



Learning from Polls During Electoral Campaigns

Lukas F. Stoetzer¹ · Lucas Leemann² · Richard Traunmueller³

Accepted: 16 November 2022
© The Author(s) 2022

Abstract

Voters' beliefs about the strength of political parties are a central part of many foundational political science theories. In this article, we present a dynamic Bayesian learning model that allows us to study how voters form these beliefs by learning from pre-election polls over the course of an election campaign. In the model, belief adaptation to new polls can vary due to the perceived precision of the poll or the reliance on prior beliefs. We evaluate the implications of our model using two experiments. We find that respondents update their beliefs assuming that the polls are relatively imprecise but still weigh them more strongly than their priors. Studying implications for motivational learning by partisans, we find that varying adaptation works through varying reliance on priors and not necessarily by discrediting a poll's precision. The findings inform our understanding of the consequences of learning from polls during political campaigns and motivational learning in general.

Keywords Pre-election polls · Electoral campaigns · Bayesian learning · Motivated reasoning

Introduction

Learning about political candidates and parties before voting is central to the foundations of democracy (Lau & Redlawsk, 2006; Lodge et al., 1995; Popkin, 2020). Citizens that use available information to learn the things that they need to know to make informed decisions are described as the democratic ideal (Bartels, 1996). For this, pre-election polls are one important source that informs voters' beliefs about

We thank Rob Johns, Patrick Kraft, Suthan Krishnarajan, Jordi Muñoz, Marco Steenbergen, and Bernhard Weßels for valuable comments on the paper. Earlier version was presented at the EPSA annual conference in Belfast 2019, the SVPW virtual annual conference 2021, the Berlin Political Behaviour Workshop 2021 and the workshop of the political methodology chair of the University of Zurich at Villa Garbald, 2019.

✉ Lukas F. Stoetzer
lukas.stoetzer@uni-wh.de
<http://lukas-stoetzer.org/>

Extended author information available on the last page of the article

uncertain political outcomes, like the support for parties, candidates, policies, or referendums. During election campaigns a stream of multiple polls creates a rich and dynamic information environment about the electorates' evolving support (Jennings & Wlezien, 2018). Through voters' learning, the information from these polls impacts their voting decisions and ultimately determines the success of candidates and parties on election day (Dahlgaard et al., 2017; Großer & Schram, 2010; Marsh, 1985; Moy & Rinke, 2012; Rothschild & Malhotra, 2014).

How do voters process and interpret information from the regularly published polls during electoral campaigns? Recent literature that studies voters' perception of polls, documents a motivated information processing (Kuru et al., 2017; Madson & Hillygus, 2020; Tsfati, 2001). This notion follows a growing literature on motivated learning, in which individuals accept information that supports their motivational aims and refute contradictory information (see e.g., Jerit & Barabas, 2012; Taber & Lodge, 2006). For example, Kuru et al. (2017) show that respondents judge a poll on policy issues as less credible when its results do not align with their pre-existing views. Similarly, Madson and Hillygus (2020) present experimental evidence in the context of the 2016 US elections that reveals that supporters of a candidate judge polls as more credible when their preferred candidate is ahead in the poll. This evidence stands at odds with the democratic ideal, as motivated citizens arrive at vastly different beliefs because they judge the available information differently. Most of these studies, however, focus on voter's judgement of polls and not on what they learn from them. In addition, they usually only consider a single poll.

In reality, voters may use information from multiple polls to inform their beliefs over the course of an election campaign. An open question, hence, is if and how these motivational biases persist in belief formation over time.

In this paper, we present a new model of how citizens learn and update their beliefs when they are introduced to new information over multiple time points. A natural starting point for how voters learn from poll results is Bayesian updating, where current beliefs are formed as a weighted combination of prior beliefs and new information (e.g., Bartels, 2002; Bullock, 2009; Gerber & Green, 1999; Guess & Coppock, 2020; Hill, 2017; Sinclair & Plott, 2012). Our Bayesian learning model is novel in that it deviates from previous Bayesian learning models in political science which almost exclusively assume (a) that both prior beliefs and evidence are normally distributed (yielding a normal posterior belief) and, importantly, (b) that citizens are trying to learn about political conditions that do not change over time (Bullock 2009). Since polling results usually refer to population shares and of course *do* change over the course of a campaign, we properly constraint citizens' beliefs distributions to the unit interval (0–1) and introduce a *dynamic* Bayesian learning model.

Our dynamic Bayesian learning model allows us to characterize the learning in terms of the *rate of adaptation* to new polling evidence. The rate of adaptation answers the question of how closely a voter updates her prior expectation to a new poll result and naturally allows for dynamic learning over the entire election campaign. This dynamic learning perspective uncovers two sources for varying adaptation rates: the perceived precision of the poll and the 'stickiness' of prior beliefs. Voters' learning depends on how accurately they judge a single poll to be in estimating actual population support, and how much information from their priors they

carry over time. The model allows studying the consequences of the two sources for learning over the entire campaign. It shows that only the combination of high stickiness to priors and very low subjective precision of the poll leads to belief persistence and reluctant updating over the entire campaign. Smaller changes wash out as more and more polls become available.

The model has important implications for our understanding of motivated learning from polls over the electoral campaign cycle. We argue that when confronted with new polling information that contradicts citizens' motivational interests, they may only reluctantly and slowly update their beliefs by either: (a) adjusting the perceived precision of the information source or (b) by relying more heavily on their prior beliefs. The perceived precision of the poll in our model speaks to the motivational perceptions of polls in prior studies (Kuru et al., 2017; Madson & Hillygus, 2020). For example, voters, and especially partisans, might judge a poll to be imprecise because it predicts their candidate losing the electoral race. The 'stickiness' of previously held prior beliefs yields an additional and novel source of motivational learning from polls. For example, partisans might more strongly carry over their prior beliefs about the chances of their candidate winning the race, if polls report unfavorable results for their candidate.

We evaluate the implications of our dynamic learning model from polls using two experiments. In the first experiment, we randomly instill different prior beliefs and sequentially present respondents with three sets of changing polling results from a hypothetical two-party electoral race. The evidence from this first experiment suggests that our learning model provides a reasonable description of how voters adjust their expectations about the electoral race to new polling information. We find that respondents update their beliefs assuming that the polls are relatively imprecise but still weigh them more strongly than their priors. With this, different prior beliefs matter, but respondents adapt to new evidence and tend to converge in their beliefs over time.

In the second experiment, we investigate if the learning process differs based on partisan motives and can result in diverging partisan beliefs about the political campaign. In this experiment, we manipulate the source (MSNBC vs. Fox) and the winning candidate of the sequence of the polls (Democratic or Republican). The results confirm the motivated processing of polling information, which is mainly due to stronger 'stickiness' of prior beliefs. We find that Democrats are particularly responsive to new information when they see their candidate winning, while in situations where their candidate is losing they are more reluctant. The results further indicate that this is not because they judge the results of favorable polls to be more precise, but rather because they are more likely to deviate from their prior beliefs. In line with prior findings, for example Kuru et al. (2017), we do not find clear source effects. This indicates that motivated learning in the context of polls seems to work through quick updating and not necessarily by discrediting the precision of the poll results due to the source.

The findings inform our understanding of the long-run consequences of learning from pre-election polls during political campaigns. One important implication of the finding that respondents weigh new poll information more strongly than their priors is that with polls being published almost daily, voters will *eventually* fully

update their beliefs about the support for candidates and parties. The initial differences in prior beliefs can wash-out if the support stays stable. However, some voters may need information from a larger number of polls to arrive at the same conclusion, especially when the information contradicts their preferences. The results support the notion that citizens may be ‘cautious’ Bayesian updaters (Hill, 2017), that adapt their ‘cautiousness’ to their motivational goals. The evaluation of poll effects on elections is important for two reasons. First, effects that work via changes in voters’ beliefs require some time to crystallize. Second, the process by which voters update their beliefs is heterogeneous, as some voters will jump to conclusions, while others are much more hesitant to form conclusions. Taken together these results have broader implications for citizens’ political behavior and ultimately the democratic process. If voters accurately process the information provided in political polls and make their political decisions accordingly, polls can play an active role in shaping democratic outcomes. If, however, citizens ignore or distort polling results according to their political needs, there is little room for polling results to inform political decisions and influence the democratic process. Our results suggest a nuanced middle ground between these two perspectives. In accordance with the democratic ideal, voters are generally responsive to political information, but retain a healthy dose of skepticism if it contradicts what they believe to be true.

Our research also provides important contributions beyond the specific context of learning from polls. First, it speaks to debates between motivated reasoning and Bayesian learning. While Bayes’ Rule is regarded as optimal to study learning from political information, a large number of studies document that political information processing is not a rational act, but prone to a host of political and other biases (e.g., Kahan, 2012; Lebo & Cassino, 2007; Leeper & Slothuus, 2014; Lord et al., 1979; Redlawsk, 2002; Redlawsk et al., 2010; Steenbergen & Howard, 2018; Taber & Lodge, 2006). Our framework demonstrates that motivated biases can be fruitfully integrated into Bayesian Learning models, by allowing learning parameters of the model, such as the perceived precision of information and the reliance on prior beliefs, to vary in line with motivational reasons. Thereby, it speaks to recent publications that discuss connections between the two (Bullock, 2009; Little, 2021). Second, and to the best of our knowledge, we present the first dynamic Bayesian learning model of its kind in political science.¹ Since our polling results of party support refer to population shares, we implement a dynamic Beta-Binomial model to properly constraint the beliefs to the unit interval and account for the fact that voter support fluctuates and changes over time. Finally, our research offers a novel experimental design and estimation procedure to study dynamic learning by integrating belief elicitation questions (Leemann et al., 2021). This permits the evaluation of additional implications, as the theoretical implications of the Bayesian learning

¹ Hill (2017) presents a Bayesian learning model of political facts with multiple rounds of new evidence but where the facts are both, discrete events (i.e., true or false) and fixed in time. Bullock (2009) discusses the implications of a dynamic Bayesian learning model using theoretical simulations, but only for the Normal-Normal case.

model not only concern the expectation, but also the variance of the beliefs (Little, 2021).

Bayesian Learning from Polling Results

To understand how voters learn from polls, it is useful to represent their beliefs using a probability distribution. This probability distribution describes both the expectation a voter has about a candidate's support (i.e., in terms of the mean) and how certain she is about this quantity (i.e., in terms of the variance). A candidate's support at a particular point during the campaign, denoted as θ_t , is a share and therefore theoretically ranges from 0 to 1. Importantly, this quantity is not fixed but dynamic, i.e., it will change over the course of a campaign and is therefore subscripted with t . Voter i 's belief about this support is represented by the probability distribution $p_i(\theta_t)$. As the belief and thereby the distribution differs between individuals, the distribution is denoted with the subscript i . The question we set out to answer is how polls y_t that publish a vote share for the candidate influences those beliefs. During the election campaign more and more polls will get available and influence beliefs sequentially.²

The Learning Model

Voters can use Bayes Rule to update their beliefs about a candidate's support. Bayes Rule formalizes this learning process as follows:

$$p_i(\theta_t|y_t) = \frac{p_i(y_t|\theta_t)p_i(\theta_t)}{p_i(y_t)}. \quad (1)$$

Before observing a poll the *prior belief* is $p_i(\theta_t)$. After observing the poll each voter can form a *posterior belief*, $p_i(\theta_t|y_t)$, by conditioning their belief on the evidence from the poll. This evidence is expressed as the subjective *likelihood* $p_i(y_t|\theta_t)$. It reflects the distribution which—in the eye of the voter—is likely to have generated the poll. To put it differently, a poll that shows the approval rate of a candidate, for example, at 55% is more likely to result from a population where 55% instead of 40% support the candidate.³

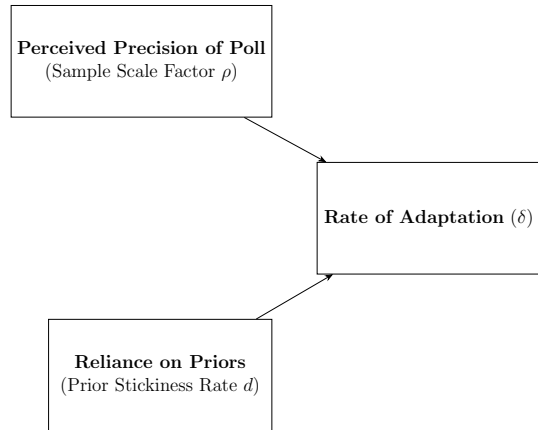
A Bayesian learning model informs different quantities about how voters learn from polls. The central question is how closely voters will adapt their expectation to the polls. We define this as the *rate of adaptation*. In particular, we are interested how strongly the prior expectation is shifted in direction of the new evidence.⁴

² Here we work with a sequence of polls that all voters receive. We discuss potential extensions and implication of selective exposure at the end of the article.

³ $p(y_t)$ is the probability of the poll that normalizes the product of the likelihood and the prior. It can be calculated by integrating over the possible vote shares for the candidate $p_i(y_t) = \int_0^1 p_i(y_t|\theta_t)p_i(\theta_t)d\theta_t$.

⁴ This is sometimes referred to as the delta rule (see e.g., Nassar et al., 2010)

Fig. 1 Reliance on priors (prior stickiness rate) and perceived precision of poll (sample scale factor) influence dynamic learning and the rate of adaptation



Denoting the expectation of the *posterior belief* as $\mu_{it} = \mathbb{E}[p_i(\theta_t|y_t)]$ we can write the adaptation equation as

$$\mu_{it} = \mu_{it-1} + \delta_{it}(y_t - \mu_{it-1}). \quad (2)$$

Here, the posterior expectation is equal to the prior expectation plus the adaptation rate δ_t times the difference between the new evidence and the prior expectation. The adaptation rate is defined between 0 and 1. A value of 0 means no adaptation to the new poll and value of 1 implies perfect adaptation.

In our Bayesian learning model the rate of adaptation depends on two aspects: (a) the perceived precision of the poll and (b) the stickiness of prior beliefs. Figure 1 shows the conceptualization that the two influence the overall adaptation rate.

The first variable depicts *how informative or precise a voter perceives the poll to be* about the vote share. Relating survey shares to the support in the population comes with several sources of uncertainty (Weisberg, 2009). A relevant part is due to the fact that only a limited amount of people are asked about their vote intention, leaving room for sampling error. The true support for a candidate can fall within a margin of error of the poll result. However, recipients of the poll results might consider alternative factors that lower or increase the information value. The sampling error is only one source of the total survey error. For example, voters could consider the error sources of individual polls next to their random sampling variability, like weighting errors, survey effects, and house-effects. Taken these into account would certainly make voters more skeptical about the poll results. But citizens may also judge the poll to be more precise in estimating the true support and mistake the survey as a nearly complete census. For all of this, it is not necessary that voters understand the conceptualization of sampling variability, but only require a subjective perception of the precision attached to the poll. So, all those aspects can be encoded in the *perceived precision of the poll*—i.e., how voters think that the true support deviates from the poll result. We can calculate this as the standard deviation of the individual likelihood: $SD_i[y_t] = \sqrt{\text{VAR}[p_i(y_t|\theta_t)]}$

Second, the rate of adaptation also depends on *how strongly voters carry over their beliefs over time*. Over an electoral campaign multiple polls will be published, and a situation a Bayesian learning model can accommodate. In essence the *posterior belief* from the past period can form the new *prior belief* that is to be revised in light of the evidence. Note, that in our case simply ‘repeating’ a standard Bayesian model multiple times would not allow us to identify learning because it would not adequately describe the dependence between prior and posterior beliefs. Instead, we need to be more explicit about how these beliefs depend on each other and evolve over time:⁵ $p_i(\theta_i|\theta_{t-1})$.

Different specifications of this process can be used to analyze the belief evolution. In general, in a dynamic setting, we expect that beliefs from the time-period before are less informative about the support today than they were yesterday. This process can be modelled in different ways. In the application below we employ a power discount model (Smith, 1979). This model discounts the fact that information was received at the last point in time, which can lead to varying reliance on the prior beliefs in the learning process.

In sum, in our conceptualization Bayesian updating steps for learning from polls depend on the perceived precision of the poll and the reliance of prior beliefs. Unlike in other applications of Bayesian learning where both are fixed (Hill, 2017), this makes it difficult to say what an ideal Bayesian updating step will look like. Each learning step depends on the perceived precision of the poll and the reliance on prior beliefs. To infer these two parameters from data we specify a particular parametric model in the next section.

Parametric Model of Bayesian Learning from Polls

The analysis of the learning process requires certain assumptions about the probability distributions to represent voters’ beliefs. We choose a set of flexible distributions that permit for a conjugate learning process: a Beta-Binomial model.⁶ This model is well-suited for our purpose as it permits us to analyze the rate of adaptation, the perceived precision of the poll, and the reliance on priors.

Let y_t define the share of support for candidate A reported in poll at time t . The poll is based on a survey of N_t respondents. We assume that the *Likelihood* function voters have in mind when evaluating the poll is a binomial distribution, where the count of respondents that support a candidate is the product of the poll’s share and the number of survey respondents $y_t N_t$.

⁵ A dynamic Bayesian model differs from the standard Bayesian model in that the posterior belief formation is now essentially a two-step process. First, a ‘forecast step’ where the prior is combined with a likely change in the event or condition of interest, and second, an ‘error-correction step’ where this forecast is then adapted or modified in light of the new evidence (see Bullock 2009).

⁶ An alternative would be a dynamic linear learning model. However, the bounds of the shares are better suited to a beta model. In addition, the beta is more flexible in that it permits for non-symmetric beliefs. A more general form might be found when employing a dynamic generalized linear model, at the cost of conjugacy.

As argued above, respondents do not need a complete understanding of the sampling process, but they need to hold a perception of the precision of the poll. To estimate this perceived precision of the poll, we rely on a particular parametrization. We multiply the sample size with a scaling parameter ρ . While this parameter indicates if the size of the sample that voters use to update their beliefs is larger or smaller than the actual sample size under random samples, it permits us to estimate varying perceived precisions of the poll. In this model, the perceived standard error of a reported poll result is given by: $\frac{1}{\sqrt{\rho}} \sqrt{\frac{y_t(1-y_t)}{N_t}}$. Therefore, $\frac{1}{\sqrt{\rho}}$ tells us if recipients perceive the standard error of the poll as smaller or larger as a random sample would suggest. For this, we assume that all voters employ the same likelihood when evaluating the poll with the only parameter ρ that estimates the perceived precision of the poll:

$$y_t N_t \sim \text{Binomial}(\theta_t \mid \rho N_t) \tag{3}$$

We further assume that beliefs about the support for the candidate are beta-distributed as candidate support is bound between 0 and 1. The posterior belief from the last update $p_i(\theta_{t-1} \mid y_{t-1}, \rho N_{t-1})$ are distributed with $\alpha_{it-1}, \beta_{it-1}$. Those shape parameters of the beta distribution guide an individual’s belief about the support for the candidate:

$$\theta_{t-1} \sim \text{Beta}(\alpha_{it-1}, \beta_{it-1}) \tag{4}$$

The beliefs are carried over to form the new prior beliefs using a power discount model (Smith, 1979):

$$p(\theta_t) = p(\theta_{t-1} \mid y_{t-1}, \rho N_{t-1})^d \tag{5}$$

where the parameter d —which we label prior stickiness rate—specifies how much of the new prior depends on the past posterior belief. It ranges from 0 to 1, where one means that the beliefs carry over and 0 that the posterior is not taken into consideration for the new period. Again we assume a common prior stickiness rate for all learners and all periods.⁷

Posterior beliefs $p(\theta_t \mid y_t, \rho N_t)$ are formed according to Bayes rule (see supplementary material (SM) A.1) and result in conjugate beta-distributed beliefs:

$$\theta_t \sim \text{Beta}(\alpha_{it}, \beta_{it}) \tag{6}$$

where $\alpha_{it} = d\alpha_{it-1} + y_t \rho N_t$ and $\beta_{it} = d\beta_{it-1} + (1 - y_t) \rho N_t$ are both a function of the prior stickiness rate, the sample scaling parameter, the poll, and the prior beliefs. Hence, all those aspects shape a voter’s learning process.

⁷ The prior stickiness rate has also been elsewhere discussed under the term ‘cognitive conservatism’, as the strong weighting of prior beliefs (Edwards, 1982). As the term conservatism is loaded with reference to ideology, we use prior stickiness rate and reliance on priors throughout the text.

In essence, the perceived precision of the poll and the reliance on prior beliefs both influence how strongly voters adapt their expectation. To see this more clearly, we can formulate the posterior expectation from the beta distributed posterior beliefs as $\mu_{it} = \frac{d\alpha_{it-1} + y_t \rho N_t}{d(\alpha_{it-1} + \beta_{it-1}) + \rho N_t}$. If the prior beliefs do not carry over, and $d = 0$, then the mean expectations are equal to the poll's reported share y_t . If the survey is relatively small $\lim_{N \rightarrow 0}$, the posterior expectation is dominated by the prior belief expectation as the prior stickiness rate cancels out: $\frac{\alpha_{it-1}}{\alpha_{it-1} + \beta_{it-1}}$. The sample scale factor influences how small the poll is perceived. Hence, if ρ value is close to zero the beliefs are dominated by the priors.

Different Learning Patterns over Time

Based on this model, we can explore how learning unfolds over multiple polls. A first relevant quantity that we can determine is how fast a voter adapts her beliefs. In the model, the prior stickiness rate and the sample scale factor constitute the rate of adaptation. Substituting the expectations in the equation 2 and solving for d gives: $\delta_{it} = \frac{\rho N_t}{d(\alpha_{it-1} + \beta_{it-1}) + \rho N_t}$. The rate of adaptation hence depends on the sample size, the prior beliefs, and the two parameters that guide the learning process. Again, if $d = 0$ the rate of adaptation is 1 and the expectations are perfectly adjusted to the evidence. If the poll is small, or recipients have a small sample scaling factor, the rate of adaptation tends towards zero $\lim_{\rho N_t \rightarrow 0} \delta_{it} = 0$. We can simplify further by assuming equal size of the evidence that informed the priors and the poll size today ($\alpha_{it-1} + \beta_{it-1} = N_t$). This way we obtain the standardized rate of adaptation:

$$\delta^* = \frac{\rho}{d + \rho} \tag{7}$$

which shows how strongly respondents adapt to new information from the poll. It has a direct interpretation. A value of 1 means that voters perfectly update their expectation to new information. A value of 50% implies an equal weight on new information and information from the priors.

This has implications on learning over campaigns. Figure 1 shows how voters in the model can closely follow the polls, neglect them altogether, or something in between for two different reasons: Either because they do not perceive the poll to be accurate or because they heavily rely on prior beliefs. The different learning behavior is guided by the prior stickiness rate (d) and the sample scale factor (ρ). Figure 2 illustrates a case of voters who start with the same prior beliefs about a race. The voters believe that it is an open race with a 50% vote share for a candidate. Subsequently, they observe three polls. All polls interview 1000 people and indicate a vote share of 60% for the candidate. The voters differ in terms of the sample scale factor (which is either low (0.1) or high (0.9)) and the prior stickiness rate (which is either low (0.1) or high (0.9)).

The illustration highlights that the voter with a low prior stickiness rate, who does not rely heavily on prior beliefs, and high sample scale factor directly updates her expectation to the polls result. This is also true for the voter with a low prior

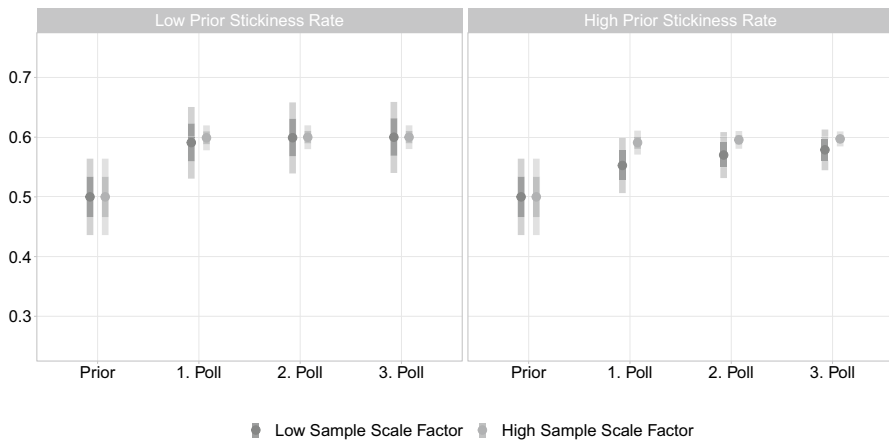


Fig. 2 The influence of the prior stickiness rate and the sample scale factor on learning from polls. The example shows three updating steps from a poll that reports a 60% vote share for a candidate based on 1000 people interviewed (hypothetical data)

stickiness rate and a low sample scale factor but the uncertainty is larger. In the case of a high prior stickiness rate, we again see that a voter with a high sample scale factor almost immediately updates the expectation. Only in the case with a high prior stickiness rate and a low sample scale factor after three polls we still see difference in the expectation and persistence of the prior beliefs. The voter here can be described as someone who judges the information from the polls to be imprecise and who puts a strong emphasis on the priors. One takeaway from this illustration is the importance of the interplay between the prior stickiness rate and precision in explaining reluctant or cautious learning. Only the combination of strong reliance on priors and low subjective precision of the poll leads to very slow learning about the expected support. A low sample scale factor only persists in the uncertainty around the expectation.⁸

Empirical Analysis

The theoretical model outlines two central aspects that influence learning over time: the perceived precision of the poll and the reliance on prior beliefs. The factors guide the way people learn from polls over political campaigns and how closely they update to new poll evidence. As it is difficult to infer the learning process from observational survey data, we conduct two survey experiments to evaluate the Bayesian learning model.⁹ The experiment in the first study simulates an electoral

⁸ Of course, these illustrative results can also be studied more formally. We provide some formalization that the expectation in the model converge to the polls in the SM A.2

⁹ SM B provides a discussion of the use of observational survey data to test the learning process.

campaign between two hypothetical parties, presents voters with a sequence of polls, and estimates the respondents' learning process.¹⁰ The experiment varies two conditions, the prior beliefs about a candidate chance and polling results. The results from the first experiment inform us about the learning process of respondents absent of any real-world cues. The second study extends this framework and includes partisan and poll-source cues. In the second study, we vary the source and the winning candidate (Democrat or Republican) to infer how partisans learn differently under the conditions. Partisan motivated learning is of particular interest because with a different learning process, beliefs about the electoral campaigns can deviate given the same information environment.¹¹

Study 1: Learning from Polls

Experimental Set-up

The first survey experiment presents a hypothetical election environment where party A and party B are competing in a district. The survey experiment provides respondents with election polls results regarding the vote share of party A and party B, based on which respondents can infer the winning chances of party A. Before respondents receive any polls, we instill prior beliefs about the vote share of party A in the district. We present respondents with 100 election outcomes of districts that are similar to the hypothetical district.¹² The survey experiment involves two groups that vary on their prior beliefs: One group that thinks that party A will win the race before seeing the first poll and one group that believes it will be a close race. More specifically, the first group sees a close race in the district, with a mean result for party A of 50%. The second group sees a clear race with a mean at 67% of party A vote share.¹³ Following the prior installment, we evaluate respondents' beliefs about the vote shares. Throughout the survey experiment, we rely on a set of questions from Manski (2009).¹⁴ In the next step, we present respondents with a sequence of

¹⁰ The sequential experimental exposure to new information has also been employed in the context of dynamic information boards (Lau & Redlawsk, 2006; Redlawsk et al., 2010). The dynamic information boards present a variety of information about a candidate including, issue positions, endorsements, background information, and preelection polls. Our experiment only presents preelection polls to study the learning process from polls.

¹¹ Replication materials for the analyses are available at: <https://doi.org/10.7910/DVN/IW4FP8>.

¹² SM C.3 shows four still frames from the GIF. Other research has successfully relied on and evaluated the same way to install priors in online surveys (Goldstein & David, 2014; Leemann et al., 2021)

¹³ For the vote share of party A in the close race, we draw 100 election results from a beta distribution with mean at 50% ($B(60, 60)$). for the clear race from a Beta distribution with mean at 67% ($B(60, 30)$).

¹⁴ This choice is based on a recent evaluation study that compares six different prior elicitation methods and shows that Manski's set of questions is best suited to elicit prior beliefs (Leemann et al., 2021). The Manski method encompasses five questions. It first asks about the "most likely" value, followed by a question about the "lower" and "upper" bound. The final two questions ask respondents about the probability that the value will be below or above the upper bound.

three polling results.¹⁵ After each poll result, respondents answer the Manski questions to elicit their beliefs about party A's vote share. This provides us with information about their posterior belief after having seen each poll and thereby allows us to infer the learning process. Respondents are further randomized into two conditions of either ascending or descending sequences of poll results for party A. Specifically, in the ascending condition, recipients see three polls which show party A support at 51%, the second poll at 54%, and the final poll at 58%. Consequently, the descending condition portrays decreasing support for party A within the polls (58%, 54%, and 51%).

Therefore, there are four experimental groups: (a) clear race prior and increasing polls; (b) clear race and decreasing polls; (c) close race and increasing polls; (d) close race and decreasing polls. Thus, some respondents receive the first polling information that aligns with their priors while others receive information that misaligns with their priors.

Statistical Modeling

We develop a statistical model for the sequence of belief elicitation questions that is based on the parametric learning model. Using the Manski elicitation method, we obtain three measures for each participant's belief distribution at each time point: The mean expectation, the lower bound, and the upper bound. In essence, our statistical models find the most likely average belief distribution for the observed answers to the Manski questions. The first model estimates the average beliefs of respondents at each time point and in each condition separately. For this, we model the measurements as a function of participants' mean beliefs and the respective quantiles (which assume to be beta distributed) and estimate the two shape parameters of the beta distribution using maximum likelihood procedure described in Leemann et al. (2021). The results of this model provide a description of the learning process without theoretical assumptions and how beliefs are informed by new polls. A second model then integrates the parametric learning model to estimate the learning process over time. For the second model, we only estimate the belief parameters of the priors freely and model the subsequent belief parameters using the theoretical learning process. This provides us with estimates of the two key parameters of the learning model, the prior stickiness rate d and the sample scale factor ρ . We can further calculate the standardized rate of adaptation. The SM A.3 describes the statistical model and the estimation procedure in detail.

¹⁵ The polls are introduced to the respondents using a small intro that is accommodated by a graphical depiction of the poll result. "In the electoral race between party A and party B, a new poll conducted in district D surveyed a total of 1000 voters." A bar chart presents the support for party A and party B, with the numerical value of support on top of each bar. The bar chart further includes the margin of error (+/- 3). See SM C.4 for the figure when polls are increasing for party A.

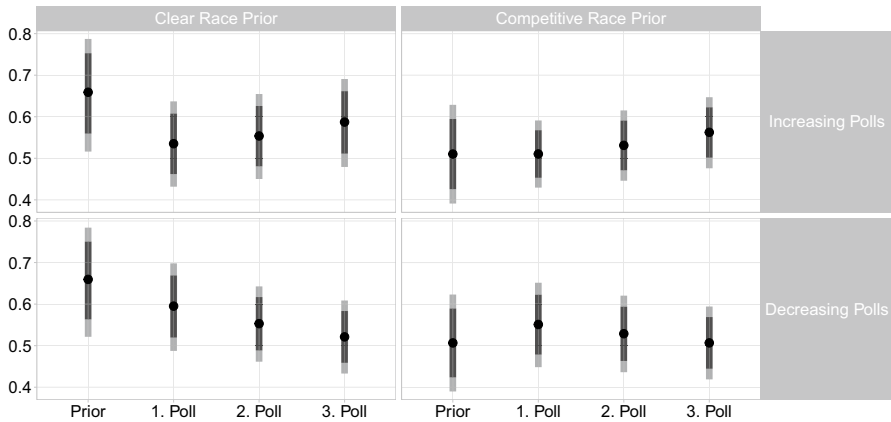


Fig. 3 Respondents’ beliefs about support for party A for different experiential scenarios and time-points. The beliefs are estimated for each time point and condition separately. The bars indicate 99% and 95% coverage of the beliefs, the points the expectations

Sample

We recruited 1388 respondents from the crowdsourcing platform Amazon Mechanical Turk (MTurk) to take part in our survey.¹⁶ The survey took place early April 2020. We follow standard practice and recruit workers on MTurk with an approval rating of more than 97% and more than 5000 HIT submissions (but see also Robinson et al., 2019). The median time for taking the survey was less than 6 mins and we paid participants \$1.10, as prescribed by the federal minimum wage. Descriptive statistics of the sample are provided in the SM D.1.¹⁷

Results

In this section we evaluate the learning process of the respondents for the different conditions. We first present results of respondents’ belief in the different conditions separately for each time point. The resulting belief distributions give a descriptive representation of the learning process. Figure 3 shows the intervals and expectation from the resulting beliefs. The figure shows that participants update their beliefs in line with the direction of the sequence of polls. More specifically, in the case of ascending poll results, the expected support for party A and B goes up, and vice versa for descending support. There is very little deviation between the prior belief and the updated expectation, with a mean absolute deviation of 0.9% percentage points. One implication of the two aspects is that the effect size of priors gets

¹⁶ MTurk can produce adequate samples, performs comparable results when compared to more established Internet surveys (Berinsky et al., 2012; Mason & Suri, 2012; Thomas & Clifford, 2017).

¹⁷ We find slightly more male participants (60%) in our sample compared to the US population, more educated (60% college degree versus 42%), but the average age aligns nicely with the average in the population (mean 40)

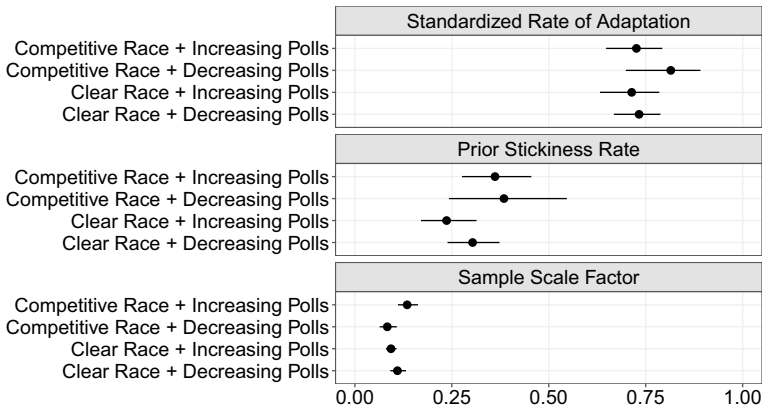


Fig. 4 Estimates for learning model from experimental study I

smaller over the three poll results in our experiment, but in some instances, it still affects belief updating (See SM D.2).

To study how the beliefs evolve we turn to the results from the parametric learning model. The estimates for the different scenarios are in Fig. 4. Across all scenarios, we estimate a standardized rate of adaptation of around 0.75, which indicates that there is similar learning behavior. This value summarizes the relatively quick adaptation rate we described above. A value of 0.75 means that one respondent of the current poll can influence the learning as much as three respondents of the previous poll.

The prior stickiness rate and the sample scaling parameter are of particular interest in understanding the origin of the learning rate. The prior stickiness rate varies around 0.30 in all four scenarios. A prior stickiness rate of 0.30 implies that respondents consider the information from a new poll, but at the same time factor in their prior beliefs. We, hence, find that participants consider their priors when updating their beliefs. With this in mind, it takes some time for the effect of priors to vanish from the belief formation about the race, which can result in the described patterns above.

We estimate the sample scaling factor to be around 0.11 across all scenarios. A sample scaling factor of 0.11 means that respondents update as if the sample size is lower than the 1000 respondents interviewed in our fictitious poll. The perceived precision of the polls is, hence, lower than the sampling error under random sampling. Calculating the value shows that participants perceive the standard error 2.78 to 3.33 larger than the standard error under a random sample. This is consistent with the idea that participants are further discounting information from polls beyond what would be reasonable if the only error source was sampling variation. Interestingly, a similar value is found when estimating the total error of a pre-election polls (see Shirani-Mehr et al., 2018).

Overall, respondents update their beliefs assuming that the polls are relatively imprecise, but still weigh them more strongly than the information from their prior. We would describe an individual with such a learning pattern as a *cautious learner*,

who is mildly skeptical about new information but still weighs it more strongly than information from their priors.¹⁸

Study II: Partisan Bias and Source Variation

Experimental Set-up

The second survey experiment follows the basic set-up of the first experiment. As before, we instill priors then show respondents three polls. However, the key difference is that we now introduce a race between a Republican and Democratic candidate in the district and provide a source for the poll. The experiment again starts the prior instilment about the political race. In this case, we show all participants the same sequence of results of a close district with a mean of 50% vote share for the Democratic and Republican candidate. We then show a sequence of poll results which indicate a) a near tie (51:49), a small majority (54:46), and a large majority (58:42).

In this experiment, we vary the source (MSNBC, Fox) and the candidate that is winning in the polls (Democrat vs Republican). The 2×2 design permits us to evaluate if partisans learn differently from polls depending on the source and which candidate is winning in the polls. In the SM subsection subsection C.4 we show the four ways in which the polls are presented: (a) Democratic winning, (b) Republican winning, (c) the poll is conducted by Fox News, and (d) the poll is conducted by MSNBC. We again elicit beliefs about the vote share of the winning candidate based using the Manski questions, after the prior inducement and after each round of polls.

To test the partisan motivated learning processes, we ask respondents about their party identification and categorize them into Democrats, Republicans, and Independents.¹⁹ The partisanship question will allow us to evaluate the learning process of Democrats and Republicans separately and analyze how they adapt their learning process in presence of ascending or descending poll results for their candidate.

Sample

Our sample consists of 2000 respondents that we recruited via Amazon MTurk. The survey took place before the 2020 presidential election, between the 28th of September and the 20th of October. The median time for taking the survey was around seven and a half minutes and we paid participants \$1.10. Descriptive statistics of the sample are again provided in the SM E.1. The sample is comparable to the first study.

¹⁸ The model further provides a reasonable description of the learning process. SM D.4 describes that parsimonious dynamic learning model fit is comparable to a model estimated for each point in time. The second study that involves partisan and source cues points in the same direction (see SM E.4)

¹⁹ We ask “Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent, or what?” and offer *strong R/D*, *weak R/D*, *leaning R/D*, and *independent* as response categories. We group *strong*, *weak*, and *leaning* into the partisan category.

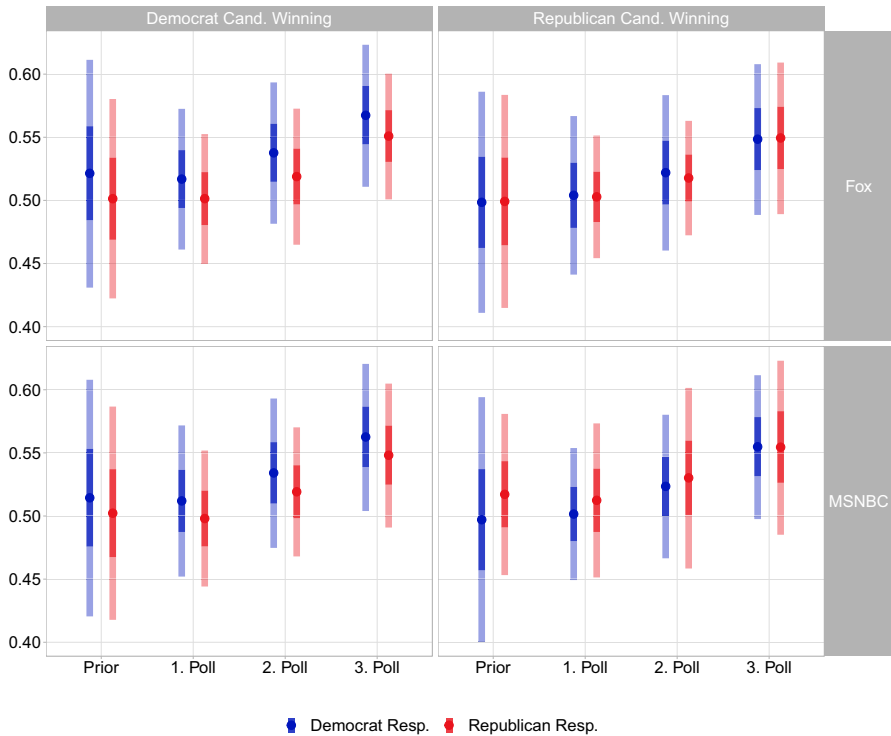


Fig. 5 Respondents' beliefs about the support for the winning candidate for different experiential scenarios and time-points. The beliefs are estimated for each time point and condition separately. The bars indicate 99% and 95% coverage of the beliefs, the points the expectations

Results

Figure 5 presents Democrat and Republican respondents' beliefs about vote shares, estimated at each time they are introduced to new poll information. The evolution of beliefs highlights some interesting patterns. First, when a Democratic candidate is winning in the polls, Republican respondents do not update their beliefs as strongly as Democratic respondents. Irrespective of the source, the expected value of Republican respondents is below the expected value of Democratic respondents (see left panel). What is more, the variance of the average beliefs is larger if the source is not partisan. For example, then the source of the poll is MSNBC, Republicans' variances are larger than when the source of the poll is Fox News. Second, the prior beliefs are biased in favor of one's own candidate. Although both observed the same sequence of prior election results, in the case of a Democrat candidate winning, Democratic respondents expect 52% and 51%, Republicans think it's 50%. The same holds in the case of Republican respondents with 52% versus 50% in the MSNBC case with a Republican candidate winning. Overall, both partisan groups actively update their beliefs with the polls. The patterns are very similar to the ones without

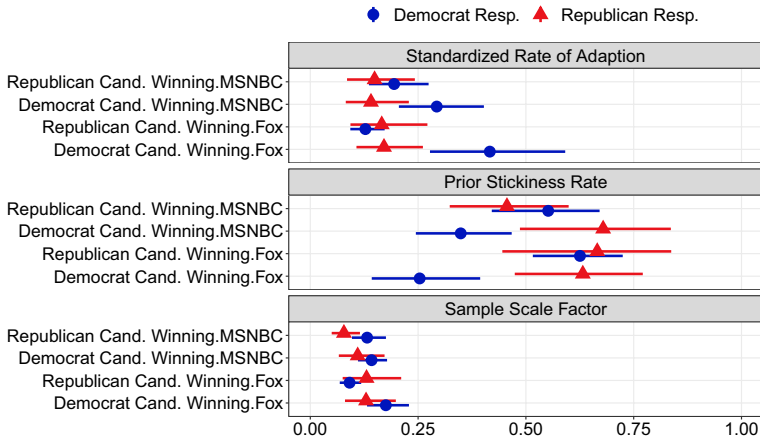


Fig. 6 Estimates for learning model from experimental study II

partisan references that we described in study 1 (competitive race before increasing polls).

We analyze the differences in learning from polls between Democrat and Republican respondents in more detail using the parameters of the learning model.²⁰ The top panel of Fig. 6 presents the estimates for the standardized rate of adaptation. We find that Democrats adopt stronger to poll results when the democratic candidate’s support is increasing. This is particularly true when the results come from Fox news. Here the standardized rate of adaptation raises from 0.13 to 0.41 for Democrats. Put differently, when the Democratic candidate is winning, a Democratic respondent weighs respondents in the Fox survey around three times more than when the Democratic candidate is losing. In the case of a poll from a MSNBC source, the increase in the standardized adaptation rate, from 0.19 to 0.29, is more subtle. The different scenarios do not affect Republican respondents’ standardized rate of adaptation, which estimate consistently around 0.15.

The difference we estimate in overall motivated learning is mostly due to a weaker reliance on prior beliefs. The estimates show a strong decrease in the prior stickiness rate for Democrats when the Democratic candidate is winning. This means that Democrats rely less on their prior beliefs when forming their posterior beliefs about the vote share of the candidates. The prior stickiness rate among Democrats is 0.25 (Fox News) and 0.35 (MSNBC) when the Democratic candidate is winning and 0.63 (Fox News) and 0.55 (MSNBC) when the Republican candidate is winning. Hence in the case where Democratic candidate sees an increase in the polls, they less strongly rely on their priors when forming beliefs based on new polls. We do not observe comparable changes among Republicans. Only in the case of an MSNBC poll, Republicans rely slightly less on their priors when the Republican candidate is winning (0.626) versus when the Democrat is winning (0.46).

²⁰ SM E.3 shows the complete parameter estimates, including the once for independent respondents.

We also estimate some motivated perceived precision of the poll. For Democrats, we estimate a sample scale factor of 0.17 when the Democratic candidate is winning, but only 0.09 if the Republican candidate is winning. This means that for a Fox News poll the perceived precision is 2.4 times the margin of error under random samples when the democratic candidate is winning but 3.3 times when the republican candidate is winning. The motivated perception that we find for the candidate that is winning does not apply to the source. It is not the case that Democrats judge Fox News polls to be generally less precise than MSNBC polls. For MSNBC polls we do not find this type of motivated perception among Democrats. Nonetheless, we do observe a subtle change among Republicans, who judge an MSNBC poll to be less precise when it reflects that the Republican candidate is winning in the polls.

Overall, the results still underline the idea that partisan identities can motivate rational learning by adapting the parameters of this process. We find that this can lead to quite different learning behavior in our respondent pool. Democrats are particularly responsive to new information when they see their candidate winning. While this is in line with previous studies of learning from polls (Madson & Hillygus, 2020), our results further indicate that the credibility of the poll is only one source for this behavior. We find that for Democrats welcome poll results lead to less reliance on prior beliefs when learning from polls. This creates different learning trajectories from the same information environment during an electoral race. With the same prior beliefs Democrats expect a vote share of 57% after having seen the three Fox News polls, Republicans expect 54%.

Why do we find these differences between Democrats and Republicans motivated learning? The results that learning can differ between the two partisans are broadly in line with the results by Baron and Jost (2019) or Morisi et al. (2019), which show that Republicans may be less willing to update with new information. But the findings provide a counter-point to Guay and Johnston (2020) that document no such partisan difference in a set of experiments on how voters interpret policy research. Also, Hill (2017) finds that Republicans and Democrats learn similarly about political facts. What differs in our experiment is the way how partisans adapt their learning process. Democrats adapt their learning process when information from the poll is favorable for their candidate, Republicans do not. Based on the research design presented in this paper we can only speculate why this might be the case. One reason might be that the hesitation of Republicans found in Baron and Jost (2019) and Morisi et al. (2019), in particular, shows in the motivated learning processes. Another reason might be the closeness of the recent US presidential election, where Biden as the Democratic candidate gained support. It could be that Democratic participants found the presented increasing support of the candidate scenario more trustworthy than Republicans.

Conclusion

This paper introduced a dynamic Bayesian learning model of how voters perceive polling information during electoral campaigns and update their beliefs about partisan or candidate strength. It reconciles Bayesian learning with motivated reasoning

by allowing for two distinct mechanisms of varying adaptation: Voters can either distrust a polling result or give undue weight to their prior beliefs (or both). The advantage of this model is that it allows us to quantify the relative importance of these two mechanisms.

In two experiments, we find evidence for the notion of cautious Bayesian learning from polls and that this cautiousness can be motivated. This finding is consistent with an account based on cautious learning (Hill, 2017) and generalizes it to a dynamic setting. Our experiments also suggest that this caution is partly driven by motivated reasoning: partisan voters differ in how quickly they incorporate new polling and discredit polls that do not conform to their preferences (Madson & Hillygus, 2020). Adaptation of the reliance on priors, however, is the stronger mechanism compared to discrediting the precision of the poll source. Still, repeated exposure to polls eventually leads to belief convergence (see SM A.2 and empirical results in Figure A5). One important implication of these findings is that a great amount of polling information may impact citizens' beliefs, sway their voting behavior, and, ultimately, affect democratic outcomes.

There are additional meaningful extensions to the learning model. Another important additional consideration is the perception of polling *bias* in the learning model. For instance, voters could perceive biased reporting by different polling sources (e.g. Fox News always reports too positively in favour of Republican support). Our model focuses on the difference in *precisions* of the polls instead of bias perception. This focus is well in line with the existing political psychology literature that focuses on the credibility of the poll (Kuru et al., 2017; Madson & Hillygus, 2020; Searles et al., 2018). Analyzing in how far learning models with bias terms yield equivalent learning outcomes over electoral campaigns is a worthwhile next step.²¹ Another important aspect is that our theoretical model, in general, allows for changing rate of adaptation within voters and over time, but our experiment does ultimately not test the possibility. In this extension, voters could perceive some polls as more precise than others, with the effect that some polls have longer-lasting effects on learning. It could also be that the reliance on priors is a function of the time between two published polls. Moreover, our model does not integrate selective exposure to poll information (Stroud, 2017). This, however, could be a driver for unequal information environments and thereby learning outcomes during election campaigns (Mummolo, 2016). Finally, our approach treats the two aspects of the learning process from polls (precision of the polls and reliance on priors) as exogenous to the learning process. But voters could in general learn about these by observing the volatility of the polls and the final results. Extended theoretical models and research designs can integrate and test this higher level of learning in the model.

The model contributes to debates on the role of pre-election polls in electoral democracies. Even if polls are a reliable source of information in electoral democracies (Jennings & Wlezien, 2018), the public debate criticizes polls for failing to predict political outcomes. The misperception might partially be due to the presentation

²¹ We present a potential extension of the our model with both bias and precision perceptions in the SM A.4.1.

of the inherent uncertainty in the polls (Westwood et al., 2020). Our experimental and theoretical contributions allow further research to evaluate the different presentations of poll results and their effects on belief formation. Studying the consequences of learning from polls for political behavior can further contribute to our understanding of concerns about polls swaying of democratic elections (Dahlgaard et al., 2017; Großer & Schram, 2010; Marsh, 1985; Rothschild & Malhotra, 2014).

Supplementary Information The online version of this article (<https://doi.org/10.1007/s11109-022-09837-8>) contains supplementary material, which is available to authorized users.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baron, J., & Jost, J. T. (2019). False equivalence: Are liberals and conservatives in the United States equally biased? *Perspectives on Psychological Science*, *14*(2), 292–303.
- Bartels, L. (2002). Beyond the Running Tally: Partisan Bias in Political Perceptions. *Political Behavior*, *24*, 117–150.
- Bartels, L. M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science*, *40*, 194–230.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, *20*(3), 351–368.
- Bullock, J. G. (2009). Partisan bias and the Bayesian ideal in the study of public opinion. *The Journal of Politics*, *71*(3), 1109–1124.
- Dahlgaard, J. O., Hansen, J. H., Hansen, K. M., & Larsen, M. V. (2017). How election polls shape voting behaviour. *Scandinavian Political Studies*, *40*(3), 330–343.
- Edwards, W. (1982). *Conservatism in human information processing* (pp. 359–369). Cambridge University Press.
- Gerber, A., & Green, D. (1999). Misperceptions about perceptual bias. *Annual Review of Political Science*, *2*(1), 189–210.
- Goldstein, D. G., & David, R. (2014). Lay understanding of probability distributions. *Judgment & Decision Making*, *9*(1), 1.
- Großer, J., & Schram, A. (2010). Public opinion polls, voter turnout, and welfare: An experimental study. *American Journal of Political Science*, *54*(3), 700–717.
- Guay, B., & Johnston, C. D. (2020). Ideological asymmetries and the determinants of politically motivated reasoning. *American Journal of Political Science*, *00*, 1–17.
- Guess, A., & Coppock, A. (2020). Does counter-attitudinal information cause backlash? Results from three large survey experiments. *British Journal of Political Science*, *50*(4), 1497–1515.
- Hill, S. J. (2017). Learning together slowly: Bayesian learning about political facts. *The Journal of Politics*, *79*(4), 1403–1418.
- Jennings, W., & Wlezienc, C. (2018). Election polling errors across time and space. *Nature Human Behaviour*, *2*, 276–283.
- Jerit, J., & Barabas, J. (2012). Partisan perceptual bias and the information environment. *The Journal of Politics*, *74*(3), 672–684.

- Kahan, D. M. (2012). Ideology, motivated reasoning, and cognitive reflection: An experimental study. *Judgment and Decision making*, 8, 407–24.
- Kuru, O., Pasek, J., & Traugott, M. W. (2017). Motivated reasoning in the perceived credibility of public opinion polls. *Public Opinion Quarterly*, 81(2), 422–446.
- Lau, Richard R., & Redlawsk, David P. (2006). *How voters decide: Information processing in election campaigns*. Cambridge University Press.
- Lebo, M. J., & Cassino, D. (2007). The aggregated consequences of motivated reasoning and the dynamics of partisan presidential approval. *Political Psychology*, 28(6), 719–746.
- Leemann, L., Stoetzer, L. F., & Trauemueller, R. (2021). Eliciting beliefs as distributions in online surveys. *Political Analysis*, 29(4), 541–553.
- Leeper, T. J., & Slothuus, R. (2014). Political parties, motivated reasoning, and public opinion formation. *Political Psychology*, 35, 129–156.
- Little, A. T. (2021). *Directional motives and different priors are observationally equivalent*. University of California.
- Lodge, M., Steenbergen, M. R., & Brau, S. (1995). The responsive voter: Campaign information and the dynamics of candidate evaluation. *American Political Science Review*, 89(2), 309–326.
- Lord, C. S., Ross, L., & Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098–2109.
- Madson, G. J., & Hillygus, S. D. (2020). All the best polls agree with me: Bias in evaluations of political polling. *Political Behavior*, 42(4), 1055–1072.
- Manski, C. F. (2009). *Identification for prediction and decision*. Harvard University Press.
- Marsh, C. (1985). Back on the bandwagon: The effect of opinion polls on public opinion. *British Journal of Political Science*, 15(1), 51–74.
- Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior Research Methods*, 44(1), 1–23.
- Morisi, D., Jost, J. T., & Singh, V. (2019). An asymmetrical “president-in-power” effect. *American Political Science Review*, 113(2), 614–620.
- Moy, P., & Rinke, E. M. (2012). *Attitudinal and behavioral consequences of published opinion polls. Opinion polls and the media* (pp. 225–245). Springer.
- Mummolo, J. (2016). News from the other side: How topic relevance limits the prevalence of partisan selective exposure. *The Journal of Politics*, 78(3), 763–773.
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, 30(37), 12366–12378.
- Popkin, S. L. (2020). *The reasoning voter*. University of Chicago Press.
- Redlawsk, D. P. (2002). Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making. *The Journal of Politics*, 64(4), 1021–1044.
- Redlawsk, D. P., Civettini, A. J. W., & Emmerson, K. M. (2010). The affective tipping point: Do motivated reasoners ever “get it”? *Political Psychology*, 31(4), 563–593.
- Robinson, J., Rosenzweig, C., Moss, A. J., & Litman, L. (2019). Tapped out or barely tapped? Recommendations for how to harness the vast and largely unused potential of the mechanical turk participant pool. *PLoS ONE*, 9, e0226394.
- Rothschild, D., & Malhotra, N. (2014). Are public opinion polls self-fulfilling prophecies? *Research & Politics*, 1(2), 2053168014547667.
- Searles, K., Smith, G., & Sui, M. (2018). Partisan media, electoral predictions, and wishful thinking. *Public Opinion Quarterly*, 82(S1), 888–910.
- Shirani-Mehr, H., Rothschild, D., Goel, S., & Gelman, A. (2018). Disentangling bias and variance in election polls. *Journal of the American Statistical Association*, 113, 1–23.
- Sinclair, B., & Plott, C. R. (2012). From uninformed to informed choices: Voters, pre-election polls and updating. *Electoral Studies*, 31(1), 83–95.
- Smith, J. Q. (1979). A generalization of the Bayesian steady forecasting model. *Journal of the Royal Statistical Society Series B (Methodological)*, 41, 375–387.
- Steenbergen, Marco R., & Howard, Lavine. (2018). Belief change: A Bayesian perspective. *The Feeling, Thinking Citizen* (pp. 99–124). Routledge.
- Stroud, N. J. (2017). Selective exposure theories. In K. Kenski & K. H. Jamieson (Eds.), *The Oxford handbook of political communication*. Oxford University Press.

- Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3), 755–769.
- Thomas, K. A., & Clifford, S. (2017). Validity and mechanical turk: An assessment of exclusion methods and interactive experiments. *Computers in Human Behavior*, 77, 184–197.
- Tsfati, Y. (2001). Why do people trust media pre-election polls? Evidence from the Israeli 1996 Elections. *International Journal of Public Opinion Research*, 13(4), 433–441.
- Weisberg, H. F. (2009). *The total survey error approach: A guide to the new science of survey research*. University of Chicago Press.
- Westwood, S. J., Messing, S., & Lelkes, Y. (2020). Projecting confidence: How the probabilistic horse race confuses and demobilizes the public. *The Journal of Politics*, 82(4), 1530–1544.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Lukas F. Stoetzer¹  · Lucas Leemann² · Richard Traunmueller³

Lucas Leemann
leemann@ipz.uzh.ch
<http://www.lucasleemann.ch>

Richard Traunmueller
traunmueller@uni-mannheim.de
<http://www.richardtraunmueller.com>

- ¹ Department of Philosophy, Politics and Economics, Witten/Herdecke University, Witten, Germany
- ² Department of Political Science, University of Zurich, Zurich, Switzerland
- ³ School of Social Sciences, University of Mannheim, Mannheim, Germany