Elections under Biased Candidate Endorsements — An Experimental Study*

Junze Sun[†] Arthur Schram[‡] Randolph Sloof[§]

October 1, 2020

Abstract

We construct an election game to study the electoral impacts of biased candidate endorsements. We derive a set of testable predictions. We test these in a laboratory experiment and find that observed election outcomes and vote shares are well predicted. We find no support, however, for our prediction that the relationship between election outcome and the endorser's bias is non-monotonic; i.e., ex ante, a candidate's winning probability will first increase and then decrease as the endorser becomes more biased towards her. Voter turnout is much less responsive to the bias than predicted. We argue that observed voting behavior can be explained, to a substantial extent, by three behavioral mechanisms: (a) distinct levels of rationality for candidate choice and turnout decisions, (b) conservative belief updating, and (c) 'partial competition neglect', where voters underestimate the correlation between the information released by an endorsement and the closeness of elections.

JEL Codes: C92, D72, D83

Key Words: biased endorsements, voting, turnout, quantal response equilibrium, experiments

*This paper supersedes the previous version entitled "Media Bias and Elections — An Experimental Study". We are grateful to Sourav Bhattacharya, Aaron Kamm, Navin Kartik, David Levine, Zara Sharif, Joel Sobel, and Leeat Yariv for insightful comments at various stages of this project. We also thank seminar participants at NYU-Abu Dhabi, the Hong Kong University of Science and Technology (HKUST), GATE-Lyon-St Etienne, the NYU-WESSI workshop in Florence, EWEBE in Bologna, BEES/BEPS in Maastricht, and the 2019 EEA-ESEM meeting, for their comments and suggestions. Financial support by the Research Priority Area Behavioral Economics of the University of Amsterdam is gratefully acknowledged. All remaining errors are ours.

†Department of Economics at the European University Institute (EUI), Florence. Email: Junze.Sun@eui.eu.

‡CREED, Amsterdam School of Economics and Tinbergen Institute, and Department of Economics at EUI, Florence. Email: A.J.H.C.Schram@uva.nl.

§Amsterdam School of Economics and Tinbergen Institute. Email: R.Sloof@uva.nl.

1 Introduction

All around the world, voters in many elections lack the information they need to properly assess the candidates on the ballot. Knowing more about the alternatives would arguably make easier the decisions they face. Of course, voters could search for such information, for example, by reading party platforms, investigating candidates' backgrounds, or visiting campaign gatherings. Such private information acquisition is costly, however, and in large elections it is unlikely that the expected benefits from the improvement in the choice made would outweigh these costs (Downs, 1957). For this reason, voters often rely on low-cost external information sources to evaluate candidates. One important source of such information is the public endorsement of a candidate by experts or other public figures. Whenever an endorsement is broadly shared, it may play an important role in shaping voters' views on candidates and on how a voter expects other voters to view the candidates. Public endorsements then potentially have a salient impact on voters' behavior and election outcomes.

Public figures who endorse candidates may well be biased, however, and such biases are often widely recognized. An important question is then, to what extent can biased endorsements affect voting behavior and election outcomes? We explore this question both theoretically and experimentally. We construct a simple election game with an informed endorser and derive both Bayesian Nash equilibria (BNE) and quantal response equilibrium (QRE) for this game. These produce theoretical predictions regarding the electoral influences of biased endorsements that we subsequently test in a laboratory experiment. In doing so, we focus on two research questions. First, how does an increase in the endorser's bias affect the election outcome and voter turnout? Second, are voters' perceptions of candidates systematically distorted by biased endorsements?

Prime examples of endorsements on public platforms are observed when large-scale media such as newspapers and TV-channels announce their support for a candidate.[1] Media, however, are often biased (Prat and Strömberg, 2013; Prior, 2013; Gentzkow, Shapiro and Stone, 2015; Puglisi and Snyder, 2015*b*). [2] If such biased media endorse a candidate, then this is precisely the kind of information to voters that we are interested in studying. Our model, however, concerns expert endorsers more generally as informed experts who possess superior information but may produce biased endorsements.[3]

[1] Despite the increasing relevance of social media in modern elections, the large-scale mass media still play the most important role (Prat, 2018; Kennedy and Prat, 2019).

[2] Media biases may occur because of various supply-side processes (Baron, 2006; Martin and McCrain, 2019; Besley and Prat, 2006; McMillan and Zoido, 2004) or be demand driven (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006). Moreover, though it is widely believed that media bias is severe in weak democracies (Djankov, McLiesh, Nenova and Shleifer, 2003), empirical evidence shows that media bias is pervasive also in well-developed countries such as the U.S. (Groseclose and Milyo, 2005; Gentzkow and Shapiro, 2010; Puglisi and Snyder, 2011, 2015*a*).

[3] While the media industry has many features, the characteristics we focus on – broad coverage and a primary

Our election game is constructed based on a more general theoretical framework developed by Sun, Schram and Sloof (2019). In the game, two candidates compete in an election under simple majority rule. These candidates differ along a quality dimension, which can be interpreted as their integrity or competence, representing attributes that all voters wish a candidate to have (Groseclose, 2001; Ashworth and De Mesquita, 2009). Aside from caring about quality, voters also have their private ideological preferences over candidates. In this way, it is possible for a voter to prefer an ideologically less-favored candidate if her quality is sufficiently superior to that of her opponent. Voters can vote for either candidate or abstain, but casting a vote is costly and these costs are private information. We assume that voters cannot directly observe candidates' qualities and have to rely on public sources such as media outlets or experts for that information.

Specifically, media outlets or other endorsers precisely observe candidates' qualities and can communicate to voters by sending public endorsements for candidates. An unbiased endorser truthfully reveals the quality-superior candidate. A biased endorser, however, may endorse a candidate even if her quality turns out to be inferior, provided that the quality difference does not exceed a certain threshold. How much quality difference the endorser will tolerate depends on the extent of its bias, which is commonly known to voters.

For example, assume that the New York Times was biased towards Hillary Clinton in the 2016 presidential elections, while Fox News was biased towards Donald Trump. We assume that voters know these biases. Media endorsements can still be informative in this environment, as long as there exist quality levels for which the New York Times would not endorse Clinton or Fox News would not endorse Trump. An endorsement allows voters to update their posterior beliefs about candidates' qualities. By affecting voters' posterior beliefs, an endorsement can affect the election outcome and voter turnout even if voters know the endorser's biases.

The key mechanism that drives the endorsers' influence in our election game is that endorsements allow voters not only to infer which candidate is more appealing in quality, but also to gauge the closeness of the election. The former information determines the election outcome, whereas the latter drives voter turnout. For instance, imagine that Fox News had endorsed Clinton in the 2016 elections. Given the publicly known Republican slant of Fox News, this endorsement would have implied that Clinton must have a much higher quality than Trump. As a result, most of the voters would have been convinced to support Clinton. Clinton would thus have obtained a much larger vote share than Trump and won the election by a significant margin. This in turn implies that the election would not end up in a close race. Because rational voters anticipate this, voter turnout would be low

source of (possibly biased) information – seem salient and pervasive. Moreover, media are a primary channel through which biased experts speak to the mass public. Finally, it is common in the media literature to model media as biased information providers (Chan and Suen, 2008; Chiang and Knight, 2011; Durante and Knight, 2012; Puglisi and Snyder, 2015*a*).

due to the "competition effect" (Levine and Palfrey, 2007).[4] Building on this mechanism, we derive precise comparative static predictions regarding the electoral influence of biased endorsements.

To test these theoretical predictions, we implemented our election game in a laboratory experiment. Aside from enabling a direct test of our hypotheses, laboratory experiments have important advantages for addressing our research questions. First, it allows us to overcome an important limitation of observational field data. Studies using the latter can only observe voting behavior conditional on the realized information released. Voting behavior conditional on other information sets is thus not observed. In contrast, exploiting the strategy method allows us to provide direct causal evidence for both *interim* (i.e., conditional on the information released by endorsements) and *ex-ante* (i.e., unconditional) impacts of the endorser's bias on voting behavior and election outcomes.[5] Second, it allows us to precisely define and induce crucial elements of the model such as candidates' quality levels, voters' ideological preferences, and the extent of an endorser's bias. Third, laboratory control enables a direct elicitation of beliefs, which allows us to rigorously identify the systematic impact of biased endorsements on voters' beliefs using Bayes' rule. Finally, one of the main goals in this paper is to test a theory. For this purpose, the environment where data are collected needs to align with the core assumptions of the theory (Schram, 2005). This is best achieved in the laboratory.

Our experimental results show that candidates' vote shares and winning prospects are much better predicted by theory than voter turnout; BNE predictions alone can explain 74.2% and 64.9% of variations in observed vote shares and election outcomes, respectively, but only 6.9% of variations in observed voter turnout. We identify two substantial deviations from theoretical predictions. First, while our theory predicts a non-monotonic relationship between the winning probability of a candidate and the endorser's bias towards her, our data show that this probability increases monotonically with the bias. Second, contrary to our theoretical predictions, observed voter turnout is systematically higher than predicted and is almost invariant to the endorser's bias. Our analysis of subjects' reported beliefs shows no evidence that their perceptions about candidates' qualities are systematically distorted by biased endorsements. Finally, we discuss the welfare implications of our results. Consistent with our theoretical predictions, aggregate welfare (as measured by subjects' average earnings) decreases in the endorser's bias and is lowest when voters are uninformed. Nevertheless, we find that voter welfare is substantially lower than predicted in treatments with low bias levels. This can be attributed to the high turnout rates observed in these treatments.

[4] The competition effect predicts voter turnout to be higher in closer elections, where voters have stronger incentives to cast costly votes.

[5] By 'interim', we mean the situation where candidates' qualities have materialized and voters have received the candidate endorsement(s). By 'ex-ante', we refer to a situation prior to the materialization of candidates' qualities and endorsements.

We explore alternative behavioral explanations for the discrepancies between observed and predicted outcomes. We find that voting behavior can be explained, to a substantial extent, by a hybrid equilibrium model that combines three behavioral mechanisms. First, voters exhibit higher levels of rationality in candidate choices than in turnout decisions. This is intuitive, because turnout decisions are cognitively more demanding than candidate choices; the former require a precise estimate of the expected benefits, while the latter only require coarse ordinal information on whether a candidate is more or less appealing than her opponent. As a result, voters are less likely to make mistakes in their candidate choices than in their turnout decisions. Second, voters update their beliefs about candidates' qualities too conservatively compared to a Bayesian. This causes voters to underestimate the actual quality difference between candidates and overestimate the closeness of the election, leading to higher turnout rates due to the competition effect. Third, voters suffer from 'partial competition neglect', that is, they underestimate the degree to which the closeness of elections is correlated with the information they received from (biased) endorsements. Such competition neglect makes voters partly ignorant of how the closeness of elections varies across information environments, leading to insufficient responsiveness of turnout to the endorser's bias.

These three mechanisms partly accommodate the irresponsiveness of turnout. They cannot, however, explain the failure of non-monotonicity. By further exploring the moderating role partisanship plays in how endorsements affect voter behavior, we provide a tentative explanation for the observed monotonic impact of endorser bias. In particular, we find that an extremely biased candidate endorsement has little impact on voting behavior of the candidate's own supporters; it neither mobilizes her supporters to turnout more nor makes them vote more often for the endorsed candidate if they do vote. In contrast, it has substantive effects on the behavior of supporters of the opposing candidate; an extremely biased endorsement not only makes them vote less, it also (albeit slightly) dissuades them from voting for their preferred candidate if they do turnout. These two effects together seem to drive the fact that a candidate's winning prospects are not harmed even if the endorsements she receives become less credible.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the election game we implement in the laboratory experiment. Section 4 conducts equilibrium analyses and formulates testable hypotheses. Section 5 explains our experimental design and Section 6 presents the results. In Section 7 we explore alternative behavioral explanations for observed deviations in subjects' behavior in the experiment. Section 8 concludes.

2 Related literature

Our paper speaks to four strands of literature. First of all, it relates to a strand of theoretical literature studying the influence of biased experts on decision making (e.g., Calvert (1985) and

Suen (2004)). In line with this literature, we model endorsers as informed experts who may provide biased advice. This literature, however, has thus far focused on environments with a single decision maker. In contrast, we study the influence of biased advice in the context of voting in a rational voter framework, where there are multiple, strategically interacting decision makers. The central insight that distinguishes our paper from preceding studies is that candidate endorsements inform voters not only with respect to which candidate to choose, but also about their turnout decisions. The strategic interaction between voters yields fundamentally different implications for the impact of biased advice than in scenarios with only a single decision maker.[6]

Second, our paper relates to a large and growing literature that empirically documents media's influence on voter behavior using observational field data (DellaVigna and Kaplan, 2007; Gerber, Karlan and Bergan, 2009; Enikolopov, Petrova and Zhuravskaya, 2011; Gentzkow, Shapiro and Sinkinson, 2011; Drago, Nannicini and Sobbrio, 2014; Adena, Enikolopov, Petrova, Santarosa and Zhuravskaya, 2015; Cagé, 2017). These studies have documented ample evidence of how media affect party vote shares and voter turnout. We complement this literature by addressing similar questions under laboratory control. This allows us to identify genuinely causal relationships between, for example, the extent of an endorser's bias and voters' reactions to the endorsement. Moreover, as discussed in the introduction, using the strategy method enables us to observe voting behavior and election outcomes under all possible information sets (those that materialized and those that did not). We can therefore directly identify not only the interim but also the ex-ante impacts of media bias.

The third strand of literature empirically identifies the persuasion effect of biased information sources.[7] The findings are mixed. Some studies find that voters are sophisticated and can filter out the influence of an endorser's bias. For instance, Chiang and Knight (2011) study the influence of media bias on voting intentions by exploiting the timing of newspaper endorsements prior to the 2004 U.S. presidential election. They find that only highly credible newspaper endorsements affect voters' beliefs. However, Cain, Loewenstein and Moore (2005) show in a controlled laboratory experiment that decision makers do not sufficiently account for the interest misalignment between themselves and informed advisors in their belief formation. We add to this literature by directly eliciting voters' beliefs in a controlled laboratory experiment. This allows us to identify the

[6] A related but different strand of literature studies the impact of dispersed information on voting behavior, taking into account the strategic interaction among voters (Feddersen and Pesendorfer, 1996, 1997, 1999). These papers demonstrate that this strategic interaction may affect how voters use information to guide their voting decisions. In particular, they identify the swing voter's curse, under which sincere voting (which is the optimal strategy in the single decision making scenario) generically fails to hold in equilibrium. Recently, Oliveros and Várdy (2015) explore how the option of abstention affects voters' media consumption decisions, taking into account the strategic interactions among the electorate. All these papers study environments where voting incurs no costs and distinct voters may be exposed to different information. In our paper, instead, we study a model in which voting is costly and voters have access to a common information source.

[7] See DellaVigna and Gentzkow (2010) for a comprehensive review.

systematic impact of biased endorsements on beliefs using a standard implication of Bayes' rule.

A final strand of literature studies voter turnout under a wide range of theoretical and empirical approaches. Most relevant for our analysis of voter turnout is the competition effect, which predicts higher voter turnout in closer elections. This is grounded in classical studies such as Downs (1957) and Riker and Ordeshook (1968) that, like ours, use a pivotal voter framework. One of the prediction that follows is that the likelihood of casting a vote is negatively correlated with the expected number of votes by others. On the other hand, information cascades (Cialdini, 2007; Coleman, 1994) or a desire for conformity (Blais and Hortala-Vallve, 2020) might make a voter more likely to turnout if she expects others to vote. In the end, it is then an empirical question whether the competition effect holds. The empirical evidence, however, is mixed. Much of the experimental literature supports a positive relationship between the closeness of elections and turnout (Levine and Palfrey, 2007; Duffy and Tavits, 2008; Großer and Schram, 2010; Agranov, Goeree, Romero and Yariv, 2018). Though various field studies also find support for the competition effect (Blais, 2000; Bursztyn, Cantoni, Funk and Yuchtman, 2017), some have recently found no relationship between closeness and turnout. An explanation often put forward for this lack of correlation is that voters have a poor perception of the closeness of an election (Enos and Fowler, 2014; Moskowitz and Schneer, 2019; Gerber, Hoffman, Morgan and Raymond, 2020). Even with a reasonable estimate of the closeness, however, the correlation between this estimate and the turnout decision is positive, but noisy (Duffy and Tavits, 2008; Gerber et al., 2020). The latter suggests limits to the rationality involved in the turnout decision.

The turnout literature points to a role for behavioral theories in the pivotal voter framework. For this purpose, we will adopt models that capture bounded rationality, such as quantal response equilibria (McKelvey and Palfrey, 1995, 1998; Goeree and Holt, 2005), and identify a novel behavioral trait – "partial competition neglect". The latter involves voters partially overlooking the correlation between information they receive (via endorsements) and the closeness of elections. While the quantal response approach allows us to capture limited rationality, partial competition neglect helps to explain poor estimates of the closeness of an election. We will argue that both mechanisms are consistent with and help to explain our empirical findings.

3 The experimental election game

The election game we implement in the laboratory is based on a general framework developed by Sun, Schram and Sloof (2019). There are 25 voters and two candidates, who compete under simple majority rule. Candidates differ in their qualities, which are commonly valued by all voters. The qualities of candidates are observable to (possibly biased) expert endorsers only. Therefore, voters must rely on information from the endorsers to infer these qualities. Voters have private and idiosyncratic ideological preferences over candidates, and voting costs that are independently drawn from commonly known distributions.

Candidates. In our experiment, the candidates are framed as two vases labeled A and B. We induce their qualities as follows. At the beginning of each election, one of the two vases is randomly chosen with equal probability and filled with x "diamonds", where x is randomly drawn from a discrete-uniform distribution on multiples of ten between 10 and 100. Let k_A and k_B be the number of diamonds allocated to vases A and B, respectively. If vase A (B) wins the election, then each voter earns k_A (k_B) "tokens" (each token is worth 2 eurocents). In this way, $k \equiv k_A - k_B$, the difference between the numbers of diamonds in vases A and B, can be interpreted as the *relative quality of A* and it follows a discrete-uniform distribution $F(\cdot)$ on $\{-100, -90, \cdots, -10, 10, \cdots, 90, 100\}$ (i.e., all multiple of tens between -100 and 100, except 0).

Voters. The electorate consists of $\mathscr{N} = 25$ voters. To induce voters' idiosyncratic ideological preferences, we use the following protocol. In every election, each voter is independently and randomly assigned to either team A or team B, with equal probability. If the label of the elected vase matches a voter's team (e.g, vase A is elected and the voter belongs to team A), then the voter receives an additional private bonus of either 20, 50 or 100 tokens. This bonus is determined randomly with equal probability and independently across voters. If instead the label of the elected vase does not match the voter's team, the voter receives no private bonus. More precisely, let v_i^A (v_i^B) denote the private bonus voter i may obtain if candidate A (B) is elected. v_i^A (v_i^B) equals zero if voter i does not belong to team A (B), and is equally likely to be 20, 50, or 100 tokens if voter i belongs to team A (B). In this way, $v_i \equiv v_i^A - v_i^B$ measures voter i's *ideological preference* and it follows a discrete-uniform distribution $G(\cdot)$ on $\{-100, -50, -20, 20, 50, 100\}$. Voter i is ideologically leaning towards candidate A (B) if $v_i > (<) 0$. The strength of her ideological preference is either *centralist* ($|v_i| = 20$), *moderate* ($|v_i| = 50$) or *extreme* ($|v_i| = 100$). In addition, for each voter i, casting a vote induces a private cost of c_i tokens, which is independently drawn from a discrete-uniform distribution $\gamma(\cdot)$ on $\{1, 2, \cdots, 15\}$ (i.e., integers between 1 and 15).

Given diamond allocation k_A and k_B, each voter i's payoff (in tokens) equals $u_A(k_A, v_i^A) = k_A + v_i^A$ if candidate A is elected, and $u_B(k_B, v_i^B) = k_B + v_i^B$ if B is elected. Therefore, from voter i's perspective, the difference in payoff from candidate A winning or losing the election equals

$$u_A(k_A, v_i^A) - u_A(k_B, v_i^B) = k_A - k_B + v_i^A - v_i^B = k + v_i \tag{1}$$

Voter i thus strictly prefers candidate A (B) to be elected if $k + v_i > (<) 0$, and the absolute payoff difference $|k + v_i|$ measures her *stake* at the election.

Expert endorsers and their biases. Throughout the election, voters never directly observe the realization of k (i.e., the relative quality of A). However, voters may receive imperfect public information about k. Let M denote the set of Expert endorsers. Following the literature, we assume that endorsers communicate to voters using a binary cutoff endorsement strategy (Calvert, 1985; Suen, 2004; Chiang and Knight, 2011). Specifically, each endorser $m \in M$ is characterized by its *bias* χ_m and sends message $s_m \in \{A, B\}$ using the following cutoff reporting strategy:

$$s(k, \chi_m) = \begin{cases} A, & \text{if } k > -\chi_m \\ B, & \text{if } k \leq -\chi_m \end{cases} \tag{2}$$

Namely, endorser m sends message A (B) if and only if k lies above (below) a certain threshold $-\chi_m$. A higher χ_m implies a lower threshold for sending message A. For this reason, message A (B) can be interpreted as an *endorsement* for candidate A (B), and χ_m measures the endorser's bias. If $\chi_m = 0$, the endorser truthfully supports the candidate with a superior quality and is thus *unbiased*. If instead $\chi_m > (<) 0$, then the endorser is said to be *A-biased* (*B-biased*) because it may support candidate A (B) even if her quality is lower than her opponent. In the experiment we set $|M| = 1$ and vary the bias χ of this single endorser to study the influence of endorser bias on voting behavior and election outcomes.[8] We distinguish between three different levels of endorser bias: *unbiased* (UB; $\chi = 0$), *weakly A-biased* (WB_A; $\chi = 55$), and *strongly A-biased* (SB_A; $\chi = 95$).

Timing. The timing of the election game is as follows. First, nature draws the quality realizations (i.e., k) and the profile of voter types (i.e., $\{(v_i, c_i)\}_{i=1}^{25}$) from their corresponding distributions. Second, observing k, the endorser sends message $s \in \{A, B\}$ using cutoff strategy (2). Third, observing s, voters simultaneously make their voting decisions. Finally, the winning candidate is elected by simple majority rule, with ties broken by a fair coin toss. All payoffs then realize.

4 Equilibrium Analyses and Hypotheses

In this section, we conduct equilibrium analyses for our election game and derive testable hypotheses regarding the electoral impacts of biased endorsers. Because this is a static game of incomplete information, we first derive the Bayesian Nash Equilibrium (BNE). Following the convention of the literature, we focus on type-symmetric BNE, where voters with the same private type (v_i, c_i) adopt the same voting strategy in equilibrium (Palfrey and Rosenthal, 1985; Levine and Palfrey, 2007).[9] As a robustness check, we also derive the logit Quantal Response Equilibria (QRE;

[8] In Online Appendix E we discuss a separate set of sessions where we set $|M| = 2$. This allows us to study the effects of the entry of new endorsers (e.g., media entry).

[9] Of course, multiple equilibria may exist. Although we have found various equilibria – which differ in the specific mixing probabilities used by different (v_i, c_i) pairs that are all indifferent between voting and abstaining – all of these are characterized by the same overall probabilities that a randomly drawn voter will vote for A or B, or will abstain. The predictions derived below hold for all these equilibria. We have not found asymmetric equilibria that possibly could lead to different overall voting probabilities. But, even if these would exist, it would arguably be rather implausible to

McKelvey and Palfrey (1995)) for our game and confirm that all our hypotheses hold under QRE as well. Formal derivations of both BNE and QRE are relegated to Online Appendix A.

4.1 Bayesian Nash Equilibrium

We assume that voters' utility function is linear in their payoffs (see (1)). In this way, any information (denoted by I) affects voters' behavior only through its influence on $E[k|I]$, the posterior expectation of candidate A's relative quality conditional on I. As a consequence, voters behave as if k is commonly known and equal to $E[k|I]$. We thus conduct the equilibrium analyses by first deriving the BNE of the election game with k commonly known to all voters, and then replacing k with $E[k|I]$. In Online Appendix A, we derive BNE numerically for all k considered in this paper.

We show that both candidate A's expected vote share and her winning probability increase monotonically in k, and equal 0.5 if $k = 0$. The intuition is straightforward; with a higher k, candidate A is more appealing in quality and able to convince more voters to support her. When $k = 0$ both candidates have equal qualities and voters are a priori indifferent between them. The expected voter turnout decreases in $|k|$, the quality difference between candidates. This is due to the "competition effect" (Levine and Palfrey, 2007): voter turnout is higher in closer elections. If $k = 0$, both candidates are expected to get equal shares of votes. In this case, the election is most likely to be a close race, where voters have the strongest incentives to vote. If instead the quality difference $|k|$ is large, then the election is expected to end in a landslide victory of the quality-superior candidate. In this situation, voters have little incentives to vote.[10]

By manipulating $E[k|I]$, endorsers like media outlets can influence voters' perceptions about the relative appeal of candidates and the closeness of elections, and therefore affect both election outcomes and voter turnout. We study the electoral impact of an increased endorser bias by considering three levels of biases $\chi \in \{0, 55, 95\}$. Let $k_s(\chi)$ for $s \in \{A, B\}$ denote the posterior expectation of k obtained by Bayes' rule, conditional on observing message s from an endorser with bias χ. Recall from equation (2) that message s is sent by the cutoff strategy: $s = A(B)$ if $k > (\leq) -\chi$. It follows that $k_A(\chi) = E[k|k > -\chi]$ and $k_B(\chi) = E[k|k \leq -\chi]$. Denote by $Pr(s|\chi)$ the ex-ante probability that the endorser with bias χ sends message s. It holds that $Pr(A|\chi) = 1 - F(-\chi)$ and $Pr(B|\chi) = F(-\chi)$. We present the endorsers' reporting strategies graphically in Figure 1 and summarize $Pr(s|\chi)$ and $k_s(\chi)$ for $\chi \in \{0, 55, 95\}$ in Table 1.

Table 1 reveals three effects of an increase in endorser bias that drive our comparative statics. *Effect I* is that the ex-ante likelihood of endorsing candidate A increases in the endorser's bias

expect that voters who are completely identical within the setup of our model (i.e. have the same private type (v_i, c_i)) use different voting strategies and are able to solve the required (equilibrium) coordination needed for it.

[10] This intuition is based on the pivotal voter framework. As discussed in Section 2, some of the literature casts doubts on the relevance of this framework for the turnout decision. We shall return to this point in Section 7, when discussing alternative behavioral mechanisms.

Figure 1: Reporting Strategies of Endorsers with Different Biases

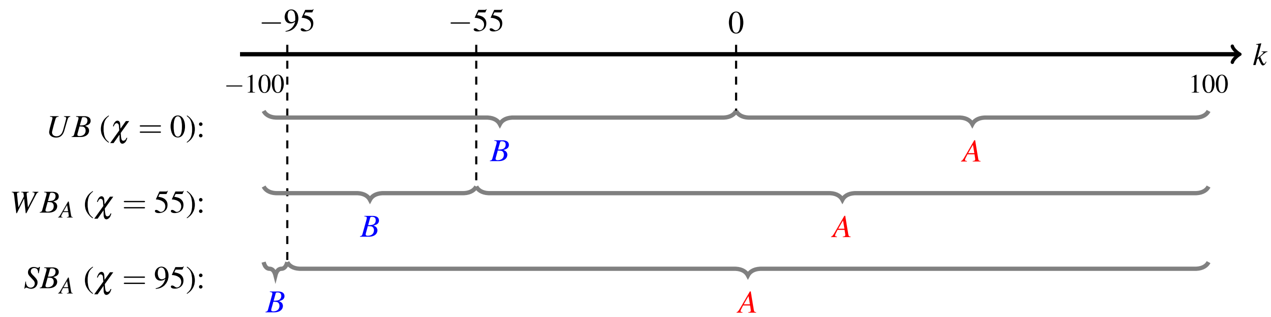

Table 1: Likelihood of Endorsements and the Rational Posterior Expectations

Endorser Bias	$\chi = 0$ (UB)		$\chi = 55$ (WB_A)		$\chi = 95$ (SB_A)	
Message	$s = A$	$s = B$	$s = A$	$s = B$	$s = A$	$s = B$
$Pr(s\|\chi)$	0.50	0.50	0.75	0.25	0.95	0.05
$k_s(\chi)$	55	-55	27	-80	5	-100

Note: $k_s(\chi)$ denotes the posterior expectation of k conditional on an endorser with bias χ sending message $s \in \{A, B\}$. It is rounded to integer values for convenience of exposure. For each χ, the law of iterated expectation holds: $\sum_{s\in\{A,B\}} Pr(s|\chi)k_s(\chi) = E[k] = 0$.

towards her. In Figure 1 this can be seen by comparing the range of k's at the various levels of biases for which A is endorsed. *Effect II* is that, as the endorser becomes more A-biased, message A becomes less credible in signaling the superiority of candidate A and induces a lower posterior expectation about the quality difference between candidates. In Figure 1 the posterior expectation about k after message A is the value of k that is at the midpoint of the range of k's for which A is sent (recall that k is uniformly distributed). With higher bias, this midpoint moves closer to $k = 0$. Finally, *Effect III* is that message B becomes more credible in signaling the superiority of candidate B and induces a higher posterior expectation about candidates' quality difference, as the endorser becomes more A-biased. In Figure 1 the endorser's threshold for sending message B becomes more stringent, and the midpoint of the range of k's for which B is sent moves further away from $k = 0$.

Figure 2 presents both the interim (i.e., conditional on message s) and ex-ante (i.e., unconditional) impacts of endorser bias. It does so for the expected vote share (panel a), winning probability of candidate A (panel b), and the expected voter turnout (panel c). Panels (a) and (b) show that, conditional on either message A or B, both candidate A's expected vote share and her winning probability decrease in endorser bias χ. This is because (i) $k_s(\chi)$ decreases in χ for both $s \in \{A, B\}$ (cf. Effects II and III), and (ii) both A's expected vote share and her winning probability are increasing functions of k. Intuitively, as the endorser becomes more A-biased, the expected relative quality of candidate A inferred from both messages A and B is lower. This dissuades some voters

Figure 2: Interim and Ex-ante Electoral Impacts of Endorser Bias

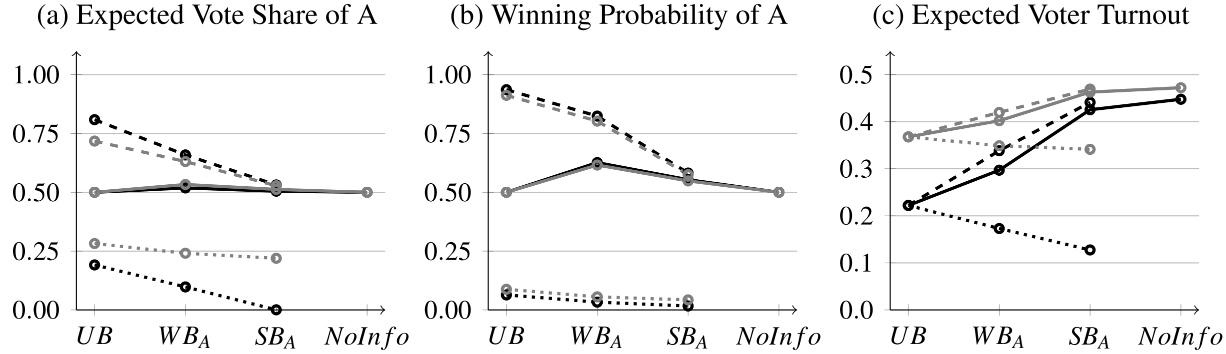

Note: In all three panels the black lines denote the BNE predictions for the interim impact of endorser bias conditional on $s = A$ (dashed), conditional on $s = B$ (dotted), and ex-ante impact of endorser bias (solid). The grey lines denote the corresponding predictions for the interim and ex-ante impact based on the logit version of QRE model, which is explained in the next subsection. On the horizontal axis of each panel, UB, WB_A and SB_A represent to three levels of endorser bias: unbiased ($\chi = 0$), weakly A-biased ($\chi = 55$), and strongly A-Biased ($\chi = 95$). $NoInfo$ refers to a scenario where voters are completely uninformed about k.

from supporting A and therefore deteriorates A's electoral prospects.

From an ex-ante perspective, candidate A's expected vote share hovers around 0.5 and is almost invariant with the endorser bias.[11] Her winning probability, however, varies non-monotonically in this bias; A's winning probability equals 0.5 under UB, 0.63 under WB_A, and 0.55 under SB_A. This non-monotonic relationship is driven by two opposite forces as the bias χ increases. On the one hand, by Effect I, the endorser supports candidate A more frequently ex-ante in an attempt to systematically manipulate the election outcome in its own favor. This force *per se* increases A's electoral prospects.[12] On the other hand, by Effects II and III, voters downwardly adjust their expectation about candidate A's relative quality conditional on either endorsement. This is because the endorsement of candidate A becomes less credible, while the endorsement for candidate B becomes more credible as bias χ increases. This force *per se* reduces A's winning chances and thus protects voters against manipulation. The net effect thus depends on which of these two opposing forces dominate. For our election game, the former force dominates as χ increases from 0 to 55, while the reverse is true as χ increases from 55 to 95.

When the endorser is unbiased (i.e., $\chi = 0$), its endorsement is highly credible and thus able to convince a majority of voters to support the endorsed candidate. Consequently, this candidate has a high chance of winning. From the ex-ante perspective, both candidates are endorsed with equal probability, so the likelihood of each candidate winning the election is 0.5 a priori. At the other

[11] We show in Online Appendix A that, for our election game, A's expected vote share predicted by BNE is an almost linear function of k and equals 0.5 if $k = 0$. Consequently, variations in A's expected vote share are almost proportional to variations in voters' posterior expectations about k. The latter are invariant from an ex-ante perspective thanks to the law of iterated expectation. This is why A's ex-ante expected vote share varies little with endorser bias under BNE.

[12] Because $k_A(\chi) > k_B(\chi)$ for all χ (cf. Table 1), Effect I *per se* increases candidate A's ex-ante winning probability.

extreme, when the endorser is strongly biased (i.e., $\chi = 95$), it almost always endorses candidate A. This endorsement, however loses credibility because it is very uninformative. As a result, voters will largely ignore this message and behave almost as if they are uninformed. For this reason, candidate A's winning chance conditional on message A sent by a strongly biased endorser is close to 0.5, the expected winning chance when voters are uninformed. These analyses yield Hypothesis 1.

Hypothesis 1. *Influence of endorser bias on the expected vote share and the election outcome:*

(a) Conditional on either message A or B, candidate A's expected vote share and winning probability are highest under UB and lowest under SB_A.

(b) Ex-ante, A's winning probability is ranked, in ascending order, by UB, SB_A and WB_A.

As for voter turnout, we can see from Figure 2.c that conditional on message A, the expected turnout increases in endorser bias while conditional on message B the reverse is true. To understand the intuition, recall that a lower expected quality difference between candidates is associated with a higher expected closeness of the election, causing higher voter turnout due to the competition effect. By Effect II (III) message A (B) induces a posterior expectation of smaller (larger) quality difference between candidates and thus implies a higher (lower) expected turnout, as χ increases.

From an ex-ante perspective, we show that voter turnout unambiguously increases in endorser bias. Moreover, voter turnout attains its maximum if voters are uninformed (i.e., in treatment $NoInfo$); in this case, voters make decisions only base on their common prior mean, which equals 0, and behave as if $k = 0$ is common knowledge. The election is then most likely to end in a close race; the probabilities of casting pivotal votes are thus highest if $k = 0$. These analyses yield Hypothesis 2.

Hypothesis 2. *Influence of endorser bias on voter turnout:*

(a) Conditional on message A, voter turnout is highest under SB_A and lowest under UB.

(b) Conditional on message B, voter turnout is highest under UB and lowest under SB_A.

(c) Ex-ante, voter turnout is ranked, in ascending order, by UB, WB_A, SB_A and $NoInfo$.

4.2 Quantal Response Equilibria

Various experimental studies have shown that a logit version of Quantal Response Equilibrium (QRE) better describes voting behavior than BNE does (Goeree and Holt, 2005; Levine and Palfrey, 2007; Großer and Schram, 2010). For this reason, we derive the logit QRE of our election game and check whether our hypotheses derived from BNE are replicated under QRE. In contrast to

BNE, QRE allows voters to make noisy best responses. The extent of decision noise is governed by the logit response parameter $\lambda \in (0, +\infty)$. As $\lambda \to 0$, noise dominates and voters' decisions are essentially purely random choices. If $\lambda \to +\infty$, then voters choose the action that maximize their expected payoff and QRE converges to a BNE. This parameterization is general enough to capture different sources of noise, such as distractions, perception biases, or miscalculations.

To determine the logit response parameter λ, we obtained the out-of-sample maximum-likelihood estimate of λ using data from a previous experiment with an electorate of size 11 conducted in the summer of 2016.[13] The estimated $\widehat{\lambda}$ equals 19.45. We then use this $\widehat{\lambda}$ to calculate QRE predictions for the candidate A's expected vote share, winning probability, and voter turnout. These predictions are presented by the grey lines in Figure 2. Regarding candidate A's expected vote share and winning probability, the BNE and QRE predictions are similar both qualitatively and quantitatively. The comparative statics predictions regarding voter turnout are also valid under QRE, though turnout levels predicted by QRE are systematically higher than the corresponding BNE predictions. In a nutshell, both Hypotheses 1 and 2 are robust under QRE.

5 Experimental design and procedure

5.1 Treatment design

In the experiment, endorsers are framed as "robots", which are computerized and communicate to voters by sending binary public messages $s \in \{A, B\}$ using cutoff strategy (2). In line with our model setup, we distinguish between an *unbiased* ($\chi = 0$; UB), a *weakly A-biased* ($\chi = 55$; WB_A) and a *strongly A-biased* robot ($\chi = 95$; SB_A). We label these robots as "ALPHA", "BETA" and "GAMMA", respectively. In addition, we introduce a control group ($NoInfo$) where no robot provides any information so that voters are uninformed about the diamond allocations. The instructions in Online Appendix D show how these treatments are presented to the subjects.

We implement a within-subject design for four distinct treatments: UB, WB_A, SB_A, and $NoInfo$. Within each session, the 25 subjects played 24 rounds of elections under each of the four treatment (a total of 96 elections per session). To facilitate direct comparisons we used for each treatment and session the same set of 24 realized parameter draws (including diamonds allocations, voters' team assignments, private bonuses and voting costs).

In each election under treatments other than $NoInfo$, we use the strategy method to elicit subjects' voting decisions contingent on each message sent by the robot. Specifically, we ask each voter "What would be your voting decision if robot X sends message s?" for both $s \in \{A, B\}$. Depending on the treatment, X can be ALPHA (for UB), BETA (for WB_A) or Gamma (for SB_A).

[13] Details about the design and analyses of the experiment conducted in 2016 are provided in Online Appendix A.2.

Each voter can choose between three options: voting for vase A, voting for vase B, and abstaining. Her contingent choice will be implemented depending on the realized message. At the end of each election, we provide subjects with information about the realized diamond allocation, as well as the total numbers of voters voting for A, voting for B, and abstaining in that election.

To minimize order effects, we randomly divided the 24 elections for each treatment into three distinct *blocks*; each block contains eight consecutive elections. With three blocks for each of the four treatment conditions this overall yielded 12 blocks of (8 consecutive) elections for each electorate. We then randomized the order of these blocks independently for each electorate using the procedure illustrated by Figure 3. Unknown to the participants, the experiment is executed in three phases. Each phase consists of 4 blocks, including one block from each of the 4 distinct treatment conditions. Within each phase, the orders of the four blocks were randomized. Aside from controlling for order effects, this allows us to investigate learning by comparing subsamples of elections in later blocks to their early counterparts. In summary, Subjects start with eight elections facing a randomly chosen treatment condition. This is followed by eight elections with a different treatment condition. Subjects experience all four distinct treatment conditions before returning to a previously faced one for another eight elections.

Figure 3: Timeline of the Experiment

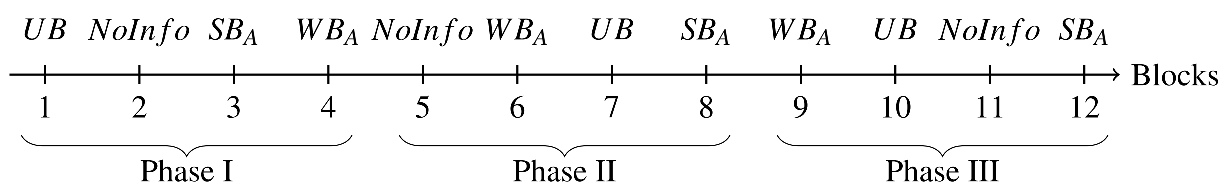

Note: Each block contains eight consecutive rounds of elections under one of the four distinct treatment conditions ($NoInfo$, UB, WB_A, and SB_A). Within each phase, the orders of treatments were randomized, independently for each electorate (the figure just illustrates three out of the 24 possible orders for each phase).

5.2 Belief elicitation

In the last election of each block in treatments other than $NoInfo$, we elicit subjects' posterior beliefs about the probability that they will receive strictly higher earnings if vase A is elected than if vase B is elected, conditional on each possible message. Specifically, we elicit these beliefs by asking "How likely do you think that vase A, if elected, pays you more than vase B, if the robot X sends message s?" for both $s \in \{A, B\}$. Depending on the treatment, X can be ALPHA, BETA or GAMMA. Subjects are required to answer these questions before making their voting decisions.

To incentivize subjects to report their beliefs truthfully, we use the choice list approach (Andersen, Harrison, Lau and Rutström, 2006). Each choice list consists of two columns. The left column contains 21 identical lotteries that pay a fixed prize of ten euros if and only if vase A yields a strictly higher monetary payoff to the subject than vase B does. The right column contains 21 lotteries that pay ten euros with probabilities varying from 0% to 100% (with increments of 5%-points). The layout of the choice list is presented in Online Appendix D.

We use the strategy method for belief elicitation. Specifically, in each round with belief elicitation, subjects are asked to fill out two choice lists, one for each possible message sent by the robot. We impose choice consistency to avoid multiple switching points and allow subjects to complete each choice list by simply making one decision.[14] After subjects have made their decisions, the relevant list (that matches the actual message sent in the subsequent election) is used to determine the lottery applied to the subject. One of the 21 choices from the chosen list is selected (with equal probability for each possible choice), and the option selected by the subject is carried out. At the end of the experiment only one out of these (nine) tasks is chosen, for each subject separately, to determine whether she earns the ten-euro prize.

5.3 Procedures

The experiment was programmed in PHP-MYSQL and conducted at the CREED laboratory of the University of Amsterdam in the summer of 2017. 150 subjects were recruited through the CREED recruitment system. In total, we have data for six sessions. At the beginning of each session, subjects read instructions on screen and have to correctly answer a set of control questions before they can start with the experiment. At the end of the session, they are invited to fill out a post-experimental questionnaire that contains questions about their decision making. The full texts of instructions, sample screen shots, control questions and post-experiment questionnaire can be found in Online Appendix D. Subjects' payments were composed of a show-up fee of ten euros, earnings from one randomly selected election in each of the 12 blocks, and from one randomly selected belief elicitation task. To avoid hedging, the elicitation task to be paid was never chosen from the set of elections selected for payment. Sessions took on average 150 minutes and the average payment was 31.3 euros (with a minimum of 18.9 euros and a maximum of 45.2 euros).

6 Results

We present our experimental results in three subsections. First, we provide an overview of aggregate voting behavior and election outcomes. This is followed by formal statistical tests for Hypotheses 1 and 2. Finally, we investigate the systematic impact of biased endorsers on voters'

[14] If a subject ticks the lottery in the left column for choice X (any number between 1 to 21), then the computer automatically ticks lotteries in the left (right) column for all choices $Y < X$ ($Y > X$). The same rule applies if the subject ticks the lottery in the right column for choice X.

beliefs about candidates' qualities. For all the analyses reported in this section, we use data from the full sample of all 12 blocks (of eight elections each).

6.1 Aggregate outcomes

Figure 4 presents the experimentally observed and theoretically predicted aggregate outcomes in the four information treatments. The left and central panels depict the interim outcomes conditional on messages A and B, respectively. The right panels show the ex-ante outcomes. The striped bars in these panels depict the predicted outcomes by a hybrid behavioral equilibrium model, which is discussed in the next section.

The top and middle panels of Figure 4 show that, except for treatment SB_A, candidate A's vote share and her winning probability are well predicted by BNE and QRE, both qualitatively and quantitatively. In treatment SB_A and conditional on message A, both equilibria predict candidate A's interim vote share and winning probability to be only slightly above 0.5, while the observed outcomes are substantially higher (cf. the top-left and middle-left panels). Consequently, A's ex-ante vote share and winning probability are also substantially higher than both equilibrium predictions under SB_A (cf. the top-right and middle-right panels).

This leads to a substantial deviation from our theoretical prediction: contrary to the non-monotonicity hypothesis 1b, candidate A's ex-ante winning probability increases monotonically as endorser bias χ increases from 0 (UB) to 95 (SB_A).

The bottom panels of Figure 4 show that the observed voter turnout deviates systematically from theoretical predictions, both qualitatively and quantitatively. On the one hand, while both BNE and QRE predict substantial impacts of endorser bias on both interim and ex-ante voter turnout, observed turnout rates are almost invariant with endorser bias. On the other hand, the observed turnout rates are systematically higher than both BNE and QRE predictions. Nevertheless, compared to the three treatments with an endorser, the observed turnout is substantially higher when voters receive no information about candidates' qualities ($NoInfo$). This result is consistent with our theory.

We next investigate the explanatory power of BNE and QRE predictions using linear regressions. Details of these analyses are reported in Online Appendix B.1. We find that BNE predictions can explain 64.9% and 74.2% of the observed variations in election outcomes and candidates' vote shares, respectively, but only 6.9% of voter turnout. For QRE predictions, these fractions are 84.9%, 77.7%, and 6.8%, respectively. These results confirm our earlier observation that vote shares and election outcomes are much better predicted by both equilibrium concepts than voter turnout is.

Finally, we discuss the welfare implications of endorser bias. In our experiment, voter welfare can be straightforwardly measured by the average earnings of voters. Theoretically, both BNE and QRE predict that voter welfare monotonically decreases in endorser bias, and is lowest in

Figure 4: Aggregate Voting Behavior and Election Outcomes

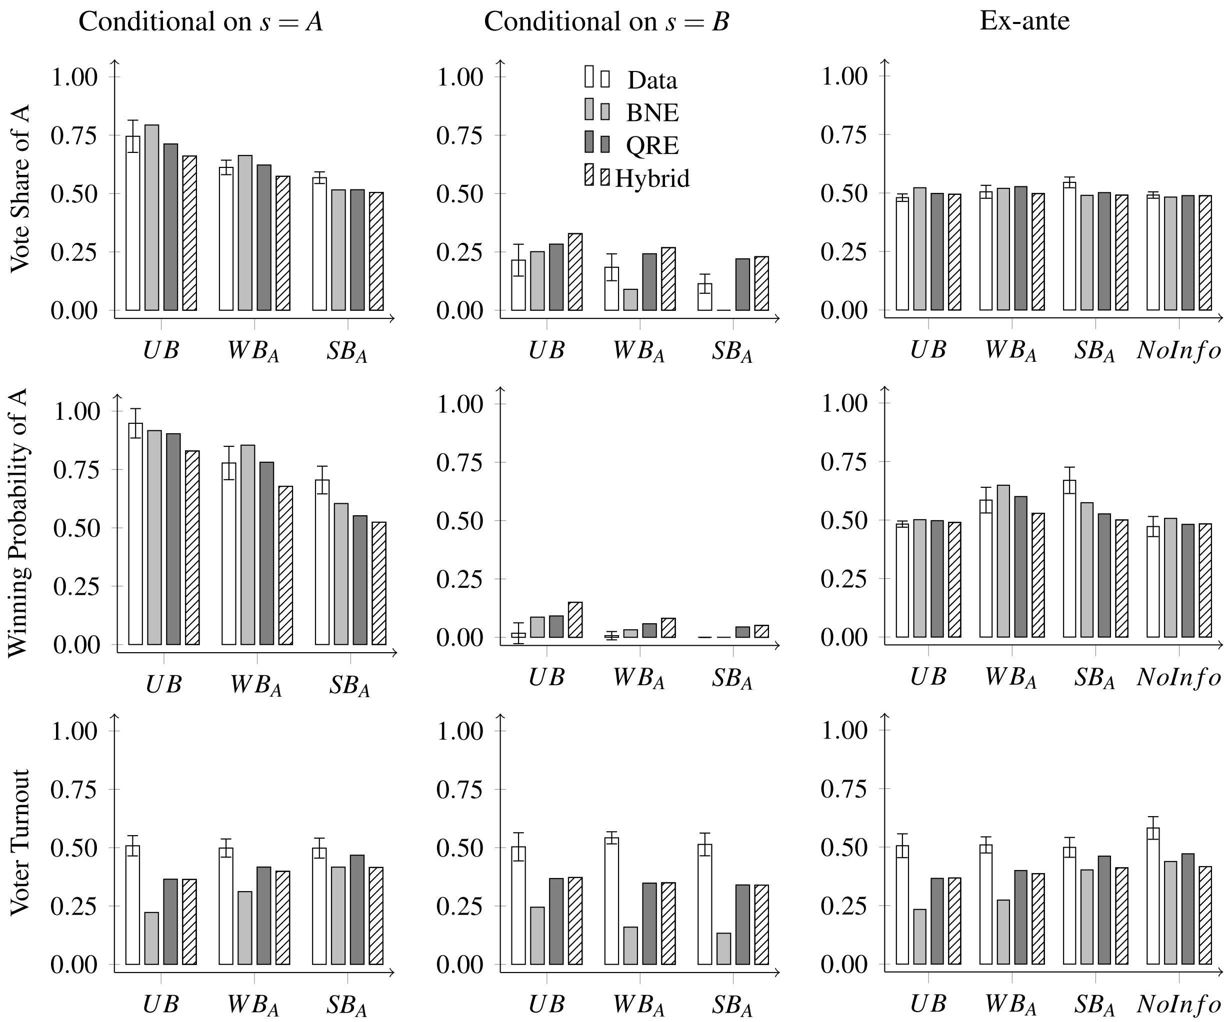

Note: In the top panels, the vote share of A is defined by the fraction of the number of votes for A relative to the total numbers of votes. In the middle panels, candidate A's winning probability equals 1 (0) if its vote share lies strictly above (below) 0.5, and it is coded as 0.5 in case of a tie. The white bars denote the observed outcomes, with 95% confidence intervals plotted. The light and dark grey bars denote the corresponding BNE and QRE predictions, respectively. The striped bars denote the predictions from our hybrid equilibrium model, which is discussed in the next section. All predictions are generated based on the realized parameter draws for the experiment.

treatment $NoInfo$. This holds for two reasons. On the one hand, the probability of electing the quality-superior candidate decreases as endorser bias rises. On the other hand, as implied by Hypothesis 2, aggregate turnout costs increase in endorser bias.

We find that subjects' average earnings indeed monotonically decrease in endorser bias and are lowest in treatment $NoInfo$, which is consistent with theoretical predictions. Nevertheless, subjects' earnings are relatively low compared to the BNE predictions when endorser bias is low (UB and WB_A), and are higher than predicted when endorser bias is high (SB_A) or when there is no information ($NoInfo$). To trace the origin of these deviations, we fist consider the losses due to electing the quality-inferior candidate. This occurs 52% of the time in the $NoInfo$ treatment, while it is predicted to be 62% by BNE. Across the three treatments other than $NoInfo$ the percentage of inferior choices is 31%, which is equal to what is predicted by BNE. We conclude that the main source of inefficiency (compared to BNE) lies in the higher turnout costs. The high and constant turnout levels that we observe (compared to the BNE predictions) imply a welfare loss that is largest for UB and WB_A, where the predicted turnout rates are substantially lower. Across treatments, this leads to the patterns of earnings that we observe.

6.2 Testing Hypotheses

In this subsection we conduct formal statistical tests to examine Hypotheses 1 and 2. All our statistical tests are two-sided even though many of our hypotheses predict treatment effect in specific directions. The results may thus be viewed as a conservative approach to testing our theoretical predictions. All statistical tests we perform are non-parametric and we use the electorate (session) as the unit of independent observation. More specifically, we use exact Fisher-Pitman permutation tests (henceforth, FPp) for pairwise comparisons in means and permutation repeated-measure ANOVA tests (henceforth, pAN) for multiple comparisons in means.

We start with Hypothesis 1, which concerns the impact of endorser bias on candidate A's expected vote share and winning probability. Both are predicted to decrease in endorser bias conditional on either message A or B (H1a). The top-left and middle-left panels of Figure 4 show that, conditional on message A, both the vote share and winning probability of candidate A indeed decrease in endorser bias. These effects are statistically significant (pAN, $p = 0.002$ for A's vote share and $p < 0.001$ for A's winning probability, $N = 6$). The top-central panel shows that A's expected vote share decreases in endorser bias conditional on message B. This effect is also statistically significant (pAN, $p < 0.001$, $N = 6$). However, candidate A's winning probability conditional on message B seems irresponsive to endorser bias (middle-central panel), and the effect is indeed statistically insignificant (pAN, $p > 0.999$, $N = 6$). This last result might be attributed to a boundary effect. As both the predicted and the observed winning probabilities of candidate A are very close to 0 for all levels of endorser bias, any treatment effects are likely to be dominated by noise.

H1b predicts that candidate A's ex-ante winning probability increases from UB_A to WB_A but decreases from WB_A to SB_A, and therefore varies non-monotonically with endorser bias. As the middle-right panel of Figure 4 shows, the observed ex-ante winning probability of candidate A monotonically increases from UB to WB_A (FPp, $p = 0.031$, $N = 6$), and from WB_A to SB_A (FPp, $p = 0.031$, $N = 6$). These results are inconsistent with the non-monotonicity hypothesis. In a nutshell, our data yield mixed support for Hypothesis 1, as summarized in Result 1.

Result 1. *Influence of endorser bias on candidate A's vote share and winning probability:*

(a) Conditional on message A, candidate A's vote share and winning probability both decrease in endorser bias. Conditional on message B, candidate A's vote share decreases in endorser bias, but its winning probability is irresponsive to endorser bias.

(b) Ex-ante, candidate A's winning probability increases monotonically in endorser bias.

Next we examine Hypothesis 2, which concerns the impact of endorser bias on voter turnout. First, H2a (H2b) predicts that conditional on message A (B), voter turnout increases (decreases) with endorser bias. As the bottom-left panel of Figure 4 shows, conditional on message A, voter turnout remains almost the same across all levels of endorser biases; the influence of endorser bias is indeed statistically insignificant (pAN, $p = 0.568$, $N = 6$). Similarly, conditional on message B (bottom-central panel), voter turnout also varies little with endorser bias; though the differences in means are marginally significant (pAN, $p = 0.057$, $N = 6$). H2c predicts that ex-ante, voter turnout is increasing in endorser bias and highest in the control group $NoInfo$. The bottom-right panel of Figure 4 suggests that ex-ante, voter turnout is indeed highest in $NoInfo$, but it responds little to endorser bias (pAN, $p = 0.544$, $N = 6$). The p values of exact Fisher-Pitman permutation tests for comparisons between $NoInfo$ and UB, WB_A, and SB_A, are all 0.031 ($N = 6$). We summarize these observations in Result 2.

Result 2. *Influence of endorser bias on voter turnout:*

(a) Conditional on message A, voter turnout is irresponsive to endorser bias.

(b) Conditional on message B, voter turnout is irresponsive to endorser bias.

(c) Ex-ante, voter turnout is highest in treatment $NoInfo$ and is irresponsive to endorser bias.

Finally, we investigate whether our conclusions are affected by learning. To do so, we conducted the same set of tests for Hypotheses 1 and 2 using only data from the last four blocks (i.e., only from Phase III in Figure 3 with one block for each treatment). We find that most conclusions remain unchanged, with only two exceptions. First, contrary to Result 1b, in the last four blocks candidate A's ex-ante winning probability no longer monotonically increases in endorser bias. Specifically, A's ex-ante winning probabilities are almost identical in WB_A and SB_A (both equal 0.64), and the difference is statistically insignificant (FPp, $p > 0.999$, $N = 6$). Even so, however, this result is still inconsistent with the non-monotonicity hypothesis 1b. Second, contrary to Result 2a, we find that in the last four blocks the expected voter turnout conditional on message A weakly increases in endorser bias, and the differences are statistically significant (pAN, $p = 0.049$, $N = 6$). Though both results suggest learning effects towards the equilibrium predictions, the effects are insufficient to create comparative static patterns that are consistent with our hypotheses.[15]

6.3 Systematic impact of biased endorsers on voters' beliefs

In this subsection, we explore the systematic impact of biased endorsers on voters' perceptions about candidates' qualities using their reported posterior beliefs about the event that they will receive strictly higher earnings if vase A is elected than if vase B is elected. According to (1), this event is equivalent to $k + v_i > 0$ for a voter with ideological preference $v_i \in \{-100, -50, -20, 20, 50, 100\}$. We can thus interpret subjects' reported beliefs as her posterior about event $k + v_i > 0$. Recall that our choice-list elicitation approach yields belief intervals of 5%-points. We use the midpoint of subjects' chosen intervals as an estimate of their reported beliefs.

Our identification strategy relies on the observation that rational voters' beliefs cannot be systematically affected by exposure to biased endorsers. Formally, let $Pr[k + v_i > 0|s, \chi]$ be the posterior belief about event $k + v_i > 0$ reported by a rational voter with ideology v_i, conditional on message s and endorser bias χ. Let $Pr[k + v_i > 0|\chi] \equiv \sum_{s \in \{A,B\}} Pr[k + v_i > 0|s, \chi] \cdot Pr(s|\chi)$ be the average posterior belief, weighted by the ex-ante probabilities of sending messages A and B when the bias is χ. It follows from Bayes' rule that $Pr[k + v_i > 0|\chi]$ is always equal to the prior $Pr[k + v_i > 0]$, which differs across v_i but is independent of χ. This allows us to identify the systematic impact of endorser bias on voters' beliefs by testing the following Hypothesis 3.

Hypothesis 3. *For all $v_i \in \{-100, -50, -20, 20, 50, 100\}$, $Pr[k + v_i > 0|\chi]$ is invariant with χ for $\chi \in \{0, 55, 95\}$.*

A violation of Hypothesis 3 indicates that voters' average posterior expectations are systematically influenced by exposure to biased endorsers. As in the previous subsection, we perform the permutation repeated-measure ANOVA tests while using the Bonferroni correction to control for the family-wise error rate at 5% level. For this analysis we exclude *inconsistent* pairs of reported

[15] A more detailed comparison of how learning affects behavior across rounds is presented in Appendix B.2. One of the things this shows is that voters' average expected earnings are higher in the last four blocks than across all blocks. Learning appears to allow subjects to better respond to information.

beliefs such that $Pr[k+v_i>0|A,\chi] < Pr[k+v_i>0|B,\chi]$.[16] Our results are not affected by such exclusion.

Figure 5: The Systematic Impact of Biased Endorsers on Voters' Reported Posterior Beliefs

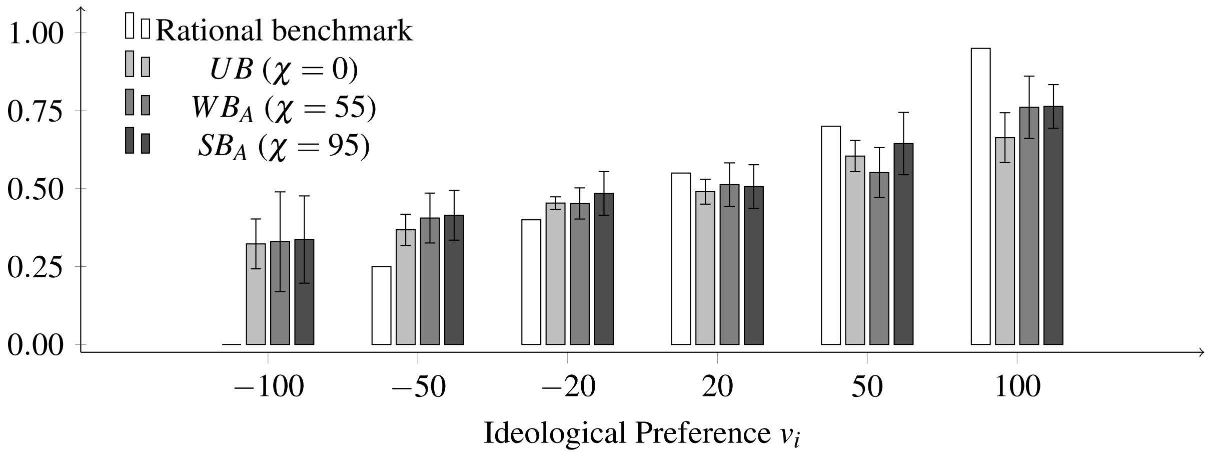

Note: The rational benchmark (white bars) equals, for each $v_i \in \{-100,-50,-20,20,50,100\}$, the objective prior $Pr[k+v_i>0]$. Other bars show the average reported posterior beliefs $Pr[k+v_i>0|\chi]$ for $\chi \in \{0,55,95\}$. 95% confidence intervals are plotted. Inconsistent pairs of reported beliefs are excluded from the analyses. Bayes' rule implies that the average posterior beliefs $Pr[k+v_i>0|\chi]$ reported by a rational voter must equal to the prior $Pr[k+v_i>0]$ (i.e., the rational benchmark).

Figure 5 shows the average posterior beliefs reported by voters with different ideological preferences $v_i \in \{-100,-50,-20,20,50,100\}$. We find that the average reported posterior beliefs about $k+v_i>0$ vary little with endorser bias for most values of v_i. Indeed, after correcting for the multiple-testing problem, we cannot reject the null Hypothesis 3 at a family-wise error rate of 5%.[17] Taken together, we find no evidence that voters' beliefs about candidates' qualities are systematically distorted by biased endorsers.

Result 3. *Voters' average reported posterior beliefs do not vary with endorser bias χ.*

It is important to note, however, that the fact that voters' beliefs are not systematically distorted by biased endorsers does not imply that voters are perfectly rational in belief formation. As Figure 5 shows, for most values of v_i voters' average reported posterior beliefs systematically deviate towards 0.5 compared to the rational benchmark (i.e., the objective prior about $k+v_i>0$). Such deviations are more pronounced for voters with moderate or extreme ideological preferences (i.e., $|v_i|=50$ or 100), whose objective priors about $k+v_i>0$ are closer to the boundaries 1 or 0. Though this result might be driven by noise (e.g., a systematic tendency to report closer to 0.5), it may equally likely be caused by non-Bayesian updating.

[16] 32% of the pairs of reported beliefs are inconsistent in this way. These pairs are evenly spread across subjects; of the nine belief elicitation tasks, only 7% of the subjects were inconsistent zero or nine times.

[17] Our conclusion is robust under other methods to control for family-wise error rates or false discovery rates, such as Holm-Bonferroni and Benjamini-Hochberg corrections, which are less stringent than the Bonferroni correction.

Until now, we have considered the unconditional posterior beliefs $Pr[k+v_i > 0|\chi]$. One might reasonably wonder whether the averaging across messages that this involves hides underlying differences. To investigate this, we also considered the posterior beliefs conditional on $s = A(Pr[k+v_i > 0|\chi,A])$ and $s = B(Pr[k+v_i > 0|\chi,B])$ Doing so shows that the comparative statics of the reported conditional posteriors are qualitatively consistent with rational updating (beliefs change in the right direction). The levels of these posterior beliefs are, however, different from the Bayesian posteriors (they do not change strongly enough). This is similar to what we observed for the unconditional posteriors in Figure 5.

7 Behavioral Explanations for Observed Discrepancies

The two main discrepancies between the theoretical predictions and the experimental results are (i) contrary to Hypothesis 1, candidate A's ex-ante winning probability increases monotonically with an endorser's bias towards her, and (ii) contrary to Hypothesis 2, voter turnout varies little with the bias.[18] These discrepancies warrant further investigation. In this section we show that subjects' observed voting behavior in our experiment can be explained, to a substantial extent, by a hybrid equilibrium model combining three behavioral mechanisms.

First, voters display a higher level of rationality in candidate choices than in turnout decisions. Arguably, this is because the latter decisions are much more cognitively demanding than the former. To make the optimal candidate choice (i.e., which candidate to vote for if one votes), a voter only needs to infer which candidate is better based on her private ideology and her belief about candidates' quality differential. To make the optimal turnout decisions (i.e., whether to cast a costly vote), the voter has to infer the chances of casting a pivotal vote and then judge whether the expected benefits from voting exceed its costs. For this reason, we expect voters to make fewer mistakes in candidate choice than in the turnout decision. This is also consistent with our finding that observed vote shares and election outcomes are much better predicted by theory than voter turnout. [19]

Second, voters do not update their beliefs about candidates' qualities rationally, even though their beliefs are not systematically distorted. It is widely documented by both laboratory experiments

[18] Of course, we may not have found all of the equilibria for our election game. We cannot, however, conceive of any equilibrium where turnout levels of 50% or more constitute best responses under SB_A and $s = B$. If an endorser is strongly biased towards A and endorses B anyway, then this implies a very high expected quality difference in favor of B. In fact, in any equilibrium, any voter who turns out will vote for B. This turns the voting game into a step-level public goods game where only one vote for B is needed to satisfy all. We cannot conceive of a Nash equilibrium that rationalizes a turnout of 50%.

[19] A full analysis of the role of beliefs about other voters' behavior would require eliciting such beliefs, as suggested by an anonymous reviewer. We believe this to be beyond the scope of our experiment. Nevertheless, the structural modeling approach that we take in this section does allow us to study how specific types of 'irrationalities' in belief formation and behavioral responses to beliefs might explain the discrpencies between the BNE/QRE predictions and observed behavior.

and field studies that belief formation systematically deviates from Bayes' rule. On the one hand, many studies suggest that people update their beliefs too conservatively (Mobius, Niederle, Niehaus and Rosenblat, 2011). In our experiment, conservatism in belief formation may lead voters to underestimate the variation of the expected quality differences, and thus also underestimated the variation of the closeness of elections, across distinct information environments. This may mitigate the predicted treatment effects of biased endorsers on voter turnout. On the other hand, people may also excessively response to information provided by biased senders, if they do not sufficiently filter out the impact of senders' biases (Cain, Loewenstein and Moore, 2005). In our experiment, this may give biased endorsers extra power in influencing the election in their favored direction. This can potentially explain the failure of the non-monotonicity hypothesis found in our experiment.

Third, voters may mis-perceive their ability to influence elections, that is, to cast a pivotal vote. We argue that voters are particularly likely to exhibit partial "competition neglect", that is, voters underestimate the degree to which the closeness of elections is correlated with the information released by endorsers.[20] Such competition neglect can lead voters to underestimate the variations in their pivotality across information environments; this would lead to insufficient responsiveness of turnout to treatment variations in our experiment.

Because the three behavioral mechanisms may lead to observably similar voting behavior, it is hard to distinguish between them by direct data comparisons. For this reason, we use a structural approach to identify their roles. To do so, we first construct a parametric hybrid behavioral equilibrium model that captures all three mechanisms, and then estimate the model parameters that best fit our experimental data. The formal model setup and estimation approach are explained in Online Appendix C. In the following two subsections, we present our estimation results and inspect the out-of-sample predictive power of our structural model.

7.1 A hybrid behavioral equilibrium model

To capture distinct levels of rationality, we construct an extensive form version of our election game in which voters make decisions in two stages and derive the *agent quantal response equilibria* (AQRE; McKelvey and Palfrey (1998)) for this model. At the first stage (turnout decision), a voter decides between casting a costly vote or abstaining. If the voter chooses to vote, then she moves to the second stage (candidate choice) and decides which candidate to vote for. We allow the voter's decisions to be affected by distinct logit response parameters: λ_p for candidate choices, and λ_t for

[20] In the experiment, the correlation between the actual closeness of elections and the information voters receive is pronounced and consistent with our theoretical predictions. More precisely, the actual fractions of close elections (i.e., elections in which the difference between the number of votes for candidate A and B is no more than one) are: 31% under $NoInfo$, 27% under $(SB_A, s = A)$, 25% under $(WB_A, s = A)$, 6% under $(UB, s = A)$, 6% under $(UB, s = B)$, 3% under $(WB_A, s = B)$, and 0% under $(SB_A, s = B)$.

turnout decisions.[21]

To account for non-Bayesian updating, we take a parsimonious approach by assuming that the electorate forms a common posterior expectation, which equals the rational posterior expectation multiplied by a parameter $\beta \geq 0$. Voters are perfect Bayesian when $\beta = 1$. If $\beta < 1$, voters are *conservative* in belief formation; their beliefs respond insufficiently to information. If instead $\beta > 1$, then voters are *excessive* in belief updating; their beliefs respond too sensitively to information.[22]

To capture competition neglect, we assume that with probability $\rho \in [0,1]$, voters infer correctly the actual probabilities of casting pivotal votes under any specific information environment; with the remaining probability $1-\rho$ voters mistakenly believes that the chances of casting pivotal votes are invariant in the information environment. Therefore, if $\rho = 0$, voters completely ignore any correlation between the actual closeness of elections and the information environment. If instead $\rho = 1$, then voters precisely infer pivotal probabilities in each information environment, as in standard BNE and QRE. With $\rho \in (0,1)$, voters partially realize the relationship between pivotal probabilities and the information environment, but insufficiently so compared to a rational agent. A higher ρ implies a lower degree of competition neglect.

By putting restrictions on model setup or parameters, we can either activate or shut down any particular behavioral mechanism. This allows us to examine the role played by each of the three mechanisms. We estimate the parameters for all possible model specifications, depending on whether each behavioral mechanism is included or not. The estimation results are presented in Table 2. Model (1) represents a standard QRE model where all the three behavioral mechanisms are excluded. In models (2) to (4), we introduce each of the three mechanism separately. In models (5) to (7), we introduce two distinct mechanisms at the same time. Finally, in model (8), we include all three behavioral mechanisms. The performance of each model in fitting our experimental data can be evaluated through the "fitting score", which re-scales the estimated log-likelihood to a range between 0% and 100%. Each 1%-point increase in the fitting score is equivalent to an increase in estimated log-likelihood by about 114.

Table 2 shows robust evidence supporting all three behavioral mechanisms. First of all, for all hybrid models allowing for distinct levels of rationality (columns (2) and (6)-(8)), the estimated $\widehat{\lambda_p}$'s are about 2.5 times higher than $\widehat{\lambda_t}$'s, which suggests that voters indeed exhibit higher levels of rationality in candidate choices than turnout decisions.[23] Second, for all models allowing for

[21] We also estimated an alternative type of extensive form AQRE model where voters decide which candidate to support at the first stage, and decide whether to abstain or to vote for the candidate they choose to support at the second stage. Using Vuong's closeness tests (Vuong, 1989), we show that this model is outperformed by the AQRE model presented here.

[22] To be consistent with Result 3 reported in the previous section, we do not allow voters' posterior beliefs to be systematically distorted by biased endorsers. Relaxing this restriction would not change our conclusions qualitatively.

[23] To examine whether our two-parameter AQRE model provides a better fit to data compared to the standard QRE, we use Vuong's closeness tests (Vuong, 1989). For all model comparisons, the p values from Vuong's closeness tests

Table 2: ML Estimation Results for the Hybrid Equilibrium Models

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Distinct Levels of Rationality	No	Yes	No	No	No	Yes	Yes	Yes
Non-Bayesian Updating	No	No	Yes	No	Yes	No	Yes	Yes
Competition Neglect	No	No	No	Yes	Yes	Yes	No	Yes
Estimated Parameters								
$\widehat{\lambda}$	15.83	-	17.51	18.29	18.46	-	-	-
$\widehat{\lambda_p}$	-	32.22	-	-	-	28.86	26.77	27.52
$\widehat{\lambda_t}$	-	8.47	-	-	-	12.01	10.97	12.46
$\widehat{\beta}$	1.00*	1.00*	0.43	1.00*	0.67	1.00*	0.51	0.74
$\widehat{\rho}$	1.00*	1.00*	1.00*	0.05	-0.06	0.06	1.00*	-0.03
Log-Likelihood	-19657	-19486	-19113	-18701	-18514	-18542	-19026	-18366
Fitting Score	70.4%	71.9%	75.2%	78.8%	80.4%	80.2%	75.92%	81.7%

Note: All models are estimated using all subjects' voting decisions from our experiment conducted in 2017. The number of observations is 25200. The fitting score (the last row), whose construction is explained in Online Appendix C, re-scales the estimated log-likelihood to a range between 0% and 100%; an 1%-point increase in fitting score is equivalent to an increase in log-likelihood by about 114. In the panel for the estimated parameters, the cells marked by asterisks are values implied by model restrictions. Specifically, all models excluding non-Bayesian updating and competition neglect assume $\beta = 1$ and $\rho = 1$, respectively.

non-Bayesian updating (columns (3), (5), (7) and (8)), the estimated $\widehat{\beta}$'s are all substantially below 1, which implies that voters update their beliefs more conservatively compared to a Bayesian.[24] Third, in all models allowing for competition neglect (columns (4)-(6) and (8)), the estimated $\widehat{\rho}$'s are all very close to 0.[25] This suggests that voters are almost completely ignorant of any correlation between the information environment and the chances of casting pivotal votes.

Aside from providing evidence for the three behavioral mechanisms, Table 2 also allows us to inspect the relative contributions of these mechanisms in explaining observed voting behavior. The comparison between models (1) and (2) shows that, relative to the standard QRE, a hybrid model including distinct levels of rationality alone improves the fitting score by 1.5%-points. By comparing models (1) and (3), we find that non-Bayesian belief updating alone improves the fitting score by 4.8%-points. By comparing models (1) and (4), we find that competition neglect alone improves the fitting score by 8.4%-points. These results suggest that, among the three mechanisms, competition neglect contributes the most in explaining observed voting behavior.

Because all three behavioral mechanisms may lead to observably similar voting behavior, it is important to understand whether the impact of any mechanism on observed voting behavior is

are less than 0.001. This suggests that, holding other behavioral mechanisms fixed, our two-parameter AQRE model outperforms the standard QRE.

[24] The p values of LR tests for null hypothesis $\beta = 1$ are below 0.001 for all these models.

[25] The p values of LR tests for null hypothesis $\rho = 1$ are below 0.001 for all these models.

already captured by other mechanisms. If this is the case, then excluding this mechanism from model (8), which include all three behavioral mechanism, should do little harm to the fitting performance. By comparing model (5)-(7) to model (8), we observe that excluding any behavioral mechanism will decrease the fitting score by more than 1.3%-points, which corresponds to a substantial drop in the estimated log-likelihood of 148. Therefore, we conclude that these behavioral mechanisms capture distinct aspects of voting behavior and are not simply substitutes.

Finally, we can relate our evidence of two of the behavioral mechanisms to previous findings in the literature, in particular on voter turnout. The higher noise levels in the turnout decision as compared to candidate choice reflects a higher level of 'irrationality'. This mirrors earlier observations that there are limits on the rationality observed in turnout decisions (Duffy and Tavits, 2008; Gerber et al., 2020). On the other hand, the high extent of competition neglect that we see supports the idea that voters are not good at using information to get a perception of the closeness that they should expect in the election (Enos and Fowler, 2014; Bursztyn et al., 2017; Moskowitz and Schneer, 2019). In this way, our experiments have helped us to develop behavioral underpinnings of regularities that have been observed elsewhere.

7.2 Out-of-Sample Predictive Power

In this subsection we inspect whether our hybrid model can explain the observed voting behavior better than BNE and QRE. To do so, we use the hybrid model to generate predictions for voters' behavior in our experimental sessions. To avoid over-fitting, all parameters are re-estimated using data from our previous experiment conducted in 2016. The predictions of the hybrid model for aggregate voting behavior and election outcomes are depicted by the striped bars in Figure 4. Regarding the aggregate vote shares and winning probabilities of each candidate, our hybrid model generates quantitatively similar predictions as with BNE and QRE. The hybrid model cannot account for the failure of non-monotonicity hypothesis. However, both interim and ex-ante turnout rates predicted by the hybrid model are much less responsive to treatment variations compared to BNE and QRE predictions. Therefore, our hybrid model partially accounts for the insufficient responsiveness of observed voter turnout.

Recall from the previous section that the failure of the non-monotonicity hypothesis is driven by the fact that both the observed vote share and winning probability of candidate A substantially exceed 0.5. Figure 6 suggests that this substantial deviation is driven by the asymmetric behavioral responses by voters from different teams. Consider how behavior changes in the transition from $NoInfo$ (panels (a) and (c)) to the case where the endorser is strongly biased and sends message A (panels (b) and (d)). The actions of voters in team A ($v_i > 0$) change very little while voters from team B ($v_i < 0$) substantially increase their vote shares for A and decrease their turnout probabilities. In other words, a very biased endorsement of A has little impact on the behavior of A's own supporters; it fails to mobilize them to turnout and does not make them vote for A (even) more. The benefit for A lies in how B's most extreme supporters respond. A very biased endorsement of A demobilizes them and even persuades more of them to vote for A. This observed effect seems contrary to the commonly held wisdom of 'confirmation bias' (Nickerson, 1998), which suggests that biased information should have stronger effects on supporters of the endorsed candidate than on the opposing supporters.[26]

The moderating role of partisanship in how voters respond to endorsements gives strongly A-biased endorsers substantially more power to influence elections than predicted, and therefore ultimately allows candidate A to win the election with a substantially higher probability than predicted in treatment SB_A when the message is A (cf. the middle-left panel of Figure 4). Such asymmetric responses can thus explain the failure of our non-monotonicity prediction. They cannot, however, be rationalized by our hybrid behavioral equilibrium model.[27] We suspect that these discrepancies might be driven by factors beyond our model, such as non-equilibrium behavior, unobserved heterogeneity across subjects, or failure of the common knowledge assumption.[28]

8 Conclusion

This paper reports results from a laboratory experiment designed to investigate the influences of biased expert endorsers (e.g., mass media) on voting behavior and election outcomes. We construct an election game based on the general theoretical framework developed in Sun, Schram and Sloof (2019) and derive both BNE and QRE for this game. These predict that ex-ante, in an environment with a single endorser the election outcome varies non-monotonically with the endorser's bias while voter turnout monotonically increases in this bias.

We subsequently test these predictions in a laboratory experiment. We find that the observed

[26] More precisely, we consider the following notion of confirmation bias: all else equal, voters may treat endorsements that advise them to vote for the a priori favored candidate more credible than endorsements that recommend the opposing candidate. If so, then upon seeing a biased endorsement for candidate A, supporters of A should deem it more reliable and respond to it more strongly than supporters of B. In the experiment we find the opposite: an extremely biased endorsement for A has a larger impact on supporters of B.

[27] Another potential mechanism that might drive subjects' asymmetric responses to information is the tendency to vote for the ex-post winner (Agranov et al., 2018). In our experiment, this tendency would imply that a voter may vote for the candidate she prefers less if she expects this candidate to win the election ex-post. Consequently, the realized vote shares of the winning candidate should be higher than their corresponding theoretical predictions. Such an effect would be strongest in UB, where the expected margin of victory is largest. Candidate A's realized vote shares should be substantially higher than equilibrium predictions after message A, while substantially lower after message B. The top panels of Figure 4, however, show no evidence of this occurring.

[28] These may include explanations such as level-k thinking, cognitive hierarchy and noisy introspection (Goeree, Holt and Palfrey, 2016; Goeree, Louis and Zhang, 2018). We leave the exploration of the roles of these non-equilibrium behavioral mechanisms for future research.

Figure 6: Observed and Predicted Aggregate Voting Behavior by Ideological Preference v_i

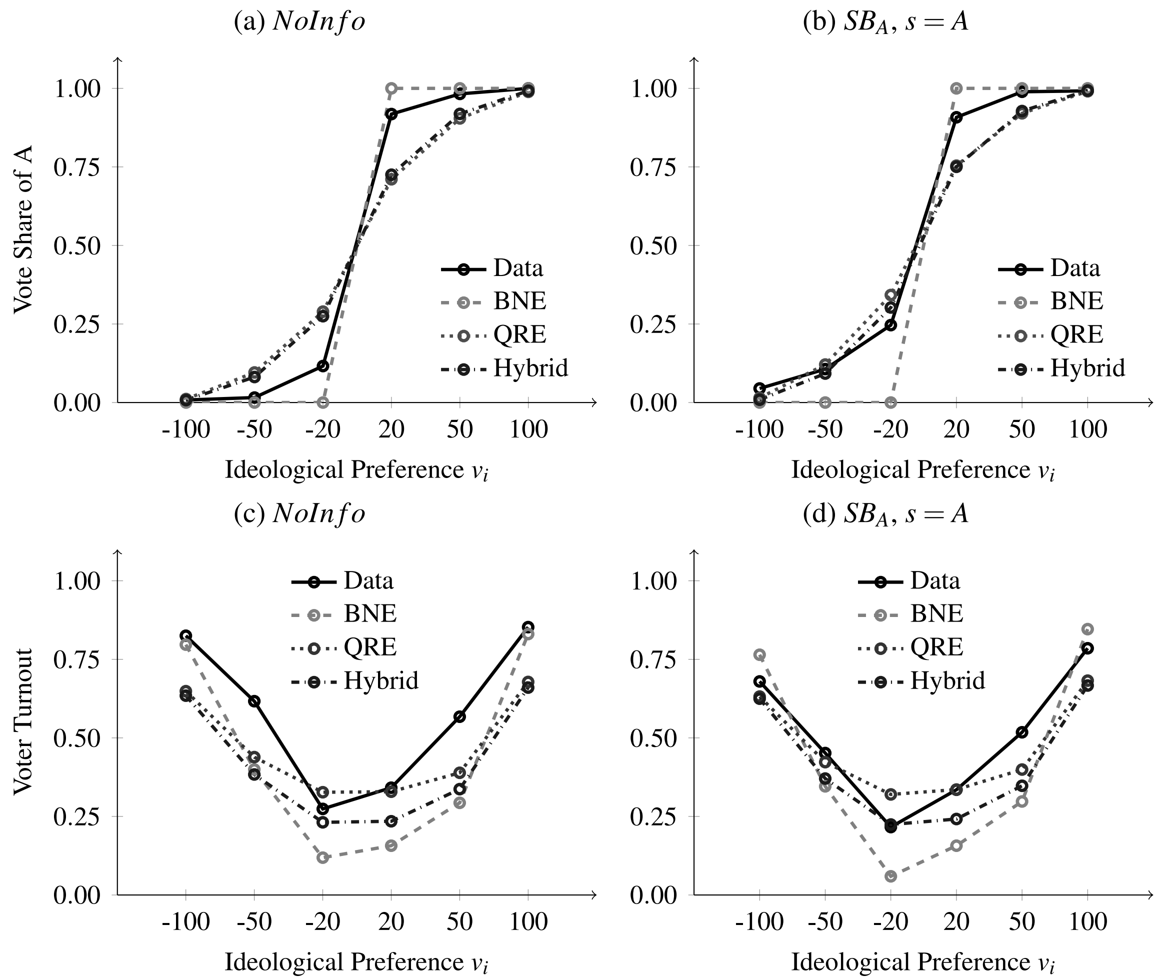

Note: Panels (a) and (b) depict the observed and predicted vote share for candidate A among voters with a given ideological preference $v_i \in \{-100, -50, -20, 20, 50, 100\}$. Panels (c) and (d) depict the observed and predicted voter turnout among voters with distinct ideological preferences. Outcomes in panels (a) and (c) are obtained in treatment $NoInfo$. Outcomes in panels (b) and (d) are obtained in treatment SB_A conditional on message A.

party vote shares and election outcomes are well predicted, albeit that the non-monotonic relationship is not confirmed. Observed voter turnout, however, is much less responsive to endorser bias than predicted. Our analyses for reported beliefs show no evidence that subjects' perceptions of candidates are systematically distorted by biased endorsers. We conclude that our data are best explained by a hybrid behavioral equilibrium model that combines (i) distinct levels of rationality for candidate choice and turnout decisions; (ii) conservative belief updating; and (iii) partial competition neglect, that is, voters underestimate the degree to which the chances of casting pivotal votes are correlated with the information environment. Among the three behavioral mechanisms, partial competition neglect plays the most important role in explaining the observed voting behavior.

Our results imply that voters are not fully irrational in their response to the information provided by biased endorsers. Nevertheless, their rationality is clearly bounded. The three bounds to rationality that we observe are intuitive and provide clear implications for understanding the world outside the laboratory. The observation that best responses are subject to noise is not surprising. In fact, our estimate of the noise parameter suggests that the noise vis-à-vis pure best response is limited. Similarly, the fact that belief updating is not perfectly Bayesian should not come as a surprise either. In fact, our results show that the Bayesian benchmark provides a solid ground to predict how voters adjust their beliefs. Finally, the fact that we find substantial competition neglect is also intuitive. In our design (and arguably also in the world outside the laboratory), subjects face an information environment that changes regularly. Keeping track of how a particular information environment affects the competitiveness of elections appears to be a cognitively challenging task.

It is worth noting that an important difference between our experimental environment and elections in the world outside the laboratory is the existence of public opinion polls. These have been shown to be effective in informing voters about the closeness of elections and therefore influencing turnout decisions (Großer and Schram, 2010; Agranov et al., 2018; Bursztyn et al., 2017). We suspect that the presence of public opinion polls might partially alleviate the difficulty of learning the closeness of elections in constantly changing information environments. On the other hand, opinion polls are notoriously better at predicting party preferences than turnout (e.g., (Rothenberg, 2014)). The extent to which polls can then help voters overcome the problem of competition neglect thus remains an open question. We consider this to be an interesting topic to explore in the future.

All in all, we believe that our results provide a solid ground for the conclusion that voters' behavior in an environment with biased endorsers may be noisy, but that it does have a strong core of rational response. Deviations from rationality can to a large extent be understood by a set of clear behavioral mechanisms.

References

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *Quarterly Journal of Economics* 130(4):1885–1939.

Agranov, Marina, Jacob K Goeree, Julian Romero and Leeat Yariv. 2018. "What makes voters turn out: The effects of polls and beliefs." *Journal of European Economic Association* 16(3):825–856.

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau and E Elisabet Rutström. 2006. "Elicitation using multiple price list formats." *Experimental Economics* 9(4):383–405.

Ashworth, Scott and Ethan Bueno De Mesquita. 2009. "Elections with platform and valence competition." *Games and Economic Behavior* 67(1):191–216.

Baron, David. 2006. "Persistent media bias." *Journal of Public Economics* 90(1):1–36.

Besley, Timothy and Andrea Prat. 2006. "Handcuffs for the grabbing hand? The role of the media in political accountability." *American Economic Review* 96(3):720–736.

Blais, André. 2000. *To vote or not to vote? The merits and limits of rational choice theory*. University of Pittsburgh Press.

Blais, André and Rafael Hortala-Vallve. 2020. "Conformity and individuals' response to information about aggregate turnout." *Political Behavior* pp. 1–15.

Bursztyn, Leonardo, Davide Cantoni, Patricia Funk and Noam Yuchtman. 2017. Polls, the Press, and Political Participation: The Effects of Anticipated Election Closeness on Voter Turnout. Working Paper 23490 National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w23490*

Cagé, Julia. 2017. "Media Competition, Information Provision and Political Participation: Evidence from French Local Newspapers and Elections, 1944-2014." CEPR Discussion Paper No. DP12198. Available at SSRN: https://ssrn.com/abstract=3016784.

Cain, Daylian M, George Loewenstein and Don A Moore. 2005. "The dirt on coming clean: Perverse effects of disclosing conflicts of interest." *The Journal of Legal Studies* 34(1):1–25.

Calvert, Randall L. 1985. "The value of biased information: A rational choice model of political advice." *The Journal of Politics* 47(2):530–555.

Chan, Jimmy and Wing Suen. 2008. "A spatial theory of news consumption and electoral competition." *Review of Economic Studies* 75(3):699–728.

Chiang, Chun-Fang and Brian Knight. 2011. "Media bias and influence: Evidence from newspaper endorsements." *Review of Economic Studies* 78(3):795–820.

Cialdini, Robert B. 2007. *Influence: The psychology of persuasion*. Revised Edition, HarperCollins New York.

Coleman, James S. 1994. *Foundations of social theory*. Harvard university press.

DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News effect: Media bias and voting." *Quarterly Journal of Economics* 122(3):1187–1234.

DellaVigna, Stefano and Matthew Gentzkow. 2010. "Persuasion: empirical evidence." *Annual Review of Economics* 2(1):643–669.

Djankov, Simeon, Caralee McLiesh, Tatiana Nenova and Andrei Shleifer. 2003. "Who owns the media?" *Journal of Law and Economics* 46(2):341–382.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. Harper, New York.

Drago, Francesco, Tommaso Nannicini and Francesco Sobbrio. 2014. "Meet the press: How voters and politicians respond to newspaper entry and exit." *American Economic Journal: Applied Economics* 6(3):159–188.

Duffy, John and Margit Tavits. 2008. "Beliefs and voting decisions: A test of the pivotal voter model." *American Journal of Political Science* 52(3):603–618.

Durante, Ruben and Brian Knight. 2012. "Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy." *Journal of European Economic Association* 10(3):451–481.

Enikolopov, Ruben, Maria Petrova and Ekaterina Zhuravskaya. 2011. "Media and political persuasion: Evidence from Russia." *American Economic Review* 101(7):3253–3285.

Enos, Ryan D and Anthony Fowler. 2014. "Pivotality and turnout: Evidence from a field experiment in the aftermath of a tied election." *Political Science Research and Methods* 2(2):309.

Feddersen, Timothy J and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse." *American Economic Review* 86(3):408–424.

Feddersen, Timothy J and Wolfgang Pesendorfer. 1999. "Abstention in elections with asymmetric information and diverse preferences." *American Political Science Review* 93(2):381–398.

Feddersen, Timothy and Wolfgang Pesendorfer. 1997. "Voting behavior and information aggregation in elections with private information." *Econometrica* 65(5):1029–1058.

Gentzkow, Matthew and Jesse M Shapiro. 2006. "Media bias and reputation." *Journal of Political Economy* 114(2):280–316.

Gentzkow, Matthew and Jesse M Shapiro. 2010. "What drives media slant? Evidence from US daily newspapers." *Econometrica* 78(1):35–71.

Gentzkow, Matthew, Jesse M Shapiro and Daniel F Stone. 2015. Media bias in the marketplace: Theory. In *Handbook of media economics*. Vol. 1 Elsevier pp. 623–645.

Gentzkow, Matthew, Jesse M Shapiro and Michael Sinkinson. 2011. "The effect of newspaper entry and exit on electoral politics." *American Economic Review* 101(7):2980–3018.

Gerber, Alan, Dean Karlan and Daniel Bergan. 2009. "Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions." *American Economic Journal: Applied Economics* 1(2):35–52.

Gerber, Alan, Mitchell Hoffman, John Morgan and Collin Raymond. 2020. "One in a Million: Field Experiments on Perceived Closeness of the Election and Voter Turnout." *American Economic Journal: Applied Economics* 12(3):287–325.

Goeree, Jacob K and Charles A Holt. 2005. "An explanation of anomalous behavior in models of political participation." *American Political Science Review* 99(2):201–213.

Goeree, Jacob K, Charles A Holt and Thomas R Palfrey. 2016. *Quantal Response Equilibria*. Springer.

Goeree, Jacob K, Philippos Louis and Jingjing Zhang. 2018. "Noisy Introspection in the 11–20 Game." *The Economic Journal* 128(611):1509–1530.

Groseclose, Tim. 2001. "A model of candidate location when one candidate has a valence advantage." *American Journal of Political Science* 45(4):862–886.

Groseclose, Tim and Jeffrey Milyo. 2005. "A measure of media bias." *Quarterly Journal of Economics* 120(4):1191–1237.

Großer, Jens and Arthur Schram. 2010. "Public opinion polls, voter turnout, and welfare: An experimental study." *American Journal of Political Science* 54(3):700–717.

Kennedy, Patrick J and Andrea Prat. 2019. "Where Do People Get Their News?" *Economic Policy* 34(97):5–47.

Levine, David K and Thomas R Palfrey. 2007. "The paradox of voter participation? A laboratory study." *American Political Science Review* 101(1):143–158.

Martin, Gregory J and Joshua McCrain. 2019. "Local news and national politics." *American Political Science Review* 113(2):372–384.

McKelvey, Richard D and Thomas R Palfrey. 1995. "Quantal response equilibria for normal form games." *Games and Economic Behavior* 10(1):6–38.

McKelvey, Richard D and Thomas R Palfrey. 1998. "Quantal response equilibria for extensive form games." *Experimental Economics* 1(1):9–41.

McMillan, John and Pablo Zoido. 2004. "How to subvert democracy: Montesinos in Peru." *Journal of Economic Perspectives* 18(4):69–92.

Mobius, Markus M, Muriel Niederle, Paul Niehaus and Tanya S Rosenblat. 2011. Managing Self-Confidence: Theory and Experimental Evidence. Working Paper 17014 National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w17014*

Moskowitz, Daniel J and Benjamin Schneer. 2019. "Reevaluating competition and turnout in US house elections." *Quarterly Journal of Political Science* 14(2):191–223.

Mullainathan, Sendhil and Andrei Shleifer. 2005. "The market for news." *American Economic Review* 95(4):1031–1053.

Nickerson, Raymond. 1998. "Confirmation bias: A ubiquitous phenomenon in many guises." *Review of general psychology* 2(2):175–220.

Oliveros, Santiago and Felix Várdy. 2015. "Demand for slant: How abstention shapes voters' choice of news media." *Economic Journal* 125(587):1327–1368.

Palfrey, Thomas R and Howard Rosenthal. 1985. "Voter participation and strategic uncertainty." *American Political Science Review* 79(1):62–78.

Prat, Andrea. 2018. "Media power." *Journal of Political Economy* 126(4):1747–1783.

Prat, Andrea and David Strömberg. 2013. "The political economy of mass media." *Advances in Economics and Econometrics* 2:135.

Prior, Markus. 2013. "Media and political polarization." *Annual Review of Political Science* 16:101–127.

Puglisi, Riccardo and James M Snyder. 2011. "Newspaper coverage of political scandals." *Journal of Politics* 73(3):931–950.

Puglisi, Riccardo and James M Snyder. 2015*a*. "The balanced US press." *Journal of European Economic Association* 13(2):240–264.

Puglisi, Riccardo and James M Snyder. 2015*b*. Empirical studies of media bias. In *Handbook of media economics*. Vol. 1 Elsevier pp. 647–667.

Riker, William H and Peter C Ordeshook. 1968. "A Theory of the Calculus of Voting." *The American political science review* 62(1):25–42.

Rothenberg, Stuart. 2014. "Election Forecasting and Public Opinion Polls." *PS, Political Science & Politics* 47(2):336.

Schram, Arthur. 2005. "Artificiality: The tension between internal and external validity in economic experiments." *Journal of Economic Methodology* 12(2):225–237.

Suen, Wing. 2004. "The self-perpetuation of biased beliefs." *Economic Journal* 114(495):377–396.

Sun, Junze, Arthur Schram and Randolph Sloof. 2019. "A Theory on Media Bias and Elections." Tinbergen Institute Discussion Paper 2019-048/I. Available at SSRN: https://ssrn.com/abstract=3420703 or http://dx.doi.org/10.2139/ssrn.3420703.

Vuong, Quang H. 1989. "Likelihood ratio tests for model selection and non-nested hypotheses." *Econometrica* 57(2):307–333.

Online Appendices for Elections under Biased Candidate Endorsements — An Experimental Study

Junze Sun[*] Arthur Schram[†] Randolph Sloof[‡]

October 1, 2020

Contents

[*]Department of Economics at the European University Institute (EUI), Florence. Email: Junze.Sun@eui.eu.

[†]CREED, Amsterdam School of Economics and Tinbergen Institute, and Department of Economics at EUI, Florence. Email: A.J.H.C.Schram@uva.nl.

[‡]Amsterdam School of Economics and Tinbergen Institute. Email: R.Sloof@uva.nl.

A Equilibrium Analyses for the Experimental Election Game

A.1 Bayesian Nash Equilibria (BNE)

In this appendix we derive the Bayesian Nash Equilibria (BNE) for our experimental election game. Let q_A and q_B be the probabilities that an ex-ante randomly sampled voter votes for candidates A and B, respectively. Given q_A and q_B, the pivotal probabilities of a vote for A and for B, under any fixed electoral size $\mathscr{N}$, are[1]

$$\begin{aligned} Pr[PivA|q_A,q_B] =& \frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i}q_A^i\cdot q_B^i\cdot q_O^{\mathscr{N}-2i} \\ &+\frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}-1}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i+1}q_A^i\cdot q_B^{i+1}\cdot q_O^{\mathscr{N}-2i-1} \end{aligned} \tag{A.1}$$

$$\begin{aligned} Pr[PivB|q_A,q_B] =& \frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i}q_A^i\cdot q_B^i\cdot q_O^{\mathscr{N}-2i} \\ &+\frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}-1}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i+1}q_A^{i+1}\cdot q_B^{i}\cdot q_O^{\mathscr{N}-2i-1} \end{aligned} \tag{A.2}$$

where $q_O \equiv 1-q_A-q_B$ is the probability of a randomly sampled voter abstaining, and $\lfloor x \rfloor$ denotes the largest integer that is not greater than x. Suppose k is directly observed, then voters with $v_i > (<) -k$ strictly prefer candidate A (B) and will cast a vote for candidate A (B) if they vote. Since voting is costly, a voter casts her vote only if the expected benefits of voting dominate the costs, i.e., $c_i \leq |k+v_i|\cdot Pr[PivA|q_A,q_B]$ ($c_i \leq |k+v_i|\cdot Pr[PivB|q_A,q_B]$) for $v_i > (<) -k$. To simplify computation, we approximate the probability distribution of individual voting costs c_i by a continuous uniform distribution on $[0,C]$, where $C = 15$. [2]

$$q_A(k) = \frac{1}{6}\sum_{v\in V, v>-k}\min\left\{\frac{|k+v|\cdot Pr[PivA|q_A,q_B]}{C},1\right\} \tag{A.3}$$

$$q_B(k) = \frac{1}{6}\sum_{v\in V, v<-k}\min\left\{\frac{|k+v|\cdot Pr[PivB|q_A,q_B]}{C},1\right\} \tag{A.4}$$

[1] The first term on the right hand side of (A.1) is the probability of tie, and the second term is the probability that candidate A falls exactly one vote behind candidate B, under multinomial distribution (q_A,q_B). In these two cases, a vote for candidate A is pivotal and improves A's winning chance by 0.5. The interpretation is analogous for (A.2).

[2] This continuous approximation is common in the literature. We have also derived equilibria for the case where (as in the experiment) only discrete costs are possible. Though the notation and analysis is more complex, the comparative statics predictions are the same as for the continuous distribution. For ease of presentation, we therefore use the approximation here. All predictions presented in the main text (for example, in Figure 2) are, however, based on the analysis with the discrete cost distribution. More details are available upon request.

where $V \equiv \{-100, -50, -20, 20, 50, 100\}$. The BNE can be obtained by solving the system of equations (A.3) and (A.4), with the pivotal probabilities therein replaced by (A.1) and (A.2). We numerically solve for the BNE under the above parameterization for each value of k considered in this paper. Panels (a) to (c) of Figure A.1 depict candidate A's winning probability ($\pi_A(k)$), expected vote share ($VS_A(k)$), and voter turnout $T(k)$ predicted by BNE as functions of k, respectively.

Figure A.1: Geometric Analyses for the Electoral Impacts of Endorser Bias

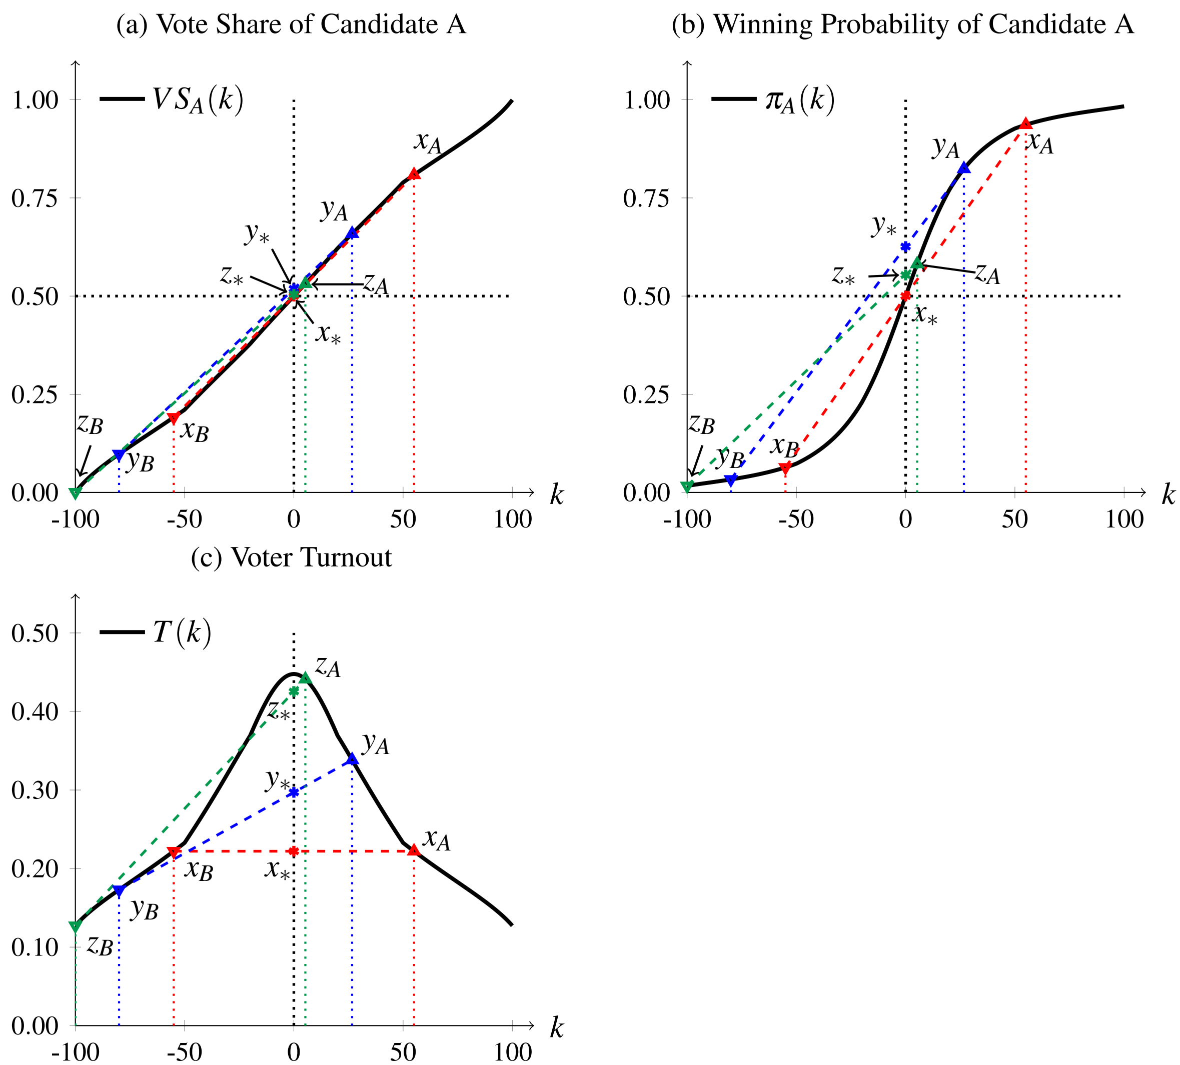

Note: Panels (a) to (c) depict candidate A's expected vote share $VS_A(k)$, winning probability $\pi_A(k)$, and the expected voter turnout $T(k)$ predicted by BNE as functions of k, respectively. Endorser's bias increases from $\chi = 0$ (red nodes) to $\chi = 55$ (blue nodes) to $\chi = 95$ (green nodes). In all three panels, w_A, w_B and w_*, $w \in \{x, y, z\}$, represent the corresponding equilibrium outcome conditional on message A, message B and unconditionally (ex-ante), respectively. The dashed line segments represent the sets of all convex combinations of w_A and w_B, for $w \in \{x, y, z\}$.

The interim and ex-ante impacts of endorser bias on candidate A's winning probability can be derived geometrically through Figure A.1b. We start with $\chi = 0$ (UB). From Table 1 in the main text we know that $k_A(0) = 55$ and $k_B(0) = -55$. As a result, candidate A's winning probability conditional on message A (B) equals $\pi_A(55)$ ($\pi_B(-55)$), represented by the red node x_A (x_B). Ex-ante, candidate A's winning probability is a convex combination of $\pi_A(55)$ and $\pi_B(-55)$, and therefore must lie on the red dashed line segment connecting x_A and x_B. By the law of iterated expectations, the posterior means must average back to the prior mean, which is 0. Geometrically, this implies that A's ex-ante winning probability can be represented by the red node x_*, the intersection of segment x_Ax_B and the vertical line $k = 0$.

A similar exercise can be done for $\chi = 55$ (WB_A), yielding interim winning probabilities conditional on message A and B represented by the blue nodes y_A and y_B, respectively, and the ex-ante winning probability represented by the blue node y_*. Likewise, for $\chi = 95$ (SB_A) the interim winning probabilities conditional on message A, B and the ex-ante winning probability are represented by the green nodes z_A, z_B and z_*, respectively. The interim impacts of an increase in endorser bias on candidate A's winning probability can thus be straightforwardly captured by the movement from x_s to y_s to z_s, for message $s \in \{A,B\}$. The ex-ante impacts of increasing endorser bias is represented by the movement from x_* to y_* to z_*.

As is evident from Figure A.1b, conditional on either message A or B, candidate A's expected winning probability decreases in endorser bias χ; $x_s > y_s > z_s$ for both $s \in \{A,B\}$.[3] Ex-ante, candidate A's winning probability varies non-monotonically with χ; $y_* > z_* > x_*$. Taken together, these results give rise to Hypothesis 1 formulated in the main text.

The same geometric approach is applied in Figure A.1a to derive the impact of endorser bias on candidate A's expected vote share. Similar to the analysis for the election outcome, we observe that conditional on either message A or B, candidate A's expected vote share decreases in bias χ; $x_s > y_s > z_s$ for both $s \in \{A,B\}$. From the ex-ante perspective, the impact of endorser bias on A's expected vote share is almost negligible; x_*, y_* and z_* are all hovering around 0.5 and are similar in magnitude. As argued in the main text, this is because $VS_A(k)$ is almost a linear function of k and equals 0.5 if $k = 0$. Consequently, variations in A's expected vote share are almost proportional to variations in voters' posterior expectations about k. The latter are invariant from the ex-ante perspective thanks to the law of iterated expectation. For this reason, the ex-ante impact of increasing endorser bias on A's expected vote share is almost negligible.

Figure A.1c presents the impact of endorser bias on voter turnout. Conditional on message A, voter turnout increases in endorser bias ($z_A > y_A > x_A$) while conditional on message B the reverse is true ($x_B > y_B > z_B$). From an ex-ante perspective, we observe that the expected voter turnout increases unambiguously in endorser bias ($z_* > y_* > x_*$). Finally, voter turnout attains its maximum if voters are uninformed; in this case, voters behave as if $k = 0$ is common knowledge. As is evident from Figure A.1c, $T(k)$ is highest when $k = 0$. These results give rise to Hypothesis 2.

[3] It is implicitly understood that whenever we compare nodes in Figures A.1b to A.1c, we are comparing their values on the vertical axis.

A.2 Quantal Response Equilibria (QRE)

In this appendix, we derive the standard quantal response equilibria (QRE; McKelvey and Palfrey (1995)) for our election game. Denote any voter i's expected utility from voting for candidate A, B or abstaining, conditional on diamond allocation k and private type (v_i, c_i), by $EU^A(k, v_i, c_i)$, $EU^B(k, v_i, c_i)$ and $EU^O(k, v_i, c_i)$, respectively. Without loss of generality, we normalize $EU^O(k, v_i, c_i)$ to zero for all i's. With this normalization, we can interpret $EU^\omega(k, v_i, c_i)$ as the difference between the expected utilities from voting for ω and abstaining, where $\omega \in \{A, B\}$. Following the cost-benefit reasoning in Appendix A.1, we have[4]

$$EU^A(k, v_i, c_i) = \frac{(k + v_i) \cdot PivA(k) - c_i}{100} \tag{A.5}$$

$$EU^B(k, v_i, c_i) = \frac{-(k + v_i) \cdot PivB(k) - c_i}{100} \tag{A.6}$$

where $PivA(k)$ and $PivB(k)$ are pivotal probabilities of a vote for candidate A and B, respectively, when k is common knowledge. In contrast to BNE, QRE allows voters to make errors in their voting decisions. This can be modeled by adding a stochastic error term to the expected utilities, which yields:

$$\begin{aligned} &EU^A(k, v_i, c_i) + \frac{\varepsilon_A}{\lambda} \\ &EU^B(k, v_i, c_i) + \frac{\varepsilon_B}{\lambda} \\ &EU^O(k, v_i, c_i) + \frac{\varepsilon_O}{\lambda} = \frac{\varepsilon_O}{\lambda} \end{aligned}$$

where ε_d for $d \in \{A, B, O\}$ denotes i.i.d. noise and $\lambda > 0$ is a parameter that captures the relative weight of noise in the perceived expected utility. In a QRE (for the normal form game), a voter will still choose the action that yields the highest expected utility, but this is now a stochastic event. For example, she will vote for A if $EU^A(k, v_i, c_i) + \frac{\varepsilon_A}{\lambda} > EU^B(k, v_i, c_i) + \frac{\varepsilon_B}{\lambda}$ and $EU^A(k, v_i, c_i) + \frac{\varepsilon_A}{\lambda} > \frac{\varepsilon_O}{\lambda}$ hold simultaneously. Or equivalently, $\varepsilon_B - \varepsilon_A < \lambda\left(EU^A(k, v_i, c_i) - EU^B(k, v_i, c_i)\right)$ and $\varepsilon_O - \varepsilon_A < \lambda EU^A(k, v, c)$. Specification of the distribution functions of ε_A, ε_B, ε_O then yields the probabilities that the voter will vote for A or B, or will abstain. We derive the logit QRE, where ε_d is assumed to follow the extreme value type I distribution for all $d \in \{A, B, O\}$. In this case, the probabilities of choosing each action d, denoted by p^d, are given by:

$$p^A(k, v_i, c_i; \lambda) = \frac{e^{\lambda EU^A(k, v_i, c_i)}}{e^{\lambda EU^A(k, v_i, c_i)} + e^{\lambda EU^B(k, v_i, c_i)} + 1} \tag{A.7}$$

[4] We divide the obtained difference in expected payoffs by 100 to make sure that k, v_i and c_i are all re-scaled to a range between 0 and 1. This is for ease of computation.

$$p^B(k,v_i,c_i;\lambda) = \frac{e^{\lambda EU^B(k,v_i,c_i)}}{e^{\lambda EU^A(k,v_i,c_i)} + e^{\lambda EU^B(k,v_i,c_i)} + 1} \tag{A.8}$$

$$p^O(k,v_i,c_i;\lambda) = \frac{1}{e^{\lambda EU^A(k,v_i,c_i)} + e^{\lambda EU^B(k,v_i,c_i)} + 1} \tag{A.9}$$

Using (A.7) to (A.9), we can compute the probabilities of a randomly sampled voter voting for A and B, $q_A(k|\lambda)$ and $q_B(k|\lambda)$, by

$$q_A(k;\lambda) \equiv E_{v,c}[p^A(k,v,c;\lambda)] = \int_v \int_c p^A(k,v,c;\lambda) d\gamma(c) dG(v)$$
$$q_B(k;\lambda) \equiv E_{v,c}[p^B(k,v,c;\lambda)] = \int_v \int_c p^B(k,v,c;\lambda) d\gamma(c) dG(v)$$

where $G(v)$ and $\gamma(c)$ are the CDFs of private ideology v_i and voting costs c_i, respectively. In our experimental election game, both v_i and c_i are independently drawn for each voter i. v_i follows a discrete uniform distribution $G(\cdot)$ over $\{-100,-50,-20,20,50,100\}$ (denote each element by v^r, $r = 1,2,\cdots,6$), and c_i follows a discrete uniform distribution $\gamma(\cdot)$ on $\{1,2,\cdots,C\}$ (denote each element by c^s, $s = 1,2,\cdots,C$). Parameter C is an integer and it specifies the upper bound of voting costs. Let $\mathbb{P}^d(k;\lambda)$ be a $6\times C$ matrix with the (r,s) element $\mathbb{P}^d_{r,s}(k;\lambda) = p^d(k,v^r,c^s;\lambda)$, for each $d \in \{A,B,O\}$. We can then simplify the above expressions to the following formulas ($\vec{e}_{1\times 6}$ and $\vec{e}_{C\times 1}$ are vectors of ones):

$$q_A(k;\lambda) = \frac{1}{6} \times \frac{1}{C} \cdot \vec{e}_{1\times 6} \cdot \mathbb{P}^A(k;\lambda) \cdot \vec{e}_{C\times 1} \tag{A.10}$$

$$q_B(k;\lambda) = \frac{1}{6} \times \frac{1}{C} \cdot \vec{e}_{1\times 6} \cdot \mathbb{P}^B(k;\lambda) \cdot \vec{e}_{C\times 1} \tag{A.11}$$

The pivotal probabilities, $PivA(k)$ and $PivB(k)$, depend on both $q_A(k|\lambda)$ and $q_B(k|\lambda)$, as well as the electorate size $\mathcal{N}$. These pivotal probabilities can be calculated by the formulas below (based on the multinomial distribution):[5]

$$\begin{aligned} PivA(k) =& \frac{1}{2}\sum_{i=0}^{\lfloor \frac{\mathcal{N}}{2} \rfloor} \binom{\mathcal{N}}{i}\binom{\mathcal{N}-i}{i} q_A(k;\lambda)^i q_B(k;\lambda)^i q_O(k;\lambda)^{\mathcal{N}-1-2i} \\ &+ \frac{1}{2}\sum_{i=0}^{\lfloor \frac{\mathcal{N}}{2} \rfloor} \binom{\mathcal{N}}{i}\binom{\mathcal{N}-i}{i+1} q_A(k;\lambda)^i q_B(k;\lambda)^{i+1} q_O(k;\lambda)^{\mathcal{N}-1-2i} \end{aligned} \tag{A.12}$$

[5] (A.12) and (A.13) are obtained from (A.1) and (A.2) by replacing q_d with $q_d(k|\lambda)$ for all $d \in \{A,B,O\}$.

$$
\begin{aligned}
PivB(k) =&\frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i}q_A(k;\lambda)^i q_B(k;\lambda)^i q_O(k;\lambda)^{\mathscr{N}-1-2i} \\
&+\frac{1}{2}\sum_{i=0}^{\lfloor\frac{\mathscr{N}}{2}\rfloor}\binom{\mathscr{N}}{i}\binom{\mathscr{N}-i}{i+1}q_A(k;\lambda)^{i+1} q_B(k;\lambda)^i q_O(k;\lambda)^{\mathscr{N}-1-2i}
\end{aligned}
\tag{A.13}
$$

Equations (A.5) to (A.13) formulate a fixed-point problem and we can solve for the QRE predictions numerically for any pair of k and λ, given model primitives $G(\cdot)$, $\gamma(\cdot)$ and $\mathscr{N}$. Some model primitives are different between our experiments conducted in 2016 and 2017. Specifically, for our 2016 experiment, the upper bound on voting costs C equals 20 and the electorate size $\mathscr{N}$ equals 11. For our 2017 experiment, $C = 15$ and $\mathscr{N} = 25$. In addition, instead of using the strategy method as in our 2017 experiment, in all treatments with a single endorser we released the public messages to voters before they made voting decisions in our 2016 experiment. As a result, for each round of election, we only observed voters' decisions conditional the realized message.

We use the maximum likelihood (ML) method explained in Appendix C to estimate the logit response parameter λ based on the experiment conducted in 2016. The estimated $\widehat{\lambda}$ equals 19.45. We then use this out-of-sample estimate to generate QRE predictions for our experimental sessions in 2017 (cf. Figure 1 in the main text). Similar to the derivation of BNE, these QRE predictions are obtained by replacing k with the rational posterior expectations of k conditional on the public information voters have.

B Additional Experimental Results

In this appendix we present two sets of additional results regarding our experiment.

B.1 The Explanatory Power of BNE and QRE

In this appendix we evaluate the explanatory power of BNE and QRE predictions. To do so, we regress the observed outcomes on their corresponding predictions:

$$Y_{is} = \beta_0 + \beta_1 Pred_{is} + X_s \gamma + Pred_{is} \times X_s \delta + \varepsilon_{is} \tag{B.1}$$

In (B.1), Y_{is} and $Pred_{is}$ are the observed and predicted aggregate outcome (A's vote share, A's winning probability, or voter turnout) of election i in electorate s. X_s is a vector of standardized electoral background variables (average age, fraction of males, and fraction of students major in economics or business) for electorate s. γ and δ are vectors of coefficients for these electoral backgrounds. The error terms ε_{is} are assumed to be independent across electorates (standard errors are clustered at the electorate level). Note that the interaction terms in (B.1) allow for the impacts of theoretical predictions on observed outcomes to be different across distinct compositions of the electorate. We estimate (B.1) for both BNE and QRE predictions. These estimation results are summarized in Table B.1.

Columns (1) and (5) show that BNE predictions alone can explain 72.3% and 58.6% of the variations in candidate A's observed vote shares and winning probabilities, respectively. Columns (9), however, shows that only 6.5% of the variations in observed voter turnout can be explained by BNE predictions. Compared to BNE, QRE explains higher fractions of variations in the observed election outcomes and vote share (cf. columns (3) and (7)), but it explains approximately the same fraction of observed variations in voter turnout (cf. column (11)). Controlling for electoral background variables does not substantially increase the explanatory power of either BNE or QRE predictions for observed vote shares and election outcomes, but does so for voter turnout by 6.2%-points. Therefore, we conclude that BNE and QRE have strong explanatory power for party vote shares and election outcomes, but much less explanatory power for voter turnout.

Finally, we check for learning effects by conducting the same regression (B.1) separately for sub-samples in blocks 1-4, 5-8 and 9-12. We find little evidence of learning for party vote shares and the election outcomes. The explanatory power of BNE (QRE) for voter turnout, however, increases from 3.9% (3.2%) in blocks 1-4 to 10.8% (13.4%) in blocks 9-12. Though these results indicate some learning effects towards the equilibrium predictions, the amount of variation in observed turnout that can be explained by BNE predictions still remains low.

Table B.1: The Explanatory Power of BNE and QRE Predictions

Dep. Variable	Vote Share of A				Winning Probability of A				Voter Turnout			
	BNE		*QRE*		*BNE*		*QRE*		*BNE*		*QRE*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Pred	0.713***	0.713***	1.245***	1.245***	0.826***	0.826***	1.138***	1.138***	0.200***	0.200***	0.458***	0.458***
	(0.037)	(0.007)	(0.067)	(0.009)	(0.012)	(0.008)	(0.019)	(0.008)	(0.022)	(0.020)	(0.059)	(0.045)
Age		0.054***		0.098***		-0.004		0.006***		0.034		0.029
		(0.009)		(0.011)		(0.004)		(0.001)		(0.020)		(0.038)
Males		−0.031***		−0.054***		−0.001		−0.002***		-0.008		0.033
		(0.004)		(0.005)		(0.002)		(0.000)		(0.009)		(0.017)
Econ		0.058***		0.109***		−0.006		0.001		0.023		0.034
		(0.012)		(0.014)		(0.004)		(0.001)		(0.024)		(0.045)
$Pred \times Age$		−0.115***		−0.203***		−0.029***		−0.051***		0.020		0.027
		(0.010)		(0.015)		(0.005)		(0.007)		(0.031)		(0.068)
$Pred \times Males$		0.057***		0.104***		0.003		0.003		0.010		0.070*
		(0.005)		(0.007)		(0.003)		(0.003)		(0.014)		(0.030)
$Pred \times Econ$		−0.126***		−0.230***		−0.031**		−0.050***		0.005		-0.024
		(0.011)		(0.016)		(0.005)		(0.011)		(0.034)		(0.077)
Constant	0.134***	0.134***	−0.130***	−0.130***	0.064***	0.064***	−0.055***	−0.055***	0.466***	0.466***	0.339***	0.339***
	(0.019)	(0.007)	(0.033)	(0.009)	(0.004)	(0.003)	(0.002)	(0.001)	(0.017)	(0.015)	(0.030)	(0.027)
#.Obs.	996	996	1,008	1,008	1,008	1,008	1,008	1,008	1,008	1,008	1,008	1,008
R^2	0.742	0.753	0.849	0.861	0.649	0.650	0.777	0.779	0.069	0.131	0.068	0.132

Note: Electoral backgrounds are standardized within the sample. Standard errors are clustered at the session level and reported in parentheses. Statistical significance: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

B.2 More Results on Learning Effects

In this appendix we present further results about learning effects in our experiment. To do so, we compare observed election outcomes and voter turnout across rounds separately for the first, second, and third phase of the experiment. Each phase consists of four blocks (covering the four treatments of the experiment) and each block consists of eight elections under the same treatment condition (cf. Figure 3 in the main text). Using this split into three phases, we performed three types of analyses.

First, we conducted the tests of our hypotheses using data from the last phase only (cf. the last paragraph of Section 6.2). Overall, the analyses show that, rather than voting mindlessly, subjects seem to gain experience and behave more 'rationally' in later blocks; the aggregate turnout rates and election outcomes are closer to the BNE predictions. Learning appears to mitigate the distortion of the election outcome caused by a strongly biased endorsement, driving the observed election outcomes closer to our predicted non-monotonic relationship. Moreover, as explained in the final paragraph of Section 6.2, most conclusions regarding hypotheses testing remain unchanged in the subsample consisting of data from the last four blocks. There are only two exceptions, both in the direction of confirming the comparative statics predicted by our theory.

Second, we considered phase-by-phase the treatment comparisons of winning probability and turnout. The results are presented in Figure B.1. The top panels show how the (interim and ex-ante) observed winning probabilities of candidate A vary over the treatments – and how these compare to BNE predictions – in the different blocks. Consistent with results based on the full sample (cf. Figure 4 in the main text), these charts show that in all stages the election outcome is reasonably well predicted for treatments UB and WB_A.

Substantial deviations from the theoretical (BNE) predictions appear mostly under treatment SB_A and conditional on message A being sent (compare the solid and dashed red lines). As explained in Section 6.1 when discussing Figure 4, the monotonic relationship between A's observed ex-ante winning probability and the endorser's bias is mostly driven by these deviations. As is evident graphically, such deviations are stronger in earlier blocks than later blocks. In particular, the top-right chart suggests that the strictly monotonic relationship between A's ex-ante winning probability and the endorser bias disappears if we focus on data from the last four blocks. This is because in these blocks the observed winning probability under SB_A and message A is much closer to the BNE prediction than before. As a consequence, A's ex-ante winning probabilities are almost identical under WB_A and SB_A (both equal to 0.64; the difference is not statistically different). This suggests again that, if anything, learning tends to push behavior in the direction of our predicted non-monotonicity hypothesis.

The bottom panels of Figure B.1 depict how the (interim and ex-ante) observed turnout rates vary with treatment conditions in different phases. Here we observe partial learning effects. On the

Figure B.1: Aggregate Voting Behavior and Election Outcomes

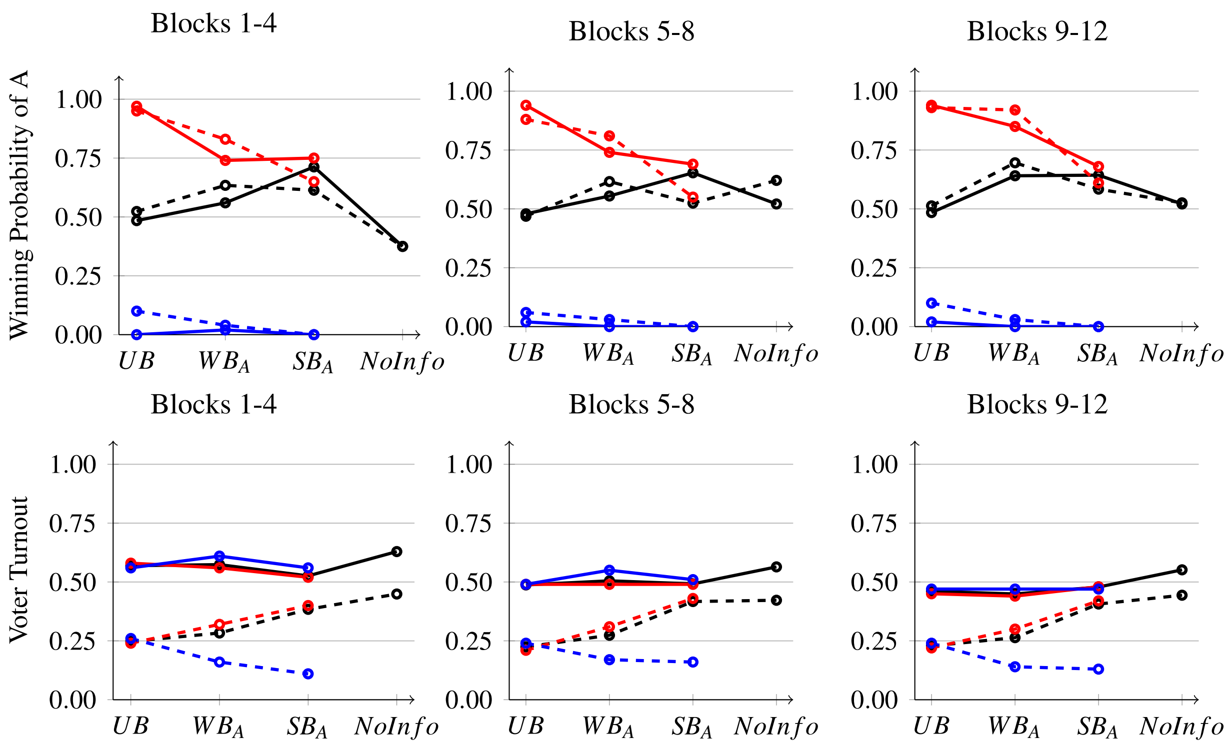

Note: In all the six panels, the dashed lines depict the electoral outcomes predicted by Bayesian Nash Equilibria (BNE) and the solid lines depict the actual electoral outcomes observed in the experiment. The red (blue) lines denote the interim outcomes conditional on message A (B), while the black lines denotes the ex-ante outcome; the average of interim outcomes weighted by the likelihood of sending messages A and B. We omit QRE predictions in these charts because they are qualitatively similar to BNE predictions.

one hand, turnout substantially drops in later blocks (suggesting that over-voting is mitigated by learning). On the other hand, however, observed (both interim and ex-ante) turnout rates remain largely irresponsive to endorser bias even in the last four blocks. This suggests a persistent deviation that is not easily mitigated by learning. This is one of the results that motivated us to explore alternative behavioral channels explained in Section 7 of the main text.

In our third analysis of the three phases we examined subjects' payoffs (as measured by their expected earnings) in the election. This allows us to determine whether learning helps subjects to better respond to candidate endorsements. Consistent with a positive learning effect, we observe that, for all the four treatments, subjects' average earnings are higher in the last four blocks than across all blocks. This confirms the learning effects we observed above.

C The Hybrid Behavioral Equilibrium Model

In this appendix we formally introduce the hybrid behavioral equilibrium model outlined in the main text and then explain the estimation methods.

C.1 The Formal Model Setup

Let τ denote an *information environment*, which summarizes the public information voters obtain in each treatment. Let Γ be the set of all information environments studied in our experiment:

$$\Gamma = \Big\{ NoInfo, (UB,A), (UB,B), (WB_A,A), (WB_A,B), (SB_A,A), (SB_A,B) \Big\} \tag{C.1}$$

Our hybrid equilibrium model extends the standard logit QRE model (presented in the previous appendix) in three ways. First of all, for each information environment τ, it allows voters' common posterior expectation about the state, denoted by $\widehat{k_\tau}$, to deviate from the corresponding Bayesian posterior k_τ. As explained in the main text, we assume that there exists some $\beta \geq 0$ such that

$$\widehat{k_\tau} = \beta \cdot k_\tau \tag{C.2}$$

Second, for each information environment τ, it allows voters' perceived pivotal probabilities, $\widehat{PivA}_\tau$ and $\widehat{PivB}_\tau$, to deviate from their actual values. More specifically, we assume that there exists a $\rho \in [0,1]$ such that

$$\widehat{PivA}_\tau = \rho \cdot PivA_\tau + (1-\rho) \cdot \overline{PivA} \tag{C.3}$$

$$\widehat{PivB}_\tau = \rho \cdot PivB_\tau + (1-\rho) \cdot \overline{PivB} \tag{C.4}$$

In (C.3) and (C.4), $Piv\Omega_\tau$ is the actual pivotal probability of a vote for candidate $\Omega \in \{A,B\}$ in equilibrium, which is derived below. $\overline{Piv\Omega} = \sum_{\tau\in\Gamma} f_\tau \cdot Piv\Omega_\tau$ for $\Omega \in \{A,B\}$, where f_τ denotes the expected frequency of encountering each information environment τ.[6] If $\rho = 0$, voters mistakenly believe that the probabilities of casting pivotal votes are constant across all information environments, but they correctly estimate these pivotal probabilities on average. If $\rho = 1$, voters precisely infer pivotal probabilities in each information environment, as in standard BNE and QRE. With $\rho \in (0,1)$, voters partly realize the relationship between pivotal probabilities and the information environment, but insufficiently so compared to a rational agent.

Notice that both non-Bayesian updating and partial competition neglect can affect a voter's

[6] In our experiment, these frequencies are $f_{NoInfo} = \frac{1}{4}$, $f_{(UB_A,A)} = f_{(UB_A,B)} = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8}$, $f_{(WB_A,A)} = \frac{1}{4} \cdot \frac{3}{4} = \frac{3}{16}$, $f_{(WB_A,B)} = \frac{1}{4} \cdot \frac{1}{4} = \frac{1}{16}$, $f_{(SB_A,A)} = \frac{1}{4} \cdot \frac{19}{20} = \frac{19}{80}$ and $f_{(SB_A,B)} = \frac{1}{4} \cdot \frac{1}{80} = \frac{1}{320}$.

judgement about the expected utility from voting. Contrary to (A.5) and (A.6), in our hybrid equilibrium model any voter i's expected utilities from voting for A and B are

$$\widehat{EU^A}(\tau, v_i, c_i) = \frac{(\widehat{k_\tau} + v_i) \cdot \widehat{PivA}_\tau - c_i}{100} \tag{C.5}$$

$$\widehat{EU^B}(\tau, v_i, c_i) = \frac{-(\widehat{k_\tau} + v_i) \cdot \widehat{PivB}_\tau - c_i}{100} \tag{C.6}$$

When $\beta = \rho = 1$, $\widehat{EU^A}(\tau, v_i, c_i)$ and $\widehat{EU^B}(\tau, v_i, c_i)$ are equivalent to, respectively, (A.5) and (A.6) formulated in Appendix A.2.

Finally, our hybrid model allows voters to display distinct levels of rationality for candidate choice and turnout decisions. We model this by an extensive form version of our election game, where voters are assumed to adopt a two-stage decision-making process. At the first stage (S1), voter i decides whether to vote or to abstain. If she decides to vote in S1, then the game moves to the second stage (S2), where the voter has to decide which candidate to vote for. We allow voters to possess distinct noise parameters for decisions at different stages. Let λ_t and λ_p denote the noise parameters for stage S1 (turnout decision) and S2 (candidate choice), respectively. In the spirit of McKelvey and Palfrey (1998), we derive the logit *agent quantal response equilibrium* (AQRE) for this extensive version of our election game.[7] We do so by backward induction. In S2 (i.e., conditional on turnout), voter i can only choose between candidate A and B. In the logit AQRE, these choice probabilities are given by

$$p^A(\tau, v_i, c_i; \lambda_p) = \frac{e^{\lambda_p EU^A(\tau, v_i, c_i)}}{e^{\lambda_p EU^A(\tau, v_i, c_i)} + e^{\lambda_p EU^B(\tau, v_i, c_i)}} \tag{C.7}$$

$$p^B(\tau, v_i, c_i; \lambda_p) = \frac{e^{\lambda_p EU^B(\tau, v_i, c_i)}}{e^{\lambda_p EU^A(\tau, v_i, c_i)} + e^{\lambda_p EU^A(\tau, v_i, c_i)}} \tag{C.8}$$

Voter i's expected utility from voting is then

$$EU^{Vote}(\tau, v_i, c_i; \lambda_p) = p^A(\tau, v_i, c_i; \lambda_p) \cdot EU^A(\tau, v_i, c_i) + p^B(\tau, v_i, c_i; \lambda_p) \cdot EU^A(\tau, v_i, c_i) \tag{C.9}$$

[7] We also estimated another type of extensive form AQRE model where voters at the first stage S1 decide which candidate to support, and at the second stage S2 decide whether to abstain or to vote for the candidate they choose to support in S1. Using Vuong's closeness tests (Vuong, 1989), we find that this model is outperformed by the current AQRE model presented in this appendix.

In S1, the voter's probabilities to vote (p^T) and to abstain (p^O) are

$$p^T(\tau, v_i, c_i; \lambda_p, \lambda_t) = \frac{e^{\lambda_t EU^{Vote}(\tau, v_i, c_i; \lambda_p)}}{e^{\lambda_t EU^{Vote}(\tau, v_i, c_i; \lambda_p)} + 1} \tag{C.10}$$

$$p^O(\tau, v_i, c_i; \lambda_p, \lambda_t) = \frac{1}{e^{\lambda_t EU^{Vote}(\tau, v_i, c_i; \lambda_p)} + 1} \tag{C.11}$$

Therefore, *a priori* the probabilities that voter i votes for candidate A and B are

$$p^A(\tau, v_i, c_i; \lambda_p, \lambda_t) = p^T(\tau, v_i, c_i; \lambda_p, \lambda_t) \cdot p^A(\tau, v_i, c_i; \lambda_p) \tag{C.12}$$

$$p^B(\tau, v_i, c_i; \lambda_p, \lambda_t) = p^T(\tau, v_i, c_i; \lambda_p, \lambda_t) \cdot p^B(\tau, v_i, c_i; \lambda_p) \tag{C.13}$$

Akin to the derivation of QRE in Appendix A.2, let $\mathbb{P}^d(\tau; \lambda)$ be a $6 \times C$ matrix with its (r, s) element $\mathbb{P}^d_{r,s}(\tau; \lambda) = p^d(\widehat{k_\tau}, v^r, c^s; \lambda)$, for all $d \in \{A, B, O\}$.[8] Following the same logic under (A.10) to (A.11) in Appendix A.2, we can derive for this two-stage AQRE model the probabilities of a randomly sampled voter voting for A and B, respectively, in information environment τ by

$$q^A(\tau; \lambda_p, \lambda_t) = \frac{1}{6} \times \frac{1}{C} \cdot \vec{e}_{1\times 6} \cdot \mathbb{P}^A(\tau; \lambda_p, \lambda_t) \cdot \vec{e}_{C\times 1} \tag{C.14}$$

$$q^B(\tau; \lambda_p, \lambda_t) = \frac{1}{6} \times \frac{1}{C} \cdot \vec{e}_{1\times 6} \cdot \mathbb{P}^B(\tau; \lambda_p, \lambda_t) \cdot \vec{e}_{C\times 1} \tag{C.15}$$

Finally, the pivotal probabilities $PivA_\tau$ and $PivB_\tau$ can be obtained from equations (A.1) and (A.2) by replacing q_d with $q^d(\tau; \lambda_p, \lambda_t)$ for $d \in \{A, B, O\}$:

$$\begin{aligned} PivA_\tau =& \frac{1}{2} \sum_{i=0}^{\lfloor \frac{\mathscr{N}}{2} \rfloor} \binom{\mathscr{N}}{i} \binom{\mathscr{N}-i}{i} q^A(\tau; \lambda_p, \lambda_t)^i q^B(\tau; \lambda_p, \lambda_t)^i q^O(\tau; \lambda_p, \lambda_t)^{\mathscr{N}-1-2i} \\ &+ \frac{1}{2} \sum_{i=0}^{\lfloor \frac{\mathscr{N}}{2} \rfloor} \binom{\mathscr{N}}{i} \binom{\mathscr{N}-i}{i+1} q^A(\tau; \lambda_p, \lambda_t)^i q^B(\tau; \lambda_p, \lambda_t)^{i+1} q^O(\tau; \lambda_p, \lambda_t)^{\mathscr{N}-1-2i} \end{aligned} \tag{C.16}$$

$$\begin{aligned} PivB_\tau =& \frac{1}{2} \sum_{i=0}^{\lfloor \frac{\mathscr{N}}{2} \rfloor} \binom{\mathscr{N}}{i} \binom{\mathscr{N}-i}{i} q^A(\tau; \lambda_p, \lambda_t)^i q^B(\tau; \lambda_p, \lambda_t)^i q^O(\tau; \lambda_p, \lambda_t)^{\mathscr{N}-1-2i} \\ &+ \frac{1}{2} \sum_{i=0}^{\lfloor \frac{\mathscr{N}}{2} \rfloor} \binom{\mathscr{N}}{i} \binom{\mathscr{N}-i}{i+1} q^A(\tau; \lambda_p, \lambda_t)^{i+1} q^B(\tau; \lambda_p, \lambda_t)^i q^O(\tau; \lambda_p, \lambda_t)^{\mathscr{N}-1-2i} \end{aligned} \tag{C.17}$$

Equations (C.1) to (C.17) formulate a fixed-point problem that we can numerically solve for any given set of model parameters $\{\lambda_t, \lambda_p, \beta, \rho\}$.

[8] Recall from Appendix A.2 that v^r denotes the r'th element of set $\{-100, -50, -20, 20, 50, 100\}$ and c^s denotes the s'th element of set $\{1, 2, \cdots, 15\}$.

C.2 Estimation Methods

Each of the models above predicts a set of decision probabilities $p^d(k,v,c|\Lambda)$ for each decision $d \in \{A,B,O\}$, conditional on (k,v,c) and the set of parameter(s) Λ. Λ equals $\{\lambda_p, \lambda_t, \beta, \rho\}$ for models concerning the distinct levels of rationality, while it equals $\{\lambda, \beta, \rho\}$ for models excluding this mechanism. Moreover, models excluding non-Bayesian updating and competition neglect assume $\beta = 1$ and $\rho = 1$, respectively. We use the maximum likelihood (ML) method to estimate the model parameters. The log likelihood function aggregates voting decisions in each election by each individual, and it is given by

$$\mathscr{L}_N(\Lambda) = \sum_{i=1}^{N}\sum_{t=1}^{T} \ln p^{d_{it}}(k_t, v_{it}, c_{it}|\Lambda) \tag{C.18}$$

where i is the index for individual subjects and t is the index for elections.

Let $\Theta \equiv \{(k,v,c)|\exists i,t \text{ such that } (k,v,c) = (k_t, v_{it}, c_{it})\}$ be the set of all combinations of (k,v,c) that appeared in the experiment and let N_θ be the number of observations for any $\theta \in \Theta$. Moreover, let f_θ^d denote the observed frequency of decision $d \in \{A,B,O\}$ for each $\theta \in \Theta$. We can then rewrite (C.18) in a more compact way as

$$\mathscr{L}_N(\Lambda) = \sum_{\theta\in\Theta}\left(N_\theta \cdot \sum_{d\in\{A,B,O\}} f_\theta^d \cdot \ln p^d(\theta|\Lambda)\right) \tag{C.19}$$

The maximum likelihood estimator (MLE) is then given by the maximizer of function (C.19). We denote the obtained MLE by $\widehat{\Lambda}$. The maximal log likelihood then equals $\mathscr{L}_N(\widehat{\Lambda})$. We construct a "fitting score" to measure the performance of the hybrid models in a range between 0% to 100% using the approach from Goeree, Louis and Zhang (2018). Specifically, we compare $\mathscr{L}_N(\widehat{\Lambda})$ to an upper bound $\overline{\mathscr{L}_N}$ and a lower bound $\underline{\mathscr{L}_N}$, which are constructed as follows:

$$\overline{\mathscr{L}_N} \equiv \sum_{\theta\in\Theta}\left(N_\theta \cdot \sum_{d\in\{A,B,O\}} f_\theta^d \cdot \ln f_\theta^d\right)$$
$$\underline{\mathscr{L}_N} \equiv \sum_{\theta\in\Theta}\left(N_\theta \cdot \sum_{d\in\{A,B,O\}} f_\theta^d \cdot \ln \frac{1}{3}\right) = -N\ln 3$$

The upper bound $\overline{\mathscr{L}_N}$ is obtained by setting the predicted choice probabilities exactly equal to the observed frequencies, which yields a global maximum of the log likelihood function. Hence, $\mathscr{L}_N(\widehat{\Lambda}) \leq \overline{\mathscr{L}_N}$ holds generically. The lower bound $\underline{\mathscr{L}_N}$ is obtained by assuming that voters just decide in a uniformly random manner, making each choice with probability one third regardless of (k,v,c). In principle, $\mathscr{L}_N(\widehat{\Lambda})$ could be lower than $\underline{\mathscr{L}_N}$; this would suggest that $\widehat{\Lambda}$ performs even worse than a purely random choice model.[9] The fitting score is then constructed as

$$R_{MLE}(\widehat{\Lambda}) = \frac{\mathscr{L}_N(\widehat{\Lambda}) - \underline{\mathscr{L}_N}}{\overline{\mathscr{L}_N} - \underline{\mathscr{L}_N}} \tag{C.20}$$

$R_{MLE}(\widehat{\Lambda})$ is a linear transformation of $\mathscr{L}_N(\widehat{\Lambda})$ to a zero-one scale, provided that $\mathscr{L}_N(\widehat{\Lambda}) \geq \underline{\mathscr{L}_N}$. The estimation results for our experiment conducted in 2017 is summarized in Table 2 of the main text.

[9] In the standard QRE model presented in Appendix A.2, $\mathscr{L}_N(\widehat{\Lambda}) < \underline{\mathscr{L}_N}$ is impossible because the random choice model is nested by setting $\lambda = 0$. In the two-stage AQRE model, however, a random choice model might perform better because it is not nested; no pairs of (λ_t, λ_p) can generate uniformly random choices. To see this, note that even by setting $\lambda_t = \lambda_p = 0$, a voter would abstain with probability $\frac{1}{2}$ and vote for each candidate with probability $\frac{1}{4}$, which is different from a uniform distribution.

D Instructions and Sample Screen Shots

[In this section we reproduce the instructions for subjects in sessions with $|M| = 1$. While reading instructions, subjects were free to refer to previous pages. Below we provide the exact texts for the main instructions, instructions for the belief elicitation tasks and post-experiment questionnaire, respectively.][10]

[Instruction, control questions and a short summary]

Welcome to the experiment! [Screen 0]

Thank you for participating in this experiment on group decision making. In this experiment, you can earn money. The amount of money you earn depends on the decisions you and the other participants make and on random events.

Before the experiment starts you will receive detailed instructions. These instructions are simple, and if you follow them carefully you may earn a considerable amount of money. The money will be paid to you in cash at the end of the experiment. We ensure that your final earnings remain confidential: we will not inform other participants of your final earnings. All of your decisions will be recorded anonymously. Nobody will be able to link any specific decisions to your name.

In today's experiment you can earn "tokens". The conversion rate is such that 50 tokens = 1€, so for each token you receive 2 eurocents. On top of this you will receive a participation fee of 10€.

During the experiment, please do not communicate with other participants. If you have any questions, please raise your hand and the experimenter will come to your table to answer your questions in private.

[10] Comments within brackets "[]" were not included in the instructions.

Elections, Groups, and Vases [Screen 1]

This experiment consists of a series of elections. In each election, there are 25 voters, including you and the other 24 participants. There are two candidates in each election. These are represented by two vases, **A** and **B** (see Figure 1). In each election, all voters will be randomly assigned to either team **A** or team **B**. Each voter is equally likely to be assigned to either team. Assignment to a team is important, because you will receive a bonus if your team wins the election. This is explained next.

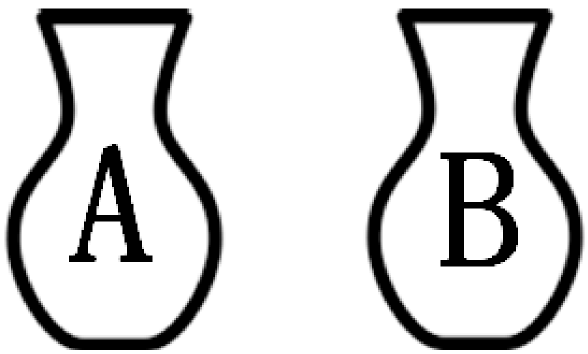

Figure 1: Two candidate vases, **A** and **B**

Bonus [Screen 2]

A voter will gain a bonus of X tokens if the vase belonging to his/her team wins the election. Thus, if you are in team **A**, you will gain a bonus of X tokens if vase **A** is elected, but no bonus if vase **B** is elected. Similarly, if you are in team **B**, you will gain a bonus of X tokens if vase B is elected, and no bonus if vase A is elected.

The size of the bonus X will be independently drawn for each voter. It equals either 20, 50, or 100 tokens. Each of these values is equally likely. Which team you are in and the size of your bonus is determined independently from other voters, and will differ from one election to another. In each election, you will be told your team and the size of the bonus X before you make your voting decision.

Additional earnings: Diamonds [Screen 3]

Aside from making money from the bonus, you can also earn money if the vase that was chosen contains so-called "diamonds". This works as follows.

At the beginning of each election, the computer will randomly select (with equal probability) one of the two vases and fill it with k "diamonds" (see Figure 2).

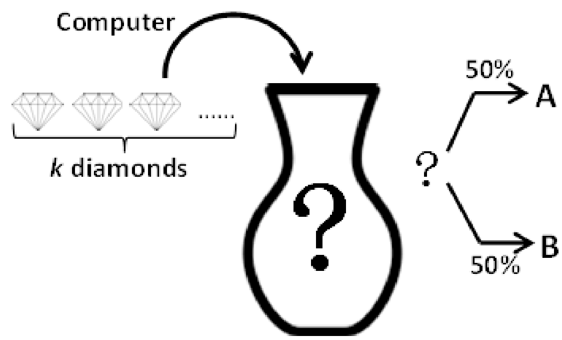

Figure 2: One random vase filled with k "diamonds"

The number of diamonds k is randomly determined and may change from one election to another. More specifically, k is equally likely to be 10, 20, 30, 40, 50, 60, 70, 80, 90, or 100 (so any multiple of 10 between 10 and 100). The other vase gets no diamonds. If the vase containing the diamonds wins the election, all voters will gain these k diamonds as tokens (so each diamond is worth one token).

In each election, the computer will independently determine which vase to fill as well as with how many diamonds. Therefore, which vase contains diamonds, and how many diamonds there are, will likely change from one election to another. You will never be directly informed about which vase has been filled and how many diamonds it contains.

Voting for a vase [Screen 4]

In each election, you have to decide whether or not to vote, and if so, for which vase. If you decide to vote, you pay a voting cost, which is equally likely to be 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, or 15 tokens (so any integer between 1 and 15). These voting costs are drawn independently across voters as well as across elections. Therefore, your voting costs are unrelated to those of other voters and may differ from election to election.

To recap: Before making your voting decision, you will be informed of which team you are in, the size of your bonus X, and your own voting costs. You will not know other voters' team, bonus or costs, nor will you receive any direct information about the allocation of diamonds.

After all voters have made their decisions, the election outcome will be determined by majority rule: the vase that obtains more votes wins the election. If there is a tie, the computer will randomly determine the winning vase by a coin toss.

Earnings [Screen 5]

Your earnings from an election are as follows:

$$\begin{aligned} \text{Earning } &= \text{bonus } X \text{ (if the winning vase belongs to your team)} \\ &\quad + \text{number of diamonds in the winning vase } - \text{voting costs (if you vote)} \end{aligned}$$

For example, suppose that in an election the computer chooses vase **B** and fills it with 40 diamonds. You are in team **A** with a bonus size of $X = 50$, and your voting costs equal 7. You decide to vote for vase **A**. The outcome of the election reveals that 7 voters vote for **A**, 8 voters vote for **B**, while all the remaining voters abstain. Vase **B** thus wins the election because it has more votes than **A**. Your earnings then equal:

$$0 + 40 - 7 = 33 \text{ tokens}$$

You do not get your bonus $X = 50$ tokens since the winning vase (**B**) does not belong to your team (**A**). The amount of 40 comes from the value of the diamonds in the winning vase (**B**). Finally, you pay a cost of 7 tokens since you decided to cast a vote.

Information about diamonds [Screen 6]

As already noted, you will never be directly informed about which vase is filled and how many diamonds it contains. However, you may receive **indirect** information about this from a "robot". The robot knows which vase was chosen and with how many diamonds it was filled.

In some elections, a robot may provide public information about the diamonds. It does so by sending either message **A** or message **B** to all voters (so all voters get the same message).

The robot comes in three different types: ALPHA, BETA, and GAMMA. These types use different strategies to determine what message to send. More specifically:

- ALPHA sends message **A** if vase **A** contains any diamonds, and message **B** if vase **B** contains any diamonds.
- BETA sends message **A** if vase **B** contains 50 diamonds or fewer, and message **B** if vase **B** contains 60 diamonds or more.

- GAMMA sends message **A** if vase **B** contains 90 diamonds or fewer, and message **B** only if vase **B** contains 100 diamonds.

In some elections, the robot will be type ALPHA, in some it will be type BETA and in some it will be type GAMMA. Aside from one of these robot types, in some elections there may be no robot at all. In that case no information about the allocation of diamonds is provided.

The table below summarizes the indirect information you (and all other voters) will receive in different cases. This table will also be given to you in a printed handout. The table shows 20 different situations that may occur. Because both vases are equally likely to be chosen and the number of diamonds is equally likely to be any multiple of 10 between 10 and 100, all of these twenty situations are equally likely to occur in any given election.

	# of diamonds in vase **A**	# of diamonds in vase **B**	No Robot	ALPHA	BETA	GAMMA
1	100	0	-	A	A	A
2	90	0	-	A	A	A
3	80	0	-	A	A	A
4	70	0	-	A	A	A
5	60	0	-	A	A	A
6	50	0	-	A	A	A
7	40	0	-	A	A	A
8	30	0	-	A	A	A
9	20	0	-	A	A	A
10	10	0	-	A	A	A
11	0	10	-	B	A	A
12	0	20	-	B	A	A
13	0	30	-	B	A	A
14	0	40	-	B	A	A
15	0	50	-	B	A	A
16	0	60	-	B	B	A
17	0	70	-	B	B	A
18	0	80	-	B	B	A
19	0	90	-	B	B	A
20	0	100	-	B	B	B

Note: The cells marked Red indicate cases where the robot sends message **A**, while the cells marked Blue indicate cases where the robot sends message **B**.

Elections in "blocks" [Screen 7]

The experiment consists of **12 blocks**, each containing **8 elections**. This makes **96 elections** in total.

In each block, one of the four scenarios "No Robot", "robot ALPHA", "robot BETA", or "robot GAMMA" applies. At the beginning of each block you are informed about the scenario that applies in that block. This means that at the start of each block we will tell you which robot (if any) will send messages in the 8 elections in this block.

In elections without a robot, you are directly asked to make voting decisions without any information about diamonds allocation. In elections with a robot providing information, the voting procedure works as follows.

Before receiving the actual message sent by that robot, you are asked to make voting decisions for each possible message. Specifically, you will be asked:

- What would be your voting decision if robot T sends message **A**?
- What would be your voting decision if robot T sends message **B**?

where we replace 'T' by either 'ALPHA', 'BETA' or 'GAMMA'. For each question, you have to make a decision among voting for vase **A**, voting for vase **B**, or abstaining. After all participants have made their decisions, the message actually sent by the robot will be revealed and all participants' corresponding voting decisions for that message will be carried out to determine the outcome for that election.

Overall payment [Screen 8]

At the end of the experiment, we will randomly choose one election from each block to calculate your earnings from elections.

In addition, before the last election in each block with a robot, we will invite you to make an estimate that gives you a chance to win a potential prize of 10 €in a lottery. You will receive instructions for these estimation tasks when you get there. We will randomly choose one of these estimation tasks for payment. The election where the estimation task is chosen will NOT belong to the elections selected to calculate your election earnings. This means that you cannot get both election earnings and the lottery prize from the same election.

Therefore, your overall payment consists of your participation fee (10 €), the earnings from 12 randomly selected elections (one from each block), and a prize of 10 €if you win the lottery belonging to the randomly selected estimation task. There are no training elections, so you start deciding for money from the very beginning.

[Screen 9: Control Questions, with correct answers in red]

To check your understanding of these instructions, please answer the following questions:

Question 1. [True/False] Each election consists of 25 voters. You will play all the elections in this experiment with the same group of participants.	[True]	[False]
Question 2: [True/False] It is possible that vase A contains 30 diamonds while vase B contains 60 diamonds in the same election.	[True]	[False]
Question 3: [True/False] Vase A is more likely to contain any diamonds than vase B.	[True]	[False]
Question 4: [True/False] Every voter is equally likely to be assigned to either team A or team B. Which team you will be assigned to has nothing to do with others.	[True]	[False]
Question 5: [True/False] If the winning vase belongs to your team, you will gain a bonus of X tokens. X is equally likely to be 20, 50, or 100 tokens. The size of your bonus may be different from those of others.	[True]	[False]
Question 6: [True/False] If you decide to vote, you have to pay a voting cost, which is equally likely to be any integer between 1 and 15 tokens. Your voting costs may be different from those of others.	[True]	[False]
Question 7: [True/False] You are informed about which vase contains diamonds and how many diamonds are there before an election starts.	[True]	[False]
Question 8: [True/False] If there is a robot providing information about diamonds allocation, all voters will receive the message sent by the robot.	[True]	[False]
Question 9: [True/False] If robot BETA sends message A, it is possible that vase A contains no diamonds at all.	[True]	[False]
Question 10: Suppose that in an election vase A contains no diamonds, and vase B contains 90 diamonds. What message will robot GAMMA send to all voters?	[A]	[B]

Question 11: Suppose that in an election vase A contains 50 diamonds. You are in team A, and the size of your bonus is 50. Your voting costs equal 10 and you decide to abstain. The election outcome is such that there are 8 voters voting for A, 5 voters voting for B, while the remaining 12 voters abstain. What are your earnings in this election, in terms of tokens? [100] tokens.

Summary [Screen 10, also distributed in printed handout]

Electorate of 25 voters: In this experiment, you will play a series of elections, with a group of 25 voters. You will be matched with the remaining 24 other participants to form the election group, and play all the elections with the **same** group of participants.

Team assignment: In each election, you are equally likely to be assigned to team **A** or to team **B**. You will gain an individual bonus if the winning vase belongs to your team.

Bonus: The size of the bonus is equally likely to be **20, 50 or 100 tokens**, and is determined independently for each voter. So the size of your bonus will likely be different from others.

Diamonds: In each election, each vase is equally likely to contain some diamonds. The number of diamonds is equally likely to be **any multiple of 10 between 10 and 100 tokens**. If the vase containing diamonds wins the election, each voter obtains all the diamonds in it.

Voting: As a voter, you can choose to vote for either vase, or to abstain. If you decide to vote, you will pay a voting cost, which is equally likely to be **any integer between 1 and 15 tokens**. Your voting costs may be different from others.

Election rule: The vase with the higher number of votes wins the election. If there is a tie, the winning vase will be determined by a coin toss, so either vase wins with probability 50%.

Information from "robots": You will not be directly informed about the allocation of diamonds. However, one of the three robot types, ALPHA, BETA, and GAMMA, may provide relevant information using the message strategies specified in the handout.

Election in "blocks": The experiment consists of 12 blocks, each containing 8 elections. In each block, there may either be no robot and thus no information about the diamonds allocation, or one of the three robots types (ALPHA, BETA, GAMMA) will send a public message to all voters before voters make their decisions.

Estimation tasks: Before the last election in each block with robots, we will invite you to make an estimate that gives you a chance to win a potential prize of 10 €in a lottery. We will randomly choose one of these estimation tasks for payment. You will receive instructions for these estimation tasks when you get there.

Overall payment: Your overall payment consists of your participation fee (10 €), the earnings from 12 randomly selected elections (one from each block), as well as a prize of 10 €if you win the lottery belonging to the randomly selected estimation task. The random selection is such that the paid estimation task does NOT belong to one of the 12 selected elections.

[Instructions for the estimation task (treatment WB_A under $|M| = 1$)]

Estimation task (with potential prize 10 €)

Before the next election starts, we kindly ask you to estimate how likely it is that vase **A** pays *strictly* more to you than vase **B**, depending on the messages sent by robot BETA. Vase **A** pays you strictly more than vase **B** if you can earn more money (without considering voting cost) under the victory of vase **A** than under the victory of vase **B**. *If both vases give the same amount, Vase **A** does not give you strictly more than Vase **B***. For example, suppose you are in team **B**, with a bonus size of 20 tokens. Then vase **A** pays you strictly more than vase **B** only if vase **A** contains 30 diamonds or more. This is because you can earn at least 30 tokens under the victory of vase **A**, while only 20 tokens under the victory of vase **B**. *Note that whether vase **A** pays you strictly more than vase **B** does not depend on whether you vote or not, nor does it depend on the actual election outcome.*

Depending on your answers, you may win a prize of 10 €in a lottery. Specifically, you will be asked to make choices by filling out the 3 lists below: (These are sample lists for illustrations, so please do not fill them out now)

This list asks you to assess the chance that vase **A** pays you strictly more than vase **B** when robot BETA sends message **A**.

How likely do you think that vase **A** pays you strictly more than vase **B**,
if robot **BETA** sends message **A**?

	Option 1	Which option do you prefer?		Option 2
Choice 1	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 0%
Choice 2	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 5%
Choice 3	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 10%
Choice 4	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 15%
Choice 5	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 20%
Choice 6	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 25%
Choice 7	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 30%
Choice 8	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 35%
Choice 9	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 40%
Choice 10	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 45%
Choice 11	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 50%
Choice 12	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 55%
Choice 13	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 60%
Choice 14	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 65%
Choice 15	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 70%
Choice 16	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 75%
Choice 17	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 80%
Choice 18	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 85%
Choice 19	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 90%
Choice 20	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 95%
Choice 21	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 100%

OK

This list asks you to assess the chance that vase **A** pays you strictly more than vase **B** when robot BETA sends message **B**.

How likely do you think that vase **A** pays you strictly more than vase **B**, if robot **BETA** sends message **B**?				
	Option 1	**Which option do you prefer?**		**Option 2**
Choice 1	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 0%
Choice 2	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 5%
Choice 3	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 10%
Choice 4	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 15%
Choice 5	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 20%
Choice 6	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 25%
Choice 7	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 30%
Choice 8	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 35%
Choice 9	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 40%
Choice 10	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 45%
Choice 11	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 50%
Choice 12	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 55%
Choice 13	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 60%
Choice 14	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 65%
Choice 15	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 70%
Choice 16	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 75%
Choice 17	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 80%
Choice 18	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 85%
Choice 19	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 90%
Choice 20	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 95%
Choice 21	Win10€ if vase **A** pays strictly more to you than vase **B**	○	○	Win10€ with probability 100%
		OK		

Only one of the lists you fill out will be used to determine your participation in the lottery. Which list is chosen depends on the actual message robot BETA sends in this election. If the robot sends message **A**, then the first list will be used. If it sends message **B** then the second list will be used. One of the 21 choices from the chosen list will be randomly selected to determine the lottery you will take part in.

So how should you decide which options to tick? That depends on how likely you think it is that vase **A** will pay you strictly more than vase **B**, depending on the message sent by robot BETA. For instance, suppose you think that there is a 63% chance that vase **A** will pay you strictly more. If you want to have the highest chance of winning the 10 €prize, then you should tick **Option 1** on any choice where **Option 2** offers you less than a 63% chance of winning the prize, and you should tick **Option 2** on any choice where the **Option 2** offers you a greater than 63% chance of winning the prize. *Our payment procedure is designed such that your expected earnings are highest if you provide your most accurate estimate.*

You will be asked to do the same estimation task in the last election of every block where robots are providing information. At the end of the experiment one of these estimation tasks will be randomly selected for payment. If an election is selected to decide the lottery, then it will not be selected to determine your election earnings. Therefore, you cannot get both election earnings and the lottery prize from the same election.

Questionnaire [Post-experiment Questionnaire]

Please fill in this short questionnaire carefully. Your feedback will be very helpful for our analysis. Thank you!

Age:	[Textbox to fill in age]	
Gender:	[Radio] Male [Radio] Female	
I study at:	[Choose from list]	
Please indicate on a scale from 1 (fully disagree) to 7 (fully agree) to what extent you agree with the following statements:		
1. In deciding whether or not to cast a vote, I paid careful attention to the voting costs.		[List: 1 to 7]
2. In deciding whether or not to cast a vote, I paid careful attention to the amount of the individual bonus.		[List: 1 to 7]
3. In the blocks with robots, my decision which vase to vote for depended on the messages of the robots.		[List: 1 to 7]
Please rank how your voting decisions rely on the messages from robots in the three scenarios from 1 to 3: (1: rely most; 3: rely least)		1: [List] 2: [List] 3: [List]
Please rank how often do you vote in the four scenarios in the experiment from 1 to 4: (1: most often; 4: least often)		1: [List] 2: [List] 3: [List] 4: [List]
Please rank how competitive do you think the elections are under the four scenarios in the experiment from 1 to 4: (1: most competitive; 4: least competitive) Note: An election is competitive if the vote shares are close for both vases.		1: [List] 2: [List] 3: [List] 4: [List]
Please rank how likely do you think your single vote may change the election outcome under the four scenarios in the experiment from 1 to 4: (1: most likely; 4: least likely)		1: [List] 2: [List] 3: [List] 4: [List]
Did other participants vote more often or less often than you expected? Why?		[(Optional) Textbox here]

[A sample screen shot of the decision interface]

[A sample screen shot of the result interface][11]

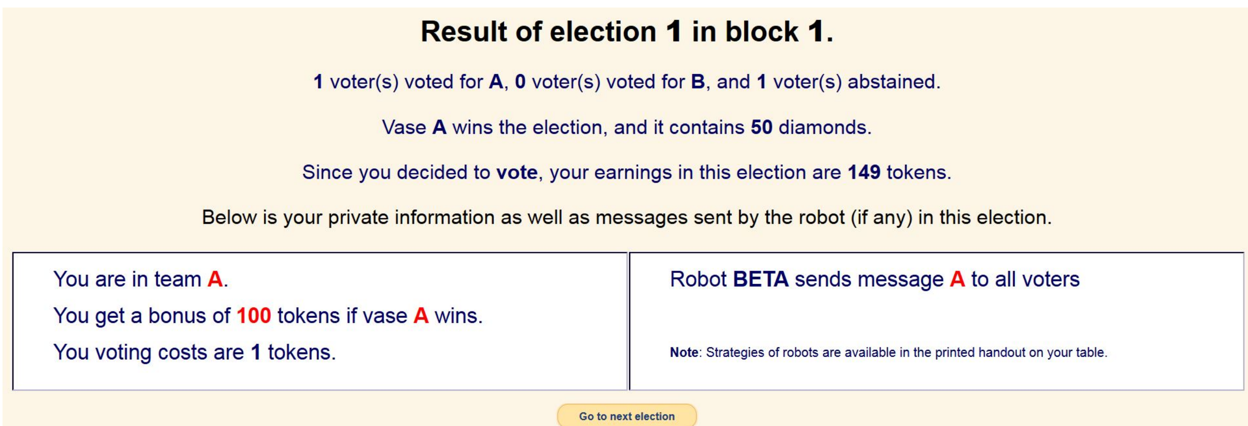

[11] The sample result interface is based on a simplified testing program with 2 voters only. In the actual experiment the numbers of voters voting for A, for B and abstaining always sum up to 25, the actual electorate size.

E Electoral Impacts of Endorser Competition

In this appendix we study the electoral impacts of increased competition between endorsers. We do so by increasing the number of endorsers $|M|$ from 1 to 2. For this purpose, we designed and implemented additional experimental sessions with $|M| = 2$ endorsers. For these sessions we introduce an additional endorser which is *weakly B-biased* ($\chi = -55$; WB_B) and labeled as robot "DELTA". Denote the biases of the two endorsers by χ_1 and χ_2, respectively. We consider three treatments with combinations of biases (χ_1, χ_2): $(0,55)$ (labeled as (UB, WB_A)), $(55,-55)$ (labeled as (WB_A, WB_B)), and $(95,0)$ (labeled as (SB_A, UB)). In these treatments, voters observe two public messages (one from each endorser). Both endorsers' reporting strategies take the cutoff structure characterized by equation (2) in the main text.

In line with our experiment reported in the main text, for these additional sessions with $|M| = 2$ we also adopted a within-subject design for distinct treatments: $NoInfo$, (UB, WB_A), (WB_A, WB_B), and (SB_A, UB). Other aspects of the experimental design and procedures are also similar; we used the strategy method to elicit voting decisions, applied randomization by blocks to minimize order effect, and elicited voters' beliefs using the choice list approach in the last round of election in each block for treatments other than $NoInfo$.[12] In total six experimental sessions with $|M| = 2$ endorsers were conducted at the CREED laboratory of the University of Amsterdam in the summer of 2017. For each session 25 subjects were recruited. These sessions took on average 200 minutes, and the average payment was about 40.1 euros (with a minimum of 28.2 euros and a maximum of 53.1 euros).[13]

The remainder of this appendix is organized as follows. In subsection E.1 we derive both BNE and QRE predictions for the scenario with $|M| = 2$. Based on these equilibrium analyses, we formulate testable hypotheses concerning the ex-ante impact of introducing a second endorser on election outcomes and voter turnout. The experimental results are subsequently reported in subsection E.2.

E.1 Equilibrium Analyses and Hypotheses

In this subsection we derive BNE and QRE predictions. Let $k_X(\chi_1, \chi_2)$ and $Pr[X|\chi_1, \chi_2]$ denote the posterior expectations of k and the ex-ante probability of sending message combination $X \in \{(A,A), (A,B), (B,B)\}$, respectively, conditional on the endorsers' biases χ_1 and χ_2. Let $\chi_+ \equiv \max\{\chi_1, \chi_2\}$ and $\chi_- \equiv \min\{\chi_1, \chi_2\}$. It follows from the cutoff reporting strategy (cf. equation

[12] Complete instructions for these experimental sessions are available upon request.

[13] Because sessions with $|M| = 2$ took on average 50 minutes longer than sessions with $|M| = 1$, we paid each subject in sessions with $|M| = 2$ an extra amount of 10 euros (unexpected and not pre-announced) at the end of the experiment as a compensation.

(2) in the main text) that

$$Pr[(A,A)|\chi_1,\chi_2] = 1 - F(-\chi_-), \text{ and } k_{(A,A)}(\chi_1,\chi_2) = E[k|k > -\chi_-]$$
$$Pr[(A,B)|\chi_1,\chi_2] = F(-\chi_-) - F(-\chi_+), \text{ and } k_{(A,B)}(\chi_1,\chi_2) = E[k|-\chi_+ < k \leq -\chi_-]$$
$$Pr[(B,B)|\chi_1,\chi_2] = F(-\chi_+), \text{ and } k_{(B,B)}(\chi_1,\chi_2) = E[k|k \leq -\chi_+]$$

Following the logic of Appendix A, conditional on the endorsers' biases (χ_1,χ_2) and realized message combination $X \in \{(A,A),(A,B),(B,B)\}$, both BNE and QRE predictions can be obtained by replacing k by the corresponding posterior expectation $k_X(\chi_1,\chi_2)$. Table E.1 summarizes all $Pr[X|\chi_1,\chi_2]$, $k_X(\chi_1,\chi_2)$, and both the interim and ex-ante outcomes predicted by BNE and QRE for all combinations of (χ_1,χ_2).

Table E.1: Rational Posterior Expectations of k and Theoretical Predictions under $|M| = 2$

Endorsers' Biases (χ_1,χ_2)	$(0,55)$			$(55,-55)$			$(95,0)$		
Message Combination X	(A,A)	(A,B)	(B,B)	(A,A)	(A,B)	(B,B)	(A,A)	(A,B)	(B,B)
$Pr[X\|\chi_1,\chi_2]$	0.50	0.25	0.25	0.25	0.50	0.25	0.50	0.45	0.05
$k_X(\chi_1,\chi_2)$	55	−30	−80	80	0	−80	55	−50	−100
BNE predictions									
Expected vote share of A	0.81	0.32	0.10	0.90	0.50	0.10	0.81	0.21	0.00
Winning probability of A	0.94	0.15	0.03	0.97	0.50	0.03	0.94	0.07	0.02
Expected voter turnout	0.22	0.32	0.17	0.17	0.45	0.17	0.22	0.23	0.13
QRE predictions									
Expected vote share of A	0.71	0.36	0.25	0.75	0.50	0.25	0.71	0.30	0.23
Winning probability of A	0.91	0.18	0.06	0.94	0.50	0.06	0.91	0.10	0.04
Expected voter turnout	0.38	0.42	0.37	0.37	0.48	0.37	0.38	0.39	0.36

Note: For all (χ_1,χ_2), the law of iterated expectation is satisfied: $\sum_{X\in\{(A,A),(A,B),(B,B)\}} Pr[X|\chi_1,\chi_2] \cdot k_X(\chi_1,\chi_2) = 0$. In generating QRE predictions, we again used the logit response parameter $\widehat{\lambda} = 17.97$ obtained from an out-of-sample estimation based on corresponding sessions in our 2016 experiment.

Using these equilibrium predictions, we study the ex-ante impact of introducing a second endorser on candidates' expected vote shares, the election outcome and voter turnout. As $|M|$ increases from 1 to 2, we denote the biases of the existing and the newly introduced endorsers by χ_1 and χ_2, respectively, and consider the aforementioned three combinations of biases, (UB, WB_A), (WB_A, WB_B) and (SB_A, UB). Given any pair (χ_1,χ_2), we identify the ex-ante electoral impact of introducing a second endorser with bias χ_2 by comparing outcomes from the one-endorser scenario with bias χ_1 to the two-endorser scenario with bias combination (χ_1,χ_2). Figure E.1 presents BNE predictions for these comparisons for candidate A's expected vote share and winning probability, as well as for voter turnout.[14]

[14] QRE predictions are qualitatively similar. More details are available upon request.

Figure E.1: The Electoral Influence of Introducing a Second Endorser

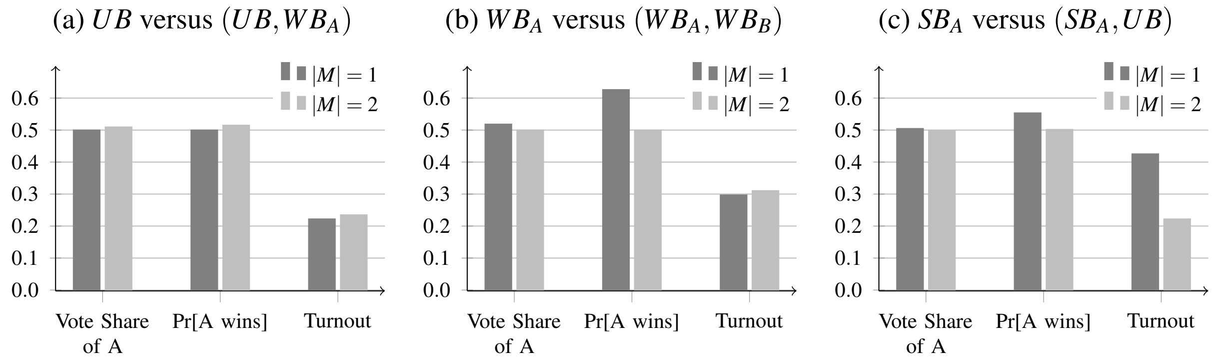

Note: Panels depict the ex-ante influence, predicted by BNE, of introducing a second endorser on candidate A's expected vote share, winning probability (Pr[A wins]) and voter turnout.

The three panels of Figure E.1 show that introducing a second endorser has very little impact on candidates' expected vote shares from the ex-ante perspective. Regarding the election outcome, we show that introducing a endorser can systematically increase A's winning probability if and only if the second endorser is more A-biased than the existing endorser, i.e., $\chi_2 > \chi_1$. In our election game, this implies that A's ex-ante winning probability increases in the transition from UB to (UB, WB_A), and decreases in the transitions from WB_A to (WB_A, WB_B), and from SB_A to (SB_A, UB). All these predictions are confirmed in Figure E.1. These observations yield Hypothesis E.1.

Hypotheses E.1. *Influence of introducing a second endorser on the election outcome:*

(a) Ex-ante, A's winning probability is higher under (UB, WB_A) than under UB.

(b) Ex-ante, A's winning probability is lower under (WB_A, WB_B) than under WB_A.

(c) Ex-ante, A's winning probability is lower under (SB_A, UB) than under SB_A.

As for voter turnout, Figure E.1 generates three predictions. These are summarized in Hypothesis E.2.

Hypotheses E.2. *Influence of introducing a second endorser on voter turnout:*

(a) Ex-ante, voter turnout is higher under (UB, WB_A) than under UB.

(b) Ex-ante, voter turnout is higher under (WB_A, WB_B) than under WB_A.

(c) Ex-ante, voter turnout is lower under (SB_A, UB) than under SB_A.

E.2 Experimental Results

The aggregate ex-ante electoral impact of introducing a second endorser are presented in Figure E.2. As is evident from the top and middle panels, except for the comparison between SB_A and (SB_A, UB), both BNE and QRE predictions correctly capture, qualitatively and quantitatively, the ex-ante impacts of introducing a second endorser on party vote shares and election outcomes. The bottom panels of Figure E.2 show that, as in sessions with $|M| = 1$, observed turnout rates are rarely affected by the entrance of a second endorser and they are again systematically higher than both BNE and QRE predictions. It is worth noting, however, that these results need not be entirely inconsistent with our theoretical predictions. This is because, except for the comparison between SB_A and (SB_A, UB), the ex-ante impacts of introducing a second endorser on voter turnout predicted by BNE and QRE are both quantitatively negligible.

In what follows we conduct formal statistical tests for Hypotheses E.1 and E.2. Specifically, we use exact Fisher-Pitman permutation tests (henceforth, FPp) for pairwise comparisons of differences in means between sessions with different $|M|$.[15]

Hypothesis E.1a predicts that introducing a weakly A-biased endorser will increase candidate A's ex-ante winning probability, if the preexisting endorser is unbiased. We indeed observe a slight increase in A's winning probability (from 48.3% to 50.7%) in the top-left panel of Figure E.2; the difference is statistically significant (FPp, $p = 0.041$, $N = 12$). Hypothesis E.1b predicts that, if the preexisting endorser is weakly A-biased, then introducing a weakly B-biased endorser will decrease A's winning probability. Indeed, we observe a sizable decrease (from 58.5% to 47.6%) in the top-central panel of Figure E.2, and the difference is statistically significant (FPp, $p = 0.009$, $N = 12$). Finally, Hypothesis E.1c predicts that, if the preexisting endorser is strongly A-biased, then introducing an unbiased endorser will decrease A's winning probability. Again, we observe a sharp decrease (from 66.9% to 52.7%) in the top-central panel of Figure E.2, and the difference is statistically significant (FPp, $p = 0.002$, $N = 12$). Note, however, that given the realized parameter draws for the experiment, both BNE and QRE predictions actually violate Hypothesis E.1c; they predict small treatment effects in the other direction, if any. These findings are summarized in Result E.1.

Result E.1. *Influence of introducing a second endorser on candidate A's ex-ante winning probability:*

(a) A's ex-ante winning probability is significantly higher in treatment (UB, WB_A) than in UB.

(b) A's ex-ante winning probability is significantly lower in treatment (WB_A, WB_B) than in WB_A.

[15] Contrary to the statistical tests for sessions with $|M| = 1$ reported in the main text, this is a between-subject comparison because the number of endorsers differs across sessions.

Figure E.2: Ex-ante Electoral Impacts of Introducing a Second Endorser

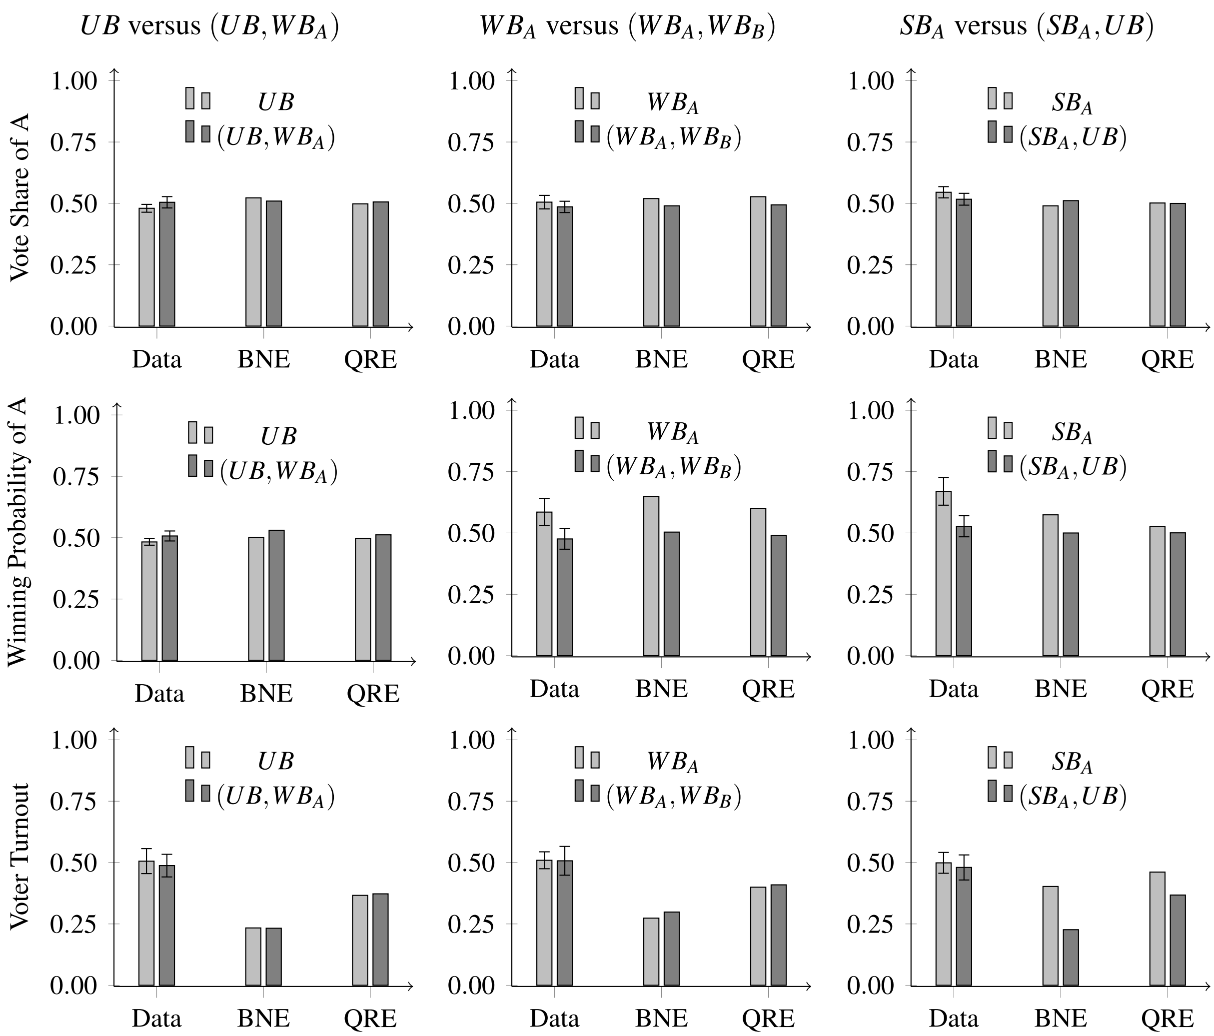

Note: Bars show candidate A's ex-ante vote share (top panels), winning probability (middle panels), and voter turnout (lower panels) and the corresponding predictions by BNE and QRE. 95% confidence intervals are plotted for the observed outcomes. All theoretical predictions are generated based on the realized parameter draws for the experiment.

(c) *A's ex-ante winning probability is significantly lower in treatment* (SB_A, UB) *than in* SB_A.

Hypothesis E.2a predicts that if the preexisting endorser is unbiased, introducing a biased endorser leads to an increase in voter turnout. Contrary to Hypothesis E.2a, our results (bottom-left panel of Figure E.2) show a slight decrease in turnout (from 50.6% to 48.8%), although this effect is statistically insignificant (FPp, $p = 0.500$, $N = 12$). Hypothesis E.2b predicts that if the preexisting endorser is weakly A-biased, then introducing an weakly B-biased endorser with an identical degree of bias will increase voter turnout. Contrary to HE.2b, ex-ante voter turnout varies little with the presence of an extra endorser (from 50.9% to 50.7%) as shown in the bottom-central panel of Figure E.2. The difference is statistically insignificant (FPp, $p = 0.931$, $N = 12$). Finally, Hypothesis E.2c predicts that if the preexisting endorser is strongly biased (as in SB_A), then introducing an unbiased endorser will decrease voter turnout. Although we indeed observe a slight decrease (from 49.9% to 48.0%, see the bottom-right panel of Figure E.2), this effect is again statistically insignificant (FPp, $p = 0.455$, $N = 12$). Overall, we find no evidence that introducing a second endorser affects voter turnout from the ex-ante perspective. These findings are summarized in Result E.2.

Result E.2. *Influence of introducing a second endorser on voter turnout:*

(a) *Ex-ante, voter turnout does not vary significantly from treatment* UB *to* (UB, WB_A).

(b) *Ex-ante, voter turnout does not vary significantly from treatment* WB_A *to* (WB_A, WB_B).

(c) *Ex-ante, voter turnout does not vary significantly from treatment* SB_A *to* (SB_A, UB).

In a nutshell, we find that introducing a second endorser systematically pushes the election outcome in the direction predicted by theory, yet it has little systematic impact on voter turnout.

References

Goeree, Jacob K, Philippos Louis and Jingjing Zhang. 2018. "Noisy Introspection in the 11–20 Game." *The Economic Journal* 128(611):1509–1530.

McKelvey, Richard D and Thomas R Palfrey. 1995. "Quantal response equilibria for normal form games." *Games and Economic Behavior* 10(1):6–38.

McKelvey, Richard D and Thomas R Palfrey. 1998. "Quantal response equilibria for extensive form games." *Experimental Economics* 1(1):9–41.

Vuong, Quang H. 1989. "Likelihood ratio tests for model selection and non-nested hypotheses." *Econometrica* 57(2):307–333.