

Forecasting runoff elections using candidate evaluations from first round exit polls

Peter Selb*, Michael Herrmann, Simon Munzert, Thomas Schübel, Susumu Shikano

Department of Politics and Public Administration, University of Konstanz, Germany

A B S T R A C T

Keywords:

Runoff elections
Two-round elections
Exit polls
Feeling thermometers

We draw attention to a simple yet underappreciated way of forecasting the outcomes of elections involving two rounds of voting: surveying the voters' candidate evaluations in first round exit polls, poststratifying the sample proportions of reported votes to official first round election returns, and redistributing the votes for eliminated competitors according to their supporters' lower-order preferences among the viable alternatives in round two. We argue that the approach is likely to outperform standard pre-election surveys, due to its better coverage and reduced measurement error, and the possibility of correcting for sample selection. We set out the practical details of the method and demonstrate its usefulness by employing a recent German mayoral election as an empirical case. Thirteen candidates were competing in the first round, while there were six candidates in the decisive second round. The runoff result was forecast two weeks in advance with an average absolute error of less than one percentage point.

1. Introduction

Election forecasting has become a thriving industry in many countries (see, for example, the 2008, 2010, and 2012 special issues of this journal). It provides party elites, candidates and voters with information about the likely outcome of an election, allowing them to adjust their campaign strategies and voting behavior accordingly. Election forecasts also provide the major ingredients of horse race style campaign coverage by the news media. Essentially, three main methodological approaches have proven successful over the years. The oldest and most widely used method is vote intention surveys of the eligible population, which are frequently conducted prior to an election by polling organizations such as Gallup, IPSOS, and IFOP. A second approach uses regression analysis to predict future outcomes based on past outcomes and other correlates of government popularity, such as the state of the economy,

which are known ahead of an election (Erikson & Wlezien, 2008; Lewis-Beck & Tien, 2008). A third and more recent approach relies on the use of exchange prices from political stock markets, in which a number of individuals trade bets on candidates' success, to predict the outcome of the election (Wolfers & Zitzewitz, 2004).

We contend that there is a fourth possible approach to election forecasting – and one that should be preferred to pre-election surveys – if an election involves two rounds of voting. In a two-round election, each voter casts one vote and the candidate with the majority of votes is elected. If no candidate wins a majority (as often occurs), a runoff election is held (often, but not necessarily, between the top two finishers) a few weeks after the first round of voting. Two-round (henceforth, runoff) elections are used around the world for the election of presidents in 82 countries,¹

* Corresponding author.

E-mail address: peter.selb@uni.kn (P. Selb).

¹ Afghanistan, Algeria, Angola, Argentina, Armenia, Austria, Azerbaijan, Belarus, Benin, Bolivia, Brazil, Bulgaria, Burkina Faso, Cape Verde, Central African Republic, Chad, Chile, Colombia, Costa Rica, Croatia, Cyprus, Côte d'Ivoire, Djibouti, Dominican Republic, East Timor, Ecuador, Egypt,

legislative assemblies in 19 countries,² and mayors and local councils in various places.³

One seemingly straightforward method of forecasting runoff elections from first round outcomes is to transfer the votes of all withdrawn candidates to the remaining competitor who is perceived to be closest to the original vote, in ideological terms. This approach is based on the heroic assumption that preference orderings are constant among the voters of eliminated candidates, and has often failed utterly in the past (see Jérôme, Jérôme, & Lewis-Beck, 1999). A superior alternative is to survey the voting intentions in hypothetical runoff scenarios. However, many survey respondents seem to be overwhelmed by such hypothetical items, and sometimes survey researchers even neglect to include the relevant runoff scenarios in the survey, as was the case, for example, with the unexpected runoff between Chirac and Le Pen in the 2002 French Presidential Elections (see Miguet, 2002). Our approach to forecasting the second round (or runoff) election works as follows:

- (1) conduct exit polls among first round voters;
- (2) ask for evaluations of all candidates on the ballot (instead of voting intentions in hypothetical second round scenarios), as well as first round vote choices;
- (3) poststratify the sample distribution of reported voting behavior to actual first round election returns, in order to account for potential selectivity; and
- (4) forecast the runoff by redistributing the votes for eliminated competitors according to their supporters' lower-order preferences among the viable candidates.

We believe that our method significantly outperforms standard opinion polls in forecasting these elections, and the following sections explain why.⁴ We then report on an application that testifies to the usefulness of our method in forecasting a mayoral election which involved thirteen candidates in the first round, and six in the decisive second round.

El Salvador, Finland, France, Gabon, Gambia, Georgia, Ghana, Guatemala, Guinea, Guinea-Bissau, Haiti, Indonesia, Iran, Kazakhstan, Kenya, Kyrgyzstan, Liberia, Lithuania, Macedonia, Madagascar, Maldives, Mali, Mauritania, Mongolia, Montenegro, Mozambique, Namibia, Nicaragua, Niger, Nigeria, Palau, Peru, Poland, Portugal, Republic of the Congo, Romania, Russian Federation, Sao Tome and Principe, Senegal, Serbia, Seychelles, Sierra Leone, Slovakia, Slovenia, Sudan, Tajikistan, Tanzania, Togo, Turkmenistan, Uganda, Ukraine, Uruguay, Uzbekistan, Yemen, Zimbabwe (Institute for Democracy and Electoral Assistance (IDEA) unified database, at: www.idea.int/uid/, retrieved September 2012).

² Bahrain, Belarus, Central African Republic, Comoros, Cuba, France, Gabon, Haiti, Iran, Kiribati, People's Republic of Korea, Mali, Mauritania, Montserrat, Republic of the Congo, Togo, Turkmenistan, Uzbekistan, Vietnam (IDEA database, www.idea.int/uid/).

³ For instance, Brazil (Sakurai & Menezes-Filho, 2008); Germany, Austria, Greece, Italy, Netherlands (Magre & Bertrana, 2007); or Romania (Stroschein, 2011).

⁴ Stock markets and statistical modeling could be used to forecast runoff elections as well, but neither approach seems superior a priori. For instance, stock markets require relatively large numbers of traders in order to capitalize on accuracy-through-aggregation phenomena, which limits their application to high salience elections. On the other hand, model-based approaches have been criticized for being atheoretical, tautological and often overfitted (see van der Eijk, 2005).

2. Exit polls versus pre-election surveys

Exit polls are traditionally used for projection purposes, but they have also been used to improve our understanding of voter choice against the backdrop of factors such as partisanship (Swyngedouw, 2001) or attention to election news (Druckman & Parkin, 2005), or even to detect electoral fraud (Scheuren & Alvey, 2008).⁵ As a prediction tool, exit polls are of public and media interest, rather than scientific interest, since they only anticipate the actual election result by a few hours, reducing their predictive value to baubler and pastime for election-day news coverage. In runoff elections, however, they also offer an opportunity to generate predictions for future rounds of voting, thereby outperforming standard pre-election surveys in terms of accuracy. This is due not only to timing, but also to the fact that exit polls are affected by sampling/nonresponse and measurement errors to a lesser extent than pre-election surveys. Potential sources of survey error for exit polls and pre-election surveys are given in Fig. 1 and discussed in what follows.

2.1. Coverage error

The target population for any survey-based election forecast is the actual voters in the upcoming runoff election, which is, of course, unobservable at the time when the survey is being conducted. In order to hit such a moving target, exit polls utilize the actual voters in the first round election as the sampling frame. This sampling frame covers the target population to the extent that those who vote in the runoff are those who turn out in round one of the election. It seems plausible to assume that this will be a good match, given what we know about the relative weight of stable personal versus situational determinants of voter participation (see Plutzer, 2002). In contrast, common pre-election surveys utilize different sampling frames – typically a list of telephone numbers, or at best a list of persons who are eligible or registered to vote – including many individuals who are unlikely to vote in either round one or the runoff. Thus, both approaches face a potential coverage error, but the problem seems to be much graver for standard pre-election polls.⁶

2.2. Nonresponse error

Getting a representative sample of actual voters from a given sampling frame seems more difficult in the pre-election setting, due to the greater likelihood of non-response error. First, polling agencies typically use random digit dialing to contact respondents over wire phone

⁵ On the methodology of exit polls, see Levy (1983).

⁶ Coverage error in exit polls might also occur as a result of the fact that one voter group – the group of absentee voters – is, by construction, excluded from the poll. While this might pose a problem for forecasts of the first round, it is hard to see why this group should exhibit distinct preference patterns, once the distribution of first round votes is taken into account through poststratification. Therefore, this should not have a negative effect on our proposed forecasting procedure.

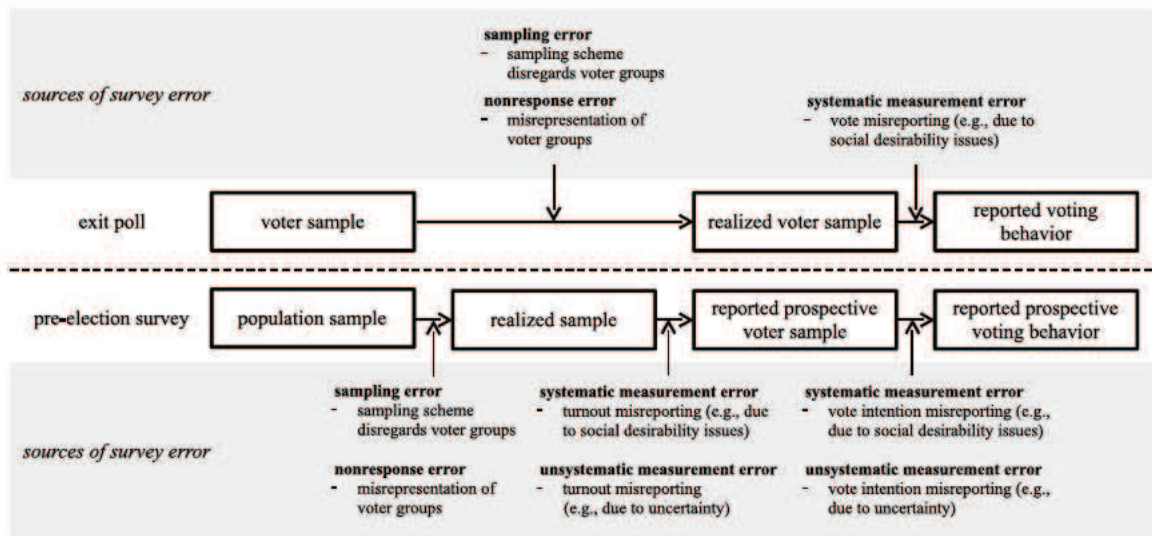


Fig. 1. Sources of survey error in exit polls and pre-election surveys.

connections. As a result, some segments of the target population will be missing from the sample due to unavailability over the phone, especially younger people with cell-phone-only connections (Keeter, 2006; Mokrzycki, Keeter, & Kennedy, 2009). Note that this problem remains essentially the same for this format, regardless of the poll's timing (before or after the first round). Second, certain voter groups (e.g., extremist voters) might disqualify themselves from the sample by refusing to participate (unit nonresponse). With exit polls, on the other hand, reaching the voters is straightforward. Unit nonresponse is a problem in exit polls, too, but it can be alleviated by post-stratifying the sample vote shares to first round results (see below).

2.3. Measurement error

In addition to nonresponse error, pre-election polls also face the problem of identifying actual voters. Sorting voters from non-voters⁷ a priori (but also a posteriori in a non-exit poll setting) introduces measurement error (Levy, 1983). It has often been observed that a non-negligible share of actual non-voters in the sample misreport their voting behavior, contributing to higher levels of self-reported turnout than actual turnout (e.g. Belli, Traugott, Young, & McGonagle, 1999; Selb & Munzert, 2013; Presser, 1990). If the reported voting of actual non-voters (i.e., turnout misreports) differs systematically from that of the actual voters in the sample, the forecasts from pre-election surveys will be biased.

In addition to turnout misreports, vote misreports may also lead to biases. Both pre-election surveys and exit polls

face problems of misreporting due to social desirability issues (systematic measurement error). However, the hypothetical nature of vote intention (as opposed to reported vote) surveys introduces an additional source of uncertainty that is absent from exit polls: unsystematic measurement error, which increases the error variance in the forecast. Furthermore, there is less reason to expect systematic measurement error in self-administered exit polls (which is the traditional mode; see Scheuren & Alvey, 2008) than in pre-election polls conducted via telephone: as Bishop and Fisher (1995) show, the self-administered mode reduces refusal rates and leads to more honest self-reports of voting than a face-to-face setting involving an interviewer.

3. Candidate evaluations versus vote intentions

Exit polls on the occasion of first round elections in the US and France often employ counterfactual questions like "what if candidate A had not been on the ballot" (Wattenberg, 1995), or offer hypothetical party choice combinations for the second round (Jérôme et al., 1999). Instead, we suggest using common feeling thermometer questions to ask the respondents about their likes and dislikes of individual candidates (see Weisberg & Rusk, 1970). Feeling thermometers are closely related to vote choices (Druckman & Parkin, 2005), and have frequently been used as substitutes for the vote when the latter was not available (see, e.g., Abramson, Aldrich, Paolino, & Rhode, 1995). We think that feeling thermometer questions offer some major advantages over runoff vote intentions.

First, surveying vote intentions in different scenarios may easily lead to inconsistent or *intransitive* preferences induced by measurement error (Brady & Ansolabehere, 1989). Such data are very difficult to interpret and analyze. Second, collecting thermometer scores for all candidates who have the potential to enter the runoff is relatively easy, even with 10 or more first round contenders. In contrast,

⁷ This is done either explicitly, by asking the respondent whether (or at what probability) she is planning to cast a vote at the upcoming election, or implicitly, by adding an "I will not vote" option in the hypothetical vote choice question.

surveying all of the hypothetical runoff scenarios in such situations quickly becomes infeasible. For example, with 10 first round candidates where only the two strongest are admitted to the runoff, there are $\binom{10}{2} = 45$ possible runoff scenarios. In the majority-plurality runoff case discussed in our empirical application, the number of possible runoffs increases to $\binom{10}{10} + \binom{10}{9} + \dots + \binom{10}{2} = 1,013$. Of course, researchers may opt to only include a pre-selection of supposedly plausible runoff scenarios in the survey, but, as the 2002 French Presidential election has taught us, this strategy runs the risk of missing the actual scenario.⁸ Lastly, candidate thermometer scores allow the researcher to evaluate every possible runoff scenario in advance, to see how the votes would be divided if some candidate(s) were to withdraw from the race. Data collected in such a way could therefore provide an important support for strategic coordination among party elites and voters.

4. Poststratification to population controls

At the very least, forecasting future election outcomes using opinion polls requires that the survey sample is (or can be made, for example by weighting) representative, in that it reflects the current distribution of preferences among voters, without systematic error. Such would not be the case if the supporters of certain candidates were less willing than others to participate in the survey and to reveal their vote intentions. For example, according to Jérôme et al. (1999), the underrepresentation of supporters of the National Front in pre-election survey samples, and their reluctance to truly report their supposedly socially undesirable preferences, forms one of the obstacles to accurate survey-based forecasts of the French Presidential elections.

An important aspect in the forecasting of runoff elections is that the 'true' distribution of candidate preferences at the time of the exit poll is actually co-evolving with the first round election result. These figures can then be used to adjust sample-based estimates of second round election outcomes for potential sample-population deviations in the current preference distribution.

5. Underlying assumptions

Like any other approach to election forecasting, the proposed method rests on a set of assumptions which are easier to defend under some conditions than under others. In this section, we will make these assumptions explicit and highlight the circumstances under which they seem more or less plausible.

- (1) *The sample is representative of the first round electorate.* This essentially requires that both the contacting by the interviewers and the respondents' willingness to participate in the survey are independent of the respondents' candidate evaluations. The former could be

⁸ This profusion of hypothetical runoff scenarios is also the reason why we cannot provide a direct comparison of the two data collection strategies in the empirical example which follows.

warranted by systematic sample selection rules (e.g., "approach each XXth person leaving the polling station"). Furthermore, representativeness with regard to the election outcome in round one can be achieved by weighting. Under poststratification to first round election results, the assumption of independent sampling is conditional on the first round candidate choice, and is therefore relaxed.

- (2) *The composition of the electorate remains the same between the ballots.* Both first round voters who abstain from the runoff and previous non-voters who participate in the second round are not systematically different from those voting on either ballot with respect to their candidate evaluations. This assumption has a higher face validity in elections where the turnout rates do not differ dramatically between the ballots. In addition, candidate evaluations allow for the identification of likely future abstainers from among first round voters (see the application below).
- (3) *The voters' candidate evaluations remain stable between the ballots.* This assumption seems to be most plausible if the interval between election dates is short, the campaign intensity is low and the candidates are well known before the first ballot, so that the voters do not learn a great deal about the candidates' qualities over the course of the campaign.
- (4) *If more than two candidates enter the runoff, the researcher inevitably has to make an assumption about the prevalence of strategic voting, since the supporters of weaker candidates have an incentive to vote strategically for one of the frontrunners.* We suggest robustness checking by redistributing the votes of potential strategic voters according to their preferences among the frontrunners (see the application below). Note that this is not possible with either hypothetical runoff scenarios or vote intention surveys conducted at the first round of voting. Furthermore, strategic voting is not an issue for most runoff elections, as only the two most successful candidates are allowed to enter the second round.

In summary, we have to keep these assumptions in mind when conducting forecasts based on the first round electorate. Evidently, though, they apply to pre-election surveys equally, unless the survey period is extended until shortly before the election. However, calling an event ahead of time ultimately entails taking some risks (cf. Lewis-Beck, 2005), and forecasting the runoff result on the day of the first round of the election clearly seems riskier when based on pre-election surveys.

6. Application: the 2012 mayoral election in Konstanz, Germany

Konstanz is a medium-sized university town with approximately 100,000 inhabitants, and is located in Baden-Württemberg, Germany. Mayoral elections in Baden-Württemberg are held every 8 years under a majority-plurality runoff system, where each voter casts a single vote for a candidate. If no candidate receives an absolute majority of votes in the first round, a second ballot is

held (usually two weeks after the first) in which a plurality of votes suffices to win the election. No candidate is automatically eliminated after the first round; in practice, however, many weak contenders subsequently withdraw. Due to the lack of automatic elimination, parties face little incentive to coordinate on joint candidates in the first round, particularly if no incumbent with a majority appeal is running for re-election. Consequently, the number of candidates competing in the first round is usually high, as was the case in July 2012, where thirteen candidates stood for office, six of which subsequently entered the runoff. To complicate matters further, most political parties were reluctant to endorse a particular candidate clearly, making it very difficult to single out, in advance, the candidates who were most likely to be leading after the first round of voting.

6.1. Data collection

Our forecasts make use of exit poll data collected at the first round election on July 1, 2012.⁹ We used a one-page questionnaire which included only the thermometer questions for the 13 candidates, plus the current vote choice. The interviews were conducted in a self-administered mode by students from the Department of Politics and Public Administration of the University of Konstanz, between 10.00 a.m. and 2.00 p.m.¹⁰ We selected our sampling points with care, to account for the city's political geography. The sampling points were stratified by urban district and polling station size. Konstanz consists of 15 urban districts, which, in turn, are divided into polling stations for election administration. We first selected 11 districts based on both known and assumed socio-demographic characteristics (i.e., age, income, family status).¹¹ From each district we then selected the largest polling station in terms of eligible voters and turnout in the previous mayoral election.¹² Financial limitations prevented us from including all of the polling stations in our sample (our overall budget was about 2000 Euros). Also, for reasons of cost efficiency, our interviewers were instructed to approach each voter as they were leaving the polling station. A systematic sampling rule would have been more appropriate, in order to ensure independent observations, for example, by preventing spouses from filling in the questionnaire.

⁹ Voters were asked to complete a paper-and-pencil questionnaire just after they had left the polling station. The questionnaire asked for their first round vote choice, as well as a general evaluation of each of the candidates on a 7-point scale, ranging from 1: "I do not like (candidate name) at all" to 7: "I like (candidate name) very much".

¹⁰ Bush and Lieske (1985) provide evidence that different groups vote at different times, leading to a temporal variation in reported party support, even on election day. However, the time interval we chose apparently did not affect the sample's representativeness.

¹¹ In the end, we had to give up one polling station due to the temporary unavailability of the interviewers. However, omitting this polling station had no significant influence on the results.

¹² Using the polling station returns from the previous election, we calculated a notional overall election outcome based solely on our sample of 11 polling stations, and found that it matched the official result of the previous election closely.

Table 1
First round vote shares by candidate at the 2012 mayoral elections: official results and sample distribution (percentages).

Candidate	Official result	Sample
Seeliger	20.1	21.8
Reiser	26.8	26.7
Linz	0.4	–
Neumann	1.0	0.9
Grossmann	0.5	0.3
Springer	0.1	0.2
Luthle	1.9	2.2
Zylla	14.3	17.3
Tartsch	4.6	2.5
Burchardt	25.9	24.9
Urban	0.5	0.5
Kaltenbach	2.3	1.1
Buchczyk	1.1	1.6
Others (write-in votes)	0.6	–

6.2. Poststratification

Based on the respondents' reported votes ($N = 899$), we reweighted our sample to reflect the actual first round vote proportions. Poststratification was conducted separately for the four leading candidates and for the remaining candidates. The sample vote shares for the 'big four' candidates were first weighted to reflect their district-level vote shares, after which the results were reweighted to reflect the overall result (including the off-sample districts and absentee votes). We did this in order to account for potential cross-district differences in the preference orderings among candidates, such as could occur, for example, if a candidate was particularly well appreciated among the voters in his or her home district. Sample vote shares for the remaining candidates were only poststratified with respect to the overall result, as the number of respondents who reported having voted for some of the smaller candidates was so low in some districts that the district-specific sample vote share of these candidates could not be estimated in any remotely reliable manner. Indeed, in our sample we do not have even one respondent who reported having voted for Thomas Linz (see Table 1). Hence, we had to recalculate the official election results excluding Linz before poststratifying.

7. Forecasting

Our forecasts of the runoff election held on July 15, 2012, are based upon the respondents' candidate rankings.¹³ Comparing the available individual ratings across all of the remaining candidates allowed us to derive each respondent's candidate ranking. Inevitably, citizens who rated all of the runoff candidates equally could not be considered in the forecast. The same is true for citizens

¹³ In a few cases, we revised the candidate ratings. On the questionnaire, the "don't know" (DK) option was offered right after the best rating-option. The response patterns of a very few respondents (2% of the sample) strongly suggest that they were confused by this ordering, since they rated the candidate they had voted for as DK, and all other candidates either worse than best or not at all. To correct for this, we replaced their DK answers with the best rating.

Table 2
Candidate vote shares in the runoff election: official result and forecasts.

Candidate	Official result	Scenario 1 forecast				Scenario 2 forecast			
		Sample raw	Abs. error	Sample poststratified	Abs. error	Sample raw	Abs. error	Sample poststratified	Abs. error
Burchardt	39.1	36.9	2.2	37.3	1.8	39.6	0.5	39.9	0.8
Reiser	31.9	28.2	3.7	29.9	2.0	30.2	1.7	31.1	0.8
Seeliger	27.6	27.7	0.1	26.6	1.0	30.2	2.6	28.9	1.3
Neumann	0.6	2.3	1.7	2.0	1.4	0.0	0.6	0.0	0.6
Urban	0.4	3.8	3.4	3.2	2.8	0.0	0.4	0.0	0.4
Springer	0.1	1.0	0.9	1.0	0.9	0.0	0.1	0.0	0.1
Mean abs. error			2.0		1.6		1.0		0.7
N			769		672		761		667

who gave their best rating to more than one remaining candidate (first-preference tie). Our procedure assumes that such voters would abstain in the second round.¹⁴

Scenario 1: Assuming that all citizens interviewed will choose their most preferred candidate in the runoff election, the distribution of first preferences with regard to the remaining candidates should provide proper forecasts of the second round result. Table 2 provides evidence supporting this assumption: in no case are the deviations from the actual electoral candidate shares greater than 3 percentage points.

Scenario 2: The forecasts can be improved further by making an additional plausible assumption: instrumentally motivated respondents who show a first preference for one of the three candidates (Neumann, Urban, Springer) who were highly unlikely to gain a reasonable share in the runoff due to the first round result, will vote strategically, i.e., support their favorite among the three most promising candidates. As can be seen from Table 2, the deviations from the electoral candidate ratios turn out to be even smaller: the average absolute difference between our forecast vote shares and the corresponding actual electoral support is only a minuscule 0.7 percentage points.¹⁵

8. Conclusion

In this paper, we have proposed an alternative procedure for forecasting election outcomes in two-round voting systems. Our procedure consists of the following steps: surveying the voters' candidate evaluations in first round exit polls, poststratifying the sample proportions based on the reported votes and official first round election results, and redistributing the votes of eliminated competitors according to their supporters' lower-order preferences among the candidates who are standing for election in the

runoff. Applying the procedure to the 2012 mayoral election in Konstanz allowed us to forecast the runoff result with negligible errors.

Our forecast's accuracy notwithstanding, we cannot rule out the possibility that specific characteristics of the election might have worked in our favor. For one, the runoff period was rather short (two weeks), and the campaign intensity was modest compared to national races. These factors might have helped to keep voters' candidate evaluations stable between the ballots. Furthermore, the turnout rate was pretty constant across election rounds, with turnouts of 42% and 44.5%, respectively, which lends some prima facie support to the central assumption that the composition of the electorate in rounds one and two was similar. Finally, the division between viable and hopeless runoff candidates was fairly clear-cut. This might explain why our forecast was robust to different assumptions regarding strategic voting.

On the other hand, the mayoral election also posed a challenge to our procedure, as it allowed not only the top two but all candidates from the first round to enter the runoff. In effect, we had to forecast the vote shares for six candidates, three of whom were expected to finish in the double digits. While the forecast performed slightly better for three candidates than for all six, the difference was minimal. Replications will be necessary to see how this simple method of forecasting runoff elections performs in other settings and relative to other approaches. Prominent runoff elections like the French Presidential elections suggest themselves as cases in point.

One potential drawback of the method compared to previous approaches such as pre-election surveys and model-based forecasts is that the lead time of our approach is limited to the runoff period, which is only one to several weeks in most contexts. However, even such a short lead time might be useful from the viewpoints of campaign strategists and voters. For instance, in settings where more than two candidates are allowed to enter the runoff, party elites have to decide which candidates should actually run, and those who voted for eliminated or hopeless candidates have to choose a viable alternative. Our approach provides the data necessary to allow them to make informed choices.

References

- Abramson, P. R., Aldrich, J. H., Paolino, P., & Rhode, D. W. (1995). Third-party and independent candidates in American politics: Wallace, Anderson, and Perot. *Political Science Quarterly*, 110, 349–367.

¹⁴ Alternatively, we could have considered these cases in the forecast by redistributing their votes among the remaining tied candidates instead, which would of course imply that these voters (which make up 22% of our sample) will turn out in round two and choose at random among the remaining tied candidates. Doing so results in a very similar forecast.

¹⁵ Of course, we could also account for sampling uncertainty, for example, by using resampling methods to arrive at confidence intervals for these estimates. In a sense, however, we would consider this as window dressing, given the fact that sampling uncertainty is only one – and perhaps not the most important – of the sources of uncertainty surrounding our forecasts (see Section 5).

- Belli, R., Traugott, M. W., Young, M., & McGonagle, K. (1999). Reducing vote overreporting in surveys—social desirability, memory failure, and source monitoring. *Public Opinion Quarterly*, 63, 90–108.
- Bishop, G. F., & Fisher, B. S. (1995). Secret ballots and self-reports in an exit-poll experiment. *Public Opinion Quarterly*, 59, 568–588.
- Brady, H., & Ansolabehere, S. (1989). The nature of utility functions in mass publics. *American Political Science Review*, 83, 143–163.
- Bush, R. J., & Lieske, J. A. (1985). Does time of voting affect exit poll results? *Public Opinion Quarterly*, 49, 94–104.
- Druckman, J. N., & Parkin, M. (2005). The impact of media bias: how editorial slant affects voters. *Journal of Politics*, 67, 1030–1049.
- Erikson, R. S., & Wlezien, C. (2008). Leading economic indicators, the polls, and the presidential vote. *PS: Political Science and Politics*, 41, 703–707.
- Jérôme, B., Jérôme, V., & Lewis-Beck, M. S. (1999). Polls fail in France: forecasts of the 1997 legislative election. *International Journal of Forecasting*, 15, 163–174.
- Keeter, S. (2006). The impact of cell phone noncoverage bias on polling in the 2004 presidential election. *Public Opinion Quarterly*, 70, 88–98.
- Levy, M. R. (1983). The methodology and performance of election day polls. *Public Opinion Quarterly*, 47, 54–67.
- Lewis-Beck, M. S. (2005). Election forecasting: principles and practice. *British Journal of Politics and International Relations*, 7, 145–164.
- Lewis-Beck, M. S., & Tien, C. (2008). The job of President and the jobs model forecast: Obama for '08? *PS: Political Science and Politics*, 41, 687–690.
- Magre, J., & Bertrana, X. (2007). Exploring the limits of institutional change: the direct election of mayors in Western Europe. *Local Government Studies*, 33, 181–194.
- Miguet, A. (2002). The French elections of 2002: after the earthquake, the deluge. *West European Politics*, 25, 207–220.
- Mokrzycki, M., Keeter, S., & Kennedy, C. (2009). Cell-phone-only voters in the 2008 exit poll and implications for future noncoverage bias. *Public Opinion Quarterly*, 73, 845–865.
- Plutzer, E. (2002). Becoming a habitual voter: inertia, resources, and growth in young adulthood. *American Political Science Review*, 96, 41–56.
- Presser, S. (1990). Can changes in context reduce vote overreporting in surveys? *Public Opinion Quarterly*, 54, 586–593.
- Sakurai, S. N., & Menezes-Filho, N. A. (2008). Fiscal policy and reelection in Brazilian municipalities. *Public Choice*, 137, 301–314.
- Scheuren, F., & Alvey, W. (2008). 2004: Fraudulent election? In F. Scheuren, & W. Alvey (Eds.), *Elections and exit polling (Chapter 3)*. Hoboken, NJ: Wiley.
- Selb, P., & Munzert, S. (2013). Voter overrepresentation, vote misreporting and turnout bias in postelection surveys. *Electoral Studies*, 32, 186–196.
- Stroschein, S. (2011). Demography in ethnic party fragmentation: Hungarian local voting in Romania. *Party Politics*, 17, 189–204.
- Swyngedouw, M. (2001). The subjective cognitive and affective map of extreme right voters: using open-ended questions in exit polls. *Electoral Studies*, 20, 217–241.
- van der Eijk, C. (2005). Election forecasting: a skeptical view. *British Journal of Politics and International Relations*, 7, 210–214.
- Wattenberg, M. P. (1995). Why Clinton won and Dukakis lost: an analysis of the candidate-centered nature of American party politics. *Party Politics*, 1, 245–260.
- Weisberg, H. F., & Rusk, J. G. (1970). Dimensions of candidate evaluation. *American Political Science Review*, 64, 1167–1185.
- Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *The Journal of Economic Perspectives*, 18, 107–126.