when candidates are legally required to register. It also marks the point in which pollsters can legitimately consider all candidates in the same scenario. Before candidates register, pollsters can only speculate, and cannot use closed-ended questions. Indeed, before then, pollsters may ask about candidates that eventually do not register, or not ask about candidates that eventually do register. In 2009 we consider polls from September through December, and in 2013 we consider polls from August through November.

6 Results

Figure 1 shows the adjusted poll prediction for the first round of the 2009 election. These are the results derived from Equation 5. It shows the percentage of support for each of the four candidates that ran for election. It

⁷Likely voter models are used by pollsters to ascertain which respondents are probably going to turn out for an election. These models can be understood as the most sophisticated prediction of a pollster.

shows that the order in their preferences did not significantly vary between October and December.

Piñera showed little variation in his vote intention, consistently polling between 40 and 45 percent. Based on the adjusted polls, we estimate that his average vote intention was 43.7% during September-October, 43.4% during October-November and 42.9% during November-December. Frei was the runner-up, polling between 28 and 32 percent. Based on the adjusted polls, we estimate that his average vote intention was 28.4% during September-October, 27.8% during October-November and 27.6% during November-December. Enríquez-Ominami followed in the race, with preferences ranging between 18 and 22 percent. Based on the adjusted polls, we estimate that his average vote intention was 23.2% during September-October, 23.3% during October-November and 22.5% during November-December. Finally, Arrate moved between 1 and 6 percent. Based on the adjusted polls, we estimate that his average vote intention was 4.5% during September-October, 5.3% during October-November and 6.8% during November-December.




Figure 1: Adjusted Polls for the First Round of 2009

Figure 2 shows the Bayesian model with the results for the first round of the 2009 election. These are the results derived from Equation 6. It uses each of the adjusted poll predictions and tests the likelihood they represent the true value of public opinion. The percentage of support for each of the candidates in days that do not have an adjusted poll prediction were esti-

mated with a MCMC random walk. The lines represent the estimated vote proportion for each candidate and the shaded areas portray the credibility interval (analogous to the confidence interval). The symbols represent the adjusted polls. It shows that the estimated vote share for Piñera remained stable throughout the campaign. He began campaigning with a vote intention of 45% in September and finished with a vote intention of 44% in December. Our final estimate for Piñera was 43.4% (he obtained 44.0% of the vote on election night). The precision of our estimate for Piñera was widest in early September and late November.

It also shows that the estimated vote share for Frei and Enríquez-Ominami converged in the middle of the election cycle. While Frei started with a vote intention around 30%, and Enríquez-Ominami started with a vote intention around 20%, they both approached 25% in mid-October. Our final estimate for Frei was 28.4% (he obtained 29.6% of the vote on election night), and our final estimate for Enríquez-Ominami was 22.1% obtained 20.2% of the vote on election night). As with Piñera, the precision for our estimates for both Frei and Enríquez-Ominami were widest in early September and late November. Finally, the figure shows that the estimated vote share for Arrate increased at a slow but steady rhythm. He moved from 3% upon registration towards 7% on election night. Our final estimate for Arrate was 6.0% (he obtained 6.2% of the vote on election night). The precision of our estimate for Arrate was the narrowest of the four candidates.




Figure 2: Bayesian Inference for the First Round of 2009

Figure 3 shows the adjusted poll prediction for the first round of the 2013 election. It shows the percentage of support for four of the nine candidates that ran for election. It shows that the order of preference between the top two candidates did not change in the final 100 days. However, there were important variations for the bottom two candidates.

Bachelet showed the largest variation in her vote intention, polling anywhere between 38 and 55 percent. Based on the adjusted polls, we estimate that her average vote intention was 47.2% during August-September, 47.4% during September-October and 45.7% during October-November. Matthei was the second favorite candidate, polling between 20 and 28 percent. Based on the adjusted polls, we estimate that her average vote intention was 22.4% during August-September, 23.1% during September-October and 24.6% during October-November. Marco Enríquez-Ominami followed in the race, with preferences ranging between 8 and 12 percent. Based on the adjusted polls, we estimate that his average vote intention was 10.1% during August-September, 9.3% during September-October and 8.4% during October-November. Parisi moved closely to Enríquez-Ominami, with preferences also ranging between 8 and 12 percent. Based on the adjusted polls, we estimate that his average vote intention was 12.1% during August-September, 12.3% during September-October and 12.6% during October-November.




Figure 4 shows the model with the results for the first round of the 2013 election. As explained above, it uses each of the adjusted poll predictions, and simulates the likelihood they are the true value of public opinion. It shows that the estimated vote share for Bachelet varied throughout the campaign. She began with a vote intention of 45% in August before climbing to 50% in October and falling to 44% in November. The precision of our estimate for Bachelet was widest in late September and early October. Our final estimate for Bachelet was 46.7% (she obtained 46.7% of the vote on election night).

Figure 4 also shows that the estimated vote share for Matthei remained stable throughout the campaign. She moved between 20% and 25%. As with Bachelet, the precision of our estimate for Matthei was also widest in late September and early October. Our final estimate for Matthei was 23.2% (she obtained 25.0% of the vote on election night). Finally, the figure shows that the vote share for Enríquez-Ominami and Parisi were not significantly different. Both, started and finished with around 10% each. The precision of our estimate for Enríquez-Ominami and Parisi were widest in late September and early October. Our final estimate for Enríquez-Ominami was 9.1% (he obtained 10.8% of the vote on election night), and our final estimate for Parisi was 12.4% (he obtained 10.1% of the vote on election night).




Figure 4: Bayesian Inference for the First Round of 2013

7 Discussion

In this article we developed a two-stage model based on a public opinion approach to forecast elections and track public opinion in Latin America. We tested the model with data from Chile. We showed that election night forecasting and campaign variation tracking is possible in developing democracies that present methodological challenges such as large multi-party systems or a lack of data. We showed that a model that could both mitigate bias and smooth polls over time was a suitable approach. Indeed, we showed that even though we do not produce an outright prediction for election night, in every case our final estimated vote share for a candidate was within two percentage points of the final result.

Our contribution is to have advanced the understanding of forecasting elections and tracking public opinion in Latin America. This is significant because of two reasons. First, because we further the understanding of forecasting and tracking in multi-party systems. Hitherto, most of the literature has focused on two-party systems (or on the margin of victory between the top two candidates). Second, because we advance the understanding of forecasting and tracking in environments with biased or incomplete information. Until now, most of the literature has focused on countries that counter this problem with large collections of data.

One advantage of our model is that it will systemically output more accurate election night forecasts than individual polls. This is true only when bracketing occurs (which normally does), and if the final adjustments is made close to the election. While not technically comparable, if someone is concerned only with the outcome of the election, our model would be the “best guess”. Our final estimation (roughly equivalent to an election night forecast) will inevitably have lower average and total error than the polls it uses as input. In the unlikely case that bracketing does not occur, our total error will be larger, but our average error will still be small. In worst case scenario, our prediction will not improve accuracy, but it will not diminish it either.

One limitation of our model is that it necessarily relies on pollsters. If pollsters are biased, our prediction will be biased. While forecasters in industrialized democracies have numerous polls at their disposal, and can legitimately assume that on the aggregate level they are not biased, this presents a serious limitation in developing democracies. This is the case, for example, when bracketing does not occur. This flaw is apparent in our Bayesian model for the first round of 2013 (see Figure 4), where pollsters consistently predicted that Enríquez-Ominami would come in fourth, only to be proven wrong on election night.

Given the advantages and limitations of our model, we suggest that future literature should consider testing the model in different environments, such as larger multi-party systems with even less available data. We also

suggest future literature should focus on developing a model that can take into account different sources of information. Our model should eventually be adjusted to incorporate more information, such as aggregate economic measures and political variables. Furthermore, adding data on the characteristics of undecided voters could eventually render a more accurate model, since it can help refine both final predictions as well as variations during the campaign.

References

- Abramowitz, Alan. 2012. "Forecasting in a Polarized Era: The Time for Change Model and the 2012 Presidential Election." *PS: Political Science and Politics* 45:618–619.
- Aichholzer, Julian and Johanna Willmann. 2014. "Forecasting Austrian National Elections: The Grand Coalition Model." *International Journal of Forecasting* 30:55–64.
- Alesina, Alberto, John Londregan and Howard Rosenthal. 1993. "A Model of the Political Economy of the United States." *American Political Science Review* 87(1):12–33.
- Armstrong, Scott. 2001. Combining Forecasts. In *Principles of Forecasting: A Handbook for Researchers and Practitioners*, ed. Scott Armstrong. Norwell, MA: Kluwer.
- Auberger, Antoine and Eric Dubois. 2005. "The Influence of Local and National Economic Conditions on French Legislative Elections." *Public Choice* 125(3-4):363–383.
- Bartels, Larry M and John Zaller. 2001. "Presidential Vote Models: A Recount." *PS: Political Science and Politics* 34:9–20.
- Bean, Louis H. 1948. *How to Predict Elections*. New York: Knopf.
- Bélanger, Éric and Jean-Francois Godbout. 2010. "Forecasting Canadian Federal Elections." *PS: Political Science and Politics* 43(4):691–699.
- Bellucci, Paolo. 2010. "Election Cycles and Electoral Forecasting in Italy, 1994–2008." *International Journal of Forecasting* 26(1):54–67.
- BolsaElectoral. 2015. "BolsaElectoral."
URL: <http://www.bolsafen.cl/>
- Brooks, Stephen and Andrew Gelman. 1998. "General Methods for Monitoring Convergence Of Iterative Simulations." *Journal of Computational and Graphical Statistics* 7:434–455.
- Bunker, Kenneth. 2014. "The 2013 Presidential and Legislative Election in Chile." *Electoral Studies* 33(4):346–348.
- Campbell, James E. 2008. "Editor's Introduction: Forecasting the 2008 National Elections." *PS: Political Science and Politics* 41(04):679–682.
- Campbell, James E. 2012. "Forecasting the Presidential and Congressional Elections of 2012: The Trial-Heat and the Seats-In-Trouble Models." *PS: Political Science and Politics* 45:630–634.

- Campbell, James E and James C Garand. 2000. *Before the Vote: Forecasting American National Elections*. Thousand Oaks: Sage.
- Campbell, James E and Kenneth A Wink. 1990. "Trial-heat Forecasts of the Popular Vote." *American Politics Quarterly* 18:251–269.
- Dolado, Juan J, Jesús Gonzalo and Laura Mayoral. 2003. "Long-Range Dependence in Spanish Political Opinion Poll Series." *Journal of Applied Econometrics* 18(2):137–155.
- Dubois, Eric and Christine Fauvelle-Aymar. 2004. *The French Voter: Before and After the 2002 Elections*. Basingstoke: Palgrave Macmillan chapter Vote Functions in France and the 2002 Election Forecast, pp. 205–230.
- Durand, Claire, André Blais and Mylène Larochelle. 2004. "The Polls in the 2002 French Presidential Election: An Autopsy." *Public Opinion Quarterly* 68(4):602–622.
- ElectionForecast. 2015. "ElectionForecast."
URL: <http://www.electionforecast.co.uk/>
- Erickson, Robert S and Christopher Wlezien. 2008. "Are Political Markets really Superior to Polls as Election Predictors." *Public Opinion Quarterly* 72:190–215.
- Erickson, Robert S and Christopher Wlezien. 2012. "The Objective and Subjective Economy and the Presidential Vote." *PS: Political Science and Politics* 45:620–624.
- Erickson, Robert S, Joseph Bafumi and Bret Wilson. 2001. "Was the 2000 Presidential Election Predictable?" *PS: Political Science and Politics* 34:815–819.
- Fair, Ray C. 1978. "The Effect of Economic Events on Votes for President." *The Review of Economics and Statistics* 60(2).
- Fernandez-i Marin, Xavier. 2011. "Estimacio del Suport Latent als Partits Politics: Eleccions al Congrés dels Diputats 2011: Enquestar les Enquestes."
- Fisher, Stephen D, Robert Ford, Will Jennings, Mark Pickup and Christopher Wlezien. 2011. "From Polls to Votes to Seats: Forecasting the 2010 British General Election." *Electoral Studies* 30(2):250–257.
- FiveThirtyEight. 2015. "FiveThirtyEight."
URL: <http://www.fivethirtyeight.com/>
- Gelman, Andrew and Donald Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences." *Statistical Science* 7(4):457–472.

- Gelman, Andrew, John Carlin, Hal Stern and Donald Rubin. 2004. *Bayesian Data Analysis*. second ed. Boca Raton: London: Chapman & Hall.
- Gibson, Rachel and Michael S Lewis-Beck. 2011. "Methodologies of Election Forecasting: Calling the 2010 UK "Hung Parliament"?" *Electoral Studies* 30(2):247–249.
- Gott, J Richard and Wesley N Colley. 2008. "Median Statistics in Polling." *Mathematical and Computer Modeling* 48:1396–1408.
- Graefe, Andreas, Scott Armstrong, Randall Jones and Alfred Cuzán. 2014. "Combining Forecasts: An Application to Elections." *International Journal of Forecasting* 30:43–54.
- Jackman, Simon. 2000. "Estimation and Inference via Bayesian Simulation: An Introduction to Markov Chain Monte Carlo." *American Journal of Political Science* 44(2):375–404.
- Jackman, Simon. 2004. "Pooling the Polls Over and Election Campaign." *Australian Journal of Political Science* 40(4):499–517.
- Jackman, Simon. 2009. *Bayesian Analysis for the Social Sciences*. Chichester: New York: Wiley.
- Jerôme, Bruno and Véronique Jérôme. 2000. "The 1998 French Regional Elections: Why so much Political Instability?" *Electoral Studies* 19(2):219–236.
- Jérôme, Bruno, Véronique Jérôme and Michael S Lewis-Beck. 1999. "Polls Fail in France: Forecasts of the 1997 Legislative Election." *International Journal of Forecasting* 15(2):163–174.
- Jones, Randall J. 2002. *Who will be in the White House?: Predicting Presidential Elections*. New York: Longman.
- Jones, Randall J. and Alfred G. Cuzán. 2008. "Forecasting US Presidential Elections: A Brief Review." *Foresight: The International Journal of Applied Forecasting* 10:29–34.
- Kernell, Samuel. 2000. "Life Before Polls: Ohio Politicians Predict the 1828 Presidential Vote." *PS: Political Science and Politics* 33:568–574.
- Larrick, Richard and Jack Soll. 2006. "Intuitions About Combining Opinions: Misappreciation of the Averaging Principle." *Management Science* 52(1):111–127.
- Lebo, Matthew J and Helmut Norpoth. 2007. "The PM and the Pendulum: Dynamic Forecasting of British Elections." *British Journal of Political Science* 37(1):71.

- Lewis-Beck, Michael S. 1991. "French National Elections: Political Economic Forecasts." *European Journal of Political Economy* 7(4):487–496.
- Lewis-Beck, Michael S and Charles Tien. 2012. "Election Forecasting for Turbulent Times." *PS: Political Science and Politics* 45:625–629.
- Lewis-Beck, Michael S and Mary Stegmaier. 2000. "Economic Determinants of Electoral Outcomes." *Annual Review of Political Science* 3(1):183–219.
- Lewis-Beck, Michael S, Richard Nadeau and Éric Bélanger. 2004. "General Election Forecasts in the United Kingdom: A Political Economy Model." *Electoral Studies* 23(2):279–290.
- Lewis-Beck, Michael S and Tom W Rice. 1984. "Forecasting Presidential Elections: A Comparison of Naive Models." *Political Behavior* 6(1):9–21.
- Lewis-Beck, Michael S and Tom W Rice. 1992. *Forecasting Elections*. Washington: Congressional Quarterly.
- Linzer, Drew A. 2013. "Dynamic Bayesian Forecasting of Presidential Elections in the States." *Journal of the American Statistical Association* 108(501):124–134.
- Lock, Kari and Andrew Gelman. 2010. "Bayesian Combination of State Polls and Election Forecasts." *Political Analysis* 18:337–346.
- Magalhães, Pedro, Luís Aguiar-Conraria and Michael Lewis-Beck. 2012. "Forecasting Spanish Elections." *International Journal of Forecasting* 28:769–776.
- Mainwaring, Scott. 1998. "Party Systems in the Third Wave." *Journal of Democracy* 9(3):67–81.
- Martin, Elizabeth A., Michael W. Traugott and Courtney Kennedy. 2005. "A Review and Proposal for a New Measure of Poll Accuracy." *Public Opinion Quarterly* pp. 342–369.
- Mosteller, Frederick. 1948. "The Pre-Election Polls of 1948." *Bulletin 60, Social Science Research Council* pp. 54–80.
- Mughan, Anthony. 1987. "General Election Forecasting in Britain: A Comparison of Three Simple Models." *Electoral Studies* 6(3):195–207.
- Nadeau, Richard and André Blais. 1993. "Explaining Election Outcomes in Canada: Economy and Politics." *Canadian Journal of Political Science* 26(04):775–790.
- Nadeau, Richard and Michael S Lewis-Beck. 2001. "National Economic Voting in US Presidential Elections." *Journal of Politics* 63(1):159–181.

- Nadeau, Richard, Michael S Lewis-Beck and Éric Bélanger. 2009. "Election Forecasting in the United Kingdom: A Two-Step Model." *Journal of Elections, Public Opinion and Parties* 19(3):333–358.
- Nadeau, Richard, Michael S Lewis-Beck and Éric Bélanger. 2010. "Electoral Forecasting in France: A Multi-Equation Solution." *International Journal of Forecasting* 26(1):11–18.
- Norpoth, Helmut and Thomas Gschwend. 2010. "The Chancellor Model: Forecasting German Elections." *International Journal of Forecasting* 26(1):42–53.
- Pasek, Josh. 2015. "Predicting Elections: Considering Tools to Pool the Polls." *Public Opinion Quarterly* 79(2):594–619.
- Pickup, Mark and Richard Johnston. 2007. "Campaign Trial Heats as Electoral Information: Evidence from the 2004 and 2006 Canadian Federal Elections." *Electoral Studies* 26(2):460–476.
- Plummer, Martyn. 2012. *rjags: Bayesian Graphical Models using MCMC*. R package version 3-9.
URL: <http://CRAN.R-project.org/package=rjags/>
- PrincetonElectionConsortium. 2015. "PrincetonElectionConsortium."
URL: <http://election.princeton.edu/>
- R Core Team. 2012. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.
URL: <http://www.R-project.org/>
- Rhode, Paul W and Koleman S Strumpf. 2004. "Historical Presidential Betting Markets." *Journal of Economic Perspectives* 18:127–141.
- Rosenstone, Steven J. 1983. *Forecasting Presidential Elections*. New Haven: Yale University Press.
- Sanders, David. 1995. "Forecasting Political Preferences and Election Outcomes in the UK: Experiences, Problems and Prospects for the Next General Election." *Electoral Studies* 14(3):251–272.
- Servel. 2015. "Servel."
URL: <http://www.servel.cl/>
- Stoltenberg, Emil. 2013. Bayesian Forecasting of Election Results in Multi-party Systems. Master's thesis University of Oslo.

- Strauss, Aaron. 2007. *Florida or Ohio? Forecasting Presidential State Outcomes Using a Reverse Random Walk*. Paper presented at Princeton Political Methodology Seminar.
- ThreeHundredEight. 2015. "ThreeHundredEight."
URL: <http://www.threehundredeight.com/>
- Toro, Sergio and Juan Pablo Luna. 2011. "The Chilean Elections of December 2009 and January 2010." *Electoral Studies* 30(1):226–230.
- TresQuintos. 2015. "TresQuintos."
URL: <http://www.tresquintos.com/>
- Wleizen, Christopher. 2003. "Presidential Election Polls in 2000: A Study in Dynamics." *Presidential Studies Quarterly* 33:172–186.
- Wright, Malcolm J., David P. Farrar and Deborah F. Russell. 2013. "Polling Accuracy in a Multiparty Election." *International Journal of Public Opinion Research* .