Estimating Coalition Majorities during Political Campaigns based on Pre-election Polls

In multi-party systems, politicians, voters, and political pundits often speculate about potential coalition governments based on current poll results. Their interest particularly centers around the question whether a specific coalition has enough public support to form a parliamentary majority. In this research note, we present a Bayesian Dynamic Multinomial-Dirichlet model to estimate the probability that a coalition will find enough public support to form a parliamentary majority. An application to German federal elections from 1994-2017 and comparisons with alternative methods underscore the value of this approach.

Keywords: Pre-election Opinion Polls; Coalitions; Multi-party Elections; German Federal Elections

The majority of democracies are governed by coalitions of political parties (Hobolt and Karp, 2010). The prospect of a coalition government often leads news media and the general public to speculate about the options for any coalition majorities, during the electoral campaign. Three related questions often arise: 1) Will a specific coalition find enough support to form a parliamentary majority? 2) Does the current coalition government hold enough public support to defend its incumbency and 3) What alternative coalitions could replace the current government? For researchers interested in the dynamics of public opinion in multi-party elections, current methods provide relatively little insight on how to address these important questions.

In this research note, we present a measurement strategy on how to trace probabilities that coalitions will find enough public support to form a majority. We introduce a dynamic Multinomial-Dirichlet model specifically tailored to the situation of tracing public support for multiple parties from pre-election polls.[1] Our dynamic model takes newly arriving results from pre-election polls and updates the estimate of latent support for different parties at the current point in time. This strategy permits us to simulate up-to-date coalition majority probabilities for any combination of parties. By further considering institutional features, such as electoral thresholds, the resulting probabilities provide a meaningful snapshot of public opinion towards various governing options, over the entire electoral cycle.

An application of the model to the German federal elections illustrates the value of this approach. The majority-probabilities draw informative timelines, accurately reflecting the competition between the two major German political coalition camps (the conservative-liberal coalition and the social-green coalition) from 1994 to 2013. Analyzing additional coalition options in previous elections further accentuates desirable

[1] Previous models of pre-election polls focus on the U.S. with political competition between two parties (see e.g. Rigdon et al., 2009; Jackman, 2014; Lock and Gelman, 2010; Linzer, 2013). Only recently have researchers turned their attention to the status of pre-election polls in multi-party elections (see e.g. Walther, 2015; Jennings and Wlezien, 2016; Munzert et al., 2017; Hanretty et al., 2016; Stoetzer et al., 2018).

dynamics of the measurement. For a coalition that later manages to win the majority of seats in parliament, the probability of success steadily increases over time. Furthermore, we show that our measurement is more reliable than adding up poll-averages or estimating a series of 'one-party-at-a-time' dynamic linear models to approximate majorities.

This research note makes a central methodological contribution to the study of public opinion in multi-party elections. By discussing a simple, yet powerful, dynamic Multinomial-Dirichlet Model it introduces an alternative to the dynamic linear models commonly applied in stduides of the dynamics in trial heat polls (see e.g. Jackman, 2014; Pickup and Johnston, 2007; Walther, 2015). In contrast to dynamic linear models, our model considers the inherent constraint that the support for all parties should sum to one: a factor especially meaningful in multi-party elections and when estimating coalition majorities. Consequently, our new measurement will be particularly useful in applied researchers intrinsically interested in the state of public opinion towards coalition-governing options.

1 A Dynamic Multinomial-Dirichlet for Poll Data

Polls entail useful information for tracing a party's support in an ongoing election campaign. Research that relies on information from polls commonly models support for each party over time, often assuming normally distributed measurement error and a random walk for the latent support (see e.g. Jackman, 2014; Pickup and Johnston, 2007; Walther, 2015). We draw on the multinomial distribution to model the number of respondents intending to vote for a party at a given time. We further suppose that the support evolves continuously over time. In this, our model follows earlier applications of dynamic time-series models used to assess qualitative information (West and Harrison, 1997; Harvey and Fernandes, 1989; Smith, 1979).

Polls administered at time-point t interview a sample of size N_t and register the number of respondents y_{kt} that intend to vote for a specific party $k \in 1, \ldots, K$. The observed number of respondents for all parties is given by a vector $\boldsymbol{Y}_t = (y_{kt}, \ldots, y_{Kt})$, which is then related to the latent support share for all party, denoted as a vector $\boldsymbol{\pi}_t = (\pi_{kt}, \ldots, \pi_{Kt})$, using a multinomial distribution:

$$\boldsymbol{Y}_t \sim Multinomial(\boldsymbol{\pi}_t, N_t). \tag{1}$$

For estimating the latent support we work with the Dirichlet distribution, as it is a conjugate prior distribution to the Multinomial. Suppose that the posterior distribution at time $t-1$ is Dirichlet distributed, defined by the parameter vector $\boldsymbol{\alpha}_{t-1} = (\alpha_{kt-1}, \ldots, \alpha_{Kt-1})$:

$$\boldsymbol{\pi}_{t-1} \sim Dirichlet(\boldsymbol{\alpha}_{t-1}). \tag{2}$$

The Dirichlet distribution implies that the latent support shares are defined on the K-simplex, accounting for the inherent constraint that the latent support has to sum to one.[2] The prior for time point t is modeled using a power discount construction (Smith, 1979), which supposes that the prior is constructed using the following equation:

$$p(\boldsymbol{\pi}_t) \propto p(\boldsymbol{\pi}_{t-1})^{\delta}. \tag{3}$$

In this formulation, δ is a discount factor that describes how much information from $t-1$ is carried-over to time point t. It's bound to lay between 0 and 1.

The formulation of the dynamic Dirichlet-Multinomial Model leads to a closed form solution to update the latent support shares to the new arriving poll information:

[2]The parameters of the Dirichlet $\boldsymbol{\alpha}_{t-1}$ provide information about the expected support, as well as the uncertainty around the expectation. Formally, the expectation is given by $E[\pi_{kt}] = \frac{\alpha_{kt}}{\sum_1^K \alpha_{kt}}$ and the variance is equal to $Var[\pi_{kt}] = \frac{\alpha_{kt}(\alpha_{0t}-\alpha_{kt})}{\alpha_{0t}^2(\alpha_{0t}+1)}$, where $\alpha_{0t} = \sum_{k=1}^K \alpha_{kt}$

$$\boldsymbol{\pi}_t \sim Dirichlet(\boldsymbol{\alpha}_t), \tag{4}$$

where $\boldsymbol{\alpha}_t = \boldsymbol{Y}_t + \delta\boldsymbol{\alpha}_{t-1}$.[3] Thus, the estimate of the latent support is a weighted combination of the current poll results and the prior latent support. How strongly the prior support is weighted depends on the discount factor δ and on the number of respondents interviewed. With a discount factor close to one, the update generally gives strong weight to prior support of parties. In this case, polls require a comparatively large sample to shift the expectation. With lower values of the discount-factor, even smaller samples can shift the expectation, which also amounts to the fact that the support easily adapts to strong changes observed in public support.

For the updating, we need to specify prior beliefs for the beginning of the timeline and a sensible estimate of the discount-factor. We choose uninformative Dirichlet priors at the beginning of the timeline by setting all $\alpha_{k0} = 1$.[4] In this article, we further propose to estimate the discount factor from the observed data using maximum likelihood.[5] The resulting estimate of the discount factor can be used to construct the latent states over time, providing estimates for each party's support over the electoral period under study.

[3]This follows from the conjugate Bayesian analysis (see e.g. Harvey and Fernandes, 1989). For completeness Appendix 1 includes the posterior calculations.

[4]This implies that in the first period with a poll, the model will update almost perfectly to the likelihood of the first poll.

[5]The estimation of the discount factor requires a timeline of observed poll results until a specific timepoint T. The likelihood function for the observed timeline is the product of individual forecast distributions which are Dirichlet-Multinomial (see e.g. Harvey and Fernandes, 1989). The log-likelihood of this is

$$lnL(\delta|\boldsymbol{D}_T) = \sum_{t=1}^{T}\left(ln(\Gamma(\alpha_{0t})) - ln(\Gamma(N_t\alpha_{0t})) + \sum_{k=1}^{K} ln(\Gamma(y_{kt} + \alpha_{kt}) - ln(\Gamma\alpha_{kt})\right), \tag{5}$$

where $\alpha_{t0} = \sum_{k=1}^{K}\alpha_{tk}$ and $\boldsymbol{D}_T$ describes all the poll results available. The log-likelihood function is maximized with respect to the only parameter δ using numerical algorithms. Please note that δ features in the Likelihood implicitly when calculating the latent states via the updating equation $\boldsymbol{\alpha}_t = \boldsymbol{Y}_t + \delta\boldsymbol{\alpha}_{t-1}$. The maximization is carried out using R's optim function.

2 Majority-generating Function

The central aim of this article is to estimate the probability that a specific coalition holds enough public support to form a parliamentary majority. Within the Bayesian framework, the task of estimating this property by using simulation is straightforward. The model generates posterior distributions of the latent support from which we can take samples to approximate the majority probabilities. For each of the draws, we check if a particular coalition of parties could secure a majority in parliament.

An important step is to consider how the electoral system translates vote shares into seat shares (Kedar et al., 2016), and how these seats can generate a majority in parliament. In our application, we define a majority-generating function that is tailored to proportional electoral systems with small dis-proportionality between seats and votes and an electoral threshold. This yields a two-step process by which we identify the coalition majority in any given draw. First, we check if one of the parties does not surpass the electoral hurdle. If a party does not pass the hurdle, it does not secure seats in parliament and is excluded from the calculation of the seat share majority. This step implies that parties do not necessarily require more than 50% of latent support, as some parties will not pass the hurdle. In the next step, the coalition requires a majority among parties in parliament. The sum of the coalition parties' seats needs to be larger than 50% of all seats allocated. The overall average of the independent draws approximates majority probabilities, which can be interpreted as a snapshot of the probability that a coalition would have attained a majority of seats, had the election been held at that point in time.

The majority-generating function can be tailored to a variety of electoral systems. For systems where seats are disproportional to vote shares, the model can be adjusted by including an additional step which translates the public support into the seat distribution. In systems in which the district results are of central importance, researchers could approximate the district results using unit-swing models (see e.g. Jackman, 2014).

This would allow district winners to be calculated for each draw from the posterior distribution, generating a seat-share distribution that takes the district expectations into account. An alternative approach would be to account for the votes-seats ratio in past elections. Linzer (2012) describes a general approach to estimating party-specific swing ratios that can be used to translate popular support into seat shares. This function can be estimated, based on results from the last election, and then used to approximate the seat shares from our model.

3 Application to the German Federal Elections

In order to gain insights into the dynamics of the mode, we apply it to pooling data on party support for the German Federal Elections from 1994-2017, as an example of a relatively stable multi-party system, to gain insights into the dynamics of the model. With its long tradition of government coalitions and almost no dis-proportionality in its electoral system, the German Federal Elections provide a particularly informative case for our model. We combine information from several German pollsters.[6] The analysis concentrates on the five major parties (The Union faction (CDU/CSU), The Social Democratic Party (SPD), Free Democratic Party (FDP), Alliance 90/The Greens (Greens) and The Left (Die Linke from 2007 on, formally named Party of Democratic Socialism (PDS)), collapsing support for other parties into a sixth category. For the 2017 election, we also consider support for the Alternative for Germany (AfD). We aggregate vote intentions from different pollsters within the same week by summing up vote intentions per party across all pollsters.

The Multinomial-Dirichlet Model depicts the evolution of support for each party in the year leading up to each election. Figure 1 presents the evolution of support (mean

[6] German's major poll agencies are Allensbach, Emnid, Forsa, Forschungsgruppe Wahlen, GMS Dr. Jung GmbH, Infratest dimap, and Institut für neue soziale Antworten (INSA). For the elections 1994-2009, we use the collected polls by (Schnell and Noack, 2014). For polls published after 2009, we employ data from the online-portal wahlrecht.de

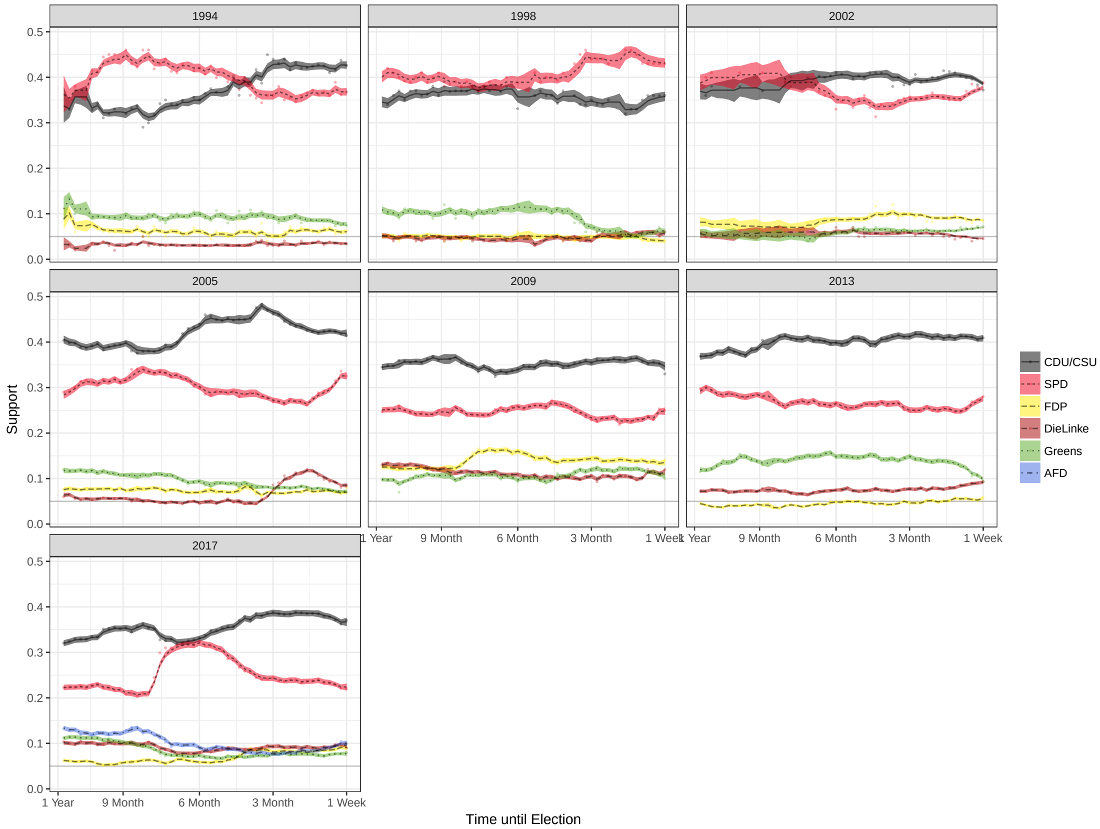

Figure 1: Latent Support for the last year of the German Federal Elections 1994-2017

and 95% credible interval) for the different elections and parties, as well as the poll results. The estimated discount-parameters are similar for the five recent elections (values between 0.3 and 0.4). This implies that the timeline closely follows the observed poll results, since with a small discount-factors the latent support is treated as almost time-independent.[7] It also means that the latent support adapts easily to jumps, as for example, experienced by Die Linke three months prior to the 2005 election. In addition, the figure shows that especially the two large parties exhibit strong volatility in the year leading to an election. For example, in 1994 the SPD started the election year neck and neck with the CDU/CSU at around 35%, then experienced an incredible increase

[7] All available polls since the previous election are used to estimate the discount factor. For three earlier elections, which have fewer polls and more weeks with missing data, the values are slightly higher (between 0.59 and 0.74). See Appendix 4 for the discount-factor estimates.

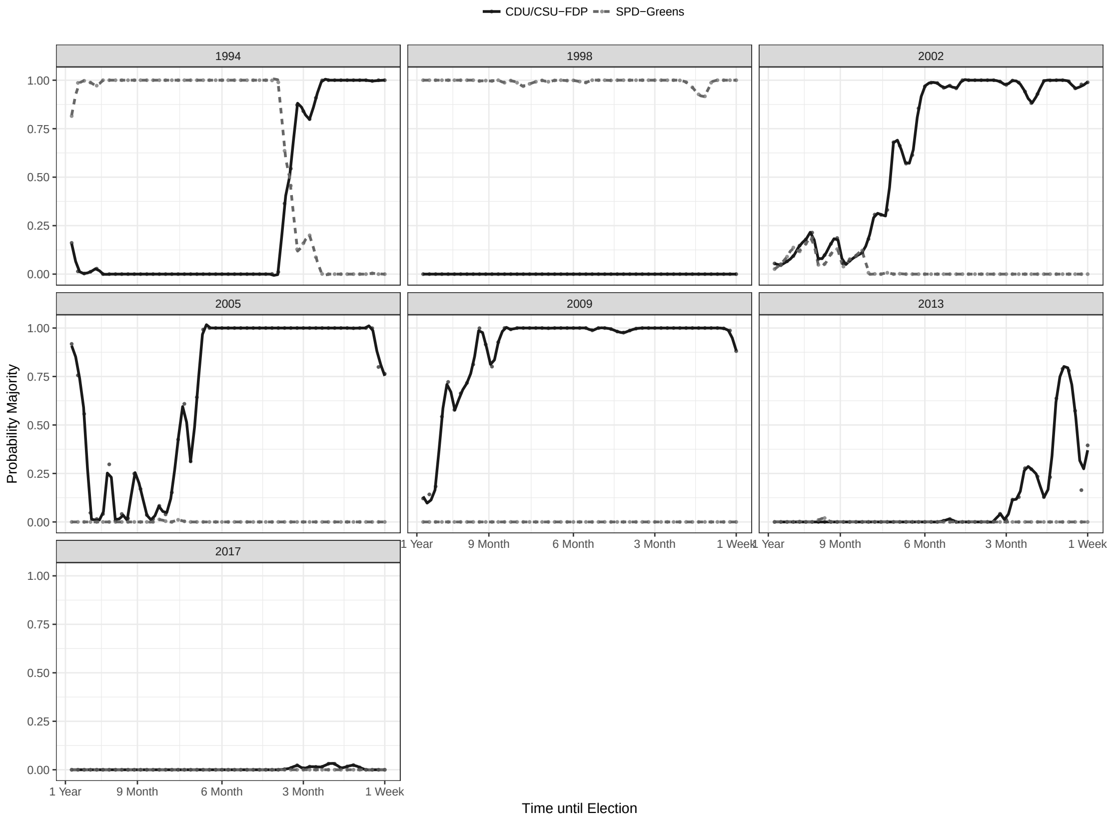

Figure 2: Estimated coalition majorities for CDU-FDP and SPD-Greens coalition for the last year before German Federal Elections 1994-2017

to almost 45%, only to fall back to their original level later on. One can also clearly observe the erratic increase in support for the SPD during the last election campaign, after the SPD announced that their chancellor candidate was Martin Schulz.

The coalition majorities present an interesting overview of coalition dynamics during electoral campaigns. Two coalitions: the CDU/CSU with the FDP, and the coalition between SPD and the Greens, as ideologically-opposed coalition camps, are of special interest. Figure 2 depicts the probability of gaining a majority for these two coalitions. At the beginning of the year before the 1994 election, pre-election polls saw the incumbent CDU/CSU-FDP coalition with zero chance of keeping their majority. Within the final nine months, however, the CDU/CSU overtook the SPD and managed to secure a clear victory with 41,4 % of the votes. This is reflected in the coalition majorities. Four months prior to the election, it became clear that the incumbent coalition would indeed find enough support, while the chances of the alternative SPD-Greens coalition ousting the Kohl-lead coalition government diminished. The CDU/CSU success did not continue in the 1998 election, in which it became clear during the electoral year that an SPD-Greens coalition would replace the CDU/CSU-FDP coalition government. The first SPD-Greens coalition under the leadership of Gerhard Schröder faced a contested election campaign in the follow-up 2002 election. Only within the last month, did the SPD manage to approximate CDU/CSU support, resulting in a final outcome in which they won with the tiny margin of 0.01% permitting them to continue their coalition government. In this case, the coalition majorities in our figure fail to show a last minute catch-up. This can be attributed to the fact that only the polls in the last week identified the SPD in front of the CDU/CSU and thus our model does not reflect the sudden jump. In the 2005 elections, the CDU/CSU rose to comfortable heights, polling around 45% (coupled with high chances of a CDU/CSU-FDP coalition) three months prior to the election, only to see a subsequent decline in support. Accordingly, the majority probability of success for the CDU/CSU-FDP dropped significantly within the last couple of weeks. Though they won the election by an unexpectedly low margin, the final outcome found no majority for either of the two camps, resulting in the first Merkel-lead Grand Coalition government. In light of this coalition, public support for small parties rose in the 2009 election cycle, giving way to clear majorities. Here again, the timeline acts as a valid forerunner in predicting a majority for the CDU/CSU-FDP coalition. For the 2013 election, the estimates underscore the small chances for either of the two camps to form a majority . Only within the last weeks, did the hopes of the CDU/CSU- FDP swell, though were unmet as the FDP, in the end, did not pass the electoral threshold. The result of the post-electoral bargaining was again a Grand Coalition with Merkel as chancellor.

In 2017, the likely entry of the AFD changed the coalition majorities from two-party coalitions, that had dominated the electoral landscape, towards three-party coalitions. Figure 2 highlights the overall low chances for the traditional constellations. Instead, the only two options with a high probability were the Jamaica coalition, between CDU/CSU-FDP and Greens, and a continuation of the Grand Coalition (See Figure 1 in the Appendix). The Jamaica coalition suffered a sharp drop in popularity during the times of the "Schulz-Effect", but over the campaign it proved to be a valid majority option. Interesting, is the fact that the probabilities for a left-alliance, between the SPD-Die Linke-Greens, briefly spiked at the height of the Schulz-Effect. This spurred media interest and led to public statements regarding this option by the parties involved.

4 Validation and Comparison to other Measures

In order to validate our measurement of majority probabilities, we consider the development over time in relation to the final size of the coalition in parliament. Our measurement should indicate high chances for those coalitions that actually win a majority of seats on election day. For coalitions that gained slightly above or slightly below 50% of the seats, the measurement should indicate high levels of uncertainty about the respective coalition winning a majority. To study this, we trace the coalition majority probability over the last year of an election campaign, looking at ideologically-connected two and three-party coalitions.[8] Figure 3 presents the results, fitting a regression model with a second order polynomial on the calculated majority probability, depending on the final seat share of the coalition.[9]

The results confirm the usefulness of our measurement in assessing governing majori-

[8] We choose coalitions that have formed at least once in any of the federal state parliaments. Two-party coalitions included in the analysis are CDU/CSU-SPD, CDU/CSU-FDP, CDU/CSU-Greens and SPD-Greens. Three-party coalitions are CDU/CSU-FDP-Greens, SPD-Greens-Die Linke, SPD-Greens-FDP.

[9] Appendix 2 shows that the relationship plotted in the Figure does hold when controlling for coalition-election fixed effects.

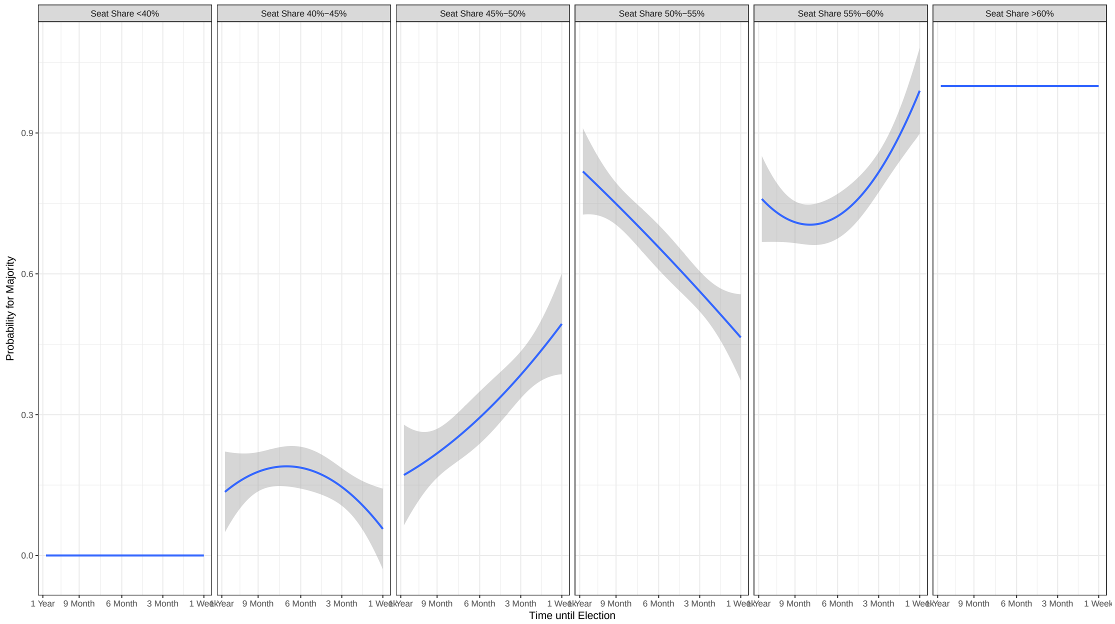

Figure 3: Evolution of coalition majority measurement within the last year for different final seat share margins

ties over time. First, for coalitions that gain a clear majority in parliament (above 60% of seats), the measurement always assigns a probability of one to the event. Conversely, for coalitions that do not come close to a majority (below 40% of seats), the probability is always zero. Second, for coalitions which secure a safe majority (above 55% and below 60% of seats), the majority probability crystallizes out over time, with increasing chances at the end of the electoral campaign. The other way round, for a coalition that in the end clearly fails to attain a majority (between 40% and 45%) the chances will decrease over time. Third, for coalitions that are closer to the margin, just below (45%-50%) and just above (50%-55%), the probabilities converge to a 50 % chance. Worth noting, is the fact that for the last-mentioned coalitions, the averages at the beginning of the year seem to be more predictive in explaining the final outcome. Thus the analysis shows that the proposed method is a valid instrument for tracing coalition majorities.

Furthermore, we compared our model to two alternative measurements of majority-probabilities. The first alternative is a measurement that can directly be derived from the polls (See Appendix 5.1). Analysis here reveals that our proposed model is a more accurate means of explaining the final results, as well as describing dynamics during the German campaigns. The second alternative measure is based on dynamic linear models. In Appendix 5.1, we discuss the standard random walk poll results (see e.g. Walther, 2015). We argue that the dynamic Multinomial-Dirichlet model should be the preferred method of estimating coalition majorities. First, we show that the party-at-a-time dynamic Linear models do not include a sum constraint on the public support, which can have consequences in the calculation of coalition majorities (especially in times of sparse polling data). Second, we find that in the application to the German Federal Elections, the dynamic Multinomial-Dirichlet model is more accurate in explaining the final outcomes.

5 Discussion

This paper describes a method of tracing public support for coalitions of political parties based on pre-election polls. It outlines a Bayesian dynamic Multinomial-Dirichlet model for such polls, by means of which, the probability that a coalition will find majority support can be estimated. An application to the German Federal Elections of 1994-2017 illustrates that the method is useful in tracing coalition options during electoral campaigns.

The majority probabilities provide a snapshot of public opinion regarding the coalition majorities, at a given time prior to the election. Although this is likely to be the information on which voters and parties base future political action, and therefore of intrinsic interest in a study of public opinion, the estimates should not be confused with a forecast of coalition governments. The dynamic process of the Multinomial-Dirichlet model can generally be used to foresee future support for coalition majorities. The expected latent support stays stable, e.g. if a party is expected to hold 40% support today, it is expected to hold 40% in the future, but the variance around this expectation increases heavily. This variance restricts the usefulness of projections regarding coalition majority probabilities. For this alternative purpose, future research could use the model to consider how the predictive power of polls evolves over time (Hanretty et al., 2016; Jennings and Wlezien, 2016), further integrating structural covariates in order to more accurately forecast the election result (Linzer, 2013; Lewis-Beck and Dassonneville, 2015; Munzert et al., 2017; Stoetzer et al., 2018).

References

Chris Hanretty, Ben Lauderdale, and Nick Vivyan. *Electoral Studies*, 41:239–243, 2016.

A. C. Harvey and C. Fernandes. Times Series Models for Count Data or Qualitative Observations. *Journal of Business & Economic Statistics*, 7(4):316–337, 1989.

Sara B Hobolt and Jeffrey A Karp. Voters and coalition governments. *Electoral Studies*, 29(3):299–307, sep 2010.

Simon Jackman. The Predictive Power of Uniform Swing. *PS: Political Science & Politics*, 47(02):317–321, 2014.

Will Jennings and Christopher Wlezien. The Timeline of Elections: A Comparative Perspective. *American Journal of Political Science*, 60(1):219–233, 2016.

Orit Kedar, Liran Harsgor, and Raz A. Sheinerman. Are Voters Equal under Proportional Representation? *American Journal of Political Science*, 60(3):676–691, 2016.

Michael S. Lewis-Beck and Ruth Dassonneville. Comparative election forecasting: Further insights from synthetic models. *Electoral Studies*, 39:275–283, 2015.

Drew A. Linzer. The relationship between seats and votes in multiparty systems. *Political Analysis*, 20(3):400–416, 2012.

Drew a. Linzer. Dynamic Bayesian Forecasting of Presidential Elections in the States. *Journal of the American Statistical Association*, 108(501):124–134, 2013.

Kari Lock and Andrew Gelman. Bayesian combination of state polls and election forecasts. *Political Analysis*, 18(3):337–348, 2010.

Simon Munzert, Lukas F. Stoetzer, Thomas Gschwend, Marcel Neunhoeffer, Sebastian Sternberg, and Steffen Zittlau. Zweitstimme.org. Ein strukturell-dynamisches Vorhersagemodell für Bundestagswahlen. *Politische Vierteljahresschrift*, 2017.

Mark Pickup and Richard Johnston. Campaign trial heats as electoral information: Evidence from the 2004 and 2006 Canadian federal elections. *Electoral Studies*, 26(2): 460–476, 2007.

S. E. Rigdon, S. H. Jacobson, W. K. Tam Cho, E. C. Sewell, and C. J. Rigdon. A Bayesian Prediction Model for the U.S. Presidential Election. *American Politics Research*, 37: 700–724, 2009.

Rainer Schnell and Marcel Noack. The Accuracy of Pre-Election Polling of German General Elections. *Methods, Data, Analyses*, 8(1):5–24, 2014.

J.Q. Smith. A Generalization of the Bayesian Steady Forecasting Model. *Journal of the Royal Statistical Society . Series B (Methodological)*, 41(3):375–387, 1979.

Lukas F. Stoetzer, Thomas Gschwend, Marcel Neunhoeffer, Simon Munzert, and Sebastian Sternberg. Forecasting Elections in Multi-Party Systems: A Bayesian Approach Combining Polls and Fundamentals. *Political Analysis*, First View, 2018.

Daniel Walther. Picking the winner(s): Forecasting elections in multiparty systems. *Electoral Studies*, 40:1–13, 2015.

Mike West and P Jeff Harrison. *Bayesian Forecasting and Dynamic Models, Bayesian Forecasting and Dynamic Models.* 1997.