



Bayesian Forecasting of Electoral Outcomes with new Parties' Competition

José Garcia Montalvo
Omiros Papaspiliopoulos
Timothée Stumpf-Fétizon

Decemberr 2018

Barcelona GSE Working Paper Series

Working Paper n° 1065

Bayesian forecasting of electoral outcomes with new parties' competition¹

José G. Montalvo²
ICREA-A & Department of
Economics and Business (UPF) and
Barcelona Graduate School of
Economics

Omiros Papaspiliopoulos
ICREA & Department of Economics
and Business at UPF and Barcelona
Graduate School of Economics

Timothée Stumpf-Fétizon
Warwick University

¹We would like to thank David Rossell and two anonymous referees for their helpful comments. This work was supported by the Spanish Ministry of Economy and Competitiveness [grant number ECO2017-82696P] and through the Severo Ochoa Program for Centres of Excellence in R&D [grant number SEV-2015-0563] and the Government of Catalonia (ICREA-Academia, 2017SGR-616).

²José G. Montalvo: Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona. Telephone: (+34) 93 542 2509. E-mail: jose.garcia-montalvo@upf.edu

Abstract

We propose a new methodology for predicting electoral results that combines a fundamental model and national polls within an evidence synthesis framework. Although novel, the methodology builds upon basic statistical structures, largely modern analysis of variance type models, and it is carried out in open-source software. The methodology is largely motivated by the specific challenges of forecasting elections with the participation of new political parties, which is becoming increasingly common in the post-2008 European panorama. Our methodology is also particularly useful for the allocation of parliamentary seats, since the vast majority of available opinion polls predict at the national level whereas seats are allocated at local level. We illustrate the advantages of our approach relative to recent competing approaches using the 2015 Spanish Congressional Election. In general the predictions of our model outperform the alternative specifications, including hybrid models that combine fundamental and polls' models. Our forecasts are, in relative terms, particularly accurate to predict the seats obtained by each political party.

Keywords: *Multilevel models, Bayesian machine learning, inverse regression, evidence synthesis, elections*

1 Introduction

Forecasting in social sciences is a challenging endeavor. Probably one of the most challenging exercises in this respect is the forecasting of election results. Most of the literature on election forecasting, including its methodological underpinning, has focused on two-party political systems, a “winner-take-all” system for the Electoral College and democracies with a long history of past elections. Instead, in this paper we develop a methodology most appropriate for elections with little historical data for some competing parties, including the case of parties entering the electoral competition for the first time, under a D’Hondt system for allocation of parliamentary seats, and where the vast majority of available opinion polls predict at the national level whereas the seats are allocated at local level.

The scientific approach to electoral forecasting relies mostly on four alternative methodologies: the statistical modeling approach based on fundamentals; the use of polls, either voting intention surveys or party sympathy surveys; the use of political prediction markets based on bets for the candidates; and combination of other methods, or hybrid models¹. The statistical modeling approach, also referred as structural approach, consists of predicting election results from historical and socioeconomic data. An example is the simple “bread and peace” model of Hibbs (2008)². In stable political systems it is known that national election votes are highly predictable from

¹Recently there have been attempts to use social media to predict elections. Using Twitter has been found to be a poor forecasting strategy (Gayo-Avello, 2012). Murthy (2015) shows that tweets are more reactive than predictive. Wang et al. (2015) uses a Xbox gaming platform to show a new methodology to forecast elections in the context of very non-representative survey data.

²There are recent examples of fundamental models used in the context of Spanish political elections. For instance Balaguer-Coll et al. (2015) show that an increase in local public spending increases the likelihood of being reelected.

fundamentals³ while polls are highly variable but contain useful information, specially close to the election day. The aggregation of polls and the use of betting markets are also classical approaches to electoral forecasting. Recently there has been an increasing interest on hybrid models which combine the outcomes of several methods. The most popular hybrid approach is the synthesis of a fundamental model and polls, e.g. Lewis-Beck (2005). Graefe (2015) averages the results of pollsters, prediction markets, experts (journalists and scholars) and quantitative models to produce a combined forecast for the 2013 German election⁴. Lewis-Beck & Dassonneville (2016) and Lewis-Beck et al. (2016) present the canonical structure of this type of models. In this case the fundamental model is a regression on GDP and government popularity. This model is then synthesized with the median of polls, using a second regression, in order to predict the national level result.

Our methodology is also hybrid but it is tailored to situations where there is little historical data to apply existing hybrid methods and where elections are determined by seats won at local level, hence the national average is not that predictive of the party's representation in the parliament. A further reality particularly relevant to the European electoral landscape is that there is limited or no polling at local level. To put things in perspective, after the beginning of the financial crisis many new parties were created in European countries to capitalize on the discontent of voters with the policy reaction to the economic crisis. Dennison & Pardijis (2016) identify 45 "insurgent" parties in Europe, many of them just a few years old, that come across the political spectrum from extreme left to extreme right. Insurgent parties held 1,329 seats in 27 EU countries in 2016, which correspond to 18.3%

³e.g. Gelman & King (1993).

⁴For an application to the US elections see Graefe et al. (2014).

of the total seats of their parliaments. The political landscape in Spain is complicated by the existence of numerous political parties with non-trivial representation in certain parts of the country (the so-called nationalist parties, e.g. in Catalonia or the Basque country), the fact that only a handful of elections have taken place since the restoration of democracy in the country in 1977 after decades of dictatorship, and that electoral polling is not as extensive as in older democracies (e.g. the USA or the UK). Moreover, as in most of the countries, polling is hardly available at higher spatial resolutions than national. However, by far the biggest challenge in the 2015 elections is that two new political parties ended up taking more than 30% of the parliamentary seats when they had no political representation in the previous parliament, which makes the Spanish election a good example of the challenges of forecasting electoral outcomes with new parties' competition.

Our approach is to learn the national average for each party primarily by published polls and use a fundamental model to learn how this national average distills down to local level. In order to deal with little or no historical data for some competing parties, we use a fundamental model of voting intention *trained* on “deep” microdata obtained in the form of pre-electoral surveys. In Spain, these are carried out by the government-sponsored research center CIS, and allow us to estimate the relationship between geographical and demographic characteristics, and voters' choice. The downside of these datasets is that the sample size in some provinces is very low, and the sample might not be representative. We address these issues using post-stratification⁵ based on census data. This model is synthesized with a polling model, which is effectively computing averages of published polls but, at the same time, correcting for potential sources of bias, such as house

⁵See e.g. Chapter 14 of Gelman & Hill (2007).

effects, the varying quality of polling methodologies, as well as time-trending that takes place as the election times approaches. Due to the absence of long historical data the synthesis is not done with a regression, but rather using a Bayesian evidence synthesis approach.

It is easiest to understand what that approach amounts to in the following way: the fundamental model produces simulations of local results for each party; these are transformed to local seats using d'Hondt method; the local results are aggregated at national level for each simulation; each of these simulations then receives a weight which corresponds to how close the implied national average is to that predicted by the polls models; then each implied national seat aggregated allocation is given the corresponding weight and weighted averages are computed to form predictions. We set up the fundamental model parameters so that the implied predictive distribution for national average is fairly flat relative to that obtained from the polling model, hence the fundamental model is useful for learning how the national result distills down to local level and for capturing correlations at that level.

Our approach has close links with recent works in electoral studies. Both the fundamental and the polling model are multilevel regression models. Park et al. (2004) use a multilevel regression model and post-stratification to obtain state level estimates from national polls⁶. Lock & Gelman (2010) use a Bayesian model to obtain a combination of polls with forecasts from fundamentals. They merge a prior distribution, obtained from previous election results, with polls to generate a posterior distribution over the position

⁶Multilevel structures are also relevant for some fundamental models. For instance Elinder (2010) shows that regional unemployment is a factor in the support for national governments.

of each state relative to the national popular vote⁷. The objective of this procedure is not to produce a forecast for the national vote but to develop a methodology that separates national vote from states' relative positions which can be very valuable for individual state forecasts.

The article is organized as follows. Section 2 introduces the challenges for forecasting electoral results in the context of emerging parties describing the Spanish electoral system, and the situation leading to the Congressional elections of 2015. The choice of this example does not compromise the general applicability of our methodology. One of the challenges of forecasting electoral results is related with the fact that the allocation of seats is very different from the proportion of votes at the national level. Section 3 describes our methodology, starting with the consideration of the fundamental model. Section 3 includes also the description of the polls model and the synthesis of both, fundamental and polls models. Section 4 applies the methodology proposed in Section 3 to the Spanish Congressional election of 2015 introduced. Section 5 contains an evaluation of the forecasting accuracy of our model compared with some alternative models proposed recently. Finally, section 6 presents the conclusions.

2 The Spanish 2015 Congressional Election

Since the end of the dictatorship in 1977 Spanish politics was characterized by the alternation in government of two political parties: PP (popular party, conservative) and PSOE (socialists); see Table 1 for the main contenders and their characteristics. Some other small and regional parties also participated in the elections but the two largest parties accounted for 75% to 85% of

⁷For the national popular vote they use the model of Hibbs (2008).

Code	Party	Ideology	2011 Result
PSOE	Partido Socialista Obrero Espanol	Center-left	0.288
PP	Partido Popular	Right-wing	0.446
Pod	Podemos (including IU)	Left-wing	N/A
C's	Ciudadanos	Center-right	N/A

Table 1: Spanish parties active at the national level in the 2015 elections.

the vote. In the 2015 Electoral Campaign there were at least four large parties because of the emergence of two new national parties: Podemos (radical left) and C's (Ciudadanos, liberal). Podemos and C's had no seats in previous Spanish parliaments⁸, whereas in the 2015 elections they ended up with 69 and 40 respectively, out of 350 in total. This structural change is one of the most challenging issues in predicting the results of the 2015 Spanish Congressional Election using standard time series regressions and, in general, in any electoral contest where the emergence of new and large political parties change the electoral environment with respect to previous elections⁹.

The dissatisfaction of a sizeable part of the population with the measures of austerity applied initially by the PSOE government since 2010 led to a popular demonstration that occupied the center of Madrid during several weeks. This social movement was named 11-M since their assemblies began May 11, 2011. In March 11, 2014 this movement crystallized into a new political party named Podemos, which soon had the support, in polls, of 15% of the

⁸Podemos did not even exist at that time.

⁹Another challenging situation for electoral forecasting in the Spanish context took place in 2004 when a terrorist attack took place in Madrid during the last week before the electoral date when no polls are allowed to be run. See J. G. Montalvo (2011). The timing of terrorist attacks to democratic elections has been frequent in recent years in European countries. Obviously, the strategic timing of elections can also be triggered by good economic conditions or business cycle peaks. For a recent reference see Voia & Ferris (2013)

likely voters. Podemos was initially marketed as the Spanish Syriza¹⁰. The leaders of Podemos came mostly from Political Science university departments. Some of them had been members of anti-capitalism parties in the radical left position of the spectrum. Although in their program for the first election they competed in, the European elections of 2014, they included the repudiation of public debt and the nationalization of many industries, their positions evolved later as to avoid extreme policies and try to escape from the radical left tag that they had from the beginning. Podemos ended up in coalition with IU (Izquierda Unida), the old communist party.

In addition, the conservative policies of PP, the corruption associated with conservative politicians and the lack of internal regeneration in the party led to the birth of a new liberal party called Ciudadanos (C's). This party was founded in 2006 but was initially geographically concentrated in Catalonia.

Both Podemos and C's appear in the CIS¹¹ polls as of July 2014. In contrast with Podemos, the support for Ciudadanos was only 0.9% in early 2014, but built up quickly. In July of 2015 polls showed a tie between these two new political contenders while the sum of the two largest political parties has gone down to 50%. Figure 1 depicts the strength of different parties by province in the 2014 European elections.

The primary challenge from a modeling perspective is that Podemos and Ciudadanos have not inherited their electorate from a distinct previous political movement. On the contrary, they are cannibalizing parties with similar ideologies. The following sections describe the modeling alternatives considered to generate a predictive method for the 2015 Spanish Congress-

¹⁰Syriza, or the Coalition of the Radical Left, is the Greek party that won the 2015 legislative election.

¹¹Center for Sociological Research (CIS) a publicly sponsored institution that runs the official polls. See Appendix A1 for the description of the data.

sional election and the difficulties imposed by the emergence of these new political parties.

The Spanish government is appointed by the *Congreso de los Diputados* which consists of 350 representatives. Each of the 52 Spanish provinces elects its own representatives from its seat contingent according to the local electoral outcome. Thus, as in US presidential elections, the popular vote at the national level is not decisive. Therefore, the notion of "local level" corresponds to "province level" in the Spanish electoral system, and we will use these two terms interchangeably in the rest of the article.

Furthermore, the allocation of seats at the province level is proportional, as opposed to the *winner-takes-all* rule that most US states apply in presidential elections. The seat allocation is determined by the *D'Hondt method* and is most easily understood in terms of the equivalent *Jefferson method*, which we may frame in terms of finding the market-clearing point in the market for seats¹².

The Jefferson method is used to find the "price" in votes per seat at which the "demand" for seats by parties equals the available budget. Thus, a simple iterative algorithm consists of increasing the price per seat until the aggregated demand for seats equals the fixed supply. Then, each party obtains the number of seats it can afford at the equilibrium price.

Since seats are an indivisible good, a party may just fall short of being able to buy an additional seat, with the remainder going to waste. This will occur in every province a party runs in. Thus, given a fixed national vote, it is preferable to have a geographically concentrated electorate. This applies

¹²Udina & Delicado (2005) use data on Spanish elections to show the forecast bias of pre-electoral polls when they convert votes into seats using D'Hondt's rule.

to the regionalist parties in Catalonia and the Basque country.

In the Spanish case, there is an additional rule which states that parties must obtain at least 3% of votes in a given province to take part in the allocation. Otherwise, their votes are disregarded. This acts as an additional penalty on smaller parties whose electorate is spread out across the nation.

3 The proposed methodology

This section provides a high-level description of the methodology we propose for predicting electoral outcomes in presence of strong emerging parties. First we specify a fundamental model in the context of emerging parties. This part of the method is based on predicting voting intention on a local level based on survey and census data. Second, we present a methodology to aggregate polls. This approach is based on averaging polls while controlling for house effects and other biases. Finally, we propose a hybrid model that synthesizes the other models. Estimation of these models, out of sample predictions for the 2015 Spanish elections, and comparisons - empirical and conceptual - with alternative fundamental, polls and hybrid models are deferred to Section 4.

3.1 Fundamental model with emerging parties

The basic characteristics of our fundamental model are driven by the following considerations:

- It should return predictions of voting intention at local level, so that they can then be turned into predictions of seat allocation at local level.

- Since these local level predictions will then be aggregated at national level, it is statistically far more efficient (and less prone to biases) to aggregate probabilities computed at local level and turn them into point forecasts at national level, as opposed to providing point estimates at local level and then aggregate those. Effectively we are computing the expectation of a non-linear function of voting intentions at local level, and the exchange of function and expectation matters. Working with probabilities at local level also allows us to capture important correlations between outcomes at the different local units¹³. We therefore adopt a probabilistic model of voting intention at local level, effectively a type of logistic regression.
- Voter choice is fundamentally not binary in the 2015 Congressional Election by contrast, for instance, with the US Presidential Election or previous Spanish Congressional Elections. Therefore, binary choice models are insufficient.
- The drastic change of the political scene and the emergence of strong new parties renders historical models insufficient for prediction since there is little or no data to train them on.

To forecast the territorial distribution of sentiment we use data on individual respondents in pre-electoral surveys. In Spain, these are carried out by the government-sponsored research center CIS¹⁴. These allow us to estimate the relationship between geographical, demographic characteristics and voters' choice. The downside of these datasets is that the sample size in some provinces is small, leading to noisy estimates. Furthermore, their sample may be biased, and in any case our results should depend as little as possible

¹³e.g. Silver (2017)

¹⁴See the Appendix A1.

to any possible bias due to non-representative sample the CIS survey might be subject to.

We correct for both these issues by stratifying the respondents into disjoint "strata"; each "stratum" is a combination of different categorical variables, e.g. "man" (in the variable "sex"), in the age group 36-55 (in the variable "age group"), with tertiary schooling (in the variable "educational level"), employed (in the variable "employment status") who lives in a small community (in the variable "community size") in the province of Albacete (in the variable "province"). Say there are N such strata (there are precisely 8424 in our specific application); and let n be a specific stratum. We model the survey response counts \mathbf{s}_n of a stratum n , which is the vector of counts in that stratum for the votes to each of the available parties, through a multinomial distribution:

$$\mathbf{s}_n | \boldsymbol{\theta} \sim \text{Multinomial}(\boldsymbol{\mu}_n) \quad (1)$$

where $\boldsymbol{\mu}_n$ is the vector of probabilities that a person belonging to such stratum votes for each of the available parties. From the most recent census data, we estimate $w_{i,n}$, the frequency of people in province i that belong to stratum n ; we then estimate the vector of probabilities to vote for each available party in province i as the weighted average

$$\sum_{n=1}^N \boldsymbol{\mu}_n w_{i,n}.$$

This formula stems from the basic decomposition:

$$\begin{aligned} & \text{Prob}[\text{vote party } l \mid \text{province } i] = \\ & \sum_n \text{Prob}[\text{vote party } l \mid \text{stratum } n] \text{Prob}[\text{stratum } n \mid \text{province } i] \end{aligned}$$

This approach is known as post-stratification, see e.g. Park et al. (2004).

The model we use for μ_n is a multinomial logistic regression. In the case of two competing parties it becomes exactly a logistic regression model for the probability of voting for one of the two parties given stratum. For the multi-party reality we are interested in, let the vector μ_n contain elements $\mu_n(l)$, which is the probability of voting party l , among $l = 1, \dots, L$ competing parties. Then, the model takes

$$\mu_n(l) = \frac{e^{f_{n,l}}}{\sum_{m=1}^L e^{f_{n,m}}}$$

where $f_{n,l}$ is a linear combination of dummy variables for the different levels of the different categorical variables that define the stratum:

$$f_{n,l} = \alpha_l + \sum_k \beta_{(k,j_k[n],l)} \quad (2)$$

where $j_k[n]$ is the level of factor k that corresponds to stratum n for party l . The Appendix A2 and A3 contains details on the model and provides the multi-level formulae that define the model rigorously¹⁵.

Therefore, we fit a main effects ANOVA model where each level of every categorical factor gets a different parameter. Additionally, we allow these parameters to differ for each party. The abundance of parameters calls for

¹⁵We follow the standard practice of setting all coefficients of the pivot category (“other parties”) to 0 for simpler interpretation.

some type of regularization, and we opt for a Bayesian multilevel approach here, whereby the parameters associated with a factor are drawn from a common prior; see Appendix A3¹⁶.

3.2 Polls model

3.2.1 An explanatory ANOVA polls model

Polls are published from a few months before until shortly before the elections¹⁷ and give prediction of voting sentiment for each of the parties at national level. The simplest possibility to aggregate polls into a single prediction would be just to average the latest period (one week, two weeks, one month). This local averaging might be carried out using overlapping or non-overlapping windows of time. Forecasting can then be done only under the assumption that there is not going to be a change in public opinion from that time period to the election day.

This local averaging implicitly assumes that polls around a period in time are independent and identically distributed around the true voting sentiment. However, this assumption is unlikely to hold for a variety of reasons¹⁸:

- Polls by the same pollster may exhibit the same systematic bias across elections. For example, some pollsters are subject to political influence, which may lead them to systematic bias. This is known as house effect¹⁹.

¹⁶Stegmueller (2013) concludes that when using multilevel models the Bayesian approach is more robust and generates more conservative tests than the frequentist approach.

¹⁷In Spain a week before, but in Andorra it is allowed to publish polls regarding the Spanish elections up to a day before.

¹⁸In fact Shirani-Merh et al. (2018), in their analysis of 7,040 polls, show that there is a substantial election-level bias and excess of variance with respect to the calculated using the standard random sampling assumption.

¹⁹e.g. Silver (2017), Shirani-Merh et al. (2018)

- Polls preceding the same election may suffer from systematic bias across pollsters. This may be due to common methodological flaws and pollsters manipulating their polls to conform with the fold. We will call this an election effect²⁰.
- Some pollsters' methodology may be superior, leading to lower error variance. Additionally, polls are carried out on samples of varying size²¹.
- Subsequent polls may be trending up or down. We will refer to this as trending.²²

We return to those later, after we have estimated our proposed model in section 4, to show the evidence our data provide for each of those.

We can formalize these components. Let \mathbf{p}_k denote a poll's predictions. Recall that in a multi-party system we have a vector of predictions, one for each competing party. Poll k takes place at some time $t[k]$ ²³, and let $\mathbf{v}_{t[k]}$ be the election result corresponding to poll k , i.e., the result of the election which this poll refers to. As earlier, let $\mathbf{p}_k(l)$ and $\mathbf{v}_{t[k]}(l)$ refer to the predictions and actual result for the l th party. We build a multi-level analysis of variance model for decomposing the error $\mathbf{p}_k(l) - \mathbf{v}_{t[k]}(l)$ as the sum of four terms:

- a time-invariant bias of the pollster that has produced the poll (house effect)
- a pollster invariant bias that applies to each election separately (elec-

²⁰Silver (2017)

²¹Shirani-Merh et al. (2018)

²²See Lock & Gelman (2010) for evidence of trending close to election day, and Linzer (2013) for a stochastic trending model.

²³This is standard multilevel notation, see Appendix A2 for details.

tion effect)

- a linear trend in time, with a coefficient that is allowed to vary across elections but is common to all pollsters (trending)
- a poll-specific idiosyncratic error that could be due to differences in methodologies across pollsters and sampling variability.

Additionally, we learn the correlations between the idiosyncratic errors for different parties, and we allow the corresponding matrices to vary by pollster. Similarly, the effects that refer to different parties are allowed to be correlated, e.g. the house effects of a pollster for different parties. Again, the abundance of parameters calls for regularization, and again we opt for a Bayesian approach to this multilevel model. All in all, the poll errors are modeled as a multivariate Gaussian distribution, the mean component and the covariance of which are implied by the decomposition described above²⁴.

$$(\mathbf{p}_k - \mathbf{v}_{t[k]}) \sim N(\boldsymbol{\gamma}_{j[k]} + \boldsymbol{\delta}_{t[k]} + d_k \boldsymbol{\epsilon}_{t[k]}, \boldsymbol{\Sigma}_{j[k]}) \quad (3)$$

where $\boldsymbol{\gamma}_j$ is the time-invariant bias of pollster j , $\boldsymbol{\delta}_t$ is the pollster-invariant bias in election t , d_k corresponds to how many days before the election poll k was published and $\boldsymbol{\epsilon}_{t[k]}$ is the pollster-invariant strength of the trend in a given election. $\boldsymbol{\epsilon}_{t[k]}$ decays as election day approaches, but $\boldsymbol{\delta}_t$ applies to all polls until the election²⁵. As with the fundamental model we use the Bayesian multilevel paradigm to deal with the abundance of parameters and refer to Appendix A4 for details on the prior distributions we have used.

²⁴Polls of different pollsters in the same election are dependent through their dependence on the election effect, polls of the same pollster in different elections are dependent through their dependence on the common house effect, etc.

²⁵See also J. Montalvo et al. (2016). The model in Shirani-Merh et al. (2018) includes a bias for each poll that is allow to change linearly over time and a variance term that captures residual variability.

In summary, we build an ANOVA model to *explain* the errors $\mathbf{p}_k(l) - \mathbf{v}_{t[k]}(l)$. From this perspective, this is not a predictive model, rather is one to understand the importance and relative magnitude of different sources of published polls variability.

3.2.2 Turning the explanatory into a predictive polls model

The model we propose in the previous section implies a joint density for all the available polls in a given election conditional on the election result:

$$\text{Prob}[\text{available polls} \mid \text{new election result}]$$

This density is obtained through the linear transformation

$$\text{poll} = \text{result} + \text{poll error}$$

and the model for the poll error we have built already. However, what we need is the "inverse probability"

$$\text{Prob}[\text{new election result} \mid \text{available polls}]$$

which we can get by use of Bayes' theorem, as

$$\text{Prob}[\text{new election result} \mid \text{available polls}] \propto$$

$$\text{Prob}[\text{available polls} \mid \text{new election result}] \text{Prob}[\text{new election result}].$$

Therefore, to get a predictive model we need a prior model for the elections to come to be combined with the explanatory model we have built. The approach followed here is an instance of what is known as *inverse regression*,

a popular approach to predictive modeling with high-dimensional data²⁶. In the hybrid model we propose in the sequel we get the fundamental model to serve as a prior. A simpler alternative, which is good enough for the purpose of predicting national average results but not seat allocation, is to use a uniform prior on the result, in which case

$$\begin{aligned} \text{Prob}[\text{new election result} \mid \text{available polls}] &\propto \\ \text{Prob}[\text{available polls} \mid \text{new election result}], \end{aligned}$$

the latter seen as a function of "new election result". Effectively, we exploit the symmetry of the Gaussian distribution in our errors model to create the predictive model as:

$$\text{result} = \text{poll} + \text{poll error}$$

The details on how to generate predictions using this predictive model are included in the Appendix A4.

3.3 The hybrid model

The basis of our hybrid model is the conditional probability we obtained in the previous section:

$$\begin{aligned} \text{Prob}[\text{new election result} \mid \text{available polls}] &\propto \\ \text{Prob}[\text{available polls} \mid \text{new election result}] \text{Prob}[\text{new election result}]. \end{aligned}$$

We use the explanatory polls model to produce the first density and the fundamental model based on surveys to produce the second. Operationally, we carry out the following procedure:

²⁶see Cook & Ni (2005)

1. We produce a simulations of local results according to the fundamental model; let one such simulation be $\mathbf{v}_{i,s}$ for $i = 1, \dots, I$, where i indicates local district and s the simulation count.
2. For each simulation we aggregate result at national level to obtain a simulation of \mathbf{v}_s .
3. Provide weight to each simulation according to

$$\mathcal{W}_s = \text{Prob}[\text{available polls} | \mathbf{v}_s]$$

which is computed from the polls explanatory model.

4. Produce election predictions by computing weighted averages

$$\frac{\sum_{s=1}^S g(\mathbf{v}_{1,s}, \dots, \mathbf{v}_{I,s}) \mathcal{W}_s}{\sum_{s=1}^S \mathcal{W}_s}$$

where g is a function of interest of local results. We are particularly interested in the function that by using D'Hondt's method maps local level results to national level number of seats for each party. Apart from point estimates we can also produce also predictive intervals in a similar way.

In the implementation described above one wants to take S , the total number of simulations, as large as possible, basically the limiting factor is the computational budget described in Section 3.2.2. In Section 4 we show graphical and numerical evidence that the resultant inference on the national-level result is dominated by the polls model; actually we devise a metric according to which we obtain that less than 35% of the hybrid model inference is driven by the prior. In the Appendix we explain how we get the fundamental model to have minor influence on the hybrid prediction of national averages.

3.4 The computational approach underpinning the methodology

The fundamental model and the polls model are ANOVA-type models, formulated as Bayesian multilevel models as described in some detail in the Appendix. We follow a fairly standard specification of the prior distributions for these models, as for example explained in Gelman et al. (2006); Gelman & Hill (2007). We estimate these models, and produce out-of-sample predictions subsequently, as implemented in *Stan* (Carpenter et al. (2017)) an open-source platform for performing full Bayesian statistical inference, which carries out Hamiltonian Markov chain Monte Carlo sampling from the corresponding posterior distributions.

What happens behind the scenes of the procedure described for the hybrid mode is the use of *importance sampling*²⁷ to explore the distribution

$$\text{Prob}[\text{new election result} \mid \text{available polls \& survey data}],$$

and the use of a sampling-based approach to carry out the Bayesian updating²⁸.

4 Results for the 2015 Spanish congressional election

In this section we apply the methodology discussed previously to the case of the Spanish Congressional election of 2015. In stable political systems it is known that the national outcome is highly predictable from fundamentals

²⁷See Chapter 2 of Liu (2001).

²⁸Smith & Gelfand (1992)

and past results. This was the case in the Spanish political system until before 2015. For example, Figure 2 (left) plots the electoral result of the 2000 election versus that of the 1996 election for the PP (blue) and the PSOE (red) in each province of Spain, numbered according to the standard postcode coding of Spanish provinces. The picture is similar in other elections prior to 2015. The results positions of provinces relative to the national average are particularly well predictable. This manifests itself through regression lines that are almost parallel the 45° line. Historical models typically include predictors such as unemployment rates, growth of personal disposable income, lagged electoral outcomes, presidential incumbency, regional trends, presidential approval, presidential home advantage (or the corresponding adjustment for party leader home province), partisanship or ideology of each state/district, etc.²⁹

However, the usefulness and predictive ability of historical models for the 2015 Congressional Election is questionable. To start with, such a model cannot provide predictions for parties with no competitive participation in previous elections. Additionally, even when making predictions for PP and PSOE, the traditional political players in Spain, it is unlikely that the model estimated under completely different political environment would have any applicability in this new situation. Figure 2 (right) shows that the pattern observed in past elections has indeed been broken in the 2015 elections.

²⁹For instance Campbell (1992, 2008), Gelman & King (1993), Klarner (2008), Fair (2009) or Hummel & Rothschild (2014).

4.1 Learning the fundamental model

To train the fundamental model for the 2015 Spanish Congressional Election, we use the 2015 CIS pre-electoral survey³⁰. This results in a total sample of about 17,452 respondents. We drop all respondents that did not report their voting intention from the sample, which amounts to assuming that their voting intention data is *missing at random*. We include factors for the province, size of the municipality, gender, age, education and labor market activity of the respondent. The categories of each variables are described in Table A.1 of Appendix A.1. The results of the estimation of the Bayesian multinomial logit model are summarized in Figure 3.³¹

In practice, because responses may not reflect behavior at the voting booth accurately, and because of the possibility of shifts between the survey and the election date, we inflate the uncertainty about the constants α , reflecting uncertainty about the national vote. In practice, this corresponds to multiplying MCMC draws of α by 1.5³².

To interpret the estimates we should notice that:

- Since the model is overparameterised, some parameters are weakly identified. This manifests in wide marginal distributions. These overstate uncertainty since the parameters are highly correlated: once a parameter has been fixed, the uncertainty resolves.
- As we fix the parameters of the pivot party (essentially consisting of regional parties) to 0, all estimates must be interpreted *relative* to these parties. Therefore, a positive intercept estimate for PSOE implies that

³⁰This surveys already contain questions about the two new parties.

³¹We run 4 chains in *Stan* with 2000 iterations each, half of which we discard.

³²The choice of that factor is partially motivated by computational constraints that arise when combining the fundamental model with the polls.

the average respondent is more likely to vote for PSOE than regional parties.

The province effects in Figure 3 indicate that PP has the most variable territorial distribution, while Podemos is fairly constant. PSOE is strong in Andalucia and Extremadura and fairly weak in Catalonia and the Basque Country. PP has its strongest base in Castille and Murcia, but is extremely weak in Catalonia and the Basque Country.

As to the other factors, Podemos and C's are slightly more urban, while the other parties' support does not vary along that dimension. PSOE and PP mostly appeal to uneducated voters. Labour market activity is mostly irrelevant after controlling for other factors. Figure 4 illustrates point predictions of the fundamental model after post-stratification, and how they compare against the actual 2015 election results.

4.2 Learning the pollster model

We use the pollster model we have described in the 2015 Spanish elections. We work with the results of 157 electoral polls published before the Congressional Elections of 1996, 2000, 2004, 2008 and 2011. This set corresponds to the subset of polls published up to 30 days before a Congressional Election.

Exploratory analysis reveals that the uncertainty about the election result close to election day by far exceeds the sampling uncertainty³³. Rolling averages, like the ones depicted in Figure 5, do not provide a direct measure of uncertainty, which is essential to building a probabilistic model. Averaging multiple polls does not eliminate the excess uncertainty. Furthermore, we sometimes observe sharp trending close to election day, even after prolonged

³³This result is also supported by the evidence provided in Shirani-Merh et al. (2018).

periods of stability. Following (and extrapolating) the trend usually takes us closer to the election result than simple averaging. Figure 6 shows for the 2015 Spanish elections how the declared margin of error in the polls, usually given as the inverse of the square root of the sample size, tends to underestimate the true uncertainty. Furthermore, using a linear trend brings us closer to the election result.

Figures 7, 8 and 9 depict the marginal distribution of pollster bias, election bias and election trend respectively³⁴. Estimated pollster biases β are generally consistent with political expedience. For example, the pollster *Sigma dos*, which mostly provides polls for the right-leaning newspaper *El Mundo*, has a consistent bias in favor of the Popular Party. *Invymark*, the pollster selected by left-leaning TV station La Sexta, shows a consistent bias in favor of the Socialist party. By contrast the polls run by the *CIS*, the public pollster, do not show any specific party bias. Election biases δ_t are large, with pollsters collectively missing the PSOE-PP differential by 7 percentage points, calling into question the quality of Spanish polling and the predictability of Spanish elections in general. Estimated trend effects ϵ_t are large in many elections, which confirms that some trend adjustment is necessary even within the last 30 days. Finally, election biases seem to coincide in sign and magnitude with trends, especially in the 2004 elections, but we deem our sample to be too small to draw further conclusions.

4.3 Synthesised predictions

Figure 10 shows how the synthesis operates in the 2015 election. As in Bayesian updating for normal distributions, the posterior's location is a

³⁴Hyperpriors are set in accordance with Stan reference priors. We sample from the models using *Stan*, running 4 chains with 2000 iterations each, half of which we discard.

compromise between prior and likelihood while inverse variance is approximately additive. Aggregating polls substantially improves the PP and Podemos prediction even though the C's forecast does benefit less. The benefits of aggregation are limited when pollsters show correlated errors due to herding or common methodological shortcomings. Since our framework explicitly allows for such a scenario, we manage to avoid undue confidence and preserve some probability mass at the outcome. We observe that the location of the predictive distribution over the national vote is largely driven by the polls model. While it is difficult to obtain quantitative importance weights for prior and likelihood in general Bayesian models, such weights exist for the case where both are Gaussian. When approximating prior and posterior by Gaussian distributions, we find a weight of 35% for the fundamental model into forming synthesized beliefs for the national average.

Figure 11 shows the predictive seat distribution for the largest five parties. The result is close to the predictive mode for PSOE, PP and Podemos, while the result of C's is in the left tail of the predictive distribution. The predictive distribution that we generated the figure from may also be used to evaluate the probability of other event on said distribution. Examples include the probability of a party coming in first or the probability of a coalition of parties achieving a parliamentary majority. The figure exhibits some features that illustrate the benefits of our strategy to model seat allocation explicitly:

- PSOE and PP are granted more seats than their national vote forecast implies. This is due to the rural bias of the provincial seat allocation.
- There are long right tails in the marginals for Podemos and C's even though national vote forecasts are symmetric. These are a consequence

of the Hondt allocation process that delivers increasing benefits to scale.

- The seat forecasts for PSOE and PP are wider than they are for Podemos and C's even though uncertainty when predicting the national vote is similar.

Figure 12 shows point predictions versus actual results and it is directly comparable to Figure 4. Both figures reveal that the CIS survey data is miscalibrated, that is, it predicts a variance between province that is too large. The phenomenon applies to all parties and it is visible before and after post-stratification. Predictions could be recalibrated by shrinking all forecasts towards a party's national mean, but this would require that the phenomenon persists between election. Our analysis of the *electoral barometers* published by CIS every trimester confirms that the miscalibration is persistent, but unfortunately the extent of the phenomenon seems to vary from one survey to the next.

5 Predictive evaluation against alternative models

While the primary appeal of our model lies in its ability to flexibly incorporate polling data into a coherent spatio-temporal probabilistic forecast, we also intended to deliver an improved point forecast relative to more straightforward approaches. In this section we elaborate several alternative models for each component and the hybrid model, and we show the gains in predictive accuracy achieved through our methodology. In keeping with the spirit of the paper, we also evaluate separately the two components of our hybrid model, the fundamental and the polls model, in isolation against their

respective alternative models. Since many alternative models are usually limited to giving a point prediction of the national-level result of the parties, we only evaluate our model based on such point predictions although our main object of interest is the distribution of seats. Additionally, in order to provide a fair comparison with other models, we only consider the prediction of the two traditional parties, avoiding the consideration of the emergent parties. However, we should notice that our modeling strategy is mostly informed by the need to accommodate the participation of new parties in political elections.

5.1 Selecting alternative models

As discussed in the 1, the prediction of elections usually rests on the specification of a fundamental model or the aggregation of the results of electoral polls and other sources of information. Recent literature on election forecasting also resorts to some combination of prediction from fundamental variables and averages of polls. In this section we describe some popular models for predicting election, and compare their predictive performance with our proposal.

We can specify an alternative fundamental model as a regression model that predicts the party's result from the growth of GDP per capital during the year preceding, and its result in the precious election³⁵:

$$\text{result} = \beta_0 + \beta_1 \times \text{lagged result} + \beta_2 \times \text{gdp growth} + \text{residual} \quad (4)$$

³⁵This specification is similar to Lewis-Beck & Dassonneville (2016) and Lewis-Beck et al. (2016) but using the results in the previous elections instead of the government approval rate, since we want to show the predictive ability of our model not only the incumbent but also for the main challenging party.

We consider separate versions of that model for predicting the national vote share and the log number of parliamentary seats won by the incumbent party.

We construct an alternative specification for the polls' model as a linear regression model that predicts the proportion of votes for a party as a simple average of all national level polls published up to 30 days before election day³⁶:

$$\text{result} = \gamma_0 + \gamma_1 \times \text{polls average} + \text{residual} \quad (5)$$

Finally, the alternative hybrid model combines the predictors of the fundamental model and the polls in a single regression following Lewis-Beck & Dassonneville (2016) and Lewis-Beck et al. (2016):

$$\begin{aligned} \text{result} = \delta_0 + \delta_1 \times \text{lagged result} + \delta_2 \times \text{gdp growth} \\ + \delta_3 \times \text{polls average} + \text{residual} \quad (6) \end{aligned}$$

Analogously, as we already pointed out, we define a model for the main challenging party, which is either PP or PSOE in the set of available elections.

5.2 Estimating the alternative models

We estimate the parameters of the alternative models through ordinary least squares (OLS). The following numbers pertain to the incumbent model trained for predicting the 2015 election, i.e. the model that includes all elections up to 2015 in its training set. This matches the training set that we

³⁶This specification is also similar to Lewis-Beck & Dassonneville (2016) and Lewis-Beck et al. (2016) who set the predictions of the polls' model to the aggregate median vote intention. In our case the median and the mean are very similar and, since many standard models like Graefe (2015) use the mean, we decided to use the mean.

used to train our proposed model. For the fundamental model, we obtain the following parameter estimates:

$$\text{votes} = 0.306 - 0.054 \times \text{lagged votes} + 0.035 \times \text{gdp growth} + \text{residual} \quad (7)$$

$$\text{log seats} = 5.994 - 0.255 \times \text{log lagged seats} + 0.108 \times \text{gdp growth} + \text{residual} \quad (8)$$

For the polls model, we obtain the following parameter estimates:

$$\text{votes} = 0.023 + 0.943 \times \text{polls average} + \text{residual} \quad (9)$$

$$\text{log seats} = 3.846 + 3.007 \times \text{polls average} + \text{residual} \quad (10)$$

For the synthetic model, we obtain the following parameter estimates:

$$\begin{aligned} \text{votes} = & 0.429 - 0.783 \times \text{lagged votes} + 0.011 \times \text{gdp growth} \\ & + 0.649 \times \text{polls average} + \text{residual} \end{aligned} \quad (11)$$

$$\begin{aligned} \text{log seats} = & 8.058 - 0.743 \times \text{log lagged seats} + 0.062 \times \text{gdp growth} \\ & + 1.468 \times \text{polls average} + \text{residual} \end{aligned} \quad (12)$$

Initially, we check the performance of the two components of the hybrid model. First of all, as shown in table 2, our *fundamental* model outperforms the alternative fundamental model, giving errors of .042/.008 compared to .044/.251 in 2015 for the incumbent and the main challenging party (PP/P-SOE). For the 2016 election our model delivers even smaller errors than the alternative model. We should keep in mind that the primary goal of the fundamental model is to provide local information on party strength, and

Election	Party	Outcome	Alternative		Our model	
			Estimate	Residual	Estimate	Residual
2015	PSOE	0.220	0.471	-0.251	0.228	0.008
2015	PP	0.287	0.331	-0.044	0.244	0.042
2016	PSOE	0.226	0.143	0.082	0.223	0.003
2016	PP	0.330	0.443	-0.112	0.262	0.067

Table 2: Predictive accuracy of the proposed *fundamental vote* forecasting model compared to the benchmark model. Predictions are made out of sample. Both models use the same set of elections for training.

Election	Party	Outcome	Alternative		Our model	
			Estimate	Residual	Estimate	Residual
2015	PSOE	0.220	0.277	-0.057	0.213	0.007
2015	PP	0.287	0.281	0.005	0.293	-0.006
2016	PSOE	0.226	0.229	0.002	0.216	0.010
2016	PP	0.330	0.303	0.027	0.302	0.028

Table 3: Predictive accuracy of the proposed *polls vote* forecasting model compared to the benchmark model. Predictions are made out of sample. Both models use the same set of elections for training.

therefore the national level point estimates are of secondary importance. See table 2 for all the estimates and outcomes.

Secondly, the alternative *polls'* model has an error of .005/.057. in 2015, whereas our pollster model has an error of .006/.007 (PP/PSOE). Thus, the quality of point estimates is slightly superior in our model. Simple polls averages predict similarly well with errors of .012/.009. See table 3 for estimates and outcomes³⁷. In this particular case it seems that modeling the biases of the polls do not provide a large advantage over other methods but, in general, it could potentially improve substantially the forecasting. For instance, the simple average would have been inadequate in elections where polls exhibit strong trending during the last months. Our proposed polls

³⁷Similar conclusions apply for the 2016 elections.

Election	Party	Outcome	Alternative		Our model	
			Estimate	Residual	Estimate	Residual
2015	PSOE	0.220	0.607	-0.387	0.203	0.017
2015	PP	0.287	0.273	0.013	0.275	0.012
2016	PSOE	0.226	0.261	-0.035	0.212	0.014
2016	PP	0.330	0.392	-0.062	0.293	0.037

Table 4: Predictive accuracy of the proposed *hybrid vote* forecasting model compared to the benchmark model. Predictions are made out of sample. Both models use the same set of elections for training.

model accounts for such trending and should yields better point predictions in those situations.

The results of the comparison of the predictive performance of our model and the alternative hybrid model are presented in Table 4 (proportion of votes) and Table 5 (seats). Our *hybrid* model outperforms the alternative model in the 2015 and the 2016 election with respect to predicting the national vote share. The comparison of the forecast with respect to our pollster model leads to less conclusive results. In fact, in this case, our pollster model seems to work a bit better than the hybrid model in predicting the national vote. However, as we have already argued in previous sections, the most likely advantage of our methodology is related with the forecasting of parliamentary seats. Table 5 shows that our forecast using the hybrid model reduces the error down to 2/14 seats from about 17/47 (PP/PSOE) of the alternative hybrid model in the 2015 elections³⁸. The large improvement in the performance of our model in the prediction of the seats in the parliament is consistent with the objectives of the specification of our fundamental model. As we discussed previously, our fundamental models tries to get information on the geographical distribution of votes, which improves greatly

³⁸The results of our model for forecasting seats are also better than the ones produced by the alternative model in the 2016 election.

Election	Party	Outcome	Alternative		Our model	
			Estimate	Residual	Estimate	Residual
2015	PSOE	90.00	137.57	-47.57	75.47	14.53
2015	PP	123.00	105.65	17.34	125.32	-2.31
2016	PSOE	85.00	112.97	-27.97	79.72	5.28
2016	PP	137.00	165.46	-28.46	119.11	17.88

Table 5: Predictive accuracy of the proposed *hybrid seats* forecasting model compared to the benchmark model. Predictions are made out of sample. Both models use the same set of elections for training.

the prediction of seats given the very nonlinear nature of the relationship between national proportion of votes and seats allocation.

6 Conclusions

This paper proposed a methodology to forecast electoral outcomes using the result of the combination of a fundamental model and a model-based aggregation of polls. We propose a Bayesian hierarchical structure for the fundamental model that synthesizes data at the provincial, regional and national level. We use a Bayesian strategy to combine the fundamental model with the information coming from recent polls. This model can naturally be updated every time new information, for instance a new poll, becomes available. This methodology is well suited to deal with increasingly frequent situations in which new political parties enter an electoral competition, although our approach is general enough to accommodate any other electoral situation. We illustrate the advantages of our method using the 2015 Spanish Congressional Election in which two new parties ended up receiving 30% of the votes. We compare the predictive performance of our model versus alternative models. In general the predictions of our model outperform the

alternative specifications, including hybrid models that combine fundamental and polls' models. Our predictions are, in relative terms, particularly accurate to predict the seats obtained by each political party.

References

- Balaguer-Coll, M. T., Brun-Martos, M. I., Anabel Forte, A., & Tortosa-Ausina, E. (2015). Local governments' re-election and its determinants: New evidence based on a bayesian approach. *European Journal of Political Economy*, *39*, 94-108.
- Campbell, J. E. (1992). Forecasting the presidential vote in the states. *American Journal of Political Science*, 386-407.
- Campbell, J. E. (2008). Evaluating us presidential election forecasts and forecasting equations. *International Journal of Forecasting*, *24*(2), 259-271.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, *76*(1).
- Cook, D., & Ni, L. (2005). Sufficient dimension reduction via inverse regression: A minimum discrepancy approach,. *Journal of the American Statistical Association*, *100*(470), 410-428.
- Dennison, S., & Pardijs, D. (2016). *The world according to europe's insurgent parties*. European Council on Foreign Relations.
- Elinder, M. (2010). Local economies and general elections: The influence of municipal and regional economic conditions on voting in sweden 1985-2002. *European Journal of Political Economy*, *26*, Issue 2, 279-292.
- Fair, R. C. (2009). Presidential and congressional vote-share equations. *American Journal of Political Science*, *53*(1), 55-72.

- Gayo-Avello, D. (2012). No, you cannot predict elections with twitter. *IEEE Internet Computing*, 16(6), 91–94.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multi-level hierarchical models* (Vol. 1). Cambridge University Press New York, NY, USA.
- Gelman, A., & King, G. (1993). Why are american presidential election campaign polls so variable when votes are so predictable? *British Journal of Political Science*, 23(4), 409–451.
- Gelman, A., et al. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian analysis*, 1(3), 515–534.
- Graefe, A. (2015). German election forecasting. comparing and combining methods for 2013. *German Politics*, 24(2), 195–204.
- Graefe, A., Armstrong, J. S., Jones, R. J., & Cuzán, A. G. (2014). Combining forecasts: an application to elections. *International Journal of Forecasting*, 30(1), 43–54.
- Hibbs, D. A. (2008). Implications of the ‘bread and peace’ model for the 2008 us presidential election. *Public Choice*, 137(1), 1–10.
- Hummel, P., & Rothschild, D. (2014). Fundamental models for forecasting elections at the state level. *Electoral Studies*, 35, 123–139.
- Klarner, C. (2008). Forecasting the 2008 us house, senate and presidential elections at the district and state level. *PS: Political Science & Politics*, 41(4), 723–728.

- Lewis-Beck, M. (2005). Election forecasting: Principles and practice. *The British Journal of Politics & International Relations*, 7(2), 145–164.
- Lewis-Beck, M., & Dassonneville, R. (2016). Forecasting methods in europe: synthetic models. *Research and Politics*, 2(1), 1-11.
- Lewis-Beck, M., Nadeau, R., & Belanger, E. (2016). The british general election: synthetic forecast. *Electoral Studies*, 41(1), 264-268.
- Linzer, D. A. (2013). Dynamic bayesian forecasting of presidential elections in the states. *Journal of the American Statistical Association*, 108(501), 124–134.
- Liu, J. (2001). *Monte carlo strategies in statistical computing*. Springer, New York.
- Lock, K., & Gelman, A. (2010). Bayesian combination of state polls and election forecasts. *Political Analysis*, 18(3), 337–348.
- Montalvo, J., Papaspiliopoulos, O., & Stumpf-Fetizon, T. (2016). Predicting elections with emerging political parties,. *arXiv:1612.03073*.
- Montalvo, J. G. (2011). Voting after bombings: a natural experiment on the effect of terrorist attacks on democratic elections,. *Review of Economics and Statistics*, 93(4), 1146–1154.
- Murthy, D. (2015). Twitter and elections: are tweets, predictive, reactive, or a form of buzz? *Information, Communication & Society*, 18(7), 816–831.
- Park, D. K., Gelman, A., & Bafumi, J. (2004). Bayesian multilevel estimation with poststratification: state-level estimates from national polls. *Political Analysis*, 12(4), 375–385.

- Shirani-Merh, H., Rothschild, D., Goel, S., & Gelman, A. (2018). Disentangling bias and variance in election polls. *Journal of the American Statistical Association*, *113*(522), 607-614.
- Silver, N. (2017). Fivethirtyeight's pollster ratings,. URL <https://projects.fivethirtyeight.com/pollster-ratings/>.
- Smith, A., & Gelfand, A. (1992). Bayesian statistics without tears: A sampling-resampling perspective,. *The American Statistician*, *46*(2), 84-88.
- Stegmueller, D. (2013). How many countries for multilevel modeling? a comparison of frequentist and bayesian approaches. *American Journal of Political Science*, *57*(3), 748–761.
- Udina, F., & Delicado, P. (2005). Estimating parliamentary composition through electoral polls. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *168*(2), 387–399.
- Voia, M.-C., & Ferris, J. S. (2013). Do business cycle peaks predict election calls in canada? *European Journal of Political Economy*, *29*, 102-118.
- Wang, W., Rothschild, D., Goel, S., & Gelman, A. (2015). Forecasting elections with non-representative polls. *International Journal of Forecasting*, *31*(3), 980-991.

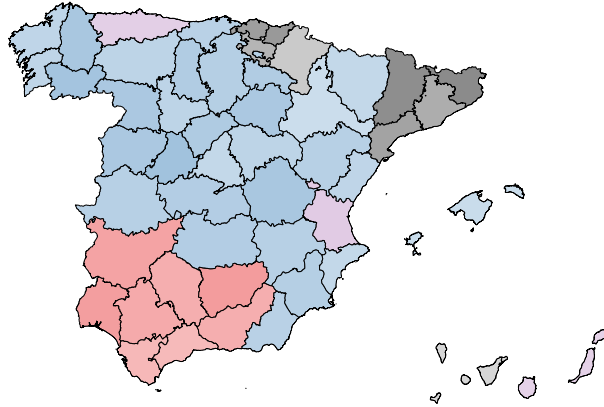


Figure 1: Map of Spanish provinces colored by strongest party in the 2014 European elections and degree of dominance, darker shades corresponding to stronger dominance. Legend: PSOE (red), PP (blue), Podemos+IU (purple), others (gray).

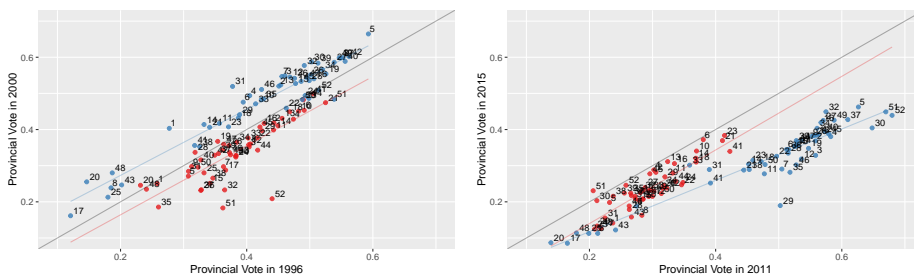


Figure 2: Scatterplot of lagged vote share vs current vote share in 2000 (left) and 2015 (right) relative to previous result, plus robust linear regression line. Legend: PSOE (red), PP (blue). In grey the 45° line. The labels refer to the INE province code.

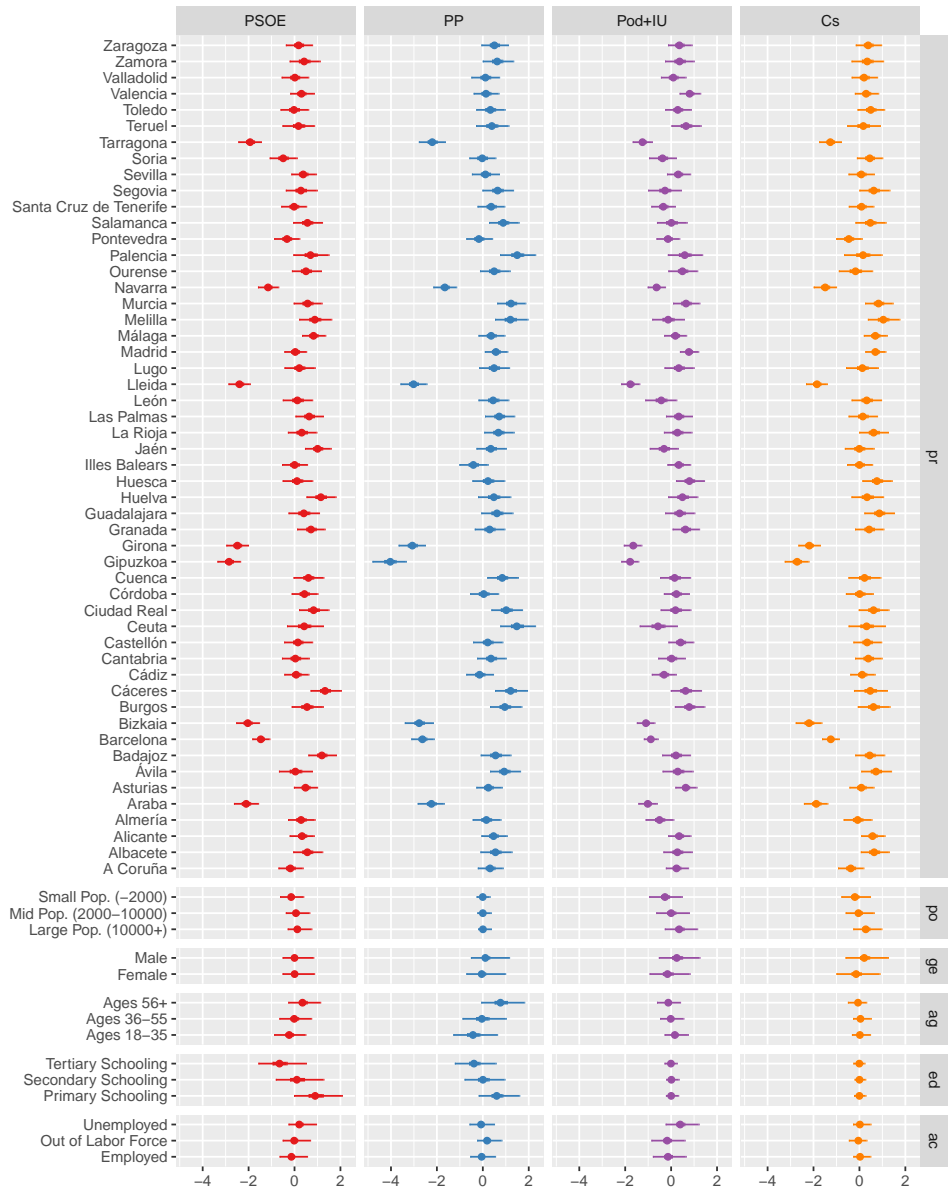


Figure 3: β_{j_i} marginal distributions (sentiment model level coefficients): median (point), 50 percent credibility interval (thick line) and 95 percent credibility interval (thin line).

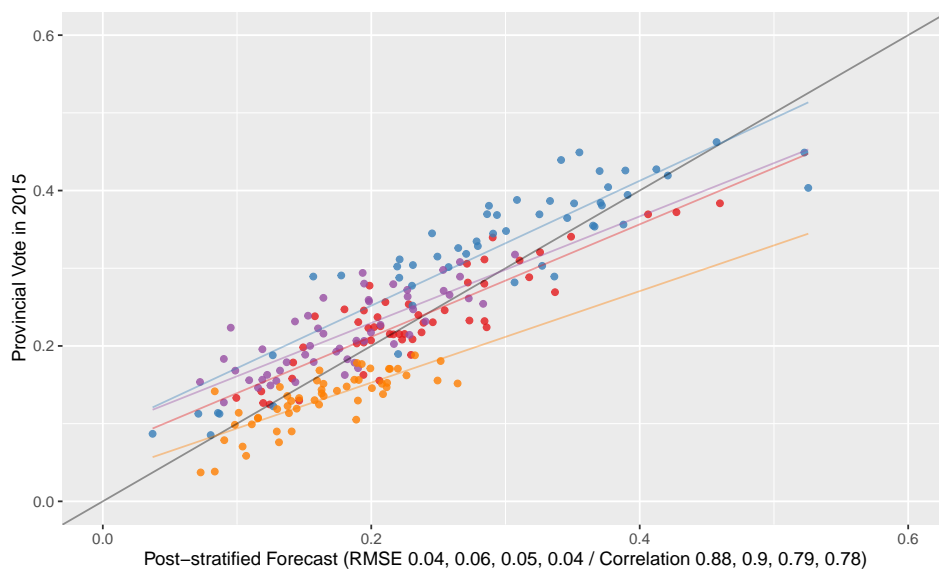


Figure 4: Scatterplot of post-stratified point estimates vs. outcomes and regression line. Statistics are listed in the usual party order. MSE is computed as the average squared difference between the mean prediction and the result over provinces. Legend: PSOE (red), PP (blue), Podemos+IU (purple), C's (orange).

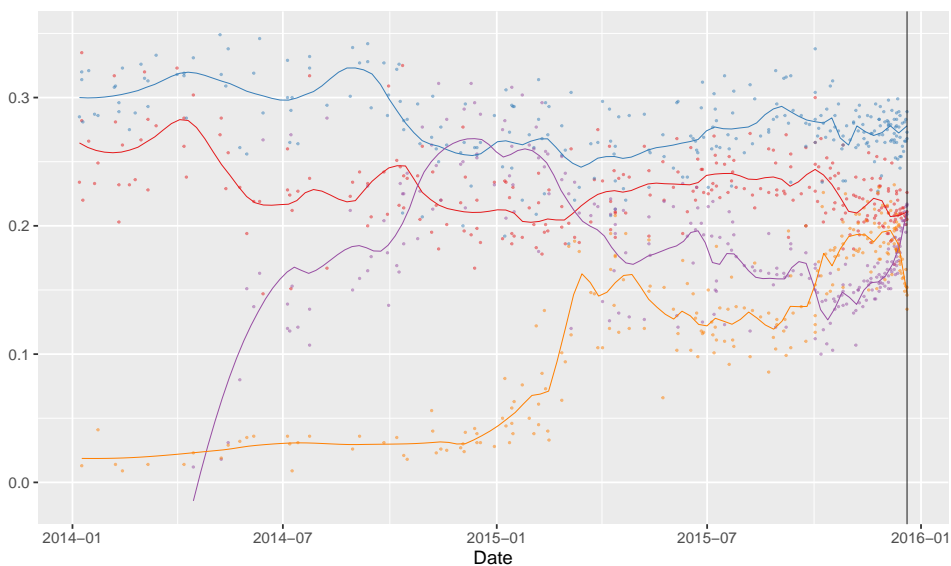


Figure 5: Polling before the general election of 2015, with LOESS smoother. Legend: Legend: PSOE (red), PP (blue), Podemos (purple), C's (orange). The election day is marked by the vertical solid line.

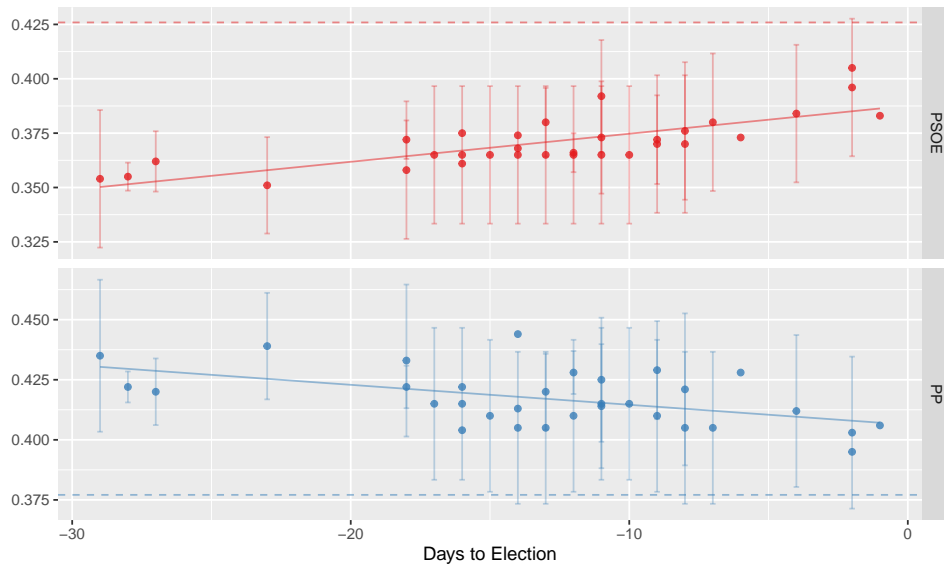


Figure 6: Polling before the general election of 2004. The solid line is a linear trend (OLS), the dashed line is the election result and the error bars correspond to the margin of error reported by the pollster.

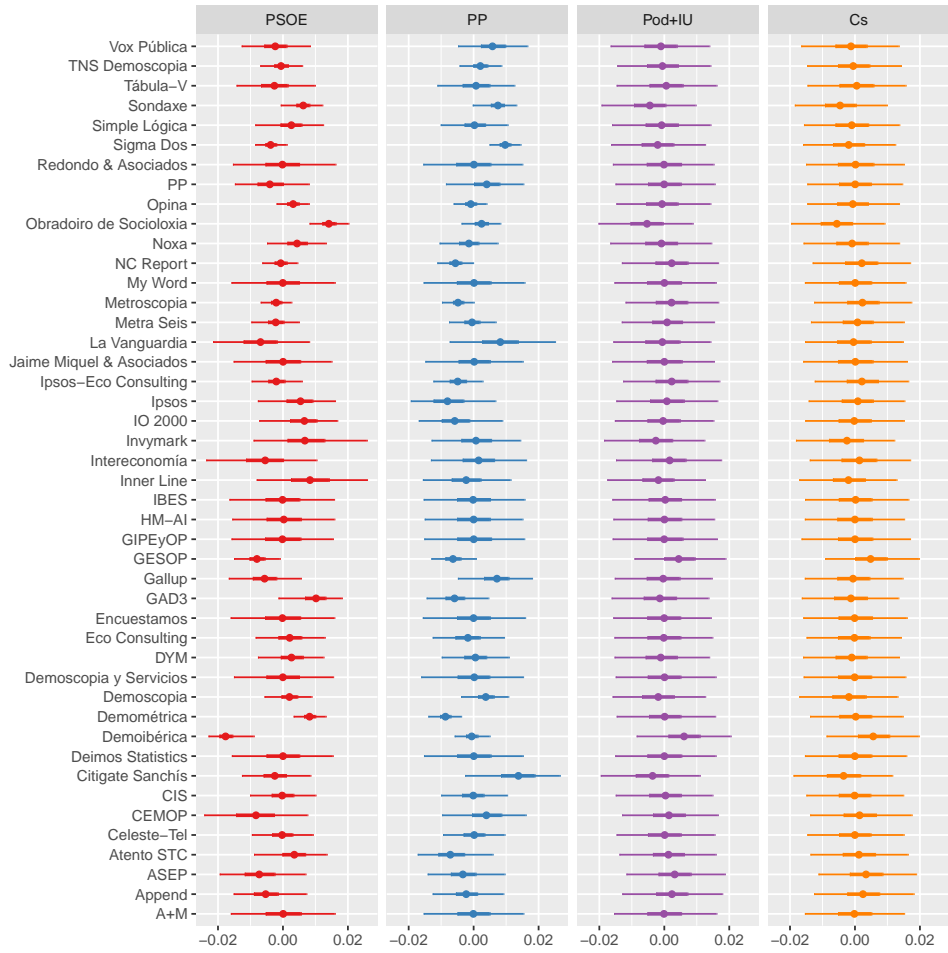


Figure 7: γ_j marginal distributions (pollster bias): median (point), 50 percent credibility interval (thick line) and 95 percent credibility interval (thin line). Positive values imply that the pollster is overestimating.

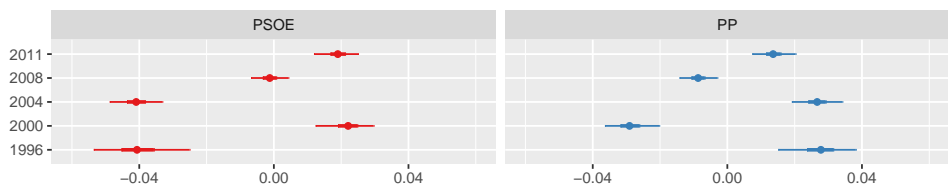


Figure 8: δ_t marginal distributions (election bias): median (point), 50 percent credibility interval (thick line) and 95 percent credibility interval (thin line). Positive values imply that the pollster is overestimating.

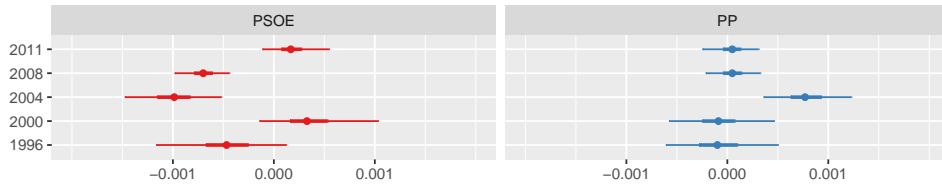


Figure 9: ϵ_t marginal distributions (election trend): median (point), 50 percent credibility interval (thick line) and 95 percent credibility interval (thin line). Positive values imply that polls are trending down.

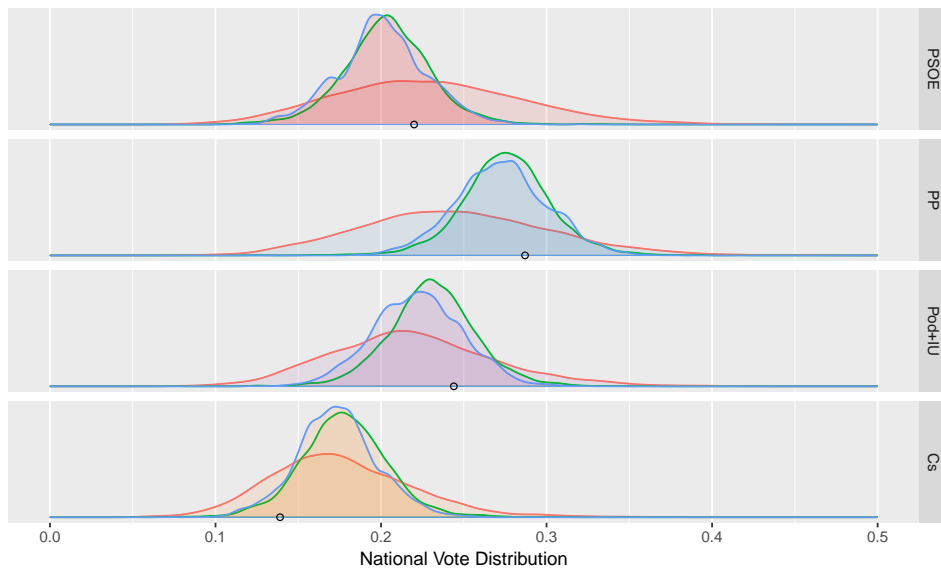


Figure 10: Predictive national vote distribution: fundamental model (red), polls model (green), synthesis (blue). The dots represent the election result.

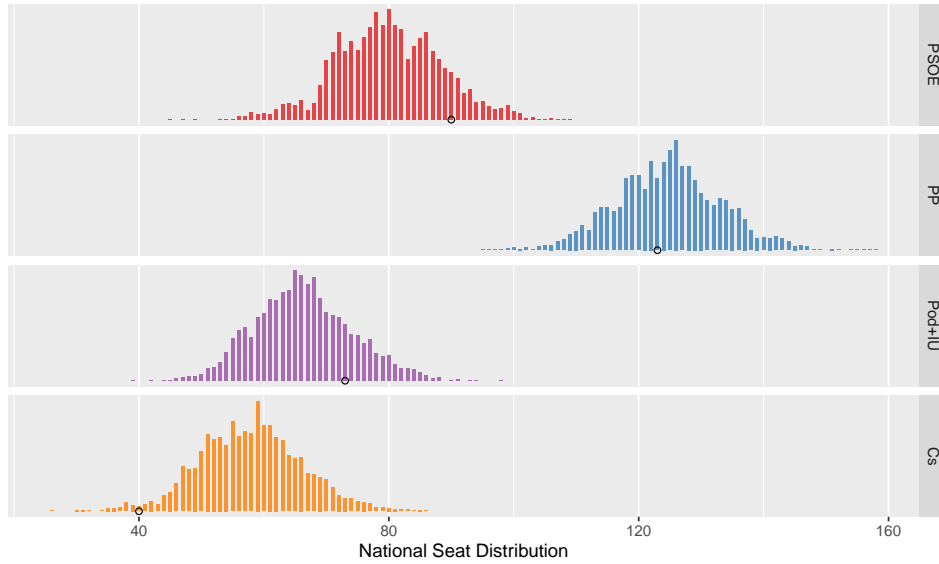


Figure 11: Predictive seat distribution and election result (black dot).

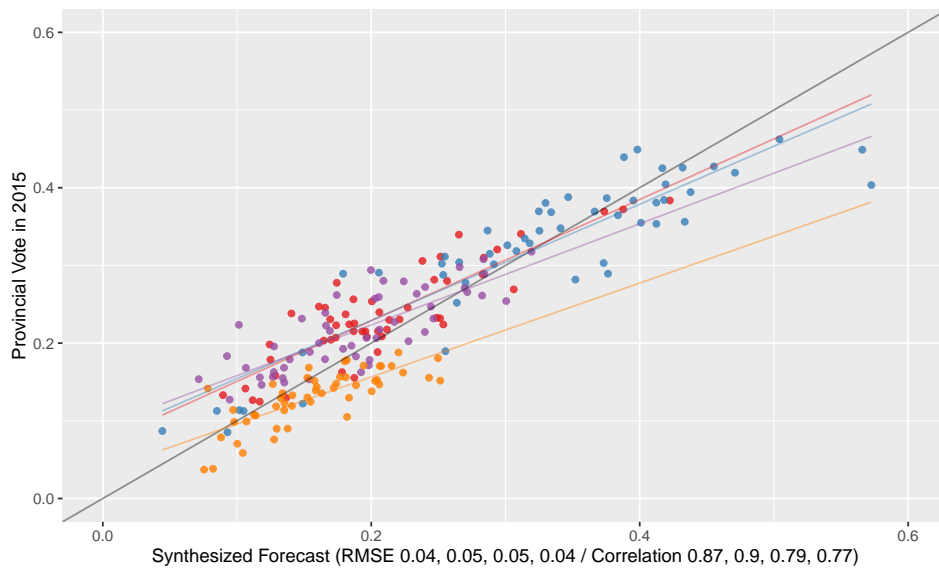


Figure 12: Scatterplot of point predictions vs. outcomes and regression line. Statistics are listed in the usual party order. MSE is computed as the average squared difference between the mean prediction and the result over provinces. Legend: Legend: PSOE (red), PP (blue), Podemos+IU (purple), C's (orange).

Appendix

A1. Data

The data on the outcome of the elections come from the Spanish Department of Home Affairs. For the fundamental model, we use the 2015 CIS (CIS is the Spanish National Center for Sociological Research) pre-electorals (CIS study number 3117). The study is openly available on <http://www.cis.es/> and includes 17452 respondents. Data was collected from October 27th to November 16th, 2015.

Factor	Code	Levels
Voting Intention	1	PSOE
	2	PP
	3	Podemos, En Comú Podem, En Marea, IU
	4	Ciudadanos
	5	Others
Province	1-52	INE Province Code
Municipality Population	1	less than 2000 inh.
	2	between 2000 and 10000 inh.
	3	more than 10000 inh.
Gender	1	Male
	2	Female
Age	1	18 to 35 y.o.
	2	36 to 55 y.o.
	3	more than 56 y.o.
Education	1	Primary or less
	2	Secondary
	3	Tertiary
Activity	1	Employed
	2	Unemployed
	3	Out of the labor force

Table A.1: Factors use in the fundamental model and their categories. These categorical features define 8424 distinct strata, or 162 distinct strata per province

To train the polls model, we use 157 polls published within 30 days of the 1996, 2000, 2004, 2008 and 2011 Congressional Elections. Furthermore, to generate predictions, we use 51 polls published within 30 days of the 2015 Congressional election.

A2. Hierarchical modeling notation

Hierarchical modeling notation is a convenient way of describing models that include a lot of categorical variables as regressors. Our hierarchical modeling notation follows the standard set by Gelman & Hill (2007) in *Data analysis using regression and multilevel/hierarchical models*.

Consider this brief explanation of the notation. Let $\{1, \dots, I\}$ index a set of observations and $\{1, \dots, J\}$ be the indices of the levels of a categorical factor. Then, the notation $j[i]$ refers to a map $\{1, \dots, I\} \mapsto \{1, \dots, J\}$ which links each observation to its respective factor level. For instance, if the factor is gender, male has index 1, female index 2 and observation 1 is female, then $j[1] = 2$.

If $\boldsymbol{\beta}$ is the vector of coefficients pertaining to the levels of some factor, we can use hierarchical modeling notation to retrieve components of that vector. In keeping with our example, $\beta_{j[1]} = \beta_2$ is the coefficient of the gender of observation 1, which is equivalently the coefficient of the female level of the gender factor.

We may express this equivalently using dummy variables, but hierarchical modeling notation tends to be more concise. For example, consider a simple regression model with one categorical factor. In dummy notation, we write $y_i = \beta_0 + \sum_j \beta_j x_{ij} + \epsilon_i$. In hierarchical modeling notation, we just write $y_i = \beta_0 + \beta_{j[i]} + \epsilon_i$.

A3. The fundamental model

In the following section, let $\boldsymbol{\theta}$ refer to the set of unknowns. We model the survey response counts $\mathbf{s}_n \in \mathbf{N}^L$ of a stratum n through a multinomial logit model:

$$\mathbf{s}_n | \boldsymbol{\theta} \sim \text{Multinomial}(\boldsymbol{\mu}_n(\boldsymbol{\theta})) \quad (13)$$

$$\boldsymbol{\mu}_n(\boldsymbol{\theta}) = \text{softmax} \left[\boldsymbol{\alpha} + \sum_k \boldsymbol{\beta}_{(k, j_k[n])} \right] \quad (14)$$

where k indexes the factors through which we define the strata (e.g. location, gender, education) and j_k indexes the possible levels of factor k (e.g. male, female, unreported for factor gender). Thus, $\boldsymbol{\beta}_{(k, j_k)}$ is the coefficient pertaining to factor k and level j_k , and $j_k[n]$ is the level of factor k that corresponds to stratum n . The operator softmax is the multivariate generalization of the logistic function.

We pool each factor's levels to a common prior:

$$\boldsymbol{\alpha} \sim \text{N}(\mathbf{0}, \mathbf{I}), \quad \boldsymbol{\beta}_{(k, j_k)} | \boldsymbol{\sigma}_k \sim \text{N}(0, (\text{diag } \boldsymbol{\sigma}_k)^2), \quad \boldsymbol{\sigma}_k \stackrel{\text{iid}}{\sim} \text{half-N}(1) \quad (15)$$

While the coefficients are identified due to the prior, we stick to the standard identifiability constraint of setting all coefficients of the residual party to zero. Then, coefficients may be interpreted as changing the response probabilities relative to the residual party.

A4. The polls model

In the following section, let $\boldsymbol{\psi}$ refer to the set of unknowns except for \mathbf{v}_t , i.e. the result of the t -th election. All vectors have dimension equal to the

number of parties minus 1. Dropping the last dimension is necessary to ensure that the distribution is non-degenerate. We assume that polls are generated by the following process:

$$\mathbf{p}_k | (\boldsymbol{\psi}, \mathbf{v}_{t[k]}) \sim \text{N}(\mathbf{v}_{t[k]} + \boldsymbol{\gamma}_{j[k]} + \boldsymbol{\delta}_{t[k]} + d_k \boldsymbol{\epsilon}_{t[k]}, \boldsymbol{\Sigma}_{j[k]}) \quad (16)$$

where $\boldsymbol{\gamma}_j$ is the time-invariant bias of pollster j , $\boldsymbol{\delta}_t$ is the pollster-invariant bias in election t , d_k corresponds to how many days before the election poll k was published, and $\boldsymbol{\epsilon}_t$ is the pollster-invariant strength of the trend in a given election. $d_k \boldsymbol{\epsilon}_{t[k]}$ vanishes as election day approaches, but $\boldsymbol{\delta}_t$ applies to all polls until the election.

We set the following priors on the random effects:

$$\boldsymbol{\gamma}_j | \boldsymbol{\psi} \sim \text{N}(\mathbf{0}, \boldsymbol{\Sigma}_\gamma), \quad \boldsymbol{\delta}_t | \boldsymbol{\psi} \sim \text{N}(\mathbf{0}, \boldsymbol{\Sigma}_\delta), \quad \boldsymbol{\epsilon}_t | \boldsymbol{\psi} \sim \text{N}(\mathbf{0}, \boldsymbol{\Sigma}_\epsilon) \quad (17)$$

If we integrate out the random effects, any two polls $k \neq k'$ have joint distribution characterized by the following mean and covariance functions:

$$\mathbf{m}(k) = \mathbf{v}_{t[k]} \quad (18)$$

$$\mathbf{C}(k, k') = 1_{(t[k]=t[k'])} (\boldsymbol{\Sigma}_\delta + d_k d_{k'} \boldsymbol{\Sigma}_\epsilon) + 1_{(j[k]=j[k'])} \boldsymbol{\Sigma}_\gamma + 1_{(k=k')} \boldsymbol{\Sigma}_{j[k]} \quad (19)$$

Accordingly, we may express the marginal polls model as a Gaussian process:

$$\mathbf{p}_k | (\mathbf{v}_{t[k]}, \boldsymbol{\Sigma}_\gamma, \boldsymbol{\Sigma}_\delta, \boldsymbol{\Sigma}_\epsilon, \boldsymbol{\Sigma}_{j[k]}) \sim \text{GP}(\mathbf{m}(k), \mathbf{C}(k, k')) \quad (20)$$

The model specified up to now defines a likelihood of polls given an upcoming

election result \mathbf{v}_{t^*} . We may complete the specification by adding a flat prior

$$p(\mathbf{v}_{t^*}) \propto 1 \tag{21}$$

thus allowing us to sample from the joint posterior of parameters and upcoming election result. Alternatively, we may use said likelihood to weight samples from some other prior over the upcoming election, e.g. our fundamental model.