

Polling bias and undecided voter allocations: US presidential elections, 2004–2016

Joshua J Bon

School of Mathematics and Statistics, University of Western Australia, Perth, WA, Australia.

E-mail: joshuajbon@gmail.com

Timothy Ballard

School of Psychology, University of Queensland, Brisbane, QLD, Australia.

Bernard Baffour

School of Demography, Australian National University, Canberra, ACT, Australia.

Summary. Accounting for undecided and uncertain voters is a challenging issue for predicting election results from public opinion polls. Undecided voters typify the uncertainty of swing voters in polls but are often ignored or allocated to each candidate in a simple, deterministic manner. Historically this may have been adequate because the undecided were comparatively small enough to assume that they do not affect the relative proportions of the decided voters. However, in the presence of high numbers of undecided voters, these static rules may in fact bias election predictions from election poll authors and meta-poll analysts. In this paper, we examine the effect of undecided voters in the 2016 US presidential election to the previous three presidential elections. We show there were a relatively high number of undecided voters over the campaign and on election day, and that the allocation of undecided voters in this election was not consistent with two-party proportional (or even) allocations. We find evidence that static allocation regimes are inadequate for election prediction models and that probabilistic allocations may be superior. We also estimate the bias attributable to polling agencies, often referred to as “house effects”.

Keywords: Undecided voters, election polls, polling bias, presidential elections, total survey error, Bayesian modelling.

1. Introduction

Timely and accurate polls are crucial in describing current political sentiment and trends. Whilst no one poll will be sufficiently precise to enable reliable election predictions, combining the results of many pre-election polls has traditionally been viewed as a way to provide accurate forecasts. However, bias at the level of the individual poll can produce systematic error in aggregate results, particularly if these biases are correlated. One important source of polling bias arises from undecided voters. For this reason, the eventual accuracy of pre-election polls is influenced by what is done to those respondents who are undecided. In the 2016 US presidential election, a large share of voters remained indecisive up until election day. When large in number, likely voters uncertain in their candidate preferences have the power to determine tight elections. Most polling firms

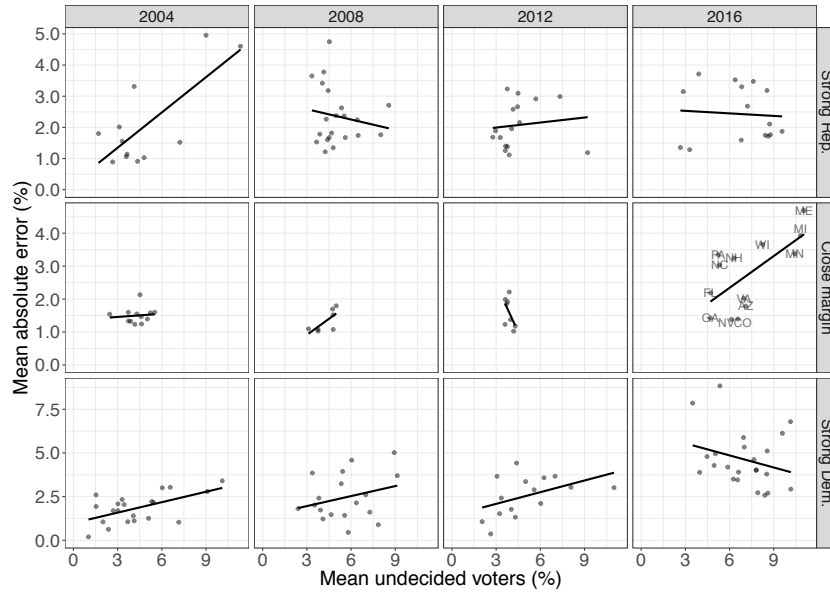


Fig. 1. Mean absolute error of state polls versus mean level of undecided voters in each state, separated by year and margin of state election victory. Polling data are polls within 35 days of US presidential elections from 2004 to 2016. “Close margin” categorises state-level elections with absolute margin $\leq 6\%$, “Strong Rep.” are races where the Republican candidate had margin $> 6\%$, and “Strong Dem.” is the remainder.

deal with undecided voters using deterministic allocation methods, the most popular being proportional or equal allocation. Static allocation methods prevent the uncertainty attributable to undecided voters from propagating through a model, and may contribute to systematic bias in the polls when the undecided voters do not split as assumed. For this reason, the use and study of probabilistic allocations methods for undecided voters is an important, but as yet under-researched area.

From a polling perspective, the 2016 US presidential election is of interest because of the public perception (and media narrative) of polling failure and the impact of high levels of undecided voters in the lead up to election day (Kennedy et al., 2017). Figure 1 shows the relationship between each state’s absolute polling error and undecided voters grouped by year (2004 to 2016) and election result margin (a strong Republican, close, or strong Democrat victory). The centre row is of most interest, as relatively high errors have the most impact in elections where the race is close. Within this subset, 2016 shows the strongest association between mean absolute error and undecided voters, indicating that the role of the undecided may have been unprecedented in the 2016 presidential election.

In this paper, we investigate the impact of undecided voters in the 2016 presidential election and presidential elections in recent years. We begin by providing background information on surveys, election polling, and undecided voters. Next, we introduce the data on national undecided voter levels to motivate our interest in the 2016 election,

followed by state-level polling data for our analysis. We present a novel extension of a recently proposed model (Shirani-Mehr et al., 2018) that allows us to quantify the bias attributable to undecided voters. We are also able to include bias attributable to certain polling agencies, often referred to as “house effects”. We show that bias in the 2016 US presidential election was critically higher than in the previous presidential elections, and a sizeable proportion of this increase can be accounted for by the high levels of undecided voters. Finally, we discuss our conclusions and recommendations from this work.

2. Surveys, election polling and undecided voters

2.1. Surveys and election polling

The accuracy (or lack thereof) of public opinion surveys and election polls has received substantial attention in recent years. All public opinion surveys, hence polls, are predicated on the assumption that citizens possess well developed attitudes on major political issues, and that surveys are passive measures of these attitudes (Converse and Traugott, 1986; Zaller and Feldman, 1992). In practice however, surveys may fail to adequately capture the sociological complexity of voting decisions and behaviour (Crossley, 1937; Gelman and King, 1993; Jacobs and Shapiro, 2005; Hillygus, 2011). Increasingly, modern public opinion surveys also have to cope with declining response rates (Keeter et al., 2000; Tourangeau and Plewes, 2013; PRC, Pew Research Center, 2012) combined with difficulties in achieving complete coverage of the population (Jacobs and Shapiro, 2005; Traugott, 2005; Leigh and Wolfers, 2006; Erikson and Wlezien, 2008). Robust evidence has demonstrated that the role of statistical uncertainty in the opinion polls has not been adequately understood (Martin et al., 2005; Sturgis et al., 2016; Hillygus, 2011; Graefe, 2014) and a failure to fully reflect this uncertainty leads to an over-statement of confidence level in predictions from survey outcomes (Erikson and Wlezien, 2008; Rothschild, 2015; Lock and Gelman, 2010).

There are a number of factors that influence the uncertainty in election polls. First, although polls play an important role in the democratic process it is becoming increasingly difficult to measure voting intention (Curtice and Firth, 2008; Keeter and Igielnik, 2016). Relatedly, there are differences between voting intentions and voting behaviour (Wlezien et al., 2013; Hopkins, 2009; Veiga and Veiga, 2004; Jennings and Wlezien, 2016). It is also generally agreed that survey respondents will not be fully representative of the entire voting public. Finally, disparities in the population result in differences in voter behaviour by geography, ethnicity, social class, gender and age, which effectively exacerbates the level of uncertainty when it comes to generalising from the sample survey to the broader population.

Understanding the mechanisms that affect polling error can provide valuable insights for better calibration and efficiency of individual polls and models derived from them. Yet undecided voters are still a relatively under-studied source of this error.

2.2. Undecided voters and election polling

Pre-election polls are typically conducted based on random sampling of likely voters who are asked for their preference among presidential candidates. Most polls record the percentage of voters who respond indecisively but do not necessarily publish this

level in their survey results (Sturgis et al., 2016). When the number of undecided voters is not reported, the polling agency may simply pre-allocate based on some judgement or decision-rule (which they may or may not report). Conversely, when the level of undecided voters is reported it can still be a challenge to incorporate these into predictive models (Hillygus, 2011; Hoek and Gendall, 1997). There is no consistent or transparent method for handling undecided voters.

To begin, we define undecided voters as individuals that are likely to vote but who have not formed a voting intention when surveyed prior to election day. The term “late-deciding” also describes these voters. Whilst similar to Kosmidis and Xezonakis (2010), our definition is restricted to likely voters as most election polls make adjustments to report results for this group only (Sturgis et al., 2016). Moreover, by necessity our definition of undecided voters must also include likely voters who have chosen not to disclose their voting preferences by stating they are undecided during a survey.

Many rule-based methods have been proposed to handle undecided voters in elections (Crespi, 1988; Daves and Warden, 1995; Fenwick et al., 1982), however some findings have indicated these assignment methods do not improve forecast accuracy (Hoek and Gendall, 1997). Simple rules for allocating undecided respondents may be adequate if the undecided voters are small in number but will likely fail when these numbers are relatively high. Additionally, any deterministic rule will not allow for variability of allocations to be modelled in predictive outcomes, which is problematic for statistical calibration. Limited research into the impact of undecided voter allocation on election poll modelling still leaves many questions about the role of indecisive voters in election polling.

The effect of undecided voters on the assessment of predictive accuracy of election polls has been considered previously (see for example, Mitofsky, 1998; Hoek and Gendall, 1997; Visser et al., 2000; Martin et al., 2005). Overall, the research has focussed on the treatment of undecided voters so that consistent accuracy measures can be defined, rather than how allocative assumptions impact polling bias. Visser et al. (2000) state that there is little published “collective wisdom” on undecided voters and better guidelines are needed, especially since excluding undecided voters was the least effective strategy in their analysis.

Investigation into undecided voter behaviour has occurred mostly in the context of election campaign assessment. For example, in US and Canada, voters who decide last minute may be more open to persuasion (Chaffee and Rimal, 1996; Fournier et al., 2004), and in the 2005 British elections, Kosmidis and Xezonakis (2010) concluded that perceived economic competence was a larger driver for the behaviour of undecided voters. However, for election outcome modellers, little can be said on the correct treatment of undecided voters for election predictions. Most notably, imputing candidate preferences for undecided voters has been found to be somewhat beneficial in Fenwick et al. (1982) whilst more recently Nandram and Choi (2008) proposed a Bayesian allocation. However, the true benefits for election forecasting is not yet clear.

Formal reports into the polling of the most recent elections in the US and UK have found some evidence of bias attributable to undecided voters. In the 2015 UK general election, Sturgis et al. (2016) report a modest, but marginal, effect from late-deciding voters, at most 1%. They assign the primary cause of the polling failure to unrepresent-

tative samples, which statistical procedures (designed to account for this) were not able to mitigate. The report into polls for the US presidential election (Kennedy et al., 2017) suggest that polls were accurate at the time they were conducted, but in some key states projection error was high due to late-deciding voters. Overall, they attribute a number of factors to polling bias in the 2016 election. In addition to late or undecided voters, they particularly emphasise over-representation of college graduates (without appropriate adjustment) and late-revealing Trump supporters – which can also manifest as larger levels of undecided voters in polls.

Based on the accuracy of the 2016 presidential election, a more sophisticated evaluation of undecided voters is necessary for predicting future elections if undecided voter levels are relatively high. This study estimates the sizeable bias attributable to undecided voters in 2016, shows evidence of undecided bias in previous presidential elections, and demonstrates that allocating undecided voters in proportion (or evenly) to the leading candidates is a poor assumption. We expect the implications of our findings will contribute to uncovering causal mechanisms that (probabilistically) determine undecided voters allocations. Unfortunately, with the available data we are not yet able to investigate these. We elaborated on the consequences for election prediction and future research in the discussion.

2.3. *Meta-analysis of polls*

Any one poll will be a snapshot of the sample collected, fraught with difficulties pertaining to sampling design, non-representativeness, and differences in methodological assumptions. For this reason, a single poll should be interpreted cautiously, even more so in situations when there is little previous experience to draw upon, for example in a referendum or, arguably, the 2016 US presidential election. Nonetheless, polling results from various sources can be compiled, compared, analysed and then interpreted using meta-analysis techniques to combine (or pool) together different polls.

Meta-polls and poll modelling can compensate for the bias and inaccuracy of individual polls, but establishing well-calibrated models require understanding the inherent problems in polls, and defensible model assumptions. In election polling, deterministic (rule-based) allocation of undecided voters is widespread (Crespi, 1988; Visser et al., 2000; Martin et al., 2005). These methods are appealingly simple and create a consistent set of data when polling organisations do not publish the number of undecided or third party voters. Undecided voters can be allocated (explicitly or implicitly) in a number of ways (Martin et al., 2005; Mitofsky, 1998). The most prevalent are:

- Splitting the undecided voters proportionately between the two leading candidates. This is equivalent (in mean) to discarding the undecided voters and normalising the two leading candidate’s voter proportions, and
- Allocating half of the undecided voters to each of the leading candidates. This is equivalent (in mean) to only reporting the margin between the two leading candidates.

Identifying the allocation procedures that polling firms use (if they do not report undecided voters) is difficult because they are averse to providing commercially sensitive information. Some meta-pollsters have published how they handle undecided voters in their

models for at least the 2016 election. FiveThirtyEight split undecided evenly between the major-party candidates (Silver, 2016), as did the Princeton Election Consortium implicitly when they used the margin between the two leading candidates (Princeton Election Consortium, 2016). In contrast, the Huffington Post used a different strategy, assuming that “one-third of undecided voters won’t vote; one-third will gravitate nationally toward either candidate; and the remaining one-third will add to this state’s margin of error” (Huffington Post, 2016a). Interestingly, the Huffington Post model is a mixture between proportional allocation, even allocation, and (imprecisely) incorporating some uncertainty from undecided voters into poll modelling.

Meta-analysis of polls is crucial to obtaining reliable predictions for elections, however, it is very difficult for these models to account for systematic bias in polls. Namely, if every poll is subject to a particular source of bias, how do we isolate and quantify its influence without external information? Investigating the role of undecided voters in polling bias will help to understand one aspect that contributed to larger polling bias in the 2016 presidential election, which may occur again.

2.4. Data

In Section 3, we examine the extent to which 2016 was an abnormal presidential election by considering the number of undecided voters relative to previous years. To investigate we summarise national polling data in US presidential elections from 2004 onwards, totalling 616 national polls. Polls from 2012 and 2016 were obtained from the Huffington Post’s Pollster API (Huffington Post, 2016b; Arnold and Leeper, 2016), data for 2008 were retrieved from an archived version of “Pollster.com” (Huffington Post, 2009), and data from 2004 were reconstructed[†] with polls available from (RealClearPolitics, 2004). Publicly available polls that reported a sample size were used. In order to model the undecided voters, the majority of polls in each election year needed to report the undecided category (explicitly or implicitly). This limited the number of elections that could be analysed to 2004, 2008, 2012, and 2016.

State level polling data from 2004 to 2016 are used to model polling bias and variance. The polls for 2012 and 2016 were obtained from Huffington Post (2016b) whilst the polls for 2004 and 2008 were retrieved from US Election Atlas (Leip, 2008). Polls were included if they occurred up to 35 days prior to their respective election. Due to the complexity of our model, state-level elections were only included if they had at least 5 polls in the dataset. Whilst other poll repositories exist for 2004 and 2008 state-level election polling data, none consistently reported undecided voter counts. In total 1,905 polls, from 129 state-level election races were analysed. State-level polling data for US presidential elections where a majority of the polls included an undecided voter category were not found for years prior to 2004 by the authors.

[†] All polls that reported a third party candidate, but did not sum to 100% of the sample were assumed to be reporting an undecided category equal in size to the remaining proportion.




Fig. 2. Mean level of undecided voters, as captured by national polls, over the 90 days prior to US presidential elections from 2004 to 2016. The number of undecided voters on each day x is the weighted average from national polls that occur within a two-week window centred at x .

3. Comparison of undecided voters in US presidential elections from national surveys

Undecided voters were much higher during the 2016 US presidential election relative to previous years. Figure 2 shows the moving average number of undecided voters over the course of presidential elections from 2004 to 2016. It can be seen that the year 2016 had a larger number of undecided voters on average and that this trend was persistent over the course of the campaign. Whilst 2016's pattern of undecided voters over time was similar to that of 2012, the major difference was that in the final week of polling the undecided voters did not continue to fall. It also appears that the 2004 and 2012 elections followed a similar pattern, however the extra variability in the 2004 election may be explained by lower numbers of polls and the reconstruction that took place. In the week prior to each election the weighted average of undecided voters was 5.1%, 3.5%, 3.9%, and 2.7% for 2016 to 2004 respectively.

The distributions of undecided voters in the months leading up to presidential elections also appear to vary over time, as seen in Figure 3. The undecided voters in 2016 and 2008 have relatively fatter tailed distributions, whilst the 2012 and 2004 elections appear to be centred between 3-4%. Undecided voter levels higher than 10% occurred more frequently in the 2016 election than 2012 and 2004, and somewhat more frequently than in 2008. The similarities of 2016 and 2008 as well as 2012 and 2004 may be partially attributable to the absence or presence of an incumbent candidate, but this is difficult to infer from only 4 elections.




Fig. 3. Histogram of undecided voters, as captured by national polls, over the course of 3 months prior to the US presidential elections 2004-2016. Each bar is relative to the number of polls from that year.

The levels of undecided voters in 2016 had an unusually high mean, differed more across states, and did not follow the final week decrease of previous elections. This finding, coupled with evidence from Figure 1, motivates an investigation into the effect of undecided voters on polling errors in the 2016 US presidential election.

4. Methods

4.1. Assessing poll bias with the total survey error framework

We use a total survey error framework (see Biemer, 2010, for overview) to analyse state polls from the US elections. The paradigm has a long history, discussed as early as Deming (1944). We refer readers to Groves and Lyberg (2010) for a thorough introduction to the literature. The total survey error paradigm attempts to account for, and assess, many sources of error with respect to user requirements. In this analysis we take the user requirement to be predictive accuracy and calibration, and our users to be either polling organisations or poll aggregators. As such, our aim is to estimate bias and variance and describe the predictive characteristics of polls to these users. In particular, we quantify the bias and variance attributable to undecided voters, and demonstrate the inadequacy of proportional (or even) allocation rules.

As we are interested in ascertaining the impact of state-level influences, we consider each state-level election from each presidential election year to be a distinct election. With regard to modelling, henceforth an election refers to a state-level presidential election from an US election year, specifically, 2004, 2008, 2012 or 2016.

Under the total survey error framework, a survey error is defined as the deviation of a survey response from its true underlying value. This error can occur either through bias or variance. The bias term captures the systematic errors which are shared by all election polls, such as shared operational practices, infrastructure and sampling frames for example. The variance term captures sampling variation, and can account for variation due to different survey methodologies, software, statistical models, or survey weighting adjustments. Therefore, the poll error which is computed through comparing the election outcomes to the predictions from multiple election polls, can be decomposed into election level bias and variance terms. We adopt this framework to ensure that the analysis of undecided voters and their role in the bias is not conflated by non-sampling errors.

On the one hand, sampling errors arise from taking a sample rather than the whole population and are usually accounted for using standard survey sampling approaches, including post-stratification (Holt and Smith, 1979), calibration (Deville and Särndal, 1992), imputation (Gelman and Carlin, 2002). On the other hand, non-sampling error is a catch-all term that refers to all other sources of error that are not a function of the sample. In theory, although a specific poll estimate may differ from the true election outcome, under favourable repeated sampling conditions polls should produce reliable estimates (Assael and Keon, 1982). However, in practice, it is well known that differences between poll results and election outcomes are only partially attributable to sampling error (Ansolabehere and Belin, 1993). Most statistical procedures to compensate for non-sampling errors assume near universal (or high) response, but this is far from the norm: the majority of election surveys have less than 10% response rates (PRC, Pew Research Center, 2012). Further, in polling there is a general difficulty in measuring voting intention and voting behaviour because polls measure the beliefs and opinions of respondents at the time of the survey, and cannot fully capture what respondents will do on election day (Bernstein et al., 2001; Silver et al., 1986; Rogers and Aida, 2014; Jowell et al., 1993). Increasingly, efforts to mitigate against this can exacerbate the inaccuracy (Ansolabehere and Hersh, 2012; Voss et al., 1995; Gelman et al., 2016; Bailey et al., 2016). This is especially true when it comes to dealing with those who are undecided, either because they are truly undecided, or are hiding extreme voting preferences (Gerber et al., 2013). The treatment of the group of polling responses that report to be undecided is therefore an important, yet relatively unstudied area of research.

4.2. A Bayesian approach to total survey error incorporating undecided voters

We use a meta-analytic approach to compare the estimates from the individual polls to the eventual election outcomes. This combines information from the various polls to produce a pooled estimate for the differences between the state-election poll means and the election outcome – which is our ground truth. The modelling framework is based on the model proposed by Shirani-Mehr et al. (2018). Their model quantifies the total error by estimating the vote share under a two-party preference, but does so by assuming proportional allocation of undecided voters. We extend their model to explicitly include undecided voter proportions. A second addition to our model is a term for house effects (or pollster-specific bias) of polling agencies. These two sources of bias do not account for all errors polls are subject to, but do provide a quantitative method for assessing

how much variability in the voting outcome is due to these sources. For other sources, the model uses state-year net bias and variance terms to capture the remaining error in the polls.

Since there are relatively few numbers of polls in some elections, a simple measure, such as root mean squared error, may yield imprecise estimates of the election-level bias. We address this by fitting a Bayesian hierarchical latent variable model (Gelman and Hill, 2007). This method pools data to determine estimates of bias and variance in states with small numbers of polls, allows bias to vary over time, and better captures the variance in excess of that expected from a simple random sample. The model “borrows strength” across states and time to estimate smoothed within state trends of both polling bias and undecided voters in each election.

With small adjustments the notation in Shirani-Mehr et al. (2018), each poll is associated with an election denoted by the index $r[i]$. Let y_i be the two-party support for the Republican candidate of poll i , n_i be the sample size, and t_i be the time at which the poll was conducted. Two-party support indicates that a proportional allocation of the undecideds voters (or scaling) has occurred, namely

$$y_i = \frac{R_i}{R_i + D_i} \quad (1)$$

where the Republican and Democratic support as measured in poll i , is R_i and D_i respectively. The time t_i is the duration between the last day the poll was conducted and the relevant election date, and is scaled to be between 0 and 1. The Republican candidate’s final two-party vote outcome is denoted by v_r . Each poll is assumed to be distributed by

$$\begin{aligned} y_i &\sim \mathcal{N}(p_i, \sigma_i^2) \\ \text{logit}(p_i) &= \text{logit}(v_{r[i]}) + \alpha_{1r[i]} + t_i\beta_{1r[i]} \\ \sigma_i^2 &= \frac{p_i(1-p_i)}{n_i} + \tau_{1r[i]}^2 \end{aligned} \quad (2)$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution parametrised by mean and variance, and the first subscript (e.g. the 1 in α_{1r}) indicates that the parameter is part of the polling model. The vote, v_r , captures the true mean of the polls allowing election level bias to be estimated by $\alpha_{1r} + t_i\beta_{1r}$ on the logit scale. Estimation on this scale ensures the estimated poll value, p_i , is bound between 0 and 1. Bias on election day is α_{1r} and the time-varying bias coefficient is β_{1r} . As for the variance, τ_{1r}^2 accounts for the excess variance above what is expected in a simple random sample.

The model is able to detect the bias in election polling at the state-year level by centring the model about the actual election outcome, whilst estimating the excess variance by anchoring the model variance at the level expected from a simple random sample. Elections with few polls are estimated by pooling the data across elections using hierarchical priors (see Table 3 in Appendix A).

Using the estimated two-party support from polls, as in (1), the model implicitly assumes that undecided voters are distributed proportionately to the major-party candidates. We relax this assumption by explicitly distributing the undecided voters in

proportion but with flexibility. To illustrate, let the undecided voters from poll i be U_i , with R_i and D_i as in (1). Scaling the polls to exclude third party candidates, we assume that

$$y'_i = \frac{R_i + \lambda U_i}{R_i + D_i + U_i} \quad (3)$$

where $0 \leq \lambda \leq 1$ allocates the undecided voters to the Republican candidate. Rather than using proportional allocation, $\lambda = \frac{R_i}{R_i + D_i}$, as is the case in model (2), we use

$$\lambda = \frac{R_i}{R_i + D_i} + \theta_i \quad (4)$$

where θ_i is an unknown bias (away from proportional allocation) which occurs at some level (i.e. poll, election or election year). Simplifying equation (3) with (4) leads to the identity

$$y'_i = y_i + u_i \theta_i \quad (5)$$

where y_i , represents the Republican two-party vote share as in (1) and $u_i = \frac{U_i}{R_i + D_i + U_i}$ represents the scaled proportion of undecided voters. This observation motivates changing model (2) to include undecided voters as an explanatory variable.

The term θ_i measures the bias away from a proportional two-party split. However, using poll-level undecided voters as an explanatory variable is problematic for two reasons; it is subject to measurement error (it is a survey estimate), and the level of undecided voters varies over time (see Figure 2). The latter issue may confound θ_i with estimates of the time-varying component of bias already in the model, β_{1r} .

To address these concerns we propose a model for the undecideds so that election day undecided voter levels can be included in model (2), rather than the undecided numbers from each poll. The model of the undecided voters is

$$u_i \sim \mathcal{N} \left(\alpha_{2r[i]} + t_i \beta_{2r[i]}, \tau_{2r[i]}^2 \right) \quad (6)$$

where α_{2r} estimates the election day level of undecided voters, the observed change in undecideds over time is accounted for by β_{2r} , and τ_{2r}^2 measures the variability within each state-level election race r . From the undecided model, α_{2r} becomes an explanatory variable measuring the election day level of undecided voters in each race. It appears in the extended model[‡]

$$\begin{aligned} y_i &\sim \mathcal{N}(p_i, \sigma_i^2) \\ \text{logit}(p_i) &= \text{logit}(v_{r[i]}) + \alpha_{1r[i]} + t_i \beta_{1r[i]} - \alpha_{2r[i]} \gamma_{g[i]} + \kappa_{h[i]} \\ \sigma_i^2 &= \frac{p_i(1-p_i)}{n_i} + \tau_{1r[i]}^2 \end{aligned} \quad (7)$$

where γ_g has replaced θ_i in (4) and is the bias attributable to mean undecided voters from the groups of state-years given in Figure 1. The 12 groups, index by g , are formed by the

[‡]Since the undecided voter proportions are between 0% and 10% approximately, we multiply α_{2r} by 10 so that it is on the same scale as t_i . This improves computational performance when estimating the model. Testing has shown negligible impact on the estimates.

cartesian product of election year (2004, 2008, 2012, and 2016) and the election result margin groups. The latter categorises how close election results were; strong Republican (margin $> 6\%$ in favour of Republican), close margin (absolute margin $\leq 6\%$), and strong Democrat (margin $> 6\%$ in favour of Democrat). These groups are in line with Figure 1 and were chosen because results with a margin greater than 6% are extremely unlikely to be affected by undecided voter allocation – in the sense that (for the majority of states) the winner collects all of the state’s Electoral College votes. Hence, the close margin group will have the most impact on the eventual outcome of the election. The groups are also a crude measure of partisan strength and influence in a state.

The γ_g coefficient is not estimated at the poll level because we now estimate (and allocate) at the state-year level of undecided voters on election day (α_{2r}). Moreover, using the 12 groups, rather than a coefficient for each state-year, occurs because of identifiability issues with α_{1r} . In addition to accounting for measurement error, using (6) also allows polls that do not report undecided voters (i.e. they have missing data) to be included in the model since only the state-level mean enters model (7).

The house effects from different polling agencies or groups are modelled by $\kappa_{h[i]}$. House effects are errors specific to a polling firm, for example they may arise from flawed survey or modelling methods or partisan prejudice. We include bias terms for 39 different polling agencies, indexed by h . Not all polls have an associated house effect, only those where at least 8 polls were in the dataset.

To ensure that the final inference is not substantially affected by the choice of prior, we specify weakly informative priors following the previous work (Shirani-Mehr et al., 2018). The hierarchical specification of the priors pull the bias and variance estimates of the poll towards the state’s average in a given election year. However, the effect is related to the number of polls in the particular election year, and the overall distribution across all polls. For states with few polls (in a given year), the estimates can be inferred from information from other polls (across state and time). For the γ_g and κ_h we specify priors with shrinkage towards zero, this ensures the probability of overestimating these effects is low. The list of priors and further explanation can be found in Table 3 in Appendix A.

We maintain a focus on the assumption of proportional allocation of undecided voters as it allows us to use polls that do not report an undecided voter number to still enter model (7). However, we rerun the model with an even split of the undecided voters to each party for a more robust analysis. The results are very similar to model with proportional allocation (see Figure 6 for example) and are discussed further in Appendix D.

Bayesian posterior sampling was conducted using the No-U-turn Hamiltonian Monte Carlo sampler (Hoffman and Gelman, 2014) implemented in `Stan` (Carpenter et al., 2017). The analysis was facilitated by the statistical coding environment and language R (R Core Team, 2017), the R interface to `stan`, `rstan` (Stan Development Team, 2018), and diagnostics were provided by `shinystan` (Stan Development Team, 2017).

5. Results

First, we compare our results from election years 2004 to 2016 with results obtained by Shirani-Mehr et al. (2018) for 2000 to 2012 using the model specified in (2) and the

Table 1. Average election-level absolute bias and average election-level standard deviation across state-elections in given year(s) from model (2). Values shown are posterior mean (s.d.) in percentage points.

	2004	2008	2012	2016	Overall	
					2004–2016	2000–2012¶
Average absolute bias	0.9% (0.11)	1.2% (0.10)	1.4% (0.11)	2.6% (0.10)	1.7% (0.06)	1.2% (0.07)
Average absolute election day bias	0.9% (0.13)	1.2% (0.12)	1.3% (0.14)	2.4% (0.13)	1.6% (0.07)	1.2% (0.08)
Average standard deviation	2.2% (0.05)	2.2% (0.04)	2.1% (0.04)	2.4% (0.05)	2.3% (0.03)	2.2% (0.04)

aforementioned paper. In Table 1, the average absolute bias§ (bias from $\alpha_{1r} + t_i\beta_{1r}$) and election day bias (α_{1r}) are both considerably higher in 2016 than in the previous three election years. The measures were at least 1.1 percentage points above previous years, having more than twice as much bias in the case of 2004 and 2008. The increased bias in 2016 can explain the increase to the overall 2004–2016 average bias compared to that of 2000–2012. The yearly averages shown in Shirani-Mehr et al. (2018) are also consistent with the results in Table 1 (2004, 2008, and 2012).

Whilst the bias in election polls seems to have played a large role in the abnormality of the 2016 election year’s polls, the average standard deviation appears to be consistent across time. The average standard deviation in 2016 was only 0.2% above the next highest year. This is not a qualitative difference given the range of values is only 2.1% to 2.4% from 2004 onwards. The consistency in average standard deviation over time lends strength to the conclusion that individual polls are subject to approximately twice as much standard deviation than what is reported (i.e. a simple random sample calculation) (Shirani-Mehr et al., 2018; Rothschild and Goel, 2016).

Second, Table 2 lists the average results from model (7) where several bias definitions are considered. Average absolute bias describes the average of election-level bias from all sources, whilst average absolute election day bias consists of all sources but the time-varying component (i.e. $t_i = 0$). Average absolute undecided voter bias and house effects consist solely of their respective bias terms, namely $\alpha_{2r[i]} \gamma_{g[i]}$ and κ_h . They are also averaged at the election-level, as is average standard deviation, the average election-level value of σ_i from model (7). Details of these calculations are in Appendix B. The estimates of bias in a given column of Table 2 will not sum to the total (average absolute bias) since they can take different signs.

The state level aggregation of undecided voters in (6), used in model (7), predicts that 3.0% to 3.8% of voters were undecided on election day between 2004 and 2012, whilst predicting 5.5% in 2016. These results are consistent with Figure 2 which suggested much higher levels of undecideds in 2016 than previous years.

The average absolute (election day) bias in Table 2 is approximately equal to those in Table 1. This indicates that both the original model, and the extended model are able to

§see Appendix B for technical definitions

¶Results from identical model using data from 2000 to 2012 (Shirani-Mehr et al., 2018).

Table 2. Average election-level absolute bias and average election-level standard deviation across state-elections in given year(s) from model (7) with assumption of proportional allocation of undecided voters

	2004	2008	2012	2016	<i>Overall</i> 2004–2016
Average absolute bias	0.8% (0.11)	1.0% (0.10)	1.3% (0.10)	2.6% (0.10)	1.7% (0.06)
Average absolute election day bias	0.8% (0.12)	0.9% (0.11)	1.3% (0.14)	2.4% (0.12)	1.6% (0.07)
Average absolute undecided voter bias	0.3% (0.17)	0.4% (0.17)	1.0% (0.29)	2.1% (0.25)	1.1% (0.11)
Average absolute house effects	0.6% (0.15)	0.4% (0.12)	0.2% (0.08)	0.2% (0.09)	0.3% (0.09)
Average standard deviation	2.2% (0.04)	2.2% (0.04)	2.1% (0.04)	2.4% (0.05)	2.2% (0.03)
Average election day undecided	3.3% (0.24)	3.8% (0.21)	3.0% (0.21)	5.5% (0.28)	4.2% (0.14)

account for approximately the same amount of bias in the US presidential election polls. However, the extended model is able to disaggregate bias into two additional sources.

The average absolute undecided voter bias was estimated to be 2.1% in 2016. This is more than twice the value in previous years. We can also calculate the average election day bias without undecided voters in 2016, which was only 1.1% (0.14). This is much closer to the total election day bias estimated in previous years (0.8-1.3%), and highlights the strong influence undecided voters had on the 2016 presidential election.

The bias attributable to undecided voters in 2004 and 2008 is very small, only 0.3-0.4%. Since the posterior estimates for the effect of undecideds on bias (γ_g) are mostly centred close to zero (see Figure 4), and the estimated level of undecided voters was relatively low (Table 2), the role of undecided voters seems minimal in these years. The aggregate effect of these factors can be seen in Figure 5. In the 2012 election a change occurs when undecided voter bias moves from 0.4% (2008) to 1.0%. The potential effect of undecided voters in this year may have been mitigated by the relatively low levels polled (only 3.0% on average).

To further elicit the role of undecided voters in 2016, Figure 4 contains the 95% and 50% credible intervals of the effect size of undecided voters on the polling bias, γ_g , on the logit scale. This parameter vector estimates the biasing effect attributable to undecided voters for twelve groups discussed in Section 4.2 and shown in Figure 1. In terms of the election outcome, we should pay most attention to the close margin category due to the winner-takes-all effect of the Electoral College system.

The credible intervals of γ_g in election years 2004 and 2008 show similar results. Both the strong Republican and close margin categories show no sign of bias from undecided voters. Evidence for polling bias in these years and groups could be mixed (changing from state to state) or inconclusive. There is weak evidence from the strong Democrat category that some bias was induced by the undecided voters (away from proportional




Fig. 4. Credible intervals of 95% (outer line) and 50% (inner line) for the effect of undecided voters on polling bias in the model (γ_g) on logit scale. A positive value indicates a bias away from proportional allocation of undecided voters in favour of the Republican candidate. A negative value is favourable for Democratic candidates.

allocation) toward the Republican candidate. In 2012, the reverse occurs – the strong Republican group has bias toward the Democratic candidate. There is some evidence that a biasing effect was present for the close margin group in 2012, but it may also be low or negligible. Figure 1 shows the mean level of undecided voters was relatively low in 2012’s close margin category, so overall undecided voters were unlikely have had a substantive effect in 2012. In general, there seems to be a skew toward the Democratic candidate in 2012, whilst the reverse pattern emerges even more strongly in 2016.

The 95% credible intervals for 2016’s close margin group range from 0.40 to 1.75 with a mean of 1.08. This is the only year (in our set) where the close margin states so clearly induced bias into election polls. Considering estimated undecided numbers averaged 5.5% on election day in 2016, this was an important source of bias in 2016 and reveals one reason why Trump underperformed in the polls. The strong Democratic group had an even higher bias in 2016 toward Trump, with mean 2.81. This is an interesting result in its own right, (as is 2012’s strong Republican category), but had little or no effect on the perception of poll failure since the binary prediction of these races are still likely to be correct, and hence unlikely to change the Electoral College predictions.

Figure 5 shows the distribution of state level average absolute bias from undecided voters on election day. This figure highlights the consequence of high numbers of undecided voters combined with their large biasing effect in 2016 (and to much less of an extent in 2012). The undecideds’ contribution to polling bias in 2016 is made up of three distinct groups, roughly translating to the aforementioned categories. The effect in 2016 was up to approximately 4%, whilst the overall effect of undecideds has been almost negligible in 2004 and 2008. The low impact of undecided voters prior to 2016 is likely due to the relatively low levels observed, combined with the possibility that undecided voters did not bias polls cohesively in previous years.

A house effect (κ_h), attributable to the bias induced by polling agencies methods and practices, is estimated for 39 such firms in the model. Overall, we find a greater number of polling agencies have high bias favouring Republican candidates than high bias towards Democratic candidates, and an overall skew favouring Republicans. However, more polls are needed to accurately assess the Democratic-leaning pollsters. The estimates of this component of bias are presented and discussed further in Appendix C.

Figure 6 compares the 2016 election day bias of the polls in each state ($\alpha_{1r} - \alpha_{2r}\gamma_g$ in the model) for the two types of static undecided voter allocations (see Appendix D for results of even allocations). The grey bars indicate the 95% credible intervals of election day bias for proportional allocation, whilst the black bars are for even allocation. The state election day biases show that neither allocations performed uniformly better than the other. Proportional allocation resulted in polls with a relatively less biased performance in Utah, Idaho, Iowa, and California for example. Whereas in Vermont, Maine, and Alaska the even allocation was performed marginally better. The relatively larger credible intervals for the even allocation can be attributed to the reduction in sample size for this model which is explained in Appendix D.

Using an even allocation of undecided voters caused no substantive change to the results above, which assumed proportional allocation. The quantitative results in Appendix D and Figure 6 show that both allocation methods were insufficient to incorporate the necessary uncertainty in undecided voter allocation. In summary, a rigid and non-




Fig. 5. Histograms of the average absolute bias from undecided voters for each state-level election, separated by year. The bias from undecided voters is the quantity $\alpha_{2r}\gamma_g$ in the model. A positive value indicates a bias away from proportional allocation of undecided voters in favour of either candidate.

probabilistic rule for allocating undecided voters was not appropriate for the 2016 US presidential election.

6. Discussion

While there have been methodological advances in election polling and modelling over time, measuring public opinion is still difficult. Polls have ‘failed’ to accurately predict winning candidates (at least according to the media) in several recent elections, such as the 2015 British election, the Scottish independence referendum, the Brexit referendum, the 2014 US House elections and the 2016 US presidential election (the focus of our paper). Our analyses revealed that one important source of bias in the polls may be undecided voters. We showed that undecided voters biased polls in the 2016 US presidential election by up to 4% in some states (all things being equal). This bias is particularly problematic given the increasing number of undecided voters observed in our analysis. We found that in 2016, 5.5% of voters were undecided on election day— up from 3-4% in previous years. Others have also found that the percentage of voters who are undecided in the final week of an election campaign is high (up to 30% in some countries) and may be increasing (Irwin and Van Holsteyn, 2008; Gelman and King, 1993; Orriols and Martnez, 2014; Sturgis et al., 2016), although Sturgis et al. (2016) note that the overall effect of late-deciding voters was modest, at most 1%, in the 2015 British general election. Given the rising prevalence of indecisive, late-revealing or late-deciding voters, and the bias they may introduce to polling, more attention needs to be given to this group when predicting election outcomes.

It is well known within the survey research community that, polls suffer from both sampling (due to the fact that information has been collected from a sample rather than everyone in the population) and non-sampling (due to the fact that there is underlying differences in voting behaviour and voting outcomes) errors. Statistical and operational adjustments compensate for sampling errors, but in reality it is the non-sampling errors that play a significant role in the discrepancies between the poll results and election outcomes. The total survey error approach provides a methodology for capturing both types of errors. We have used this approach to understand, interpret and report the various sources of error that exist in election polling. We have focused specifically on undecided voters, and provide evidence that found that there was substantial differences in the degree of undecidedness in pre-election polling in the 2016 US election.

Since the majority of polling agencies had no specific methodology to include this in their predictive models, the reported results over-estimated the lead of the Democratic candidate, Clinton, against the Republican candidate, Trump. Our results show that voters who were undecided at the time of being surveyed tended to behave differently to those who decided which party or candidate to vote for earlier. Although it is well recognised that undecided respondents contribute to polling error, there is no consensus about the inferences that can be drawn from their data (Henderson and Hillygus, 2016; Fenwick et al., 1982; Hillygus, 2011). Our research has demonstrated that a failure to adequately include them leads to inaccuracies in polling predictions. Specifically, static proportional and even allocations of undecided voters led to bias in polling of the 2016 presidential election. Because of this, we argue that undecided voter counts




Fig. 6. 95% Credible intervals for state election day bias (from true election results) in the 2016 presidential election for proportional (grey line) and even (black line) undecided voter allocation models, after accounting for house effects. A positive value indicates a bias of the polls in favour of the Republican candidate, hence an underperformance in the actual election result.

should always be reported, and probabilistic allocation of undecided voters needs to be implemented in future predictive models. This will allow the uncertainty attributable to undecided voters to propagate through these election models, and improve predictions.

The American Association for Public Opinion Research’s (AAPOR) report on US presidential election polling in 2016 (Kennedy et al., 2017) concluded that, in general, polls did not fail relative to historical standards but did under-estimate the support for a Trump presidency in some key states. The authors cited late-deciding voters as one of several explanations for Trump outperforming Clinton on election day. Our study complements AAPOR’s report by demonstrating a sizeable number of undecided voters on election day relative to previous elections and by finding a significant effect of undecided voters on the bias in election polls in 2016 using a total survey error model. The AAPOR report relied on exit poll data, whereas we use polling data prior to the election, meaning it may be possible to detect and mitigate these effects in future predictions. We also show that the role of undecided voters in 2016 was different to the 2004, 2008, and 2012 elections. In these previous elections, undecided voters did not significantly affect the races which were close, nor did they significantly contribute to overall polling bias as the number of undecideds remained low. Some bias due to undecided voters was visible in elections prior to 2016, however, the affected states were not tight races and therefore did not influence the binary outcome of the election.

Despite the novelty of our findings in ascertaining the role of undecided voters and the adoption of the total survey error framework, our analyses has some noteworthy limitations. Our models remain associational and only provides evidence to support the hypothesis that there is a relationship between the increase in undecidedness and polling accuracy. Properly understanding the underlying effects, and causal mechanisms surrounding how indecision directly influences election outcomes will require future interdisciplinary research. Figure 4 shows that there have been other election years where undecided voters have biased polls away from proportional allocation for certain groups of states. Future research could elicit similarities between these groups which may be useful for predicting the biasing effect of undecided voters.

In our analysis we would have liked to model more sources of error explicitly, particularly other sources of error highlighted in Kennedy et al. (2017). For example, accounting for over-representative sampling of college graduates in some polls would be helpful. However, we were constrained by lack of survey methodology disclosure by polling agencies that limits the number of attributes recorded for polls in our dataset.

Additionally, though our modelling allows us to quantify the errors that are left unmeasured in standard election level estimates of accuracy, we have not translated this to a model of how respondents will vote in future elections. For those wishing to predict US presidential elections, our analysis demonstrates that if undecided voter levels are high they must be included in the modelling process and that deterministic allocation is not appropriate. For example, an averaging method could model undecideds levels by state, such as in Equation (6), and allocations to candidates can be simulated (by Bayesian methods or bootstrapping say). Probabilistic allocation should at least lead to better estimates of uncertainty, and thereby model calibration. The posterior estimates of undecided allocation bias presented in this paper may also be used as a starting place for predictive modellers to incorporate this bias.

Undecided voters played a pivotal role in the 2016 US presidential election, contributing significantly to the bias observed in the polls. We have shown that static allocation methods (proportional and even) were inadequate in the 2016 election. Probabilistic allocation should be considered in future elections, but further investigation and validation of specific methods is needed especially considering the limited supplementary information provided by commercial polls.

Acknowledgement

The authors would like to thank the anonymous reviewers for their valuable comments which improved the paper.

References

- Ansolabehere, S. and T. R. Belin (1993). Poll faulting. *Chance* 6(1), 22–28.
- Ansolabehere, S. and E. Hersh (2012). Validation: What big data reveal about survey misreporting and the real electorate. *Political Analysis* 20(4), 437–459.
- Arnold, J. B. and T. J. Leeper (2016). *pollstR: R client for Pollster API*. R package version 1.4.0.
- Assael, H. and J. Keon (1982). Nonsampling vs. sampling errors in survey research. *Journal of Marketing* 46(2), 114–123.
- Bailey, M. A., D. J. Hopkins, and T. Rogers (2016). Unresponsive and unpersuaded: The unintended consequences of a voter persuasion effort. *Political Behavior* 38(3), 713–746.
- Bernstein, R., A. Chadha, and R. Montjoy (2001). Overreporting voting: Why it happens and why it matters. *Public Opinion Quarterly* 65(1), 22–44.
- Biemer, P. P. (2010). Total survey error: Design, implementation, and evaluation. *Public Opinion Quarterly* 74(5), 817–848.
- Carpenter, B., A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell (2017). Stan: A probabilistic programming language. *Journal of Statistical Software, Articles* 76(1), 1–32.
- Chaffee, S. H. and R. N. Rimal (1996). Time of vote decision and openness to persuasion. In D. C. Mutz, P. M. Sniderman, and R. A. Brody (Eds.), *Political persuasion and attitude change*, pp. 267–291. Ann Arbor: The University of Michigan Press.
- Converse, P. E. and M. W. Traugott (1986). Assessing the accuracy of polls and surveys. *Science* 234(4780), 1094–1098.
- Crespi, I. (1988). *Pre-election polling: Sources of accuracy and error*. Russell Sage Foundation.
- Crossley, A. M. (1937). Straw polls in 1936. *Public Opinion Quarterly* 1(1), 24–35.

- Curtice, J. and D. Firth (2008). Exit polling in a cold climate: the BBC-ITV experience in Britain in 2005. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171(3), 509–539.
- Daves, R. P. and S. Warden (1995). Methods of allocating undecided respondents to candidate choices in pre-election polls. *Presidential polls and the news media*, 101–19.
- Deming, W. E. (1944). On errors in surveys. *American Sociological Review* 9(4), 359–369.
- Deville, J.-C. and C.-E. Särndal (1992). Calibration estimators in survey sampling. *Journal of the American statistical Association* 87(418), 376–382.
- Erikson, R. S. and C. Wlezien (2008). Are political markets really superior to polls as election predictors? *Public Opinion Quarterly* 72(2), 190–215.
- Fenwick, I., F. Wiseman, J. F. Becker, and J. R. Heiman (1982). Classifying undecided voters in pre-election polls. *Public Opinion Quarterly* 46(3), 383–391.
- Fournier, P., R. Nadeau, A. Blais, E. Gidengil, and N. Nevitte (2004). Time-of-voting decision and susceptibility to campaign effects. *Electoral Studies* 23(4), 661–681.
- Gelman, A. and J. B. Carlin (2002). *Survey Nonresponse*, Chapter Poststratification and Weighting Adjustments. New York: Wiley.
- Gelman, A., S. Goel, D. Rivers, and D. Rothschild (2016). The mythical swing voter. *Quarterly Journal of Political Science* 11(1), 103–130.
- Gelman, A. and J. Hill (2007). *Data analysis using regression and multilevel/hierarchical models*, Volume 1. New York: Cambridge University Press.
- Gelman, A. and G. King (1993). Why are American presidential election campaign polls so variable when votes are so predictable? *British Journal of Political Science* 23(04), 409–451.
- Gerber, A. S., G. A. Huber, D. Doherty, C. M. Dowling, and S. J. Hill (2013). Who wants to discuss vote choices with others? Polarization in preferences for deliberation. *Public opinion quarterly* 77(2), 474–496.
- Graefe, A. (2014). Accuracy of vote expectation surveys in forecasting elections. *Public Opinion Quarterly* 78(S1), 204–232.
- Groves, R. M. and L. Lyberg (2010). Total survey error: Past, present, and future. *Public opinion quarterly* 74(5), 849–879.
- Henderson, M. and D. S. Hillygus (2016). Changing the clock the role of campaigns in the timing of vote decision. *Public Opinion Quarterly*, nfw027.
- Hillygus, D. S. (2011). The evolution of election polling in the United States. *Public Opinion Quarterly* 75(5), 962–981. 10.1093/poq/nfr054.

- Hoek, J. and P. Gendall (1997). Factors affecting political poll accuracy: An analysis of undecided respondents. *Marketing Bulletin, Department of Marketing Massey University* 8, 1–14.
- Hoffman, M. D. and A. Gelman (2014). The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15(1), 1593–1623.
- Holt, D. T. and T. F. Smith (1979). Post stratification. *Journal of the Royal Statistical Society. Series A (General)*, 33–46.
- Hopkins, D. J. (2009). No more Wilder effect, never a Whitman effect: When and why polls mislead about black and female candidates. *The Journal of Politics* 71(3), 769–781.
- Huffington Post (2009). State archive: National polls. <http://www.pollster.com/polls/us/>. Accessed: 2016-12-20.
- Huffington Post (2016a). 2016 president forecast. <http://elections.huffingtonpost.com/2016/forecast/president>. Accessed: 2016-12-01.
- Huffington Post (2016b). Pollster API v2. <http://elections.huffingtonpost.com/pollster/api/v2>. Accessed: 2016-12-20.
- Irwin, G. A. and J. J. Van Holsteyn (2008). What are they waiting for? Strategic information for late deciding voters. *International Journal of Public Opinion Research* 20(4), 483–493.
- Jacobs, L. R. and R. Y. Shapiro (2005). Polling politics, media, and election campaigns. *Public Opinion Quarterly* 69(5), 635–641. 10.1093/poq/nfi068.
- Jennings, W. and C. Wlezien (2016). The timeline of elections: A comparative perspective. *American Journal of Political Science* 60(1), 219–233.
- Jowell, R., B. Hedges, P. Lynn, G. Farrant, and A. Heath (1993). Review: The 1992 British election: The failure of the polls. *The Public Opinion Quarterly* 57(2), 238–263.
- Keeter, S. and R. Igielnik (2016). Can likely voter models be improved? Evidence from the 2014 U.S. House elections.
- Keeter, S., C. Miller, A. Kohut, R. M. Groves, and S. Presser (2000). Consequences of reducing nonresponse in a national telephone survey. *Public opinion quarterly* 64(2), 125–148.
- Kennedy, C., M. Blumenthal, S. Clement, J. D. Clinton, C. Durand, C. Franklin, K. M. andLee Miringoff, K. Olson, D. Rivers, L. Saad, E. Witt, and C. Wlezien (2017, May). An evaluation of 2016 election polls in the U.S. Technical report, American Association for Public Opinion Research.

- Kosmidis, S. and G. Xezonakis (2010). The undecided voters and the economy: Campaign heterogeneity in the 2005 British general election. *Electoral Studies* 29(4), 604–616.
- Leigh, A. and J. Wolfers (2006). Competing approaches to forecasting elections: Economic models, opinion polling and prediction markets. *Economic Record* 82(258), 325–340.
- Leip, D. (2008). Atlas of US Presidential elections. <http://uselectionatlas.org/>. Accessed: 2016-12-20.
- Lock, K. and A. Gelman (2010). Bayesian combination of state polls and election forecasts. *Political Analysis*, 337–348.
- Martin, E. A., M. W. Traugott, and C. Kennedy (2005). A review and proposal for a new measure of poll accuracy. *Public Opinion Quarterly* 69(3), 342–369.
- Mitofsky, W. J. (1998). Review: Was 1996 a worse year for polls than 1948? *The Public Opinion Quarterly* 62(2), 230–249.
- Nandram, B. and J. W. Choi (2008). A Bayesian allocation of undecided voters. *Survey Methodology* 34(1), 37–49.
- Orriols, L. and . Martnez (2014). The role of the political context in voting indecision. *Electoral Studies* 35, 12–23.
- PRC, Pew Research Center (2012). Assessing the representativeness of public opinion surveys. <http://www.people-press.org/2012/05/15/assessing-the-representativeness-of-public-opinion-surveys/>.
- Princeton Election Consortium (2016). About PEC and the meta-analysis (faq). <http://election.princeton.edu/faq/>. Accessed: 2016-12-01.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- RealClearPolitics (2004). 3-way race: Bush/Cheney vs Kerry/Edwards vs Nader/Camejo. http://www.realclearpolitics.com/bush_vs_kerry.html. Accessed: 2016-12-20.
- Rogers, T. and M. Aida (2014). Vote self-prediction hardly predicts who will vote, and is (misleadingly) unbiased. *American Politics Research* 42(3), 503–528.
- Rothschild, D. (2015). Combining forecasts for elections: Accurate, relevant, and timely. *International Journal of Forecasting* 31(3), 952–964.
- Rothschild, D. and S. Goel (2016). When you hear the margin of error is plus or minus 3 percent, think 7 instead. <https://www.nytimes.com/2016/10/06/upshot/when-you-hear-the-margin-of-error-is-plus-or-minus-3-percent-think-7-instead.html>. The Upshot, New York Times. Posted online 2016-10-5. Accessed: 2016-12-10.

- Shirani-Mehr, H., D. Rothschild, S. Goel, and A. Gelman (2018). Disentangling bias and variance in election polls. *Journal of the American Statistical Association*, forthcoming. <http://www.stat.columbia.edu/~gelman/research/published/polling-errors.pdf>. Accessed: 2018-02-15.
- Silver, B. D., B. A. Anderson, and P. R. Abramson (1986). Who overreports voting? *American Political Science Review* 80(02), 613–624.
- Silver, N. (2016). A user’s guide to fivethirtyeight’s 2016 general election forecast. <https://fivethirtyeight.com/features/a-users-guide-to-fivethirtyeights-2016-general-election-forecast/>. Accessed: 2016-12-01.
- Stan Development Team (2017). ShinyStan: Interactive visual and numerical diagnostics and posterior analysis for Bayesian models. <http://mc-stan.org/>. R package version 2.4.0.
- Stan Development Team (2018). RStan: the R interface to Stan. <http://mc-stan.org/>. R package version 2.17.3.
- Sturgis, P., B. Nick, C. Mario, F. Stephen, G. Jane, W. Jennings, K. Jouni, L. Ben, and S. Patten (2016). Report of the inquiry into the 2015 British general election opinion polls.
- Tourangeau, R. and T. J. Plewes (2013). *Nonresponse in social science surveys: A research agenda*. http://www.nap.edu/openbook.php?record_id=18293: National Academies Press.
- Traugott, M. W. (2005). The accuracy of the national preelection polls in the 2004 presidential election. *Public Opinion Quarterly* 69(5), 642–654. 10.1093/poq/nfi061.
- Veiga, F. J. and L. G. Veiga (2004). The determinants of vote intentions in portugal. *Public Choice* 118(3), 341–364.
- Visser, P. S., J. A. Krosnick, J. Marquette, and M. Curtin (2000). Improving election forecasting: Allocation of undecided respondents, identification of likely voters, and response order effects. *Election polls, the news media, and democracy*, 224–260.
- Voss, D. S., A. Gelman, and G. King (1995). A review: Preelection survey methodology: Details from eight polling organizations, 1988 and 1992. *The Public Opinion Quarterly* 59(1), 98–132.
- Wlezien, C., W. Jennings, S. Fisher, R. Ford, and M. Pickup (2013). Polls and the vote in Britain. *Political Studies* 61, 66–91.
- Zaller, J. and S. Feldman (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American journal of political science*, 579–616.

Table 3. Priors used in models for analysis of state polls^{||}.

Model	Component	Prior	Hyper-prior	
			Mean	Variance
Polling	Mean	$\alpha_{1r} \sim \mathcal{N}(\mu_{1\alpha}, \sigma_{1\alpha}^2)$	$\mu_{1\alpha} \sim \mathcal{N}(0, 0.2)$	$\sigma_{1\alpha} \sim \mathcal{N}_+(0, 0.2)$
		$\beta_{1r} \sim \mathcal{N}(\mu_{1\beta}, \sigma_{1\beta}^2)$	$\mu_{1\beta} \sim \mathcal{N}(0, 0.2)$	$\sigma_{1\beta} \sim \mathcal{N}_+(0, 0.2)$
	Variance	$\gamma_g \sim \mathcal{L}(0, 0.05)$	$\mu_\kappa \sim \mathcal{N}(0, 0.05)$	$\sigma_\kappa \sim \exp(1/0.05)$
		$\kappa_h \sim \mathcal{N}(\mu_\kappa, \sigma_\kappa^2)$		
		$\tau_{1r}^2 \sim \mathcal{N}_+(0, \sigma_{1\tau}^2)$	$\sigma_{1\tau} \sim \mathcal{N}_+(0, 0.05)$	
Undecided voters	Mean	$\alpha_{2r} \sim \mathcal{N}(\phi_{y[r]}, \sigma_{2\alpha}^2)$	$\phi_y \sim \mathcal{N}(0.04, 0.01)$	$\sigma_{2\alpha} \sim \mathcal{N}_+(0, 0.02)$
		$\beta_{2r} \sim \mathcal{N}(\mu_{2\beta}, \sigma_{2\beta}^2)$	$\mu_{2\beta} \sim \mathcal{N}(0, 0.02)$	$\sigma_{2\beta} \sim \mathcal{N}_+(0, 0.02)$
	Variance	$\tau_{2r}^2 \sim \mathcal{N}_+(0, \sigma_{2\tau}^2)$		$\sigma_{2\tau} \sim \mathcal{N}_+(0, 0.01)$

A. Prior distributions for undecided voter analysis

The priors used for the polling model, in Table 3, are identical to those from Shirani-Mehr et al. (2018) except for the undecided voter allocation, γ_g , and bias from house effects, κ_h , which are new to the model.

The (non-hierarchical) prior probability distribution for γ_g is Laplacian (or double-exponential) as we wish to shrink the estimates towards zero where there is insufficient evidence to suggest a true effect exists. We base this prior on Figure 1 where we observe only some categories appearing to demonstrate a relationship between error and the level of undecided voters. The undecided voters on election day, α_{2r} , is scaled in the polling model by 10, so that its value is approximately on the same scale as t_i . This is done for computational reasons, but as a result the prior scale assigned to γ_g (0.05) is comparable to the scales of other priors. We set its value to be a quarter of the original parameters prior scales to induce further shrinkage.

The house effects bias, κ_h , is specified with a hierarchical prior, similarly to the other hierarchical parameters in the model. Some polling agencies have few polls included in the sample and we would like the parameter estimates to borrow strength across the agencies – just as in the state-level election day bias parameter, α_{1r} . However, we assign an exponential distribution to the standard deviation, σ_κ to induce further shrinkage, relative to α_{1r} and β_{1r} . Implicitly, we are assuming that some pollsters actually have negligible bias due to their practices so we reflect this in our priors. The scale parameters of this prior are also a quarter of the original parameters prior scales.

The priors for the undecided model are also hierarchical where appropriate. We assign variance levels so that only weak information over the observed undecided voter proportion range (between approximately 0 and 0.2%) is provided. The prior for α_{2r} has mean, ϕ_y , that varies across year. This is included because the mean of the model is uncentered, unlike the polling model which uses v_r , and we observe different levels of undecided voters across years. The prior for ϕ_y , the mean level of undecided voters in each year, is weakly-informative but centred on 0.04. This value is the approximate mean observed in Figure 2.

^{||}The notation $\mathcal{N}_+(\mu, \sigma^2)$ denotes half-normal distribution, while $\mathcal{L}(\mu, \sigma)$ represents the Laplacian or double-exponential distribution.

Table 4. Descriptions and definitions of poll-level bias quantities used to calculate election-level bias (and averages) on percent scale.

Name	Sym.	Poll-level metric
Model (2)		
All bias	b_r	$\text{logit}(p_i) = \text{logit}(v_{r[i]}) + \alpha_{1r[i]} + t_i \beta_{1r[i]}$
Election day bias	b_r^e	$\text{logit}(p_i^e) = \text{logit}(v_{r[i]}) + \alpha_{1r[i]}$
Model (7)		
All bias	b_r	$\text{logit}(p_i) = \text{logit}(v_{r[i]}) + \alpha_{1r[i]} + t_i \beta_{1r[i]} - \alpha_{2r[i]} \gamma_{g[i]} + \kappa_{h[i]}$
Election day bias	b_r^e	$\text{logit}(p_i^e) = \text{logit}(v_{r[i]}) + \alpha_{1r[i]} - \alpha_{2r[i]} \gamma_{g[i]} + \kappa_{h[i]}$
Undecided voter bias	b_r^u	$\text{logit}(p_i^u) = \text{logit}(v_{r[i]}) - \alpha_{2r[i]} \gamma_{g[i]}$
House effects	b_r^h	$\text{logit}(p_i^h) = \text{logit}(v_{r[i]}) + \kappa_{h[i]}$

B. Mean election-level bias and variance calculations

The definitions of mean bias and variance quantities presented in Table 1, Table 2, and Figure 5 are taken from Shirani-Mehr et al. (2018) where appropriate. We restate them here, and add additional definitions for bias from undecided voters and house effects. For a given election race, the average bias, b_r , is defined as

$$b_r = \frac{1}{|S_r|} \sum_{i \in S_r} (p_i - v_r) \quad (8)$$

where S_r is the set of polls from state-level election r , $|S_r|$ denotes the size of the set S_r , and p_i is taken from either model (2) or (7) depending on which is in use. Calculating the mean of $|b_r|$ over a given set of elections determines the average election-level absolute bias over that set. The set of races, for example, could be one election year (e.g. all state-level races in 2004), or all of the elections in the dataset.

The bias on election day can be calculated by setting $t = 0$ in p_i and calculating the same quantity above. More specifically

$$b_r^e = \frac{1}{|S_r|} \sum_{i \in S_r} (p_i^e - v_r) \quad (9)$$

where $\text{logit}(p_i^e) = \text{logit}(v_{r[i]}) + \alpha_{1r[i]} - \alpha_{2r[i]} \gamma_{g[i]} + \kappa_{h[i]}$. Similarly, bias attributable to certain components or excluding components can be calculated by setting the unwanted parameters in p_i to zero. The full list of bias measures are in defined at the poll-level Table 4. As in Equations (8) and (9), the election-level bias for each race r can be calculate by taking the mean over the polls of that election.

Finally, the average absolute election-level bias, $\mu_{b,S}$, of type b , over subset of elections S , can be calculated by

$$\mu_{b,S} = \frac{1}{|S|} \sum_{s \in S} |b_s|. \quad (10)$$

The average election-level polling standard deviation is defined similarly by

$$\sigma_r = \frac{1}{|S_r|} \sum_{i \in S_r} \sqrt{\frac{p_i(1-p_i)}{n_i} + \tau_{1r}^2} \quad (11)$$

and the average over a set of elections can be calculated based on this quantity also.

Table 5. Average house effects across elections. Only those polling agencies with mean posterior greater than 0.5% are shown.

Polling agency or group	Posterior	
	mean	s.d.
ARG	0.61	0.35
CNN	0.50	0.48
Gravis Marketing	1.01	0.33
Grove Insight	-0.67	0.49
JZ Analytics / Newsmax	-0.69	0.54
Lucid / The Times Picayune	-1.23	0.49
Mason Dixon	1.27	0.29
Monmouth University	-0.51	0.55
Rasmussen	0.95	0.22
Remington Research Group / AxiomStrategies	1.64	0.35
Strategic Vision	1.27	0.31
University of Cincinnati	0.58	0.53
University of New Hampshire	-1.28	0.53
University of Wisconsin	-1.06	0.59
UPI/CVOTER	0.69	0.32
Zogby	0.60	0.46

C. Bias from house effects

The house effect for a given polling organisation is the bias attributable to their specific polling methods and practices. Unfortunately, we are only able to capture an aggregate of house polling bias, because polling agencies only release a small amount of information on the assumptions behind their polls. Of the total number of polls, 85% have an associated house effect estimated. Polls without an associated house effect are due to insufficient poll numbers from the pollster.

Figure 7 shows some interesting trends in partisanship and polling. Overall, there is a small skew toward pollsters being biased in favour of Republican candidates, and far more of these Republican-biased agencies are decisively biased. Rasmussen, Gravis Marketing, Mason-Dixon, Strategic Vision, and Remington Research Group/Axiom Strategies have estimated average bias between 0.95% and 1.64% on average (Table 5). On the other hand, the University of New Hampshire and Lucid/The Times Picayune appear to have substantial bias in favour of the Democratic candidate – over 1.2% is shown in Table 5. The wider credible intervals, relative to their Republic-bias counterparts, may be due to a lower number of polls included in the data set (18 and 13 respectively, versus between 24 and 198 polls per group).

Table 5 shows selected estimates of the overall bias, b_h , attributable to pollster h on the percentage scale. These values are calculated by

$$b_h = \frac{1}{|S_h|} \sum_{i \in S_h} (p_i^h - v_{r[i]}) \quad (12)$$

where S_h is the set of polls taken by pollster h , p_i^h is the poll-level house effect bias (see

Table 4), and $v_{r[i]}$ is the election result in race r associated with poll i .

The house effect bias of a particular agency may be unintentional and unassociated with partisanship. For example, the relative bias toward Democratic candidates by Survey Monkey could be a result of their online sampling frame which is likely skewed toward a younger demographic. Most, if not all, organisations use methods to mitigate this bias but it is not always possible when limited demographics are sampled.

The pollsters with arguably the best performance are youGov and Public Policy Polling. However, both have a relatively high number of polls included in the dataset, 63 and 106 each. The average number of polls per pollster is 41, and range of polls is 8 to 184.

D. Modelling with an even allocation of undecided voters

In order to embed an even allocation of undecided voters into model (7), we make the following changes.

- (a) We remove the polls that do not report an undecided voter number as, unlike the proportional case, we cannot infer the even allocation from the data. Following their removal, 1,725 polls remain in the dataset.
- (b) The poll represented in (1) changes so that the poll values are scaled to remove third party candidates then Republicans are allocated 50% of the undecided voters. Formally, this can be represented most succinctly by

$$y_i = \frac{R_i + 0.5 \times U_i}{R_i + D_i + U_i}.$$

- (c) The undecided voter allocation with uncertainty, λ in (4), changes to $\lambda = 0.5 + \theta_i$.

Given the updated definitions of y_i and λ , using (3) we note that the identity in (5) still holds and modelling can proceed as in the case of a proportional allocation of undecided voters.

The analysis with an even allocation of undecided voters was very similar to the original results in Section 5. In general, most measures of bias across the years slightly increased, and the average standard deviation in the model increased. This is indicative of a bias-variance trade-off in the modelling process. Overall, summary measures for elections years 2004–2016 had an increase in the bias of 0.1% (average absolute bias and election day bias), and a decrease in the standard deviation of 0.1%. Focussing on 2016, average election day bias increased by 0.2%, whilst the average standard deviation decreased by 0.2%.

The effect of undecided voters on the polling bias in the model (γ_g) is replotted in Figure 8 with credible intervals. This plot is largely unchanged and still shows that undecided voters in the 2016 election year had a considerably different effect on bias compared to previous years.

The histogram of the absolute bias from undecided voters is replotted for each state-level election in Figure 9 under the assumption of even allocation of undecided voters. It shows a similar story as before. With careful comparison to Figure 5, one can see that bias has increased for many of the states.




Fig. 7. Credible intervals of 95% (outer line) and 50% (inner line) for the house effects bias from polling organisations in the model ($\kappa_{i,j}$), on the logit scale. A positive value indicates a bias in favour of the Republican candidate, hence underperformance in the election result.




Fig. 8. Credible intervals of 95% (outer line) and 50% (inner line) for the effect of undecided voters on polling bias in the model (γ_g) , on logit scale. A positive value indicates a bias away from an even allocation of undecided voters in favour of the Republican candidate.

Table 6. Average election-level absolute bias and average election-level standard deviation across state-elections in given year(s) from model (7) with assumption of even allocation of undecided voters.

	2004	2008	2012	2016	<i>Overall</i> 2004–2016
Average absolute bias	0.8% (0.11)	1.1% (0.10)	1.4% (0.11)	2.8% (0.10)	1.8% (0.06)
Average absolute election day bias	0.8% (0.12)	1.0% (0.12)	1.4% (0.14)	2.6% (0.12)	1.7% (0.07)
Average absolute undecided voter bias	0.4% (0.19)	0.5% (0.21)	1.0% (0.30)	2.3% (0.24)	1.3% (0.11)
Average absolute house effects	0.6% (0.17)	0.4% (0.14)	0.3% (0.12)	0.2% (0.06)	0.4% (0.09)
Average standard deviation	2.1% (0.04)	2.1% (0.03)	2.0% (0.03)	2.2% (0.04)	2.1% (0.03)
Average election day undecided	3.3% (0.24)	3.8% (0.21)	3.0% (0.21)	5.5% (0.28)	4.2% (0.14)




Fig. 9. Histograms of the average absolute bias from undecided voters for each state-level election, separated by year. The bias from undecided voters is the quantity $\alpha_{2r}\gamma_g$ in the model. A positive value indicates a bias away from an even allocation of undecided voters in favour of either candidate.