

Can Biased Polls Distort Electoral Results? Evidence from the Lab and the Field¹

Aristotelis Boukouras,² Will Jennings,³ Lunzheng Li,⁴ and Zacharias Maniadis⁵

Abstract

Biased exposure of voters to the outcome of polls constitutes a risk to the principle of balanced and impartial elections. We first show empirically how modern communication (through social media) may naturally result in such biased exposure. Then, in a series of experiments with a total of 375 participants, we investigate the impact of such biased exposure on election outcomes in an environment where only a strict subset of voters has information on the quality of the two candidates. Thus, polls serve to communicate information to uninformed voters. In our treatment conditions, participants have access to a biased sample of polls' results, favouring systematically one candidate. Participants in the biased treatment conditions consistently elect the candidate favoured by polls more often than in the unbiased control conditions. Remarkably, this holds even when voters are a priori informed about the bias. Accordingly, our results indicate that – in an experimental setting at least – biased polls distort election results via two channels: (i) by distorting the information set of voters, and (ii) by providing an anchor for subjects' expectations regarding the election outcome. Overall, biased exposure distorts elections in a very robust manner.

Keywords: biased polls, candidate valence, information aggregation

JEL Classification: D72, D83

¹ 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. We would like to thank David K. Levine, Adriana Bunea, Philippos Louis, David Rahman, Joel Sobel, Erik Gahner Larsen, 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 Behaviors and Algorithms, Cyprus, and the Central European Programme on Economic Theory 2019.

² University of Leicester, Division of Economics, School of Business, email: Aris.Boukouras@le.ac.uk

³ University of Southampton, Department of Politics and International Relations, email: W.J.Jennings@soton.ac.uk

⁴ University of Southampton, Department of Economics, email: ll12g13@soton.ac.uk

⁵ University of Southampton, Department of Economics, email: Z.Maniadis@soton.ac.uk

1. Introduction

The rise of populism in western democracies over the last few years has changed the political landscape, upsetting political balances that survived for decades and bringing new forces into the forefront. Some academics attribute the phenomenon to economic factors (Funke et al., 2016; Guriev, 2018), while others to recent developments in traditional and social media (Dellavigna and Kaplan, 2007; Petrova, 2008; Boxell et al., 2017), or to cultural factors (Oesch, 2008; Van Hauwaert and Van Kessel, 2018).⁶ Regardless of the origins for its resurgence, populism has ramifications for both economic policy (Kaltwasser, 2018) and political stability. A key aspect of modern populism is distrust in democratic institutions, of which the most fundamental is elections under free media.

The role of voting-intention polls, in particular, has been under heavy criticism. Questions have been raised about the reliability of polls (Shirani-Mehr et al., 2018) and their effects on democratic elections. A key reason for such scepticism is widespread perception of poor predictive performance of polls in recent high profile elections, most notably the 2016 US presidential election and the UK general elections of 2015 and 2017.⁷ Some prominent politicians, such as Lord Foulkes in the UK and Ron Paul in the US, have even claimed that polls may be manipulated by status-quo groups in an attempt to cling to power.⁸ If people perceive polls to be biased in favour of a candidate or a party, this perception may erode trust in democratic institutions and reinvigorate populist agendas.

Thus, it is crucial to examine whether criticisms of polls, such as the ones presented above, are unfounded, or whether polls have potential to skew election results in the current political environment. If the feedback that the public receives from opinion polls is not

⁶ For broader treatments of populism, see also the papers by Boeri et al. (2018), Inglehart and Norris (2017) and Rodrik (2018). For the role of social media in the proliferation of fake news see Allcott and Gentzkow (2017).

⁷ See also Whiteley (2016) on the reasons behind the failure of polls to predict the 2015 general election in the UK.

⁸ For general economic models of such manipulation see Maniadis (2014) and Cipullo and Reslow (2019).

representative of the true preferences of the electorate, one may worry that this could distort the democratic process. For instance, in a two-party election race, imagine that poll results showing the left party ahead are more likely to be revealed to the public than poll results showing this party trailing. Then, a critical question arises: would this systematic bias in exposure to poll results affect the electoral race? How significant is such an effect and how does it depend on what voters know about the bias? In this paper, we first show that this type of biased exposure to the results of polls may arise naturally from the dynamics of modern communication (social media). We then examine the causal effect of such biased exposure on elections. We use the experimental approach and a series of robustness checks and find that biased exposure consistently leads to meaningful and systematic changes in election outcomes.

Even leaving aside the possibility of conscious manipulation, several plausible mechanisms could generate systematic bias on the feedback that citizens receive about the results of voting-intention polls (Sturgis et al., 2016). First of all, pollsters have methodological flexibility similar to other empirical scientists (Ioannidis, 2005) and if they have strong priors about who is leading, they may choose methods that verify these priors (for example, turnout adjustments). Moreover, the traditional media reveal poll results selectively, either to match the expectations of their audience (Gentzkow and Shapiro, 2010) or to simply make interesting news (Larsen and Fazekas, 2019). Finally, the voters themselves may propagate results in a biased manner, especially via social media. Our objective in this paper is to first empirically substantiate such biased feedback and then examine its causal effects on election outcomes.

It is difficult to find appropriate data to substantiate most of the mechanisms described above. For this reason, in Section 3 we shall provide empirical evidence regarding only the last-mentioned channel that mediates communication and may lead to biased propagation of poll results: online publics. We use Twitter data from the US and the UK and examine econometrically how the patterns of retweeting of news about poll results are affected by the

results themselves. We find systematic biases in the manner in which poll results are propagated via social media. In particular, there is a fundamental asymmetry between parties in the pattern of propagation. In our data, ‘good news’ about the popularity of conservative parties seem to receive less propagation than ‘bad news’, whereas the opposite is true about liberal parties.

After having empirically established the existence of biased exposure of the public to poll results, we turn to an examination of its consequences. Given the great difficulty of using observational data not only to measure the degree of bias in polls but also to examine the electoral consequences of this, we take an experimental approach to address the question. This allows us to control voters’ information both at the poll and at the election stage. We focus on an environment of two-party elections and we postulate a straightforward ‘biased rule’ according to which the revelation of poll results takes place. In particular, only the most ‘favourable’ results for a particular party (the ‘favoured party’) are revealed. This censored rule captures the essence of the idea that propagation of poll results to the public depends on the results themselves. This information pattern could ensue from any of the reasons discussed earlier: conscious manipulation, pollsters’ priors that the favoured party is ahead (coupled with methodological flexibility), traditional media wishing to match the public’s expectations, the biased way social media propagate poll outcomes, etc.

In our experiments, we observe the outcome of fifteen electoral races between the same two parties (we call them parties K and J) who field different candidates every time. The two candidates differ in their ‘valence’, and the exact valences are known to only 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 some signal about 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 addition, we conduct two robustness checks. 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), we keep the same control condition as E1, but in the treatment condition participants are informed beforehand about the (non-random) rule for selecting the two polls.

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 vote share of the favoured candidate K by an average of 20 percentage points, 11.7 percentage points and 7.3 percentage points in E1, E2 and E3, respectively. These differences are very consistent across sessions and rounds and their magnitudes are meaningful politically. Importantly, these effects do not go away as participants gain more experience.

There is some evidence in E1 that learning fails in our environment, allowing biased polls to distort democratic outcomes over prolonged periods. It seems that the self-confirming nature of biased polls limits the scope of receiving the type of feedback that would reveal the bias. Perhaps more remarkably, explicitly informing voters in E3 about the biased rule for revealing poll results does not eliminate the electoral advantage that the bias yields to the 'favoured' candidate K. This indicates that voter unawareness about the bias of polls is not the only factor that drives our results. Even when they are aware, subjects do not appear to rationally weigh the information content of polls. Instead, it seems that, in forming their expectations about the electoral results, voters use polls merely as judgemental anchors, so they

overweight the reference point that polls provide and they underweight their informational content. This interpretation is consistent with the well-known process of anchoring-and-adjustment (Tversky and Kahneman, 1974).

An analysis of the relationship between participants' beliefs (regarding the election winner) and the revealed poll results supports the aforementioned mechanisms. We find that these beliefs are highly correlated with the average revealed poll vote shares, both in the treatment and in the control setting, especially in E1 and E3. Econometric results further indicate that beliefs do not increasingly deviate from revealed poll results as time passes. This means that voters do not discard or discount poll results in later rounds, indicating that very limited learning takes place. Moreover, averaged revealed poll results are a good predictor of electoral results in the treatment condition for all three experiments, although these polls were selected in a biased manner.

Overall, there is only weak evidence that participants are either able to realise the biased nature of polls or that they can sufficiently account for it when they are informed of it. These failures lead subjects to overestimate the popularity of the favoured candidate and to vote for her more frequently. Therefore, biased polls can have a significant and robust impact on election outcomes even when the public knows or suspects the bias. Consequently, our experiments may inform the public debate on whether or not biased polls can skew behaviour in real election settings.

The rest of the paper is structured as follows. Section 2 places our findings in the relevant literature. Section 3 presents our observational study on the propagation of poll results on social media. Section 4 discusses the design of our three experiments. In Section 5 we present descriptive results of our experiments, whereas in Section 6 we provide the econometric analysis of voting behaviour. Section 7 presents a short discussion of our findings and concludes.

2. Related Literature

The effects of polls on election outcomes have been the topic of both theoretical and empirical study. This large literature contains important experimental studies, but as far as we can tell, none of them considers biased feedback on actual polls combined with opportunities for learning. Economic experiments 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 Ordershook, 1984; McKelvey and Ordershook, 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.⁹ This is closer to the approach taken in political science, where many experiments strategically manipulate the

⁹ 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 subjects 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 very strict approach, even information about other subjects' 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 subjects form beliefs about whether information is biased or not should be an open research question, not a forbidden one.

poll information that participants receive. Typically, these experiments are non-incentivised. The early study by Fleitas (1971) indicates that voting is not responsive to 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 Rothshchild 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 affected by the manipulation but behaviour not so much. As with previous experiments, rational choice theories predicting turnout do not perform very well.

The main difference between the aforementioned political science studies and ours is that these studies are not examining whether subjects are capable of understanding that manipulation is taking place and of accounting for it. In particular, in these studies participants face biased or manipulated polls only once, so they do not learn from past mistakes. Our design allows for multiple rounds of repetition so that we explore the participants' scope for learning. We believe that this is a critical aspect, as a standard approach of rational choice theory is to consider 'equilibrium' behaviour, i.e. crystallised behaviour after the effects of learning have taken place. 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. This indicates that biased polls influence voters through multiple channels. 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.

3. Bias in Online Propagation of Poll Results: Evidence from the Field

Although we mentioned several plausible mechanisms that could result in selective exposure of the public to poll results, most of them are difficult to study empirically. For instance, we do not have access to the pool of all methodologies that pollsters have in their disposal, neither are there published data on the set of poll results available to the media when they choose what news to broadcast. For this reason, we shall resort to showing that biased exposure can also ensue from the natural structure of modern political communication, namely social media.

To what extent do online publics paint a representative portrait of the existing results of opinion polls? We shall show that the nature of social networks results in a biased exposure of the public to poll results. People have various cognitive mechanisms that result in selective attention, such as negativity bias (Soroka, 2014), motivated reasoning (Taber and Lodge, 2006), cognitive dissonance (Morwitz and Pluzinski, 1996) or disproportionate responsiveness to outliers. Users of social media are also not demographically or politically representative of the general population (Mellon and Prosser, 2017), which could give rise to further biases in attention, via selective reporting. As a result, individuals attend to and, crucially, propagate to others, the results published by polling firms in a systematically biased manner.

In this analysis, we examine the biased propagation of published opinion poll estimates or trackers in the United States (US) and the United Kingdom (UK). Specifically, we consider measures of voting intentions for US Presidential elections (reported by HuffPost pollster.com) and for UK parliamentary elections (YouGov's political tracker). This enables us to assess the spread patterns of the published poll results in two different countries. Our objective is to show that in some real-life electoral races a subset of voters is exposed to poll results in a manner that systematically depends on the results of the polls themselves.

3.1. Opinion Polling in the US and UK

While opinion pollsters in the US and UK ask a wide variety of survey questions on political issues, among the most prominent measures of political attitudes are for presidential elections (in the US) and Westminster voting intentions (in the UK). These are central to depictions of the ‘horse race’ by media (Iyengar, 1991; Matthews et al., 2012). In the US, George Gallup famously introduced random sampling methods to measure national voting intentions in the 1936 presidential election. Variants of the question “If the election were held today, whom would you vote for?” have been asked regularly ever since. During the 2016 presidential election campaign there were well over 400 national opinion polls of voting intentions for Donald Trump and Hilary Clinton, yielding a steady flow of information on the election horse race.

In the UK, pollsters have been asking people about their voting intentions for Westminster parliamentary elections since 1943 (see Wlezien et al., 2013; Sturgis et al., 2016). YouGov has become one of the highest volume pollsters in the UK since their introduction of online methods in 2001, regularly fielding the question “If there were a general election held tomorrow, which party would you vote for?” During the first Cameron government, it fielded a survey almost every other day.

3.2. Opinion Polling Data on Social Media (Twitter)

HuffPost Pollster and YouGov each report their latest poll estimates via their official accounts on the social media platform Twitter (in the case of HuffPost this is polls conducted by other polling firms). This provides a regular stream of poll information that enables us to analyse patterns of selective reporting, by social media users, in an observational setting. With frequent estimations of public opinion (at least every other day), most fluctuations in poll estimates are attributable to noise due to sampling error (even where dampened by poll aggregators), and thus most users are (arguably) reacting to random short-term fluctuations,

rather than systematic trends.¹⁰ While it would in theory be possible to collect data on wider engagement with poll estimates on Twitter, this approach enables us to model a fairly stable source of poll information.

We obtained relevant tweets of poll estimates from @pollsterpolls and @YouGov using an advanced Twitter search with terms corresponding to the standard form of poll reporting used by each organisation (removing all extraneous cases from the scraped data). Details of these search terms are provided in Table 1. All the tweets report the current *level* of voting intentions for the relevant candidate or party. We calculate the change in voting intentions from the previous poll estimate in our dataset. This forms the independent variable of our analysis – the change in observed poll estimates.

Table 1. Twitter reporting of poll estimates in the US and the UK

	US – Presidential election voting intentions	UK – General election voting intentions
Choice	Trump/Clinton	Labour/Conservatives/Liberal Democrats
Pollster	All pollsters	YouGov
Start	8 September 2015	9 April 2010
End	8 November 2016	8 December 2017
Measure	Voting intention, by candidate	Voting intention, by party
N of polls	445	1,451
N of days	428	2,801
Polls per day	1.04	0.52
N of Δ in vote	444	1,450
Retweets	5,054	41,291
Source	@pollsterpolls	@YouGov
Search terms	“2016 General Election”, “Trump”, “Clinton”	“Lab”, “Con”, / “Westminster voting intentions”

Crucially, we also collected data on the number of ‘retweets’ for each tweet. This provides us with a measure of online propagation of the poll result, our dependent variable. On average, each poll estimate received 24.4 retweets, with an upward trend over time in the

¹⁰ We will show that a systematic bias in the propagation of poll results is even evident in responses to short-term noise, rather than more sizable long-term trends. With such trends we should expect such bias to play an even more important role.

number of retweets as usage of Twitter grew. Our analysis undertakes an ordinary least squares regression of the number of retweets of a given poll estimate (*Retweets*) as a function of change in candidate or party support ($\Delta Vote$). In the US we focus on change in the ‘margin’ between the candidates, i.e. the lead of Clinton over Trump. In the UK we focus on change in support for the Labour, Conservative and Liberal Democrat parties. This focus on *change* enables us to determine whether biased propagation of poll results can stem from mere short-term fluctuations, rather than structural differences between particular candidates or parties.¹¹ The estimated models therefore take the following form, where *Equation 1* refers to the US, and *Equation 2* to the UK.

$$\text{US: } \textit{Retweets} = a_0 + b_1 \Delta(\textit{Vote}(\textit{Clinton}) - \textit{Vote}(\textit{Trump})) + \varepsilon \quad (1)$$

$$\text{UK: } \textit{Retweets} = a_0 + b_1 \Delta \textit{Vote}(\textit{Con}) + b_2 \Delta \textit{Vote}(\textit{Lab}) + b_3 \Delta \textit{Vote}(\textit{LD}) + \varepsilon \quad (2)$$

The results for this analysis are reported in Tables 2 and 3. These reveal largely consistent, and also interesting, patterns (both within and across countries) in attention to poll estimates. In the US (Table 2), an one-unit increase in the Clinton-Trump lead in polls reported by HuffPost Pollster was associated with 1.0 additional retweets of the poll estimate. This might signify the partisan lean of the users of Twitter or the followers of this specific polling account, but it does hint at a selective reporting mechanism of poll estimates that fundamentally distorts voters’ (salient) information on the popularity of candidates. We should emphasise that we wish to establish empirically the existence of this distortion, and we do not wish to claim causality.

In the UK (Table 3), we see a similar pattern whereby a one-unit increase in voting intentions for the Conservative Party leads to 2.5 fewer retweets of the poll. In contrast, a one-unit increase in support for Labour leads to extra 7.3 retweets. There are no systematic

¹¹ This focus on short-term dynamics enables us to show that even between adjacent days there is a systematic bias in propagation, and in particular selective propagation occurs regardless of how popular a candidate or party is.

differences for the Liberal Democrats, at least observed during this period. Predicted values of the regression models are depicted in Figure 1. These confirm the findings: there are distinct partisan differences in the online promulgation of poll results, specifically a bias where increases in support for left-wing parties/candidates are propagated more in online platforms, whereas it is drops in that support that receive wider spread for right-wing parties/candidates. In the UK context, interestingly, the pattern is more pronounced for Labour than the Conservatives, so this is not a purely symmetrical relationship.

Table 2. Selective propagation of poll estimates of the US 2016 presidential election

	Retweets
$\Delta(\text{Clinton-Trump})$	1.021 (0.148)***
Intercept	11.353 (0.666)***
N	444
R-squared	0.10
Adjusted R-squared	0.10

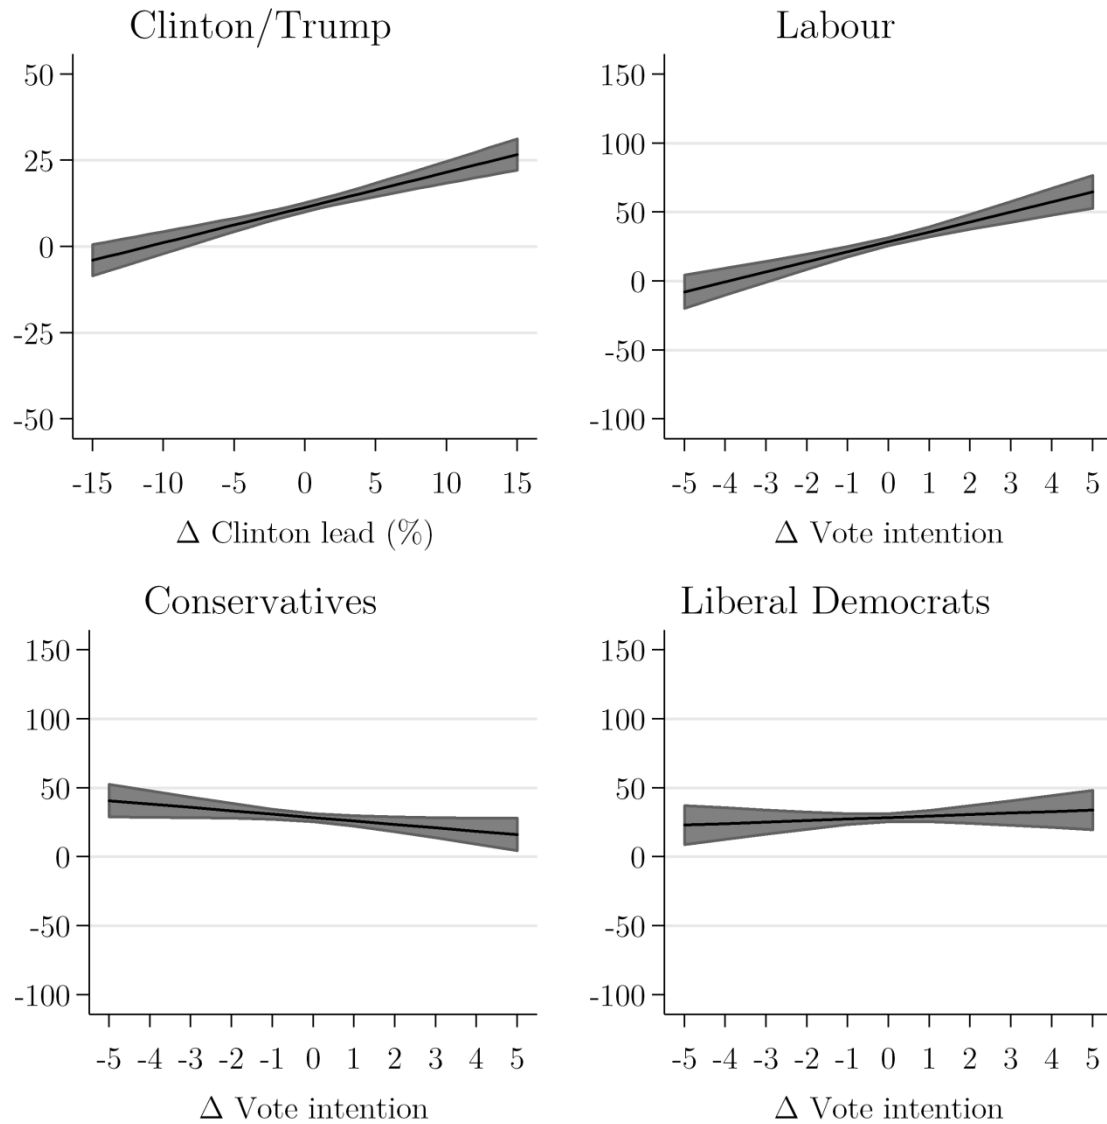
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 3. Selective propagation of poll estimates of voting intention, UK

	Retweets
$\Delta\text{Vote(Con)}$	-2.464 (1.162)*
$\Delta\text{Vote(Lab)}$	7.335 (1.187)***
$\Delta\text{Vote(LD)}$	1.110 (1.419)
Intercept	28.423 (1.480)***
N	1,450
R-squared	0.04
Adjusted R-squared	0.04

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Figure 1. Adjusted predictions (with 95% confidence intervals) of the number of retweets, by Δ Vote



We have thus shown empirically that in modern democracies the public is likely to be exposed to the results of pre-election polls in a biased manner.¹² Now we may ask: which are the implications of such a bias for democratic elections? If biased exposure skews elections,

¹² Of course, given the network structure of the social media, it is not true that the same biased sample of polls is revealed to every voter (as is the case in our experiment). For instance, if left-leaning people mostly re-tweet to other left-leaning people, right-leaning people may not be exposed much to those polls. In fact, there is evidence that “the network of political retweets exhibits a highly segregated partisan structure, with extremely limited connectivity between left- and right-leaning users.” (Conover et al., 2011). Accordingly, we are not claiming that the empirical pattern established here always generates a biased propagation pattern similar to the one used in our experiments. However, our correlational results do establish that different subsets of voters are exposed more to a certain type of results rather than another type of results. As stated before, we believe that there are other propagation mechanisms that could more plausibly lead to an information pattern similar to the one used in our experimental design.

then we should be concerned and maybe need to address this by policy changes. The problem is that establishing a causal relationship about such a complex phenomenon in the field can be very difficult. For this reason, we employ the experimental method increasingly popular in economics and political science (Palfrey, 2016). The virtue of this approach is that it can establish causality and discover general patterns of social behaviour in a controlled setting.

4. Our Experimental Environment

In general, the information conveyed by poll results can be relevant to voters for many reasons (e.g. 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, voters have bandwagon preferences, etc.). 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 meaningful for the aggregation of 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 treatments 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 manipulation (two polls out of five) is conveyed, but in a random, rather than a systematically selective, manner. 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.

In a final experiment (E3), we also 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 systematically biased (and subjects know this fact). We perform this test by replicating experiment E1 with one important modification. In particular, in the treatment condition, participants are informed explicitly about the (biased) selection rule. All experiments are described in detail in Table 4 (page 21).

4.1. 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 Figure 2. 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 rounds¹³ 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.

Each candidate’s valence is drawn at the start of every round from a uniform distribution with values between 0 and 120.¹⁴ 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 α 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, $Q_h \sim U[0,120]$.


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

¹³ The interpretation is that the two parties consistently pick candidates that share their ideological views.

¹⁴ 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.

voters. As explained earlier, it is important that the distribution of informed voters is not symmetric in the ideological spectrum. If that were the case, then the socially efficient outcome would be for uninformed voters to abstain from elections and let voting amongst 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.

Figure 2. Ideological preferences in the experimental interaction

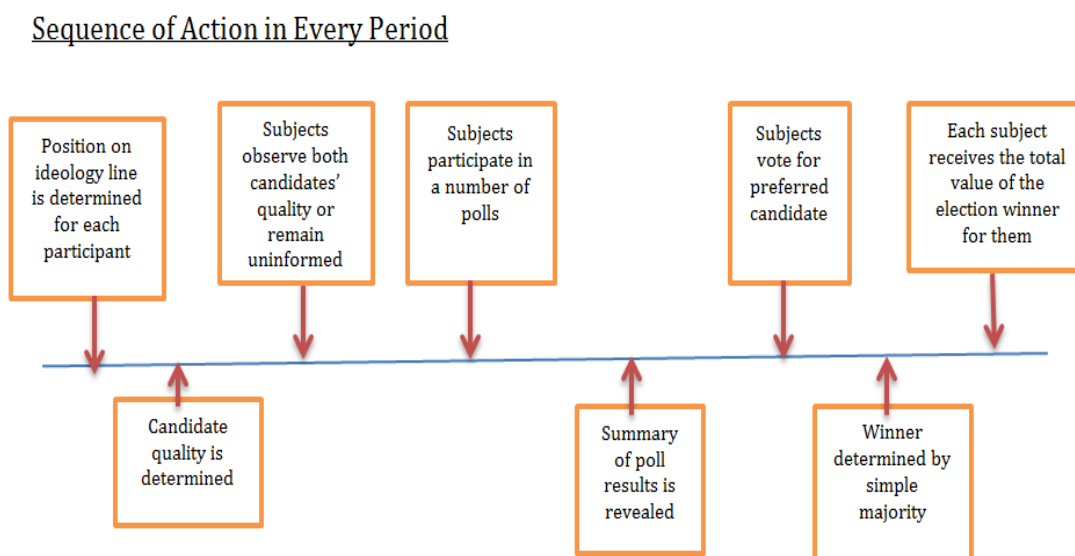


For example, assume that some uninformed voter observes substantial support in favour of K in the polls. If she perceives polls as unbiased and other subjects as rational, she will infer that K's valence is higher than J's, since some informed voters who are close to J's ideological position prefer to 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, if she is sufficiently close to the centre-left of the political spectrum.

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).¹⁵ 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.¹⁶ 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 subjects) is illustrated in Figure 3. The winner of the election is determined by simple majority, with ties broken by a 50-50 coin toss.

Figure 3. Sequence of actions in each experimental period



¹⁵ Voters choose from the following three options: 'K', 'J', and 'Prefer not to participate'. Polls do not contain information on 'non-participation', as this substantially simplifies the feedback that subjects observe about the results of polls. This is especially important, since subjects need to infer overall support for each candidate on the basis of results from multiple polls.

¹⁶ 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.

4.2. The Three Experiments

The only stage (out of those displayed in Figure 3) that differs across the two treatments in each of our three experiments is the one where “summary of poll results is revealed”. Table 4 describes our experimental design and Table 5 illustrates the information selection in the treatment environment. 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 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. Instead, they have the chance to infer such biases and agendas through experience. In our experimental setting, this is accomplished because voters can compare poll predictions with actual election results in every round. In Experiment 3, we examine the consequences of providing a priori information about the exact nature of the bias to voters. Table 4 below summarises the three experiments and the relevant ‘control’ and ‘treatment’ conditions in each one of them.

Table 4. 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.

Table 5. 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%

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 the treatment condition of all three experiments, only the two results of companies B and E would be revealed if these were the actual five sets of results.

Experiments E1 and E2 had 120 participants each,¹⁷ with eight 15-subject sessions (four control sessions and four treatment sessions).¹⁸ 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

¹⁷ 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.

¹⁸ We denote individual sessions as Ei_Cj or Ei_Tj 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, $E1_C1$ denotes the first control session in E1 and $E2_T1$ the first treatment session in E2.

30 subjects) to achieve perfect randomisation by containing one control and one treatment session, with participants being randomly allocated between the two.¹⁹

In each session, subjects read instructions from their computer screens.²⁰ 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 Figure 3. The only aspect that was not described is the ‘belief elicitation’ stage. In particular, after the release of the polls, participants were asked to state their beliefs about the vote shares of the two candidates in the elections. The information about polls took the form of a single probability distribution for each result, as shown in Table 5. Subjects’ beliefs at the elicitation stage were also described in terms of this binary probability distribution.

5. Results and Descriptive Analysis

Let us first provide an overall summary of the *primary treatment effect* across the three experiments: the rate of electoral success. Table 6 illustrates the number of rounds won by each of the two parties in the treatment and control conditions across the three experiments. As can be seen, 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. As we shall see in detail later, these differences are relatively homogeneous in their magnitude and extremely consistent in their sign, both across sessions of a given treatment and across rounds of a given session. In terms of statistical significance of the differences in individual experiments, the difference is significant in E1

¹⁹ The only three exceptions in this approach were sessions *E1_C2*, *E1_T2* and *E3_T5*, which were the only sessions of their block because of insufficient subject participation or lab capacity constraints.

²⁰ We programmed the experiments using O-tree (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 University of Newcastle and of York respectively.

(Fisher exact test with one-sided alternative) but not so in E2 and E3. Still, the differences are politically significant and very consistent, as we shall illustrate now.

Table 6. Number of elections won for each party in each treatment and results of Fisher's exact test

	E1		E2		E3	
	Control	Treatment	Control	Treatment	Control	Treatment
K	36	48	37	44	34	48
J	24	12	23	16	26	27
p-value	0.0138		0.121		0.245	

The alternative is that the number of rounds won in treatment condition is higher than in the control condition.

5.1. Experiment E1

Recall that in Experiment E1, the 15 participants in each control session voted every period after having been exposed to the results of all five polls, while in the treatment condition, the respective 15 participants were exposed to the two polls that had the greatest voting intention for the candidate of party K (but this was not explicitly stated). The most important general finding is summarised by descriptive analysis: 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. Figures 4 and 5 juxtapose the fraction of election rounds won by K and vote shares for K in treatment vs. control sessions. It is clear that the electoral performance of K is consistently better in all treatment sessions relative to any control session. K won more rounds than J in both treatment and control sessions. This is, however, to be expected since (by pure chance) there were more rounds where the randomly drawn valence for K was higher than the drawn valence for J. In fact, in 11 out of the 15 regular rounds K has higher valence than J, and in 9 of those the difference in favour of K is over 20 points.

Figure A in the appendix indicates that there is enough heterogeneity in the findings of the five polls, so that revealing a biased selection of poll results is meaningful. For almost all rounds of the treatment sessions, the vote share of K differs substantially across polls, so

selecting the ones with the highest share gives a non-representative image of the average vote share of K.

Figure 4. Fraction of rounds won by K in each experimental session of E1





Figure 5. Average vote share for K in each experimental session of E1



Furthermore, the difference in vote shares does not appear only at the average level, but also for each individual round. Figure 6 shows the average vote share of treatment and control condition that candidate K received in each round (averaging across the four sessions of each treatment). The figure indicates that ‘treatment’ rounds have consistently higher vote shares for K than ‘control’ rounds. In fact, vote shares in ‘treatment’ are higher than vote shares in

‘control’ for all rounds. This is important because it does not seem to be the case that the difference vanishes in the last few rounds. Accordingly, these data are consistent with the interpretation that participants behave as if they perceive polls in the treatment as unbiased: they do not seem to be discounting them, even after several opportunities for learning. We shall now delve deeper into this important issue.

Figure 6. Comparison of vote share round-per-round in E1





Figure 7. Relationship between elicited beliefs and poll results revealed in E1



Evidence from Beliefs

As we explained previously, after the ‘summary of polls’ stage and before elections, participants were asked to state their beliefs about the vote shares of the two candidates in the upcoming elections. We use this elicitation of subjects’ beliefs to examine whether they are in alignment with the poll information that participants received. 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 would lead to a low correlation between their beliefs and the average of the revealed poll information. However, Figure 7 shows that the correlation is clearly not larger in the control sessions relative to the treatment sessions. Figures B and C in the appendix illustrate this relationship in more detail. In particular, they juxtapose (in each round and session) the average vote share of K according to revealed polls and the analogous vote share that subjects expect according to their beliefs. In both conditions, beliefs closely follow the average vote share revealed in polls, with no discernible pattern of differences. This is consistent with the idea that participants perceive polls as unbiased in both the treatment and the control condition.

5.2. Experiment E2

Experiment E1 compared the electoral results in a ‘biased regime’, where there is a systematically biased selection of poll results revealed to the public, to a ‘full information’ regime. This full information regime is a natural benchmark to consider: the public is informed about the totality of relevant evidence for democratic decision-making. However, a weakness of this benchmark is that it provides more information than the control treatment (the results of five polls instead of two). For this reason, it is important to also employ a control condition where the amount of information is similar to the treatment (the ‘biased regime’). For this purpose, we conducted the same number of randomised blocks (four 30-subject blocks) in an additional experiment (Experiment E2) where the control treatment revealed the results of only two out of the five polls, and these two polls were chosen randomly.

Figures 8 and 9 contain the basic descriptive results from this additional experiment. The evidence points consistently to the direction observed in Experiment E1, but the treatment effects are smaller. This should not be surprising, if one considers the censored nature of the results revealed in the treatment condition. This entails that the treatment and control conditions

are closer to each other in E2 (in terms of revealed poll results), than in E1. In other words, by pure chance the poll results revealed in the treatment and the control in E2 can be close to each other or even identical, which (almost certainly) cannot be the case in the comparison between treatment and control in E1.

Figure 8. Fraction of rounds won by K in each experimental session of E2




Figure 9. Average vote share for K in each experimental session of E2




Figure 11 illustrates the vote share round-by-round. Once more, the pattern is that the vote share for K is uniformly higher in the treatment than in the control, while the difference does not seem to disappear with learning. These differences appear somewhat smaller than in E1. However, the difference is still meaningful: the number of rounds won by J in the control is nearly 50% larger than in the treatment (23 vs. 16). Overall, the consistency of the pattern indicates that the biased release of poll information has relatively robust and predictable effects on electoral results.²¹

Evidence from Beliefs

There is no a priori reason to expect that subjects in the treatment condition of E2 would behave differently than in the treatment condition of E1, since these conditions are practically identical. However, a careful inspection of Figures E and F of the appendix reveals that now some systematic patterns of discounting poll results might exist. In particular, in Figure E it seems that average beliefs in the treatment sessions tend to be lower than average revealed poll results. This means that subjects seem to somewhat discount the (inflated because of bias) advantage in favour of K presented in the polls. This pattern does not seem to hold for the control sessions, as evidenced by Figure F. Figure 10 below tends to confirm this pattern. In particular, it now appears possible that the correlation between average beliefs and revealed polls is systematically higher in the control (where information is unbiased) than in the treatment (where information is biased). However, as we shall see in Section 6, further econometric analysis cannot provide evidence for learning effects.

²¹ 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 significant differences. However, implementing this would probably be too burdensome for participants in our current experimental environment.

Figure 10. Relationship between elicited beliefs and poll results revealed in E2





Figure 11. Comparison of vote share round-per-round in E2



5.3. Experiment E3

Experiment 1 showed in a particularly robust manner that election outcomes in a ‘biased feedback’ environment (where poll feedback is systematically selected to maximise the seeming popularity of a particular party) are distorted by this reporting bias. The effect is large in political terms: compared to the ‘transparent democracy’ benchmark, where the information from all polls is revealed, the ‘biased feedback’ environment resulted to an increase of winning rate for party K (favoured by the biased feedback) by twenty percentage points. Further analysis will show that the experimental results in E1 are consistent with the notion that the systematically selected poll results become self-confirming, which prevents feedback that could have exposed the biased rule underlying them. Experiment 2 showed that the treatment effect is robust (but its magnitude is unsurprisingly smaller) when one accounts for the fact that


the ‘transparent democracy’ benchmark has more information revealed than the ‘biased feedback’ condition.

However, 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 most 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 make 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 them.

To address these concerns, we run a third experiment (E3) where the treatment condition entails using the same biased rule as in the treatment conditions of E1 and E2 but with full transparency 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 sessions of the ‘fully transparent biased rule’ and four control sessions with the ‘transparent democracy’ information environment) are illustrated in Figures 12-14. As can be seen, even in this case, the biased feedback rule seems to offer an advantage to candidate K. 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 ‘fully transparent 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 striking.

Figure 12. Fraction of rounds won by K in each experimental session of E3



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. Figure 14 additionally shows that the difference does not depend on the particular round, but that is rather sizable for every individual round. Again, the picture that emerges is that biased exposure of the public to poll results affects elections in a pretty robust manner.

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 fully cognitively discount for the bias in the polls, they only seem to use them as anchors and maybe adjust them until they reach an acceptable range. The use of such as heuristic is reasonable, given the demanding learning environment of the voting interaction.

Figure 13. Average vote share for K in each experimental session of E3

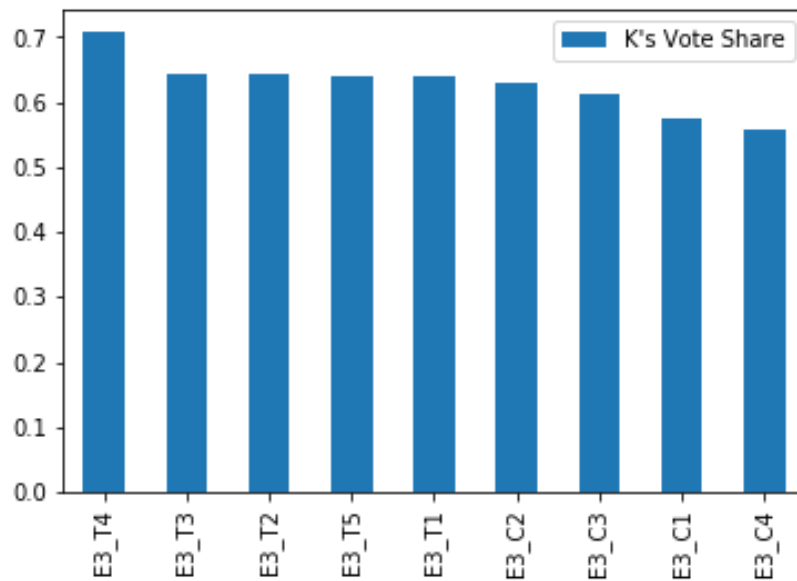



Figure 14. Comparison of vote share round-per-round in E3



Evidence from Beliefs


The comparison between average results of revealed polls and elicited beliefs becomes very interesting, especially compared to experiments E2 and E1. Figure 15 illustrates the correlation between average beliefs and revealed polls. Remarkably, no systematic pattern

emerges. In particular, it does not appear to be the case that subjects account for the bias when they form their beliefs. Beliefs are not closer to revealed poll results in the unbiased control condition than in the biased treatment condition, a finding consistent with the belief correlations of E1, but not of E2. This is very surprising, since on the one hand you have a set of poll results known to be a biased sample from the available evidence on voters' preferences, and on the other hand you have the totality of the evidence. Yet, results indicate that subjects do not appear to distrust the first set of results more than the second one, reinforcing the interpretation that polls generate judgemental anchors for beliefs and adjustment is insufficient.

5.4. Additional Descriptive Analysis

In terms of the welfare effects of biased polls, a rough measure of utilitarian welfare is the average experimental payoffs in each condition. Intuition suggests that biased polls should have a negative impact on this measure as they introduce noise in the information conveyed by polls to voters. Moreover, if voters do not discount the information of 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 7 below. 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.

Figure 15. Relationship between elicited beliefs and poll results revealed in E3**Table 7.** 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

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 (the opposite direction was not observed). In E2, while in the control condition there were 3 elections where the high-valence candidate lost, this increased to 8 elections in the treatment condition.²² 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

²² Out of all these instances, only once did J win when K had the higher valence (it happened in the control condition).

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.

It is worth emphasizing that for an overall assessment of the social, economic and political implications of our experimental results, these utilitarian welfare effects are only of complementary importance. Changing the margins of victory also has important implications for the parliamentary representation of parties and for long-run political competition, and these aspects cannot be captured by this limited analysis.

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 higher payoffs, which they can easily calculate.²³ Accordingly, if these individuals' votes deviate from 'optimal behaviour' this would indicate that the assumption of rational, money-maximising political agents is violated. Table 8 below illustrates the behaviour of informed voters. For instance, in 8.57% of the 420 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.

Nonetheless, a non-trivial fraction of decisions, slightly lower than 15% for the control 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

²³ For simplicity, we shall call 'h-voter' an informed candidate whose money-maximising choice is candidate h, where $h \in \{J, K\}$.

is not captured by monetary payoffs. Interestingly, J-voters are more likely to vote for candidate K in the treatment setting than in the control, and within the treatment setting 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.

Table 8. Behaviour of informed voters

Preferred/ Voted for	E1		E2		E3	
	Percent of total Choices		Percent of total Choices		Percent of total Choices	
	Cont.	Treat.	Cont.	Treat.	Cont.	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

‘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 525 such decisions).

It is also important to discuss the behaviour of voters at the poll stage. Figure J in the appendix provides an overall summary of the voting behaviour in the different sessions in our three experiments. The results are broken down by different status of voters (informed vs. uninformed). Certain insights can be inferred from Figure J: informed voters are more likely to participate to polls, while uninformed voters are more likely to vote for K in polls (which makes sense, since their ideologies are closer to K). Moreover, there seems to exist no significant difference between treatment and control, which is again unsurprising, since the treatment is different from the control only when voters observe poll results.

Table 9 compares the voting choice at the poll stage to the one at the actual elections.²⁴ 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, 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%.²⁵ These were mainly uninformed voters who are closer to J, but were induced to switch to K in the elections because of the treatment.

Table 9. Comparison of individuals' voting at the polls vs. the real election

	E1		E2		E3	
poll/election	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%

6. Regression Analysis

The preceding descriptive analysis shows that biased exposure to polls increases the likelihood of 'favoured' candidate K being elected. A key question is why this is happening.

²⁴ 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 9 contains information for about 70% of overall decisions.

²⁵ 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 a subject voted both in the polls and the elections) that a voter chose J in the polls but K in the elections. The higher occurrence of this in the treatment condition can be interpreted as a treatment effect.

Because of the relatively complex environment that we are studying, it is unlikely that voters use the strategic structure of the environment to predict behaviour deductively, thus we shall focus our analysis on the effects of feedback and learning on beliefs and behaviour. In particular, in E1 and E2, do subjects manage to learn that in the treatment condition the revealed polls are not a representative image of subjects' preferences at that particular time? What is the relationship between the poll information that subjects observe, their beliefs and election results? In the following, we shall employ our measures of subjects' beliefs and try to tackle these questions. The first model we estimate (Model 1) takes the following form:

$$B_t = a + b_1 P_t + b_2 T + b_3 (T * P_t) + b_4 D + b_5 (D * P_t) + e_t \quad (3)$$

We consider one round as the unit of observation, so data are at the session level. The dependent variable B_t is the average subjects' beliefs about candidate K's vote share in period t . P_t is the share of voters supporting K that can be inferred by the revealed polls in round t . For instance, in E1, treatment condition, this share is derived as the average of two polls, while in the control condition, this share is derived as the average of five polls.²⁶ T is the 'late rounds' dummy variable taking the value 0 for early rounds (rounds 1 to 10) and 1 for late rounds (rounds 11 to 15). D is the treatment dummy (1 if the session is in the treatment condition, 0 otherwise).

We are principally interested in the coefficients of the two interaction terms. A significant negative coefficient in the interaction term $P_t \cdot T$ would indicate that the degree to which revealed poll results affect beliefs weakens through time. This would be consistent with the notion that subjects distrust polls at the treatment condition (but we would not expect the same for the control condition). On the other hand, a significant negative interaction between P_t and D would imply that in the treatment condition there is a weaker relationship between

²⁶ Thus, this specification models voters as rather unsophisticated, forming inferences about each candidate's support by merely taking the average of the polls revealed to them.

beliefs and average announced poll results. Of course, we would expect such a negative interaction to exist in E3, since participants are explicitly informed about the bias.

As Table 10 indicates, the results of the model do not support the notion that subjects in E1 are able to learn and discount manipulation in the treatment condition. In particular, there is no significant interaction between the ‘late rounds’ dummy and average announced poll information, although the respective coefficients are negative in both the treatment and the control. Tables 11 and 12 show a similar pattern. Interestingly, the coefficient b_3 is positive in the control but negative in the treatment setting in both E2 and E3. However, none of this is statistically significant. On aggregate, there seems to be very weak, if any at all, evidence that subjects somewhat discount poll information in late rounds. The experimental condition also does not seem to make a difference: coefficient b_5 does not have a consistent sign across the three experiments and it is not statistically significant in any of them.

Table 10. Results of Model 1 in E1

	Pool	Treatment	Control
Avg. poll info.	0.858*** (0.029)	0.848*** (0.035)	0.856*** (0.034)
Late Rounds dummy	12.350** (6.183)	3.823 (9.519)	16.819* (9.718)
Late Rounds dummy *Avg. poll info.	-0.114 (0.076)	-0.021 (0.109)	-0.166 (0.133)
Is treatment*Avg. poll info.	-0.005 (0.045)		
Is treatment	-0.740 (3.245)		
Constant	9.185*** (1.711)	8.975*** (2.655)	9.094*** (1.895)
Observations	120	60	60
R-squared	0.945	0.931	0.942

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 11. Results of Model 1 in E2

	Pool	Treatment	Control
Avg. poll info.	0.798*** (0.021)	0.787*** (0.029)	0.815*** (0.023)
Late Rounds dummy	-4.151 (4.133)	14.000 (12.397)	-2.863 (4.684)
Late Rounds dummy *Avg. poll info.	0.080 (0.049)	-0.092 (0.133)	0.036 (0.062)
Is treatment*Avg. poll info.	0.002 (0.033)		
Is treatment	-3.657 (2.500)		
Constant	11.987*** (1.341)	8.667*** (2.150)	11.694*** (1.364)
Observations	120	60	60
R-squared	0.963	0.954	0.969

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 12. Results of Model 1 in E3

	Pool	Treatment	Control
Avg. poll info.	0.887*** (0.025)	0.847*** (0.025)	0.897*** (0.032)
Late Rounds dummy	3.268 (4.573)	7.117 (5.523)	-0.342 (8.055)
Late Rounds dummy *Avg. poll info.	0.018 (0.054)	-0.015 (0.062)	0.044 (0.100)
Is treatment*Avg. poll info.	-0.038 (0.033)		
Is treatment	-2.041 (2.442)		
Constant	5.825*** (1.605)	3.592* (1.811)	5.750*** (1.823)
Observations	135	75	60
R-squared	0.960	0.962	0.958

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Finally, with regards to variable P_t , the results of Table 12 indicate that even if the bias is known a priori, there is a very strong relationship between beliefs and average revealed poll

results. The respective coefficient b_1 is positive and statistically significant at the 1% level in all settings and for all experiments. As expected, the relationship appears stronger in the control condition (although the difference is not statistically significant).

The second model we examine (Model 2) takes the form:

$$V_t = a + b_1 P_t + b_2 D + b_3 (D * P_t) + e_t \quad (4)$$

V_t is the share of votes K received in the elections of round t . Election results are regressed on average poll information, the treatment dummy and the interaction term. This specification will allow us to examine whether average revealed poll results are good predictors of the electoral performance of K. In particular, if voters realise the bias and discount for it, we would expect that in the treatment condition polls have less influence on the final election. The results are presented in Tables 13-15. For all experiments, we find that the coefficient of the interaction term is small and not significant. Moreover, the coefficient in the treatment condition is often larger. There is clearly no evidence that revealed poll results are better predictors of election results in the control, rather than in the treatment condition.

The above finding is at face value paradoxical, since unbiased polls should be closer to election results than biased ones. However, this conclusion ignores the fact that revealed poll results might affect behaviour. The data indicate that the change in behaviour induced by polls can sometimes render the predictions of the biased sample of polls self-confirming – or, at least, nearly as good a predictor of elections as the unbiased sample. Figures 16-18 below illustrate the distribution of the differences between the average revealed poll results and the actual election results of the same round (both of these results are represented by the voting share for K).²⁷ These figures can help us assess this issue further.

²⁷ In particular, the variable $diff_PV$ in Figures 16-18 is defined as $diff_PV = P_t - V_t$.

Table 13. Results of Model 2 in E1

	Pool	Treatment	Control
Avg. poll info.	1.165*** (0.062)	1.182*** (0.066)	1.165*** (0.070)
Is treatment	-9.609 (7.287)		
Is treatment*Avg. poll info.	0.017 (0.099)		
Constant	-6.297 (3.866)	-15.907*** (5.268)	-6.297 (4.362)
Observations	120	60	60
R-squared	0.844	0.846	0.827

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 14. Results of Model 2 in E2

	Pool	Treatment	Control
Avg. poll info.	1.018*** (0.052)	1.134*** (0.064)	1.018*** (0.055)
Is treatment	-19.668*** (6.524)		
Is treatment*Avg. poll info.	0.116 (0.086)		
Constant	0.334 (3.446)	-19.334*** (5.181)	0.334 (3.656)
Observations	120	60	60
R-squared	0.854	0.845	0.856

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

A visual inspection of Figure 16 indicates that, in Experiment E1, the average vote share of K according to the revealed poll results deviates from the election vote share for K no more in the treatment sessions than in the control sessions. If anything, the variance in the deviations seems to be larger in the control. On the other hand, Figures 17 and 18 indicate that, in experiments E2 and E3, there is a pattern whereby in the treatment – but not in the control – average revealed polls systematically over-predict K's vote share.

This is important, as it might cast some light on the inability of participants (especially in E1 and E3) to account for the biased nature of poll results. Subjects who do not account for strategic incentives, but learn from experience alone (in the spirit of reinforcement learning) will have the opportunity to observe differences such as those presented in Figures 16-18 for a number of rounds. If these observed differences are not systematically greater in the treatment condition than in the control condition, we should not expect that such naïve learners would discount the ‘biased’ revealed poll results of the treatment any more than the ‘unbiased’ results of the control. This is consistent with what happens in E1.

However, as we noted regarding experiment E2, Figure 17 illustrates systematically larger disparities between revealed polls and election results in the treatment condition vs. the control condition. There is evidence that this resulted in subjects’ discounting somewhat the average revealed poll results in the treatment of E2: correlations between these poll results and elicited expectations are lower in the treatment condition (Figure 10). However, in E3, despite the fact that subjects are a priori informed of the bias in addition to the opportunity of observing a systematic overprediction of K’s vote share in the treatment (see Figure 18), there is not much evidence of discounting average poll results in the treatment condition.

Table 15. Results of Model 2 in E3

	Pool	Treatment	Control
Avg. poll info.	1.184*** (0.052)	1.139*** (0.050)	1.184*** (0.052)
Is treatment	-8.060 (5.319)		
Is treatment*Avg. poll info.	-0.045 (0.072)		
Constant	-11.461*** (3.483)	-19.521*** (4.000)	-11.461*** (3.505)
Observations	135	75	60
R-squared	0.888	0.876	0.899

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 16. Distribution of differences in the vote share of K between revealed poll results and elections, E1


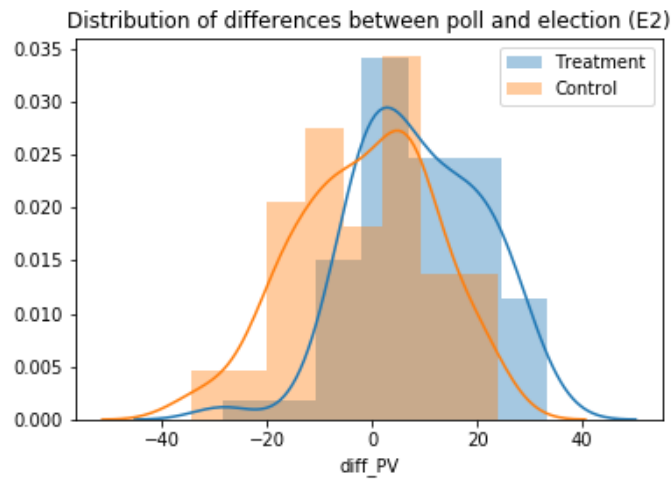


Figure 17. Distribution of differences in the vote share of K between revealed poll results and elections, E2



The third model we examine (Model 3) considers differences:

$$\Delta BP_t = a + b_1 \Delta PV_{t-1} + b_2 D + b_3 (D * \Delta PV_{t-1}) + e_t \quad (5)$$

ΔBP_t equals B_t minus P_t and ΔPV_t is P_t minus V_t . We use model 3 to explicitly examine whether there is evidence for learning. Again, we focus on reinforcement-type learners, who observe the model's variables through time. If they observe that ΔPV_{t-1} is large, this means that (in the last period) polls overestimated the performance of K relative to the election outcome. We expect that if subjects learn, this will result in subjects adjusting their

beliefs (for K's share) downwards conditional on the poll results, hence we expect a decrease in ΔBP_t . However, as Tables 16-18 indicate, the coefficients for ΔPV_{t-1} are small and not significant. Once more, we receive little evidence that subjects, in forming their beliefs, adjust for the existence of biased polls in our experimental environment.

Figure 18. Distribution of differences in the vote share of K between revealed poll results and elections, E3

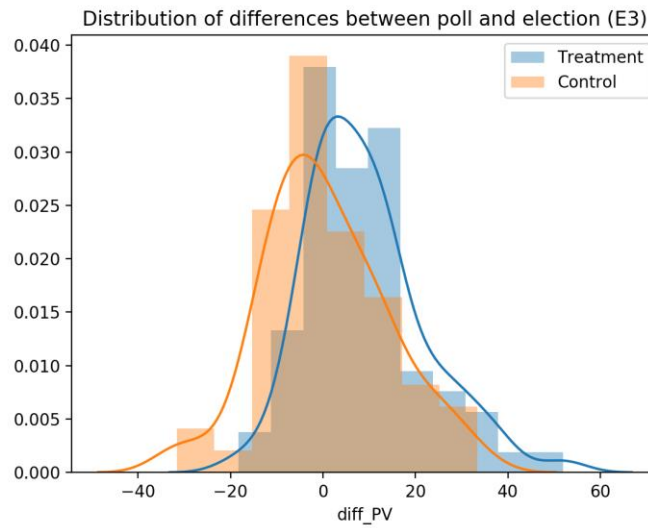


Table 16. Results of Model 3 in E1

	Pool	Treatment	Control
ΔPV_{t-1}	0.034 (0.058)	-0.086 (0.069)	0.034 (0.062)
Is treatment	-4.662*** (1.238)		
Is treatment * ΔPV_{t-1}	-0.120 (0.094)		
Constant	3.083*** (0.870)	-1.579* (0.819)	3.083*** (0.927)
Observations	112	56	56
R-squared	0.137	0.029	0.005

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 17. Results of Model 3 in E2

	Pool	Treatment	Control
ΔPV_{t-1}	0.051 (0.066)	-0.032 (0.067)	0.051 (0.072)
Is treatment	-6.826*** (1.421)		
Is treatment * ΔPV_{t-1}	-0.083 (0.099)		
Constant	1.304 (0.893)	-5.523*** (0.996)	1.304 (0.974)
Observations	112	56	56
R-squared	0.229	0.004	0.009
Standard errors in parentheses			
*** p<0.01, ** p<0.05, * p<0.1			

Table 18. Results of Model 3 in E3

	Pool	Treatment	Control
ΔPV_{t-1}	-0.031 (0.065)	-0.064 (0.060)	-0.031 (0.063)
Is treatment	-6.042*** (1.280)		
Is treatment * ΔPV_{t-1}	-0.033 (0.087)		
Constant	1.118 (0.866)	-4.924*** (0.963)	1.118 (0.842)
Observations	126	70	56
R-squared	0.217	0.016	0.004
Standard errors in parentheses			
*** p<0.01, ** p<0.05, * p<0.1			

7. Discussion and Conclusions

In this paper we examined the existence and implications of biased mechanisms that propagate the results of voting intention polls. We first established the existence of a systematically biased propagation pattern in the field, by analysing how news regarding the electoral ‘horse race’ are reproduced in social networks. The data indicate that ‘good news’ have a higher chance of being propagated in the examined network if they concern liberal

politicians than if they concerned conservatives. We then presented results from a series of experiments with majority voting where participants received information regarding poll results in a systematically selective manner. The environment we considered is a two-candidate election contest with common values (concerning candidates' valence) and no voting costs.

Our findings indicate that biased exposure to polls consistently skews the electoral outcome in a predictable way. In a very robust manner, elections that took place in the 'biased polls' environment provided an electoral advantage to the party that was favoured by the bias. This effect was smaller when in the control condition two polls were randomly revealed, as opposed to all five polls being revealed, but the direction of the effect was consistently the same. Similarly, effects were smaller when the voters were explicitly informed of the selection rule under which poll information was revealed, but the treatment effect was still politically significant and very consistent. Overall, the empirical results from E1 and E2 show limited evidence that the repeated opportunities for learning allowed voters to understand the systematic bias and account for it. The evidence from E3 indicates that it is especially the second part of this statement that matters (failing to account for the bias once one realises it).

Why is it so? We argue that a key reason is the genuinely complex environment where voting takes place. For instance, as Figure 16 indicates, pure feedback alone is unlikely to be sufficient for learning. Therefore, in Experiment 1, in terms of comparing election results to the average revealed poll predictions, the biased condition would not appear as particularly more 'suspicious' to an active learner than the unbiased one. It seems that the self-confirming nature of polls renders corrective feedback difficult. Accordingly, voters who are unable to infer from the strategic nature of the interaction, but only learn from experience, are unable to adjust their behaviour. However, in E3 subjects have a priori information about the bias, and Figure 18 indicates that, in the treatment condition, elections tend to diverge from average poll predictions in a systematic way. Despite all this, subjects fail to sufficiently discount the

revealed poll information in forming their beliefs (as evidenced by Figure 15). This indicates that even perfect a priori information in conjunction with subsequent feedback are not enough to make subjects discount the results of biased polls.

How applicable to real-world settings can results derived from our experimental environment be? We believe that our primary result, that people are unable to account for biased polls and hence such bias might robustly distort elections, is likely to generalise to the real world. Our subjects participate in fifteen elections. The number of rounds is reasonably high given the length of the typical experiment in the literature. Indeed, the number of elections that real-life voters may participate in their lifetime is not very large. If anything, the time delay in real election environments might make learning more challenging. Obviously, our stylised environment simplifies important aspects of real elections. However, it seems to us that if subjects are unable to adjust to systematic bias in a stylised environment such as this, they are unlikely to do so in more complicated real elections environments. In addition, the results of E3 indicate that even intergenerational transmission of information about the incentives, agendas and biases of information providers is unlikely to undo the electoral effects of selection and bias in revealing poll results.

In summary, it seems that both our lab and our field evidence raise considerable concerns regarding the risk that biased mechanisms of propagation of poll results may affect democratic outcomes. More evidence from the lab and the field is needed (including replications of this study) before safe policy conclusions can be made. The stakes for electoral policy are particularly high.

References

- Agranov, M., Goeree, J., Romero, J. and L. Yariv (2017): “What Makes Voters Turn Out: The Effects of Polls and Beliefs,” *Journal of the European Economic Association*, 16, 825–856.
- Allcott, H. and M. Gentzkow (2017): “Social Media and Fake News in the 2016 Election,” *Journal of Economic Perspectives*, 31, 211–36.
- Bock, O., Baetge, I., and A. Nicklisch (2014): “hroot: Hamburg Registration and Organization Online Tool”, *European Economic Review*, 71, 117-120.
- Boeri, T., Mishra P., Papageorgiou C., and A. Spilimbergo (2018): “A Dialogue Between a Populist and an Economist,” *AEA Papers and Proceedings*, 108, 191–95.
- Boxell, L., Gentzkow M., and J. M. Shapiro (2017): “Is The Internet Causing Political Polarization? Evidence From Demographics,” Tech. Rep. Working Paper No. W23258, National Bureau of Economic Research.
- Brown, K. M. and C. E. Zech (1973): “Welfare Effects of Announcing Election Forecasts,” *Public Choice*, 14, 117–123.
- Chen, D. L., Schonger, M. and C. Wickens (2016): “oTree—An Open-Source Platform for Laboratory, Online, and Field Experiments,” *Journal of Behavioral and Experimental Finance*, 9, 88-97.
- Cipullo, D. and A. Reslow (2019): “Biased Forecasts to Affect Voting Decisions?: The Brexit Case,” Working paper.
- Conover, M.D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F. and A. Flammini (2011): “Political Polarization on Twitter,” In Fifth international AAAI conference on weblogs and social media.

- DellaVigna, S. and E. Kaplan (2007): “The Fox News Effect: Media Bias and Voting,” *The Quarterly Journal of Economics*, 122, 1187–1234.
- Feddersen, T. and W. Pesendorfer (1997): “Voting Behavior and Information Aggregation in Elections with Private Information,” *Econometrica: Journal of the Econometric Society*, 65, 1029–1058.
- Fleitas, D. W. (1971): “Bandwagon and Underdog Effects in Minimal-Information Elections,” *American Political Science Review*, 65, 434–438.
- Forsythe, R., T. Rietz, R. Myerson, and R. Weber (1996): “An Experimental Study of Voting Rules and Polls in Three-Candidate Elections,” *International Journal of Game Theory*, 25, 355–383.
- Funke, M., M. Schularick, and C. Trebesch (2016): “Going to Extremes: Politics After Financial Crises, 1870–2014,” *European Economic Review*, 88, 227–260.
- Gentzkow, M. and J. Shapiro (2010): “What Drives Media Slant? Evidence from US Daily Newspapers,” *Econometrica* 78.1 35-71.
- Gerber, A., Hoffman, M., Morgan, J., and C. Raymond (2017): “One in a Million: Field Experiments on Perceived Closeness of the Election and Voter Turnout,” No. w23071. National Bureau of Economic Research.
- Goeree, J. K., and J. Grosser (2007): “Welfare Reducing Polls,” *Economic Theory* 31.1 51-68.
- Greiner, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1(1), 114-125.
- Guriev, S. (2018): “Economic Drivers of Populism,” *AEA Papers and Proceedings*, vol. 108, 200–203.

- Inglehart, R. and P. Norris (2017): “Trump and the Populist Authoritarian Parties: the Silent Revolution in Reverse,” *Perspectives on Politics*, 15, 443–454.
- Kaltwasser, C. (2018): “Studying the (Economic) Consequences of Populism,” *AEA Papers and Proceedings*, 108, 204–07.
- Klor, E. F. and E. Winter (2007): “The Welfare Effects of Public Opinion Polls,” *International Journal of Game Theory*, 35, 379.
- Larsen, E. G. and Z. Fazekas (2019): “Transforming Stability into Change: How the Media Select and Report Opinion Polls,” *The International Journal of Press/Politics*, in press.
- Maniadis, Z. (2014): “Selective Revelation of Public Information and Self-Confirming Equilibrium,” *International Journal of Game Theory* 43, 991-1008.
- Matthews, J., Pickup, M. and F. Cutler (2012): “The Mediated Horserace: Campaign Polls and Poll Reporting,” *Canadian Journal of Political Science*, 45(2), 261-287.
- McKelvey, R. and P. Ordeshook (1984): “Rational Expectations in Elections: Some Experimental Results Based on a Multidimensional Model,” *Public Choice* 44.1 61-102.
- McKelvey, R. D. and P. C. Ordeshook (1985): “Elections With Limited Information: A Fulfilled Expectations Model Using Contemporaneous Poll and Endorsement Data as Information Sources,” *Journal of Economic Theory*, 36, 55–85.
- Meffert, M. F. and T. Gschwend (2011): “Polls, Coalition Signals and Strategic Voting: An Experimental Investigation of Perceptions and Effects,” *European Journal of Political Research*, 50, 636–667.
- Mellon, J. and C. Prosser (2017): “Twitter and Facebook Are not Representative of the General Population: Political Attitudes and Demographics of British Social Media Users,” *Research & Politics* 4.3.

- Morwitz, V. and C. Pluzinski (1996): “Do Polls Reflect Opinions or Do Opinions Reflect Polls? The Impact of Political Polling on Voters' Expectations, Preferences, and Behavior,” *Journal of Consumer Research*, 23.1, 53–67.
- Oesch, D. (2008): “Explaining Workers’ Support for Right-Wing Populist Parties in Western Europe: Evidence from Austria, Belgium, France, Norway, and Switzerland,” *International Political Science Review*, 29, 349–373.
- Palfrey, T. (2016): “Experiments in Political Economy,” *The handbook of experimental economics* 2, 347-434.
- Petrova, M. (2008): “Inequality and Media Capture,” *Journal of Public Economics*, 92, 183–212.
- Plott, C. (1982): “A Comparative Analysis of Direct Democracy, Two Candidate Elections, and Three Candidate Elections in an Experimental Environment,” *Social Science Working Paper*, California Institute of Technology
- Rodrik, D. (2018): “Is Populism Necessarily Bad Economics?” *AEA Papers and Proceedings*, 108, 196–99.
- Rothschild, D. and N. Malhotra (2014): “Are Public Opinion Polls Self-Fulfilling Prophecies?” *Research & Politics*, 1, 1–10.
- Shirani-Mehr, H., Rothschild, D., Goel, S. and A. Gelman (2018): “Disentangling Bias and Variance in Election Polls,” *Journal of the American Statistical Association*, 87, 1–23.
- Sinclair, B. and C. R. Plott (2012): “From Uninformed to Informed Choices: Voters, Pre-Election Polls and Updating,” *Electoral Studies*, 31, 83–95.

- Soroka, S. (2014): “Negativity in Democratic Politics: Causes and Consequences” (Cambridge Studies in Public Opinion and Political Psychology). Cambridge: Cambridge University Press.
- Sturgis, P., Baker, N., Callegaro, M., Fisher, S., Green, J., Jennings, W., Kuha, J., Lauderdale, B. and P. Smith (2016): “Report of the Inquiry into the 2015 British General Election Opinion Polls.”
- Taber, C. S. and M. Lodge (2006): “Motivated Scepticism in the Evaluation of Political Beliefs,” *American Journal of Political Science*, 50: 755-769.
- Tversky, A. and D. Kahneman (1974): “Judgment Under Uncertainty: Heuristics and Biases,” *Science* 185.4157 1124-1131.
- Van Hauwaert, S. M. and S. Van Kessel (2018): “Beyond Protest and Discontent: A Cross-National Analysis of the Effect of Populist Attitudes and Issue Positions on Populist Party Support,” *European Journal of Political Research*, 57, 68–92.
- Whiteley, P. (2016): “Why Did the Polls Get It Wrong in the 2015 General Election? Evaluating the Inquiry into Pre-Election Polls,” *The Political Quarterly*.
- Wlezien, C., Jennings, W., Fisher, S., Ford, R. and Pickup, M., 2013. “Polls and the Vote in Britain,” *Political Studies*, 61(1_suppl), 66-91., 87, 437–442.
- .

APPENDIX 1. Additional Graphs

Figure A. Poll outcomes in treatment sessions of E1

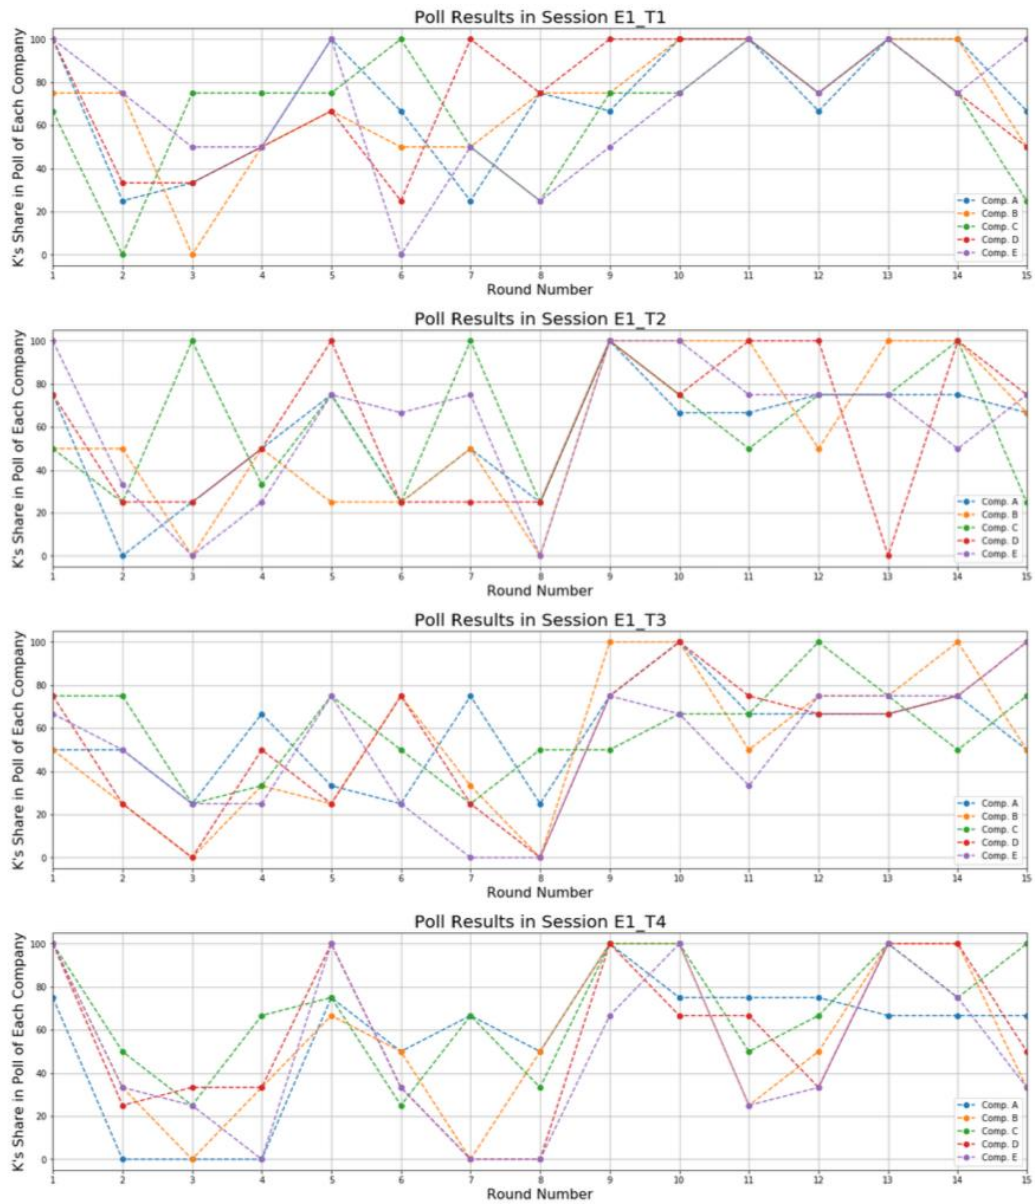


Figure B. Average Beliefs vs. Poll outcomes in control sessions of E1

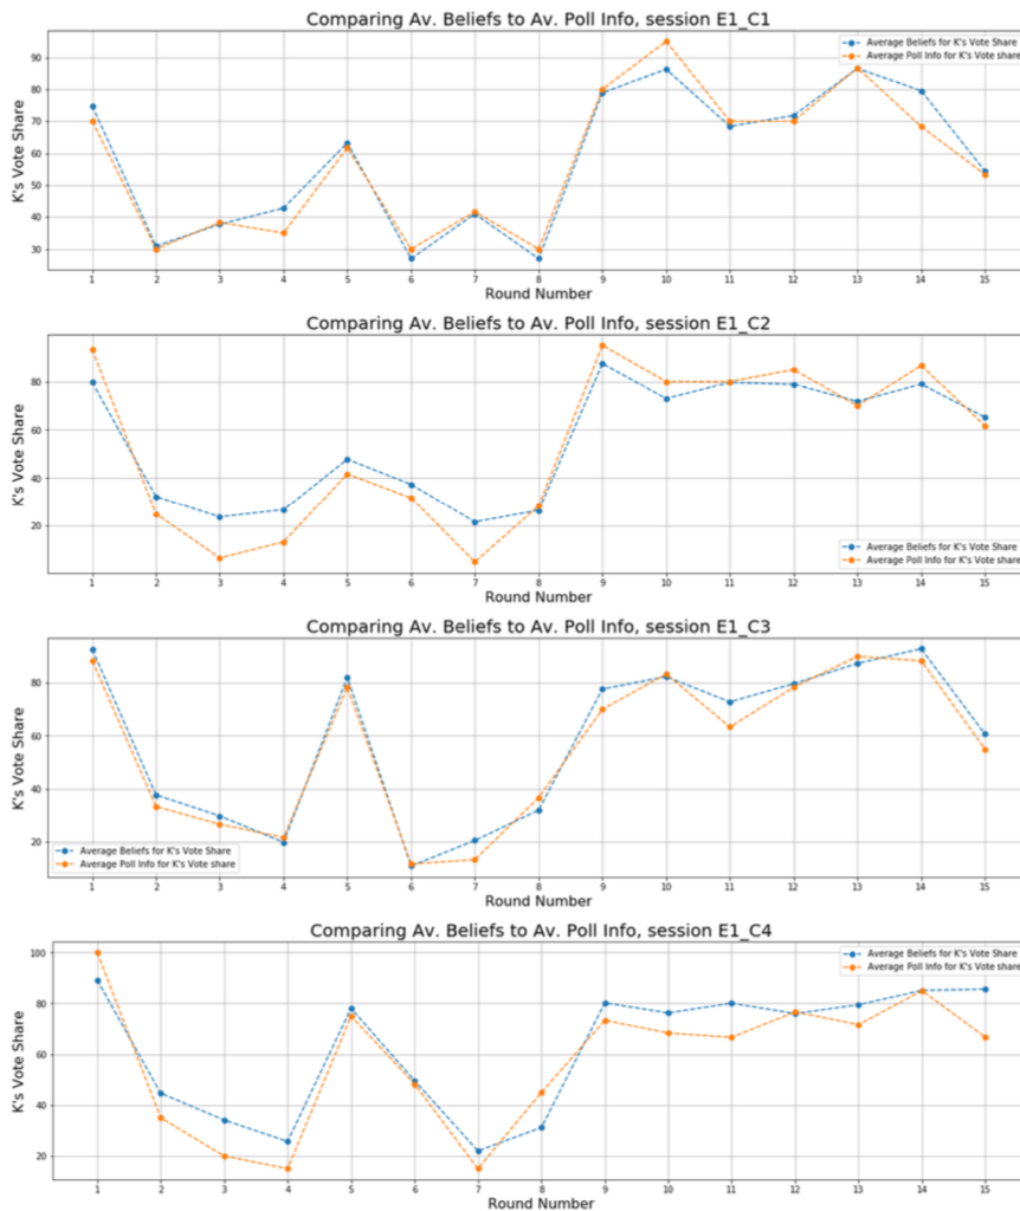


Figure C. Average Beliefs vs. Poll outcomes in treatment sessions of E1

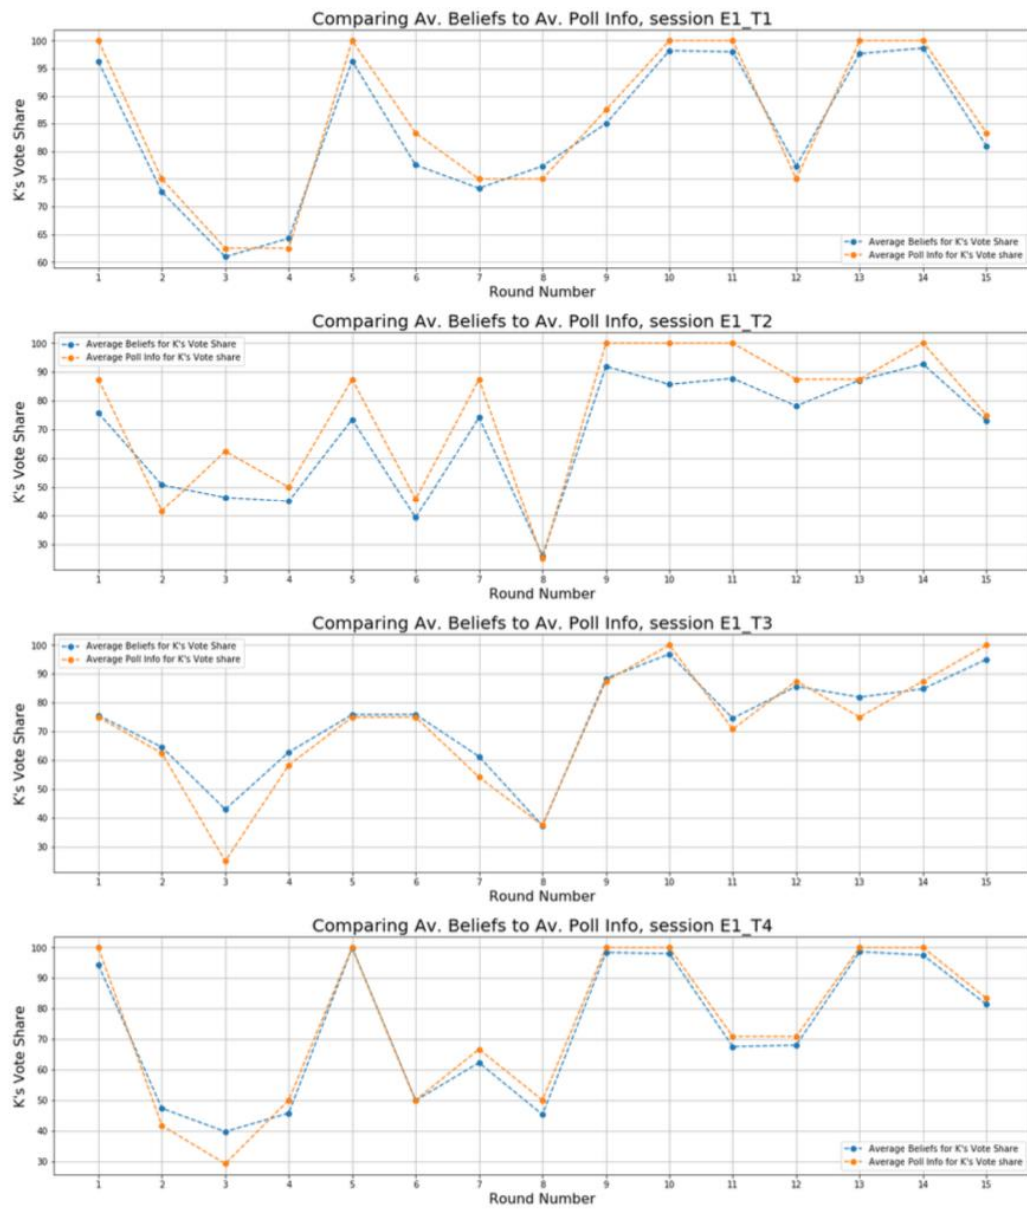


Figure D. Poll outcomes in treatment sessions of E2

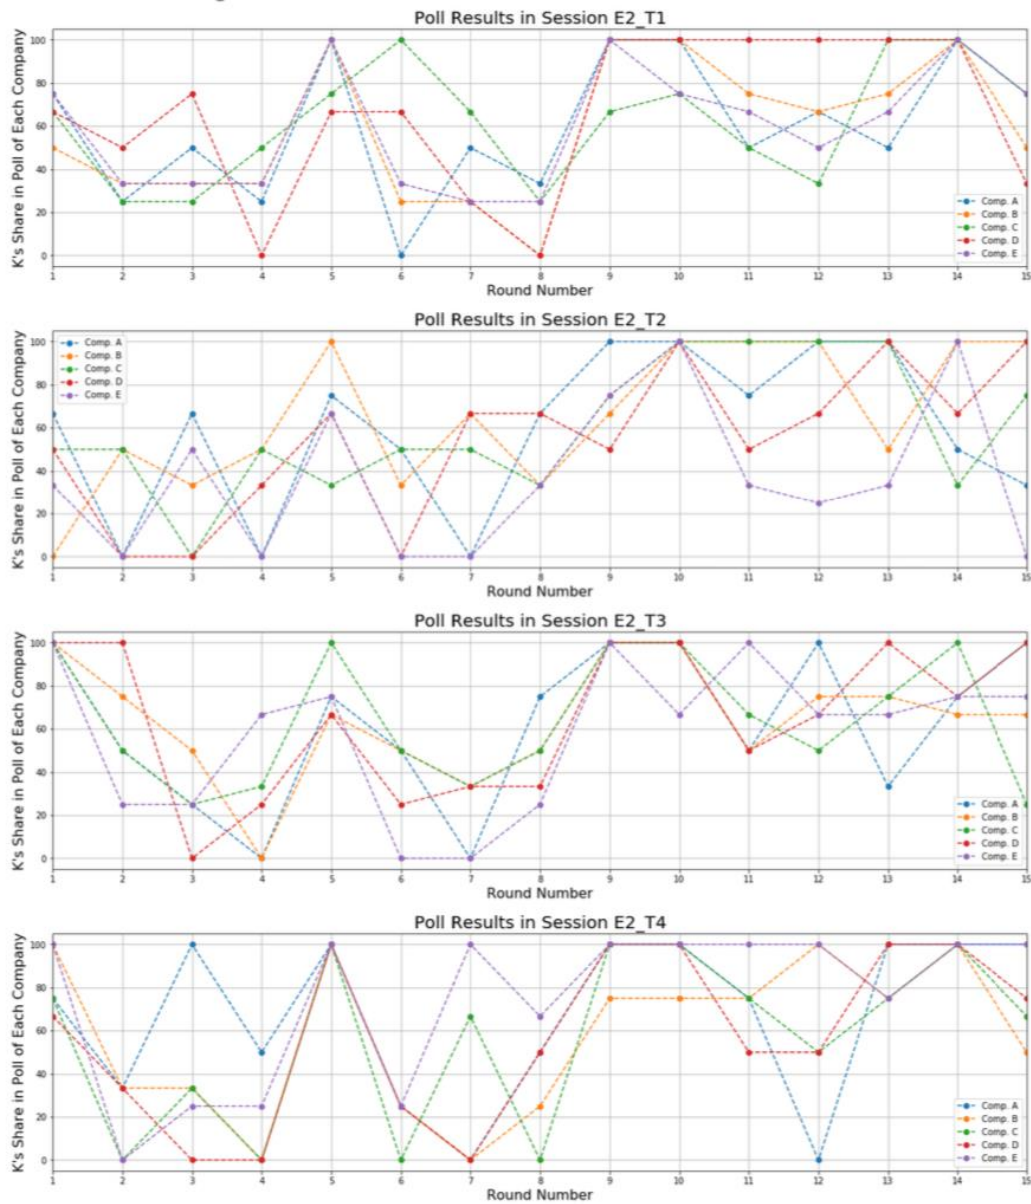


Figure E. Average Beliefs vs. Poll outcomes in treatment sessions of E2

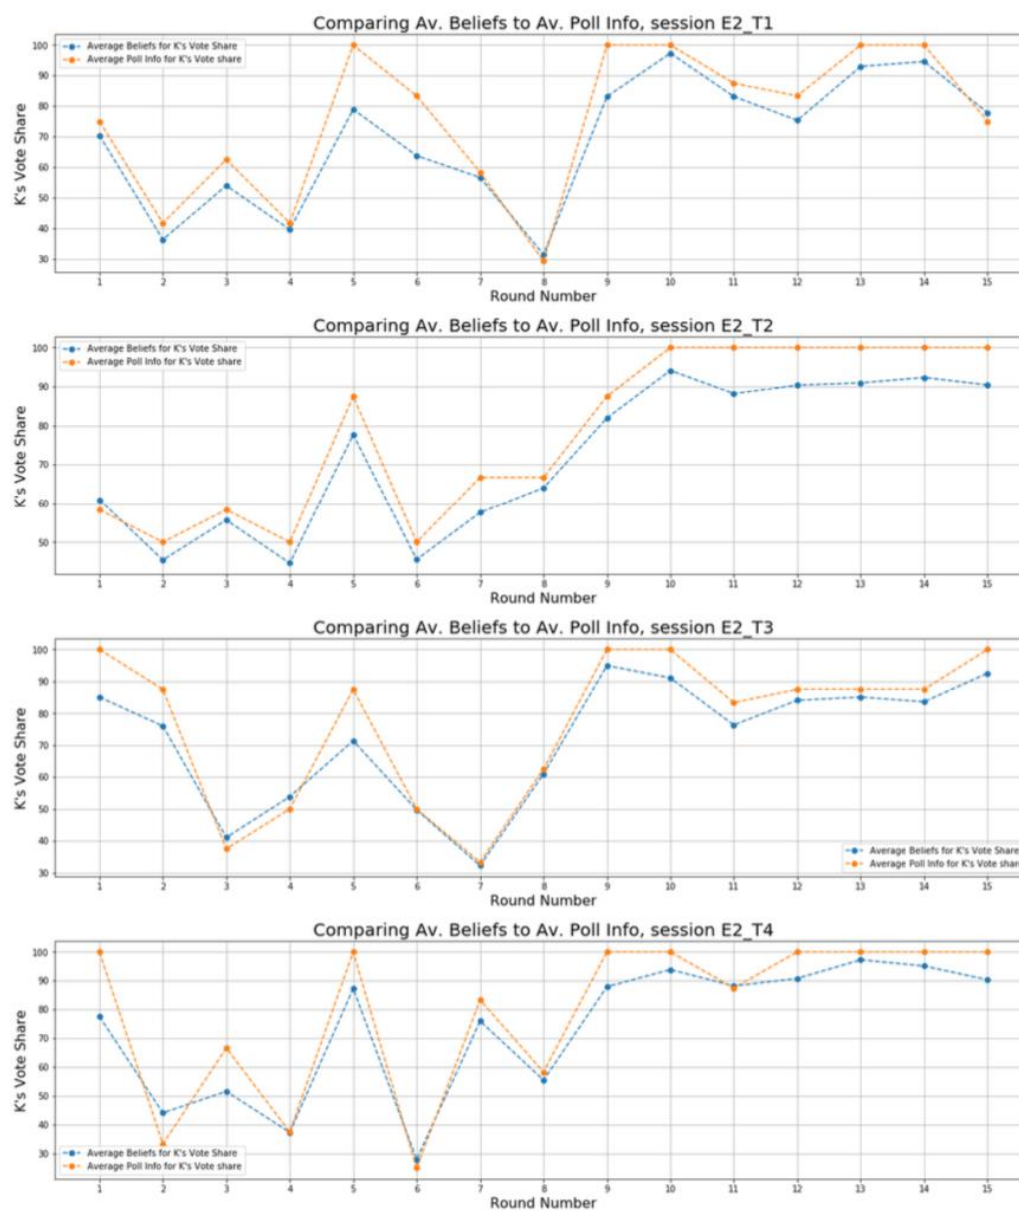


Figure F. Average Beliefs vs. Poll outcomes in control sessions of E2

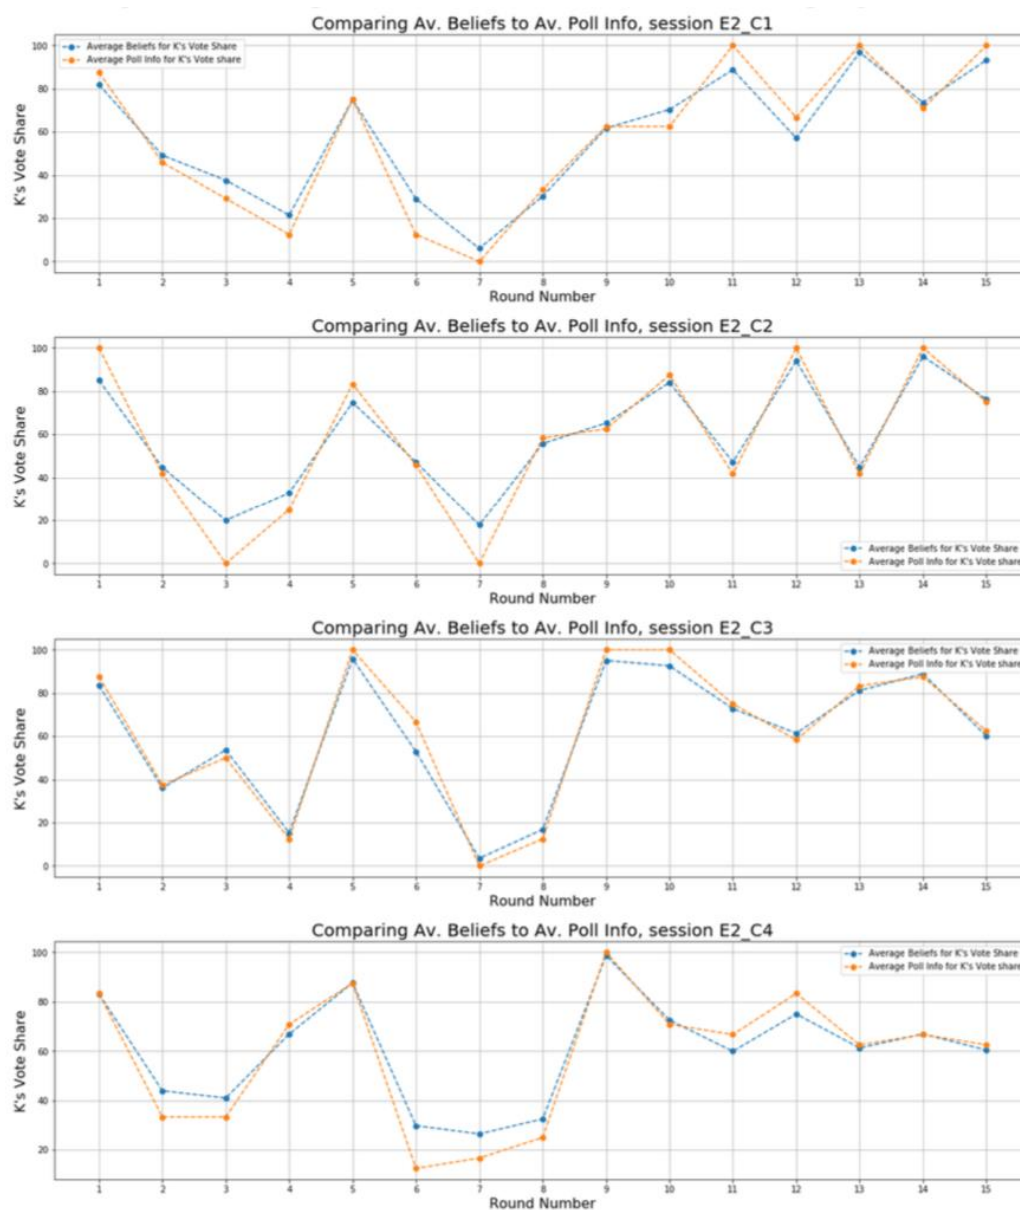


Figure G. Poll outcomes in treatment sessions of E3

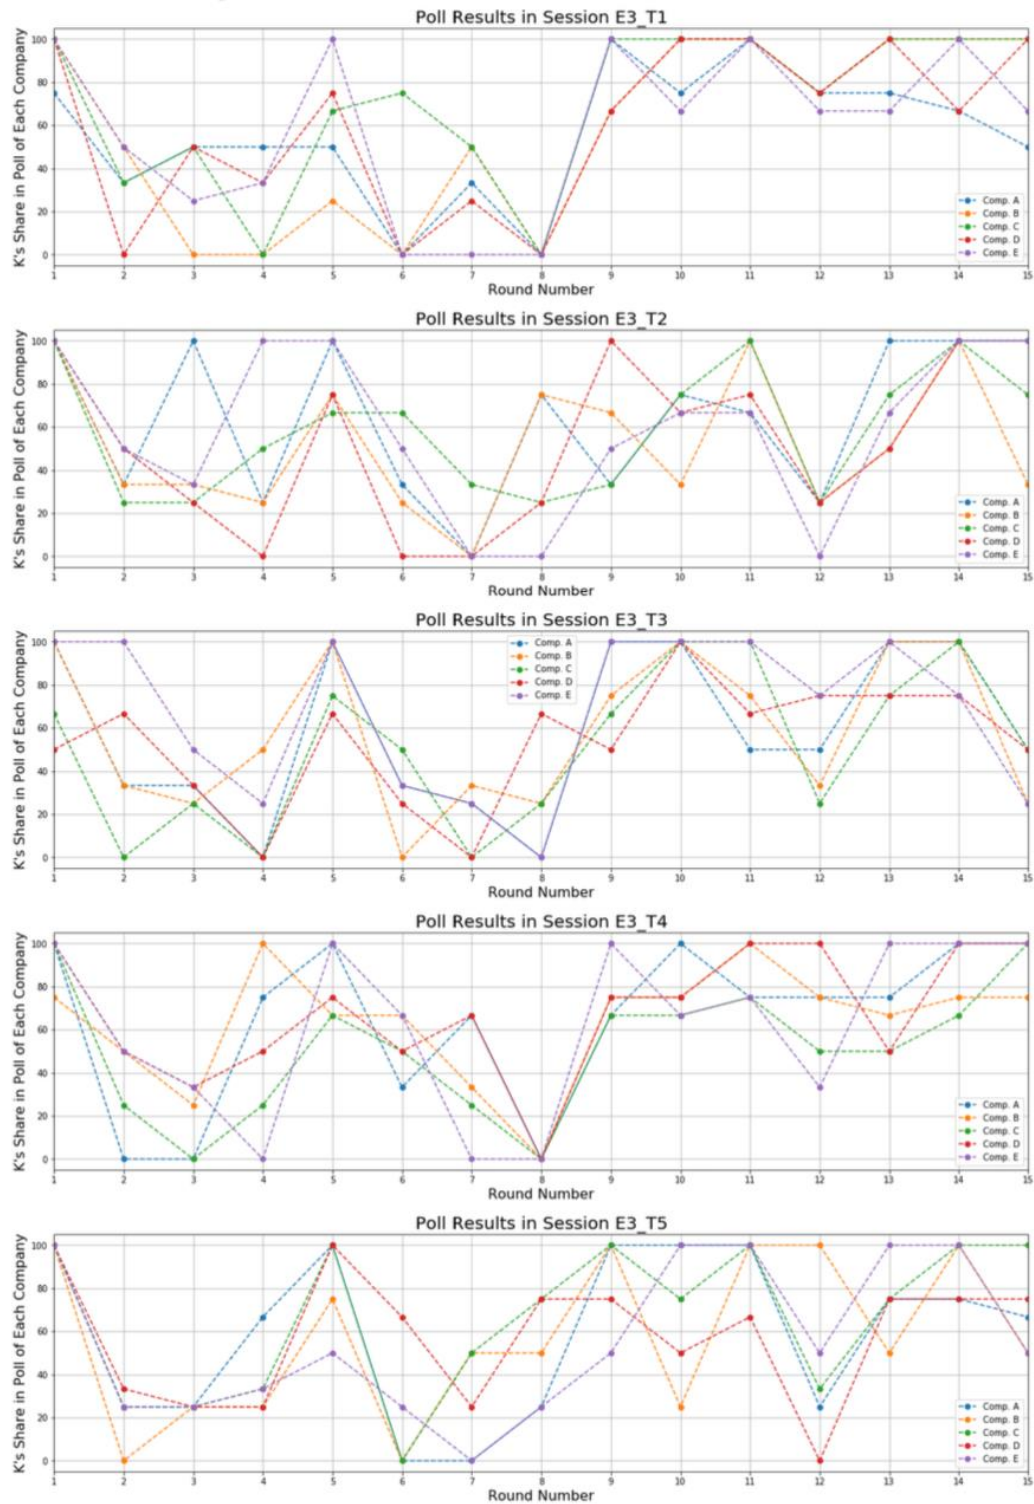


Figure H. Average Beliefs vs. Poll outcomes in treatment sessions of E3

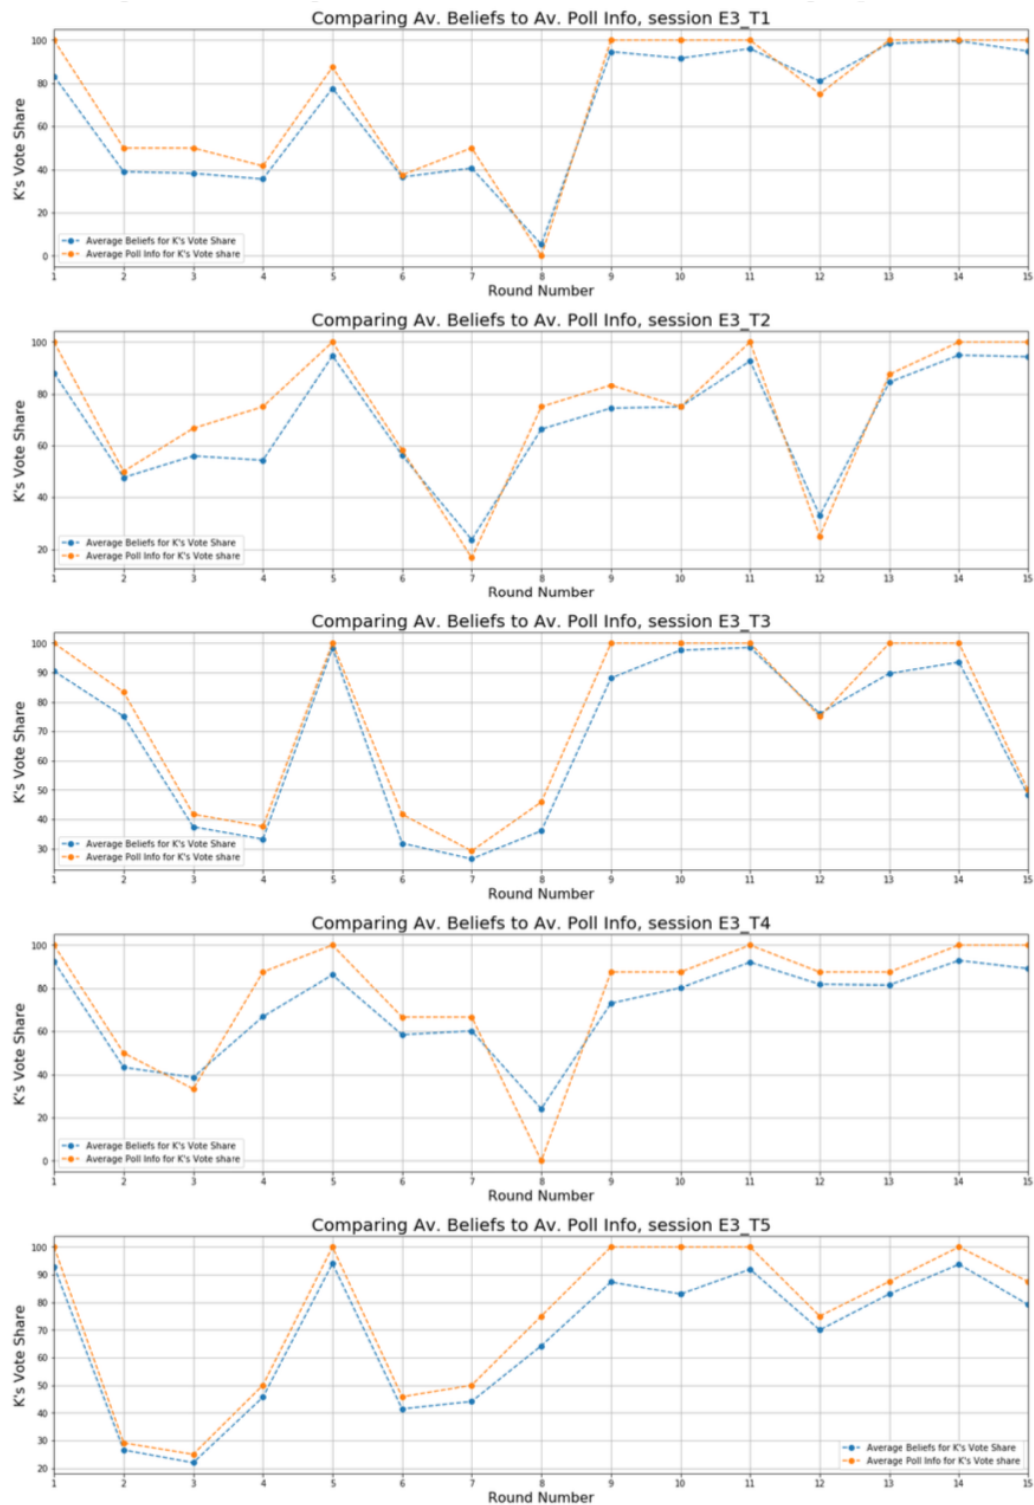


Figure I. Average Beliefs vs. Poll outcomes in control sessions of E3

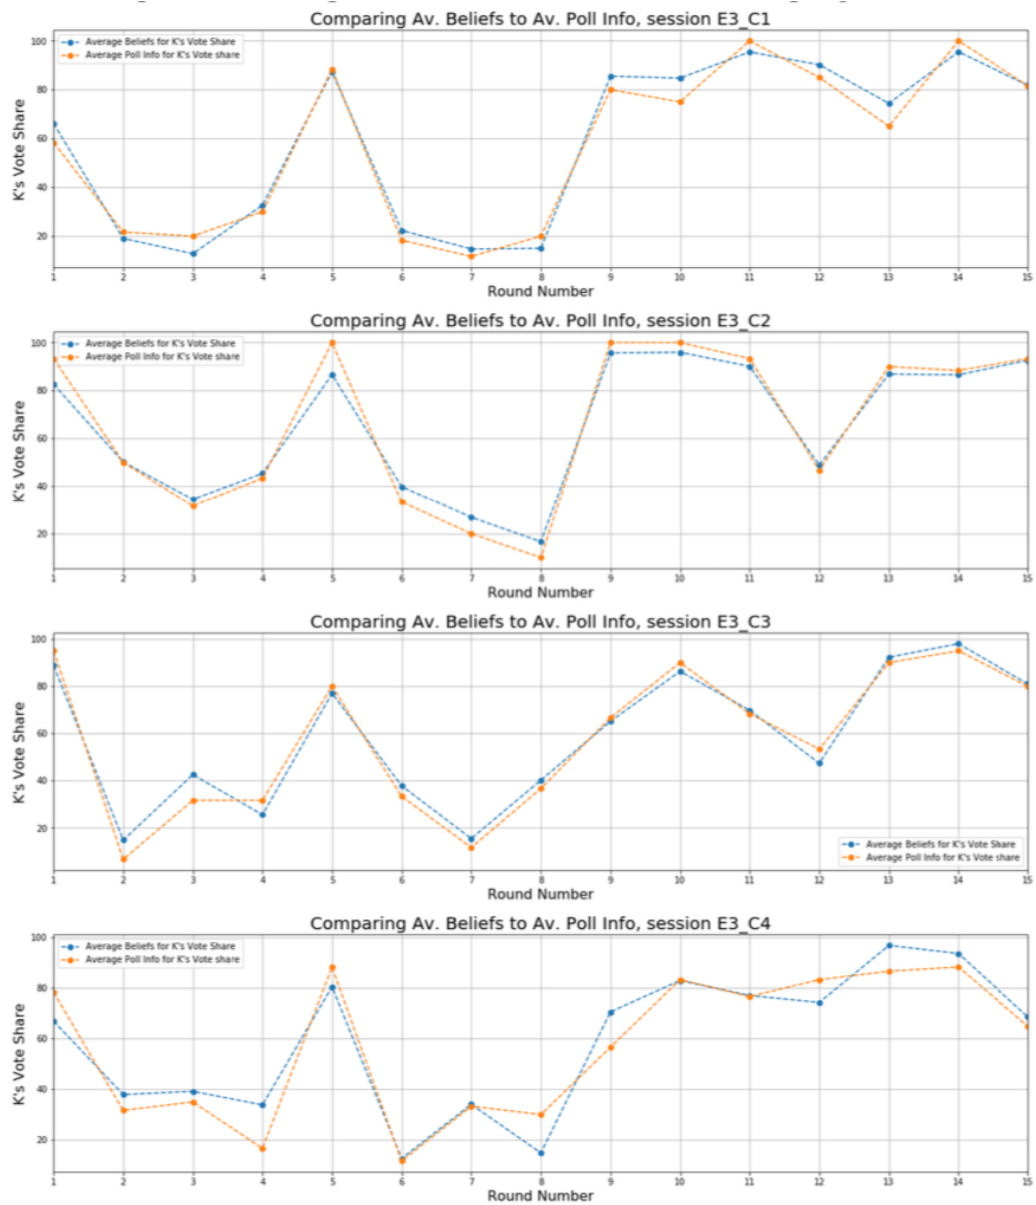


Figure J. Descriptive summary of voting behaviour at the poll stage, pooled at session level

	session	E1_C1	E1_C2	E1_C3	E1_C4	E1_T1	E1_T2	E1_T3	E1_T4
uninformed	J	26.09%	29.21%	29.67%	38.30%	17.20%	30.85%	46.67%	21.74%
	k	42.39%	47.19%	37.36%	51.06%	67.74%	53.19%	43.33%	47.83%
	non-participation	31.52%	23.60%	32.97%	10.64%	15.05%	15.96%	10.00%	30.43%
informed	J	44.05%	42.50%	44.44%	43.82%	34.94%	45.24%	38.37%	51.19%
	k	55.95%	50.00%	54.32%	55.06%	62.65%	54.76%	59.30%	47.62%
	non-participation	0.00%	7.50%	1.23%	1.12%	2.41%	0.00%	2.33%	1.19%

	session	E2_C1	E2_C2	E2_C3	E2_C4	E2_T1	E2_T2	E2_T3	E2_T4
uninformed	J	30.21%	29.21%	33.71%	22.92%	28.71%	26.32%	27.37%	29.21%
	k	43.75%	56.18%	39.33%	51.04%	48.51%	37.89%	47.37%	41.57%
	non-participation	26.04%	14.61%	26.97%	26.04%	22.77%	35.79%	25.26%	29.21%
informed	J	38.75%	45.45%	44.19%	46.91%	45.56%	46.84%	37.50%	34.94%
	k	57.50%	53.41%	55.81%	50.62%	48.89%	50.63%	62.50%	61.45%
	non-participation	3.75%	1.14%	0.00%	2.47%	5.56%	2.53%	0.00%	3.61%

	session	E3_C1	E3_C2	E3_C3	E3_C4	E3_T1	E3_T2	E3_T3	E3_T4	E3_T5
uninformed	J	34.88%	21.98%	29.03%	23.76%	30.85%	20.83%	27.84%	23.96%	32.97%
	k	50.00%	43.96%	41.94%	39.60%	38.30%	43.75%	47.42%	54.17%	56.04%
	non-participation	15.12%	34.07%	29.03%	36.63%	30.85%	35.42%	24.74%	21.88%	10.99%
informed	J	43.82%	37.18%	43.02%	48.10%	37.04%	52.33%	44.83%	42.35%	44.57%
	k	49.44%	60.26%	56.98%	50.63%	62.96%	47.67%	54.02%	56.47%	54.35%
	non-participation	6.74%	2.56%	0.00%	1.27%	0.00%	0.00%	0.00%	1.18%	1.09%