

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## European Journal of Political Economy

journal homepage: [www.elsevier.com/locate/ejpe](http://www.elsevier.com/locate/ejpe)Can biased polls distort electoral results? Evidence from the lab<sup>☆</sup>Aristotelis Boukouras<sup>a</sup>, Will Jennings<sup>b</sup>, Lunzheng Li<sup>c,\*</sup>, Zacharias Maniadis<sup>c,d</sup><sup>a</sup> University of Leicester, Division of Economics, School of Business, United Kingdom<sup>b</sup> University of Southampton, Department of Politics and International Relations, United Kingdom<sup>c</sup> University of Cyprus, Department of Economics (SInnoPSis Unit), Cyprus<sup>d</sup> University of Southampton, Department of Economics, United Kingdom

## ARTICLE INFO

Dataset link: <https://github.com/lunzheng-li/biased-poll-replication-package>

## Keywords:

Lab experiment  
Voting  
Biased polls  
Beliefs

## ABSTRACT

We introduce a new methodological approach for studying the effect of biased polls on election outcomes and apply it to a set of new experiments with 375 participants. Voters may observe and learn about the bias by playing multiple voting rounds. While in control conditions, polls are unbiased, in treatment conditions, participants view only poll results where a particular candidate's vote share is the largest. This candidate is consistently elected more often in the treatments than in the controls, because biased polls robustly distort voters' expectations about vote shares. This effect holds after eighteen election rounds, out of which the first three are practice rounds, but somewhat more weakly in our main treatment where voters are *explicitly* informed about the bias.

## 1. Introduction

Pre-election polls are an expression of voter preferences at a given point in time. However, society often uses them as predictors of election outcomes, as evinced by recent formal attempts to account for discrepancies between polls and election results (see [Sturgis et al., 2016](#)). As this formal UK inquiry shows, it is possible that in given historical periods and countries, differences exist between the average expression of voters' preferences as exhibited in the polls and election results.

We propose an experimental methodology for examining the potential effects of such discrepancies in the lab. We capture non-random patterns in the manner that voters are exposed to poll results by generating a collection of poll results in each election round, each corresponding to a random subset of participants. In turn, which of these poll results voters observe may be non-random. We apply this approach to a set of novel laboratory experiments with 375 UK participants, and we find that for a particular type of non-randomness, biased exposure to poll results may indeed affect election results and welfare, even in the main treatment where participants are fully informed about the structure of non-randomness.

We are one of the first studies to perform such an experimental examination, and for good reasons. It is difficult to measure the degree of systematic discrepancies between polls and election results, and their electoral consequences with observational data

<sup>☆</sup> We are greatly indebted to Daniel Zizzo, John Hey and Alexandros Karakostas for their support and assistance. This work was supported by BA/Leverhulme small grant SG162362 and the project "SInnoPSis" funded by Horizon 2020 under grant agreement ID: 857636. We would like to thank David K. Levine, Adriana Bunea, Philippos Louis, David Rahman, Joel Sobel, Erik Gahner Larsen, seven anonymous referees, seminar participants at the Universities of Portsmouth, Southampton, Leicester, King's College London, as well as participants at the Second Cherry Blossom Workshop on Experimental Economics in Wuhan, China; the 2018 and 2019 CRETE conferences in Tinos, Greece; the 2018 Workshop on Perceptions and Public Policies, Bar Ilan, Israel, the 2019 Cyprus Meeting on Behaviours and Algorithms, Cyprus, and the Central European Programme on Economic Theory 2019, Italy.

\* Corresponding author.

E-mail addresses: [Aris.Boukouras@le.ac.uk](mailto:Aris.Boukouras@le.ac.uk) (A. Boukouras), [W.J.Jennings@soton.ac.uk](mailto:W.J.Jennings@soton.ac.uk) (W. Jennings), [lunzheng.li@ucy.ac.cy](mailto:lunzheng.li@ucy.ac.cy) (L. Li), [maniadis.zacharias@ucy.ac.cy](mailto:maniadis.zacharias@ucy.ac.cy) (Z. Maniadis).

<https://doi.org/10.1016/j.ejpoleco.2023.102383>

Received 6 April 2022; Received in revised form 4 December 2022; Accepted 1 March 2023

Available online 5 March 2023

0176-2680/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

alone. Ideally, one should conduct a randomised controlled trial, but it would not be ethical to distort a real electoral race. An alternative approach would be to use purely hypothetical surveys and to embed the study in a real election. However, if we chose to do this, we would be unable to incentivise thoughtful behaviour with real money and we would be unable to examine the effects of repetition and learning. Summing up, we tackle an important problem with the only methodological strategy – in our view – that establishes causality and allows for incentives and feedback.

In our experiments, we observe the outcome of fifteen electoral races (plus three practice rounds) between two parties (Party K and Party J) who field different candidates in every round. The two candidates differ in their ‘valence’, and the exact valences are known to some participants (the ‘informed voters’). ‘Uninformed voters’ are only told the statistical distribution out of which the valences were drawn. Before each election, five voting-intention polls are generated by randomly sampling participants. In this manner, polls allow informed voters to provide a noisy signal regarding the valence of the two candidates. In Experiment 1 (E1), we start by comparing a biased regime – where the results of only the *two polls most favourable for one candidate* (the candidate of party K, or simply *candidate K*) are revealed – to a natural control setting, where *all five polls* are revealed. In Experiment 2 (E2), the control setting entails revealing the results of *two randomly selected polls*, rather than all five polls. Finally, in Experiment 3 (E3), whereas in the control condition all five polls are revealed, in the treatment condition participants are informed beforehand about the (non-random) rule for selecting the two polls to be revealed.

If a party’s popularity is systematically ‘inflated’ in the polls, does this result in an electoral advantage for that party? Our results suggest that this is indeed the case. Both in terms of the number of rounds that candidate K was elected and in terms of average vote share, candidate K performed better in the treatment than in the control condition in a robust manner. In particular, the biased feedback mechanism increased the average vote share of the favoured candidate K by 16 percentage points, 8.6 percentage points and 6.1 percentage points in E1, E2 and E3, respectively. These differences are consistent across sessions, and there is limited evidence that these effects go away as participants gain more experience.<sup>1</sup>

Our methodology addresses research design issues present in previous literature. Previous experiments with biased or manipulated polls exist only in the political science literature (Meffert and Gschwend, 2011; Gerber et al., 2017) and they are all conducted in one-shot election environments, which do not permit voters to infer the accuracy of polls through experience. Also, it is important to emphasise that our results run counter to the predictions of some established theories. In fact, some political scientist colleagues, drawing on the theory of psychological reactance (Brehm, 1966, 1972), predicted that the bias, once publicly revealed, will backfire against candidate K. Namely, if voters realise that polls are biased, they could perceive it as an attempt to limit their freedom on political choice and so they could vote against the polls and in favour of the election ‘underdog’. However, in our experiments we find no such evidence.

Economic experiments in the literature have examined a variety of mechanisms that can drive poll effects on elections, with neutral phrasing and a theory-testing focus. An important mechanism examined in the lab is asymmetric information among voters (McKelvey and Ordeshook, 1984, 1985; Brown and Zech, 1973; Sinclair and Plott, 2012). This experimental strand finds that polls aggregate information reasonably well, although voters exhibit some robust elements of bounded rationality. A second studied mechanism has been coordination and strategic voting in multi-candidate elections (Forsythe et al., 1996; Plott, 1982), where the evidence indicates that polls can often be instrumental in coordinating voters’ choices. An additional important mechanism is turnout under costly voting. Most studies (Klor and Winter, 2007; Agranov et al., 2017; Gerber et al., 2017) point to a failure of the standard prediction that polls discourage majority group voting and that they are welfare reducing (Goeree and Grosser, 2007), although the effects seem generally complex.

However, the economics literature is mainly focused on unbiased polls, whereas our paper is concerned with biased polls and their effects on voting behaviour.<sup>2</sup> This is closer to the approach taken in political science, where many experiments strategically manipulate the poll information that participants receive. Typically, these experiments are non-incentivised. The early study by Fleitas (1971) indicates that voting is not responsive to the quantitative information revealed in polls. Meffert and Gschwend (2011) present different versions of newspaper articles that report voter support for German parties in multicandidate elections, while Rothschild and Malhotra (2014) manipulate the ostensible public support for several important issues and examine how this affects subjects’ stated preference on the issues. These studies find that manipulation affects beliefs and moderately alters behaviour. Gerber et al. (2017) conduct large field experiments where they selectively convey poll results to manipulate the ostensible closeness of the race. Again, beliefs seem to be affected by the manipulation, but behaviour is not affected much. As with previous experiments, rational choice theories, which predict voter turnout, do not perform very well.

<sup>1</sup> Furthermore, our analysis of voting behaviour and elicited beliefs provides support for the mechanism of anchoring and (insufficient) adjustment (Tversky and Kahneman, 1974). Beliefs are highly correlated with the average vote shares displayed in the revealed polls. Moreover, average revealed poll results are a good predictor of electoral results, although these polls were selected in a biased manner.

<sup>2</sup> We suspect that at least part of the reason for this omission in the experimental economics literature is reluctance to use what can be viewed as explicit manipulation in the lab. For instance, we refer to several studies in political science that expose subjects to different poll results (sometimes fabricated) and examine how this affects their behaviour. In our experiments, we avoid this approach that would unambiguously qualify as deception and we only provide truthful information. Still, some colleagues would count as deception any omission of information, as long as participants are expected to behave differently in the presence of this information. However, most experiments where information is a treatment variable can be considered problematic under this strict definition. Moreover, according to this strict approach, even information about other participants’ behaviour, or about the research objectives, should be shared with all subjects, but of course this would sometimes jeopardise the research design. We argue that the question of whether and how people are able to identify biased information can and should be examined in the economics laboratory, and how participants form beliefs about whether information is biased or not should be an open research question, not a forbidden one.

**Table 1**  
The experimental design.

	Experiment E1	Experiment E2	Experiment E3
Treatment	The two polls (out of the five) with the greatest support for K are revealed.	The two polls (out of the five) with the greatest support for K are revealed.	The two polls (out of the five) with the greatest support for K are revealed. Subjects are a priori informed about this.
Control	All five polls are revealed.	Two out of the five polls are randomly revealed. Subjects are a priori informed about this.	All five polls are revealed.

The main difference between the aforementioned political science studies and ours is that our design allows for multiple rounds of repetition where the predictions of polls can be juxtaposed with the publicly known election results in every period. In addition, our experiments show that voters are influenced by biased polls even when they are aware that polls are biased, a test that is absent from the aforementioned papers. To the best of our knowledge, no other study has attempted to disentangle the factors driving the effects of biased polls on election outcomes. Finally, our study is conducted in a laboratory and decision-making is incentivised with real money.

The rest of the paper is structured as follows. Section 2 discusses the design of our three experiments. In Section 3 we present descriptive results of our experiments. Section 4 provides additional analysis on welfare effects and individual behaviour. Section 5 presents a short discussion of our findings and concludes.

## 2. Our experimental environment

In general, the information conveyed by poll results can be relevant to voters for many reasons (for instance, voting is costly and voters need to estimate the closeness of the race, there are multiple candidates and voters need to focus on a viable candidate, or voters have bandwagon preferences). The particular environment we choose to study here is akin to Feddersen and Pesendorfer (1997), where voters assess candidates on two dimensions, their ideological position and their intrinsic quality (valence). In our setting, there are two political parties, party K and party J, each one of which fields a candidate. We refer to the candidates' identity by the name of the political party they stand for, hence the candidates are K and J.

All voters know the closeness of the candidates' political views to their own, i.e. the ideological position of the two candidates, but they differ in their knowledge of the candidates' valence. Some voters are informed and know precisely the valence of each candidate, while the remaining are uninformed and they know only the statistical distribution out of which each valence is drawn. Moreover, in our setting informed voters are on average left-wing leaning in terms of ideological positions, while uninformed voters are on average right-wing leaning, so the voting intentions of the informed voters are not representative of the overall population. As a result, elections across the entire set of voters (not within the set of informed voters only) are needed to aggregate the electorate's preferences, while pre-election polls convey valuable information to uninformed voters by helping them make inferences about candidates' valence. In our setting, we have five voting intention polls taking place prior to each election.

Our research question, then, focuses on whether election outcomes are *affected* by giving voters a biased sample of the total information (total information in every round consists of the results from five polls), which systematically depicts the candidate of party K performing 'better' than in reality. This 'biased selection' environment constitutes our experimental *treatment manipulation*. We define the concept of 'affected' italicised above relative to two control conditions as benchmarks. Our first control (in E1) is simply an environment where the total information is released to voters. Our second control (in E2) is a setting where an equal amount of information as in our treatment condition (two polls out of five) is conveyed, but in a random, rather than a systematically selective, manner.<sup>3</sup>

In our main experiment (E3), we test whether the effect of biased polls is due to subjects perceiving the polls as unbiased (despite the feedback that they receive in every round) or due to their inability of properly inferring from feedback which is known to be systematically biased. In particular, whereas in the control of E3 the total information is released to voters, in the treatment condition, participants are informed explicitly about the (biased) selection rule. All experiments are described in detail in Table 1.

### 2.1. Voters' preferences on candidates, voter information, polls, and elections

In each experimental session, there are fifteen human voters (the two non-human 'candidates' are inactive, hence they do not vote). Voters are ordered according to their ideological positions as illustrated in Fig. 1. Voter 1 is the most left-wing voter, while Voter 15 is the most right-wing voter. The median voter is in position 8, while candidates of parties J and K are in position 6 and in position 10, respectively. Ideological positions of candidates are the same in all rounds<sup>4</sup> and all voters know it in advance. At the beginning of each round, the ideological position of each voter is randomly drawn from integers between 1 to 15 (inclusive) without replacement.

<sup>3</sup> The second control allows us to test whether the difference between observing all five polls and two selected polls is due to disparate quantities of information, i.e. observing a smaller set of polls (two instead of five), or whether it is due to the selection per se.

<sup>4</sup> The interpretation is that the two parties consistently pick candidates that share their ideological views.




Fig. 1. Ideological preferences in the experimental interaction.

Notes. This figure illustrates the distribution of preferences and information across the fifteen participants in any given experimental round. The positions on the line are occupied by different participants in every round.

Each candidate's valence is drawn at the start of every round from a uniform distribution with values between 1 and 120.<sup>5</sup> At the time of the polls and the elections, the two drawn valences are known to voters in ideological positions  $\{1, 2, 3, 5, 7, 9, 11\}$  who are the *informed voters*. The remaining voters, i.e. the ones in ideological positions  $\{4, 6, 8, 10, 12, 13, 14, 15\}$ , are the *uninformed voters*. They only know the distribution out of which the quality (valence) of the candidates is drawn.

The utility that voter  $i \in \{1, 2, \dots, 15\}$  obtains in the case where candidate  $h \in \{J, K\}$  wins the election is given by  $U_{ih} = X_i - \alpha d_{ih} + Q_h$ , where  $U_{ih}$  is voter  $i$ 's overall utility from candidate  $h$  being elected,  $X_i$  is voter  $i$ 's utility from having a candidate with the same ideological position as herself being elected, while  $d_{ih}$  is the distance between the ideological positions of voter  $i$  and candidate  $h$ .  $Q_h$  is the valence of candidate  $h$ , and  $\alpha$  is a parameter that measures the utility loss per unit of distance in ideological positions between  $i$  and  $h$ . For the purposes of our experiments, we set  $X_i = 100$  and  $\alpha = 5$  (for all voters, rounds, and sessions) and, as stated previously,  $Q_h \sim U[1, 120]$ .

After the valence is drawn for both J and K and informed voters receive this information, five polls, each inquiring four randomly chosen voters, take place. Sampled subjects are asked for whom they would like to vote in the upcoming elections (they may choose not to participate).<sup>6</sup> Given the number of drawn subjects who choose to participate, a poll reports the fraction of those in favour of K and in favour of J, respectively. For example, a poll revealing the following fractions: [25% for J, 75% for K] indicates that out of the four voters, all of whom chose to answer, three expressed support in favour of candidate K and one in favour of J.<sup>7</sup> Note that a single voter may participate in multiple polls. After the five polls are created by the above process, some subset of the results (depending on the experimental condition) is presented to all voters. The summary of each experimental round (as it was provided to participants) is illustrated in Fig. 3. The winner of the election is determined by simple majority, with ties broken by a random draw. Subjects participated in eighteen election rounds like the one described above, that is, three practice rounds and fifteen incentivised real ones.

Since some voters are uninformed about the difference in valence between the two candidates, pre-election polls can be socially valuable in this setting. In particular, they can be utilised to transmit information about the candidates' valence from informed to uninformed voters. It is theoretically important that the distribution of informed voters is not symmetric in the ideological spectrum. If the distribution was symmetric, then the socially efficient outcome would be for uninformed voters to abstain from elections and let voting among informed voters determine the election outcome. In such an environment, polls would not perform a politically valuable role, because participation of uninformed voters would not be necessary. Instead, polls are meaningful in our setting, because they aggregate information about candidate valence when the ideological preferences of informed voters do not represent the ideological preferences of uninformed voters.<sup>8</sup>

To illustrate the hypothesised inference process of participants who do not account for the bias, let us assume that some uninformed voter observes substantial support in favour of K in the polls. If she perceives polls as unbiased and other subjects as rational, she will infer that K's valence is higher than J's, since some informed voters who are close to J's ideological position

<sup>5</sup> To reduce noise across sessions, we drew these valences once and for all before the start of the first session and used the same random draws for every session and for all experiments. In rounds 1, 5, 9, 10, 11, 12, 13, 14 and 15, K has a higher valence, while J's valence is higher in rounds 2, 3, 4, 6, 7 and 8.

<sup>6</sup> Voters choose from the following three options: 'K', 'J', and 'Prefer not to participate'. Experimental poll results do not show the information on 'non-participation', as this corresponds to the salient information that voters receive in real polls, especially when multiple polls are presented. For instance, Fig. 2 illustrates the format of presentation of UK polls used by 'YouGov.co.uk'. This format of presentation is common for almost all online media appearing in an online search for 'voting intention polls', such as Financial Times tracker, 'ukpollingreport.co.uk' and 'markpack.org'. Accordingly, our experimental approach substantially simplifies the feedback that participants observe about the results of polls, while keeping in line with the presentation structure used in real life elections. This is especially important, since participants need to infer overall support for each candidate on the basis of results from multiple polls, without being cognitively overwhelmed.

<sup>7</sup> If, out of the four sampled voters, three opted to support K and one chose not to participate, the poll would be presented as 0% in favour of J and 100% in favour of K.

<sup>8</sup> However, this does not impact our experimental design, since even if the distributions of political preferences of informed and uninformed voters were identical and symmetric, a biased sample of polls (if not appropriately discounted) would still tilt the election result in the favoured candidate's direction.



Voting Intention Tracker (GB)

From 2019 General Election - Present

		Con	Lab	Lib Dem	SNP	Green	Brexit Party	Other	Con lead over Lab
		%	%	%	%	%	%	%	%
Start of Fieldwork	End of Fieldwork								
<b>2020</b>									
29/05/2020	30/05/2020	45	35	6	5	5	2	2	10
25/05/2020	26/05/2020	44	38	6	5	4	2	1	6
18/05/2020	19/05/2020	48	33	6	5	5	2	1	15
05/05/2020	06/05/2020	50	30	7	4	5	3	1	20
16/04/2020	17/04/2020	53	32	5	4	3	1	2	21
01/04/2020	02/04/2020	52	28	8	5	5	1	1	24
09/02/2020	10/02/2020	48	28	10	4	6	2	2	20
31/01/2020	02/02/2020	49	30	8	4	5	2	2	19
24/01/2020	26/01/2020	49	29	10	5	4	2	1	20
<b>DECEMBER GENERAL ELECTION</b>		44	32	12	4	3	2	3	12

Fig. 2. YouGov voting intention sample screenshot. Source: <https://yougov.co.uk>.

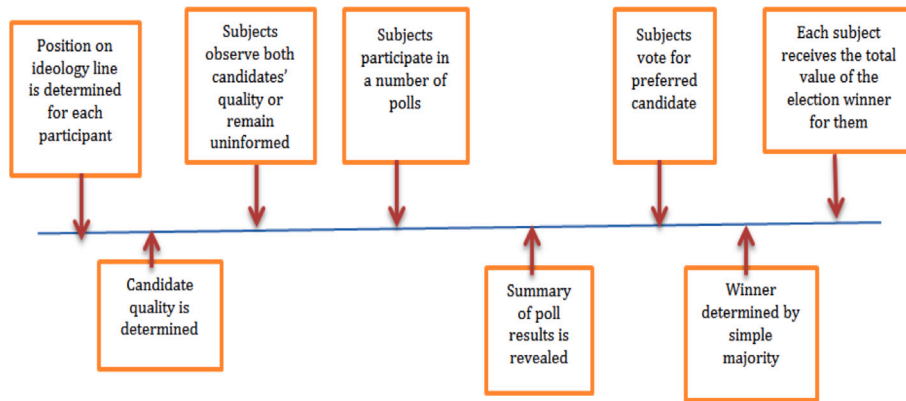


Fig. 3. Sequence of actions in each experimental round.

prefer to vote for K. These voters would do so only if K is of significantly higher valence than J. Accordingly, the uninformed voter, who observes the polls herself, infers from them the higher valence of candidate K and she may herself change her voting intention from J to K, depending on her position in the political preferences spectrum.

Note that the bias in exposure to poll results can in principle be detected through learning in all of our treatment conditions. Voters may perceive polls as unbiased in the early rounds of our experiments. However, informed voters know the true valence parameters of the two candidates and, if they are motivated primarily by pecuniary incentives, they will vote for the candidate that gives them the highest experimental payoffs. If this behaviour persists, poll results will systematically overstate candidate K’s vote share in comparison to the election results. Voters able to learn from experience should detect this systematic difference, and adjust for it in their beliefs and behaviour. In short, our research design allows for polls to alter subjects’ beliefs in favour of candidate K, but also for participants to detect the bias through experience.

2.2. Implementation

The only stage that differs across the two experimental conditions in each of our three experiments is the one where the summary of poll results is revealed (see Fig. 3). Table 2 illustrates how the information on poll results is revealed to subjects in the two experimental conditions of E1 and E3. Finding meaningful differences between ‘control’ and ‘treatment’ would indicate that biased polls can skew elections. The first benchmark (the control condition in our first experiment), which we use to judge whether ‘skewing’ takes place, is a perfectly transparent regime where all existing information (all five polls) is available to the public. This is a natural starting point. We also consider another benchmark (the control condition in our second experiment) where two out of the five polls are revealed in a random manner.

In terms of the treatment conditions, our natural point of departure (in E1 and E2) is an environment where voters observe the revealed information and have no a priori knowledge concerning how the two polls out of five are chosen to be revealed. In our view, this corresponds to many natural election environments of interest, where voters are not provided with any ‘manual’ describing the possible biases or agendas of those that reveal poll information. In our main condition (Experiment 3) we use a larger sample and examine the consequences of providing a priori information about the exact nature of the bias to voters.

**Table 2**  
Example presentation of poll results in each condition.

Treatment					
Company	B		E		
Candidate K	75%		100%		
Candidate J	25%		0%		
Control					
Company	A	B	C	D	E
Candidate K	33%	75%	25%	67%	100%
Candidate J	67%	25%	75%	33%	0%

*Notes.* There are five polling companies, A to E. The result of each company is represented in terms of the two fractions measuring support for each candidate. In the control of E1 and E3, all five results are revealed, in a format similar to the example of the table. In addition, if the above table represented an actual set of poll results, then, in the treatment condition of all three experiments, companies B and E would be revealed, since these polls yield the highest support for candidate K.

Experiments E1 and E2 had 120 participants each,<sup>9</sup> with eight 15-subject sessions (four control sessions and four treatment sessions).<sup>10</sup> Experiment E3 had 135 participants, with four control sessions and five treatment sessions. Participants in E1 and E2 were students at the University of Southampton and Newcastle Business School, and the experiments took place between May and November 2018. Participants in E3 were students at the University of York, and the experiment took place in June 2019. Our objective was for each experimental block (of 30 subjects) to achieve perfect randomisation by containing one control and one treatment session, with participants being randomly allocated between the two.<sup>11</sup>

In each session, subjects read instructions from their computer screens.<sup>12</sup> After the instructions, subjects participated in 18 rounds of play, including three practice rounds. At the end of the session, they were asked to complete a short questionnaire and were informed about their final score and monetary earnings. The core design of each round has been summarised in Fig. 3. The only aspect that was not described there is the ‘belief elicitation’ stage. In particular, after the release of information on poll results, participants were asked to state their beliefs about the vote shares of the two candidates in the forthcoming election. The information on poll results revealed to participants took the form of a single probability distribution for each result, as shown in Table 2. Participants’ beliefs at the elicitation stage were also expressed in terms of this binary probability distribution.

### 3. Results

Let us first provide an overall summary of the primary treatment effect across the three experiments: the rate of electoral success. Table 3 illustrates (for all three experiments) the number of rounds won by each candidate in each session of the two treatments. In addition, the table shows – in the parentheses – the total number of votes that each candidate received in each session. Adding up across all sessions in a given treatment, we can see that in E1 party K won 60% of all rounds in the control condition but 80% of the rounds in the treatment condition. In E2 party K won 61.6% of all rounds in the control but 73.3% of the rounds in the treatment, while in E3 party K won 56.7% of all rounds in the control but 64% of the rounds in the treatment. These differences are relatively homogeneous in their magnitude and consistent in their sign, both across sessions of a given treatment and across rounds of a given session.

We examine whether these differences are statistically significant using a Mann-Whitney test. For the three experiments, we treat each experimental session as an observation, and the continuous variable we compare across the two treatments is the number of rounds won by K in a session. As we can see from Table 3, the difference in rounds won by K is statistically significant for Experiments 1 and 2 but not for Experiment 3. We also compare the total election votes for K in a session, and the difference is statistically significant for all three experiments (p-values in the parenthesis). In general, the differences are sizable and consistent, as we shall illustrate below.

<sup>9</sup> We shall use the words ‘session’ to denote each experimental interaction among 15 subjects who vote in the same election, and ‘block’ to denote the two sessions (one control and one treatment) taking place at the same time in the lab. A block has 30 subjects.

<sup>10</sup> We denote individual sessions as  $E_iC_j$  or  $E_iT_j$  where  $i \in \{1, 2, 3\}$  denotes experiment,  $j \in \{1, 2, 3, 4, 5\}$ , denotes session, ‘C’ stands for control, and ‘T’ for treatment. For instance,  $E1C1$  denotes the first control session in E1 and  $E2T1$  the first treatment session in E2.


<sup>11</sup> The only three exceptions in this approach were sessions  $E1C2$ ,  $E1T2$  and  $E3T5$ , which were the only sessions of their block because of insufficient subject participation or lab capacity constraints.

<sup>12</sup> We programmed the experiments using oTree (Chen et al., 2016) and recruited subjects via ORSEE (Greiner, 2015) in the University of Southampton and via hroot (Bock et al., 2014) in the Universities of Newcastle and York. The full set of instructions are presented in Online Appendix A.


**Table 3**  
Number of elections won and votes received by each party across sessions and experiments, and results of Mann-Whitney U Test.

	E1		E2		E3	
	K	J	K	J	K	J
C1	9 (137)	6 (87)	9 (123)	6 (92)	9 (126)	6 (93)
C2	9 (118)	6 (97)	9 (135)	6 (85)	8 (139)	7 (82)
C3	9 (125)	6 (88)	9 (136)	6 (80)	8 (134)	7 (85)
C4	9 (131)	6 (82)	10 (126)	5 (82)	9 (123)	6 (97)
T1	15 (195)	0 (30)	10 (142)	5 (77)	9 (141)	6 (79)
T2	10 (151)	5 (72)	11 (143)	4 (70)	8 (138)	7 (77)
T3	13 (165)	2 (55)	12 (148)	3 (66)	9 (141)	6 (78)
T4	10 (150)	5 (63)	11 (163)	4 (56)	12 (158)	3 (65)
T5					10 (143)	5 (80)
P-value	0.010 (0.015)		0.018 (0.015)		0.120 (0.018)	


Notes. The numbers in the parentheses correspond to the total votes for each party in different sessions. The p-values are calculated as follows. In each session, the number of elections won by K (number of votes received by K) constitutes our continuous measure. Each of the three tests examines the null hypothesis that the probability that this continuous measure in a random control session is larger than the analogous measure in a random treatment session is equal to 0.5.



(a) Fraction of rounds won by K in each experimental session




(b) Average vote share for K in each experimental session




(c) Statistical correlation between average beliefs and average revealed polls

Fig. 4. Comparison of descriptive results across three experiments.


Notes. Fig. 4(a): In E1 and E2, the first three control sessions (C1, C2 and C3) have identical winning percentages for K, thus the relevant data points in the figure overlap.



(a) E1



(b) E2



(c) E3

Fig. 5. Comparison of vote share round-per-round in all experiments.

### 3.1. Experiments E1 and E2

In our two incomplete information experiments (E1 and E2), the treatment did offer a considerable advantage to party K. Biased exposure to polls increased both the likelihood of party K winning the election and its vote share. Figs. 4(a) and 4(b) juxtapose the fraction of election rounds won by K and vote shares for K in treatment versus control sessions. It is clear that for E1 and E2 the electoral performance of K is consistently better in treatment sessions relative to any control session.<sup>13</sup>

Furthermore, the difference in vote shares does not appear only at the average level, but also for each individual round. Figs. 5(a) and 5(b) show the vote share of candidate K in the treatment and the control condition for each round (averaging across the four sessions of each condition). The figures indicate that ‘treatment’ rounds have consistently higher vote shares for K than ‘control’

<sup>13</sup> K won more rounds than J in both the treatment and the control condition. This is to be expected since (by pure chance) in most rounds the randomly drawn valence for K was higher than the drawn valence for J. In fact, in 9 out of the 15 regular rounds K had higher valence than J, and in 7 of these the difference in favour of K was over 20 points.



rounds. In fact, in E1, vote shares in ‘treatment’ are higher than vote shares in ‘control’ for 14 out of 15 rounds, and in E2, for 13 out of 15 rounds.<sup>14</sup> This is important, because it does not seem to be the case that the difference vanishes in the last few rounds.

We will now use the elicitation of subjects’ beliefs to examine whether they are in alignment with the poll information that participants received. At this point, we need to define two measures that we shall use frequently in the subsequent analysis. First, ‘average revealed polls’ in a given round is the share of voters supporting K that can be inferred by the revealed polls in this round. For instance, in E1, in every round of the treatment condition, this share is derived as the average of two polls, while in rounds of the control condition this share is derived as the average of five polls. Second, ‘average beliefs’ in a given round will refer to the elicited expected vote share for K averaged across session participants. If participants in the treatment condition perceived polls to be biased, then they should predict different vote shares for the election than the analogous poll information revealed, and this could potentially lead to a low correlation between average beliefs and average revealed polls. Fig. 4(c) shows that the correlation is clearly not larger in the control sessions relative to the treatment sessions of E1, so there is no systematic pattern in subjects’ beliefs. Looking at the line pertaining to E2, it appears that the correlation between average beliefs and average revealed polls is systematically higher in the control (where information is unbiased) than in the treatment (where information is biased).<sup>15</sup>

### 3.2. Experiment E3: Our main condition

It may be argued that in actual democratic elections people have enough experience with the political process and the media in order to gauge the agendas and incentives of those who reveal poll information. In particular, it is likely that some voters have a strong prior about the ‘biased feedback’ rule. Accordingly, our environment in the treatment conditions of E1 and E2 might be criticised as capturing only the special case of elections with young or inexperienced voters, especially in early rounds of play. Moreover, the structure of the treatment conditions of E1 and E2 makes it difficult to pinpoint exactly the mechanism that drives the treatment effect. In particular, the effect may be either because of the inability of voters to understand that the information is selected in a systematically biased manner, or due to their difficulty in deducing information from a biased set of results even when they know the biased process that generates it.

To alleviate these concerns, in our main experiment (E3), the treatment condition entails using the same biased rule as in the treatment conditions of E1 and E2, but with full clarity about this biased rule. In particular, the instructions mentioned that: “After polls have taken place in each round, the findings of the two companies which exhibit the greatest support for candidate K will be revealed to you. All participants will observe the fraction of votes that each of the two candidates received in the polls of these two companies” and then provided an example to illustrate the biased rule. In this environment, a rational participant would observe the results of these two companies and then try to gauge information about the valence of the two candidates accounting for the selection rule underlying these results. Once more, the issue is whether subjects sufficiently discount the information (typically) in favour of K having the higher valence, and thus whether society avoids the swaying of election results due to the biased reporting rule.

The basic results of E3 (which had five treatment sessions with the ‘known biased rule’ and four control sessions with the ‘transparent democracy’ information environment) are illustrated in Figs. 4 and 5(c). As can be seen, even in this case, the biased feedback rule seems to offer an advantage to candidate K, but this advantage is somewhat smaller than in E1 and E2. In particular, the four sessions with the best electoral performance for K (as measured by the fraction of elections won) are all sessions with the ‘known biased rule’. The difference is – once more – politically meaningful: the number of rounds won by J per session in the control is about 20% larger than in the treatment (6.5 vs. 5.4). Again, it is the consistency and robustness of the effect of the biased release of poll information on electoral results that is interesting. A similar message is conveyed by examining the average vote share of K in each session. In particular, in all treatment sessions K has a higher vote share relative to any control session. Fig. 5(c) shows that the difference exists for most rounds, and that it is rather sizable when ‘ceiling effects’ are not binding.<sup>16</sup>

One interpretation of this finding is that polls create a judgemental anchor for voters’ beliefs regarding election outcomes. Voters do not seem to have the capacity to account for the bias in the polls to its full extent. Instead, they seem to use poll results as anchors, which they adjust until they reach an acceptable range for their beliefs. The use of such a heuristic is reasonable, given the complex setting and its cognitive implications for participants.

The comparison between average revealed polls and average beliefs becomes interesting, especially compared to E2 and E1. In the treatment sessions of E3, there is a weak tendency for average revealed polls to exceed average beliefs (grand means are 74.5% and 68.7%, respectively). This is not true in the control sessions (grand means are 59.7% and 60.4% for average revealed polls and average beliefs, respectively).<sup>17</sup> However, correlational analysis shows no systematic patterns (see the line corresponding to E3 in

<sup>14</sup> In our study, the qualitative effects are robust across several dimensions (experiment, session, round). However, as is common in social sciences, the exact effect sizes may depend on the context (Kessler and Vesterlund, 2015). For instance, we expect that if we were to conduct experiments with 10 polls instead of 5 (keeping other aspects of the experimental environment constant), we would likely find larger treatment effects. However, implementing this would be burdensome for participants in the current experimental environment.

<sup>15</sup> However, this pattern observed in the descriptive results is not corroborated by our more rigorous analysis. Our Online Appendix B details additional econometric analysis, in which we examine whether there are systematic patterns of learning. This analysis provides very little support for the claim that participants discounted biased polls in E2.

<sup>16</sup> In the last five rounds of E3, the treatment effect appears small. However, in these rounds the vote share of K is so high that the treatment does not have much scope for increasing it further. This is referred to as a ‘ceiling effect’ in the behavioural literature, and the small treatment effects could be an artefact of this.

<sup>17</sup> Figures C.8 and C.9 in the appendix indicate that small discounting takes place in most rounds of the treatment condition, but such discounting not discernible in the control condition.

**Table 4**  
Average payoffs in each session.

Session in E1	Average Experimental Payoffs	Session in E2	Average Experimental Payoffs	Session in E3	Average Experimental Payoffs
E1_C1	171.27	E2_C1	171.27	E3_C1	171.27
E1_C2	171.27	E2_C2	167.40	E3_C2	170.80
E1_C3	171.27	E2_C3	171.27	E3_C3	170.80
E1_C4	171.27	E2_C4	169.33	E3_C4	171.27
E1_T1	159.87	E2_T1	169.93	E3_T1	171.27
E1_T2	168.93	E2_T2	166.93	E3_T2	170.80
E1_T3	164.27	E2_T3	165.93	E3_T3	168.73
E1_T4	168.93	E2_T4	166.93	E3_T4	165.67
				E3_T5	169.27

*Notes.* These payoffs are the average individual experimental points across all rounds. In a particular round, a voter's payoffs depend on the distance of her ideological position to the winning candidate's position, and on the winning candidate's valence (see experimental instructions in Online Appendix A for details).

Fig. 4(c)). In particular, average beliefs do not seem more strongly correlated to average revealed polls in the control condition than in the treatment condition.<sup>18</sup> In conclusion, despite the fact that participants are fully informed in E3, belief adjustment is small and insufficient, so that biased polls end up affecting electoral results.

#### 4. Welfare effects and individual behaviour

##### *Welfare effects*

In terms of the welfare effects of biased polls, a rough measure of utilitarian welfare is the average experimental payoffs in each condition. One reason that biased polls should have a negative impact on this measure is that they introduce noise in the information conveyed by polls to voters. Moreover, if voters do not discount the information contained in biased polls properly, then they will tend to vote more frequently for candidate K even if he is of lower valence than candidate J.

Indeed, our findings confirm these conjectures, as can be seen in Table 4. In particular, sessions in the treatment condition were generally associated with lower payoffs per subject than sessions in the control condition. In fact, average individual payoffs across conditions were 171.3 (control) vs. 165.5 (treatment) in E1, 169.8 vs. 167.4 (respectively) in E2 and 171.03 vs. 169.15 (respectively) in E3. This disparity resulted from the fact that the high-valence candidate lost in the treatment condition more often than in the control condition.

Specifically, in the control of E1 the high-valence candidate always won. In contrast, in the treatment condition of E1 there were 12 elections where candidate J lost, despite having the higher valence (candidate K never lost when their valence was higher). In E2, while in the control condition there were three elections where the high-valence candidate lost, this increased to eight elections in the treatment condition.<sup>19</sup> In E3, in the control condition, out of 60 elections, there were two cases where K was the high-valence candidate but J won in the end. The opposite never happened. In the treatment condition, out of 75 elections, there were two times when K was the high-valence candidate but J won in the end, and four times when J was the high-valence candidate but K won in the end.

##### *Behaviour at the individual level*

It also worthwhile to provide some insights on the behaviour of informed voters. We should note that in our experiments, informed voters face an easy decision: they should simply vote for the candidate that gives them the highest payoff, which they can easily calculate. Accordingly, if these individuals' votes deviate from 'optimal behaviour' this would indicate that the assumption of rational, money-maximising political agents is violated. Table 5 illustrates the behaviour of informed voters in the election stage, depending on whether they are J-voters or K-voters.<sup>20</sup> For instance, in 8.57% of the 420 election vote decisions that informed voters made in the control condition of E1, informed voters chose candidate K although the money-maximising choice was candidate J. Similarly, in 34.05% of the 420 decisions that informed voters made in the treatment condition of E2, informed voters chose candidate J and their money-maximising choice was also candidate J. As can be seen, most decisions by informed voters are consistent with the money-maximising model.

<sup>18</sup> This lack of strong support in favour of subjects' discounting of biased polls is corroborated by the analysis of Online Appendix B, which generally finds very weak evidence for belief adjustment to biased polls, and no evidence for systematic belief adjustment mechanisms.

<sup>19</sup> Out of all these instances, only once did J win when K had the higher valence (it happened in the control condition).

<sup>20</sup> For simplicity, we shall call 'h-voter' an informed voter whose money-maximising choice is candidate  $h$ , where  $h \in \{J, K\}$ .

**Table 5**  
Behaviour of informed voters.

Preferred/Voted for	E1		E2		E3	
	Percent of total choices		Percent of total choices		Percent of total choices	
	Control	Treat.	Control	Treat.	Control	Treat.
K/J	5.24	3.33	4.29	3.81	3.10	1.90
J/J	38.33	27.86	37.86	34.05	38.57	39.05
J/K	8.57	19.76	8.81	13.81	8.81	9.14
K/K	44.52	47.38	46.43	47.38	46.38	49.52

*Notes.* This table presents the voting behaviour of informed voters in the final elections. The data are pooled across rounds and also at the experiment level. ‘Preferred’ stands for the money-maximising choice of candidate, while ‘voted for’ signifies the actual voting choice in the elections. Please note that the fractions do not add up to 100%, because abstention is allowed at the election voting stage. In total, there are 420 decisions by the seven informed voters in the four sessions of each condition of each experiment (except E3, where in the treatment condition there are five sessions and thus 525 such decisions).

**Table 6**  
Comparison of individuals’ voting at the polls vs. the final election.

Poll/Election	E1		E2		E3	
	Treatment	Control	Treatment	Control	Treatment	Control
J/J	58.00%	76.06%	72.43%	78.35%	77.50%	78.69%
J/K	40.80%	22.01%	26.75%	18.90%	20.63%	19.67%
J/A	1.20%	1.93%	0.82%	2.76%	1.88%	1.64%
K/J	4.68%	13.99%	5.40%	12.57%	5.82%	14.04%
K/K	94.55%	84.55%	94.03%	86.03%	93.32%	85.67%
K/A	0.78%	1.46%	0.57%	1.40%	0.86%	0.29%

*Notes.* The table juxtaposes voting at the poll stage with the respective vote in the elections for the same individual and the same period. Behaviour at treatment vs. control conditions is compared and data are pooled at the experiment level. ‘A’ stands for abstention in the final elections. Only the decisions of individuals who voted for some candidate at the polls are considered.

Nonetheless, a non-trivial fraction of decisions, slightly lower than 15% for the controls and ranging between 11% and 23% for the treatments, deviates from the prediction of the model of selfish money-maximising agents. A possible explanation for this behaviour is ‘bandwagon preferences’, i.e. a genuine willingness of the participants to vote for the likely winner, which is not captured by monetary payoffs. Interestingly, J-voters are more likely to vote for candidate K in the treatment condition than in the control, and within the treatment condition this type of behaviour is more common than the opposite (i.e. K-voters voting for candidate J). Thus, ‘bandwagon preferences’ are likely to be relevant, and in particular they seem to amplify the effects of biased polls.

It is also useful to discuss the behaviour of voters at the poll stage. Table 6 compares the voting choice at the poll stage to the one at the actual elections.<sup>21</sup> The table indicates that, if subjects truthfully reported voting intentions in the polls, the treatment induced some voters to switch in the direction of voting for K in the elections. Moreover, the voting pattern for those that chose K in the polls is similar across experiments: in all experiments, about 8%–10% of poll voters for K, who would otherwise depart from voting K in the elections (as indicated by behaviour in the control condition), are induced by the treatment to stick to K. However, there are significant differences across experiments in the behaviour of those who chose J at the poll stage, and these can partially account for the heterogeneity of the primary treatment effect across experiments. In particular, as we move from Experiment 1 to Experiment 2, and then to Experiment 3, the effects of treatment in inducing those that voted for J in polls to switch to K in the elections falls from 19.2% to 7.85% to about 1%.<sup>22</sup> These were mainly uninformed voters who were likely induced to switch to K in the elections because of the treatment.<sup>23</sup>

## 5. Discussion and conclusions

In this paper, we presented a novel experimental design to examine the effect of biased polls on electoral results. The environment we considered is a two-candidate election contest with common values (concerning candidates’ valence) and no voting costs. In

<sup>21</sup> Note that this table does not contain the behaviour of all subjects, since some were not randomly chosen to any poll, and some who were chosen opted not to participate. In total, Table 6 contains information for about 70% of overall election votes.

<sup>22</sup> These percentages are obtained as the difference between treatment and control in the J/K row in each experiment. Recall that the entries in this row correspond to the percentage of cases, out of all cases where someone voted both in the polls and the elections and chose J in the polls, that this voter voted for K in the elections. The higher occurrence of this in the treatment condition can be interpreted as a treatment effect.

<sup>23</sup> Tables D.1–D.3 in Online Appendix D provide an overall summary of voting behaviour at the poll stage in the different sessions of the three experiments. The results are broken down by different status of voters (informed vs. uninformed). Certain insights can be inferred from Tables D.1–D.3: informed voters are more likely to participate to polls, while uninformed voters are more likely to vote for K rather than J in the polls (which makes sense, since their ideologies are closer to K). Moreover, there seem to exist no systematic differences between treatment and control, which is again unsurprising, since the treatment is different from the control only when voters observe poll results.

particular, a strict subset of voters has information on the valence of the two candidates, and polls communicate information to uninformed voters. In our treatment conditions, participants have access to a biased sample of polls' results, favouring systematically one candidate. By allowing for multiple rounds of repetition, voters have opportunities to compare the actual electoral results with the polls. In our main condition, voters are a priori informed about the bias, but we also examine what happens when they are not.

Our findings indicate that biased exposure to polls consistently skews the electoral outcome in a predictable way. In a robust manner, elections that took place in the 'biased polls' environment provided an electoral advantage to the candidate that was 'favoured by the bias'. This effect was smaller for our main condition, where voters were explicitly informed of the selection rule under which poll information was revealed, but it was still sizable and consistent. Overall, the results show limited evidence that the repeated opportunities for learning allowed voters to understand the systematic bias,<sup>24</sup> and voters fail to account for it even if they are informed.

A possible explanation for this behaviour is the genuinely complex environment where voting takes place. In E1, in terms of comparing election results to average revealed polls, the biased treatment condition would not appear as particularly more 'suspicious' to an active learner than the unbiased control condition. However, in the treatment condition of E2 and E3, election results systematically assigned a lower vote share to K than average revealed polls.<sup>25</sup> Despite all of this, in these experiments subjects only discounted the revealed poll information weakly in forming their beliefs. This indicates that even perfect a priori information, in conjunction with subsequent feedback, are not enough to ensure that voters sufficiently discount the results of biased polls. Participants seem to anchor their beliefs on average revealed polls and insufficiently adjust for the feedback they receive.

In summary, in this study we proposed a novel experimental paradigm for examining the effects of voters being exposed to a non-random subset of available poll information. Applying this paradigm in a new set of experiments, we found some consistent – but in our main condition not very strong – evidence in favour of electoral effects of biased polls. More evidence using this paradigm in the lab and the field is needed (including replications of this study) before safe policy conclusions can be made.

### Data availability

Data and code are available at <https://github.com/lunzheng-li/biased-poll-replication-package>.

### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ejpoleco.2023.102383>.

### References

- Agranov, M., Goeree, J.K., Romero, J., Yariv, L., 2017. What makes voters turn out: the effects of polls and beliefs. *J. Eur. Econ. Assoc.* 16 (3), 825–856.
- Bock, O., Baetge, I., Nicklisch, A., 2014. hroot: hamburg registration and organization online tool. *Eur. Econ. Rev.* 71, 117–120.
- Brehm, J.W., 1966. *A Theory of Psychological Reactance*. Academic Press.
- Brehm, J.W., 1972. *Responses to Loss of Freedom: A Theory of Psychological Reactance*. General Learning Press.
- Brown, K.M., Zech, C.E., 1973. Welfare effects of announcing election forecasts. *Public Choice* 14 (1), 117–123.
- Chen, D.L., Schonger, M., Wickens, C., 2016. oTree—an open-source platform for laboratory, online, and field experiments. *J. Behav. Exp. Finance* 9, 88–97.
- Fedderson, T., Pesendorfer, W., 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* 65, 1029–1058.
- Fleitas, D.W., 1971. Bandwagon and underdog effects in minimal-information elections. *Amer. Polit. Sci. Rev.* 65 (2), 434–438.
- Forsythe, R., Rietz, T., Myerson, R., Weber, R., 1996. An experimental study of voting rules and polls in three-candidate elections. *Internat. J. Game Theory* 25 (3), 355–383.
- Gerber, A., Hoffman, M., Morgan, J., Raymond, C., 2017. One in a Million: Field Experiments on Perceived Closeness of the Election and Voter Turnout. Tech. Rep., NBER Working Paper No. w23071.
- Goeree, J.K., Grosser, J., 2007. Welfare reducing polls. *Econom. Theory* 31 (1), 51–68.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1 (1), 114–125.
- Kessler, J., Vesterlund, L., 2015. *The External Validity of Laboratory Experiments: The Misleading Emphasis on Quantitative Effects*. vol. 18, Oxford University Press Oxford, UK.
- Klor, E.F., Winter, E., 2007. The welfare effects of public opinion polls. *Internat. J. Game Theory* 35 (3), 379.
- McKelvey, R.D., Ordeshook, P.C., 1984. Rational expectations in elections: some experimental results based on a multidimensional model. *Public Choice* 44 (1), 61–102.
- McKelvey, R.D., Ordeshook, P.C., 1985. Elections with limited information: a fulfilled expectations model using contemporaneous poll and endorsement data as information sources. *J. Econom. Theory* 36 (1), 55–85.
- Meffert, M.F., Gschwend, T., 2011. Polls, coalition signals and strategic voting: an experimental investigation of perceptions and effects. *Eur. J. Polit. Res.* 50 (5), 636–667.
- Plott, C.R., 1982. *A Comparative Analysis of Direct Democracy, Two Candidate Elections, and Three Candidate Elections in an Experimental Environment*. Tech. Rep., California Institute of Technology Working Paper Series.
- Rothschild, D., Malhotra, N., 2014. Are public opinion polls self-fulfilling prophecies? *Res. Polit.* 1 (2), 1–10.
- Sinclair, B., Plott, C.R., 2012. From uninformed to informed choices: voters, pre-election polls and updating. *Elect. Stud.* 31 (1), 83–95.
- Sturgis, P., Baker, N., Callegaro, M., Fisher, S., Green, J., Jennings, W., Kuha, J., Lauderdale, B., Smith, P., 2016. *Report of the Inquiry into the 2015 British General Election Opinion Polls*. NCRM, British Polling Council, Market Research Society.
- Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. *Science* 185 (4157), 1124–1131.

<sup>24</sup> Regression results of models of learning are provided in Online Appendix B. Although all models show insignificant results, we believe that allowing the opportunity for learning is an important part of the design, and replications with larger sample sizes and more sophisticated econometrics models are needed in future research.

<sup>25</sup> See the “Opportunities for Learning” part in Online Appendix B.