Media Bias and Elections - An Experimental Study

Sun, J.; Schram, A.; Sloof, R.

Publication date
2018

Document Version
Submitted manuscript

Media Bias and Elections – An Experimental Study*

Junze Sun† Arthur Schram‡ Randolph Sloof§

This version: November 30, 2018.
Preliminary draft. Please do not cite.

Abstract

We study the impact of media bias on voters' behavior and election outcomes in a laboratory experiment. We model this interaction and derive the Bayesian Nash Equilibria. These predict for a single media that, ex-ante, an increased media bias affects candidates' winning probabilities non-monotonically and increases voter turnout. Introducing a second media can affect the election outcome and voter turnout in either direction. We test these predictions in a laboratory experiment and find that both observed election outcomes and vote shares are well predicted. Voter turnout, however, is much less responsive to media bias than predicted. We show that subjects' observed behavior can be rationalized, to a substantial extent, by a quantal response equilibrium model combined with (a) distinct noise parameters for candidate choice and turnout decisions, (b) non-Bayesian belief updating, and (c) "partial competition neglect" (i.e., voters' imperfect ability to infer closeness of elections from messages announced by media).

JEL Codes: C92, D72, D83, D91

*Financial support from the University of Amsterdam Research Priority Area in Behavioral Economics is gratefully acknowledged. Comments by seminar participants at NYU-Abu Dhabi and the Hong Kong University of Science and Technology, and at the NYU-WESSI workshop in Florence and EWEBE in Bologna are gratefully acknowledged. All remaining errors are ours.

†CREED, University of Amsterdam and Tinbergen Institute. Email: J.Sun@uva.nl.

‡CREED, University of Amsterdam and Tinbergen Institute, and Robert Schumann Center for Advanced Studies, EUI, Florence. Email: A.J.H.C.Schram@uva.nl.

§CREED, University of Amsterdam and Tinbergen Institute. Email: R.Sloof@uva.nl.

1 Introduction

Around the world, voters in many elections lack the information needed to properly assess their candidates. 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 it is unlikely that the costs would outweigh the expected benefits from the improvement in the choice made, especially in large elections (Downs, 1957).[1] For this reason, voters often rely on low-cost external sources to evaluate candidates. One of the most important sources of such information is the media coverage of campaigns, candidates, and elections. Mass media thus plays an important role in shaping voters' views on candidates and therefore potentially have a salient impact on voters' behavior and the election outcomes.[2]

Mass media would be efficiency enhancing if they provide unbiased information that both reduces the private costs of information acquisition by voters and increases the likelihood of a victory of the efficiency-maximizing candidate. However, a large strand of literature shows that media bias is pervasive and can arise from various sources.[3] From the supply side, media bias could come from the private motives of journalists (Baron, 2006), editors (Sobbrio, 2014), government officials (Besley and Prat, 2006; McMillan and Zoido, 2004) or interest groups who finance advertising (Ellman and Germano, 2009). Alternatively, media bias could also come from the demand side. For example, Mullainathan and Shleifer (2005) show that media bias can arise if media outlets choose their editorial positions to cater readers with confirmation bias. Gentzkow and Shapiro (2006) show that reputation concerns may lead media outlets to cater to readers' priors. Importantly, although it is widely believed that media bias is most severe in weak democracies (Djankov et al., 2003), empirical evidence shows that media bias is pervasive even in well-developed countries like the U.S. (Groseclose and Milyo, 2005; Gentzkow and Shapiro, 2010; Puglisi and Snyder, 2011, 2015) and voters are skeptical about the neutrality of mass media outlets (Pew Research Center, 2008). An important question is then, what are the consequences of media bias?

We address this question in this paper. We first study the impact of media bias on voter behavior and election outcomes in a simple electoral environment, and derive both Bayesian

[1] The expected benefits of being informed are decreasing in the size of the electorate because the effect of a single vote decreases. Rational voters then do not engage in costly information acquisition. This is known as *rational ignorance* (Downs, 1957).

[2] Of course, in the current era, social media also play an important role in voters' quest for information (Allcott and Gentzkow, 2017). Here we focus on more traditional mass media such as newspapers, radio, and TV channels.

[3] See Prat and Stromberg (2013) and Gentzkow et al. (2015) for comprehensive reviews.

Nash equilibria (BNE) and quantal response equilibrium (QRE) for this game. These provide theoretical predictions that we subsequently test in a laboratory experiment. In doing so, we focus on two questions. First, in an environment with a single media outlet, how does an increase in media bias affect parties' vote shares, the election outcome, and voter turnout? Second, what is the impact of media entry on those electoral outcomes? We investigate this question by studying the comparative statics of introducing a second (possibly biased) media outlet on top of the preexisting one.

Our theoretical model is an application to our experimental environment of a more general model developed in our companion paper (Sun et al. (2018), henceforth SSS18). In the model we consider an election (under simple majority rule) with two candidates, A and B, who differ along the *quality* dimension. The quality of a candidate may be interpreted as her ability or valence and is commonly valued by all voters. Apart from caring about qualities, voters have heterogeneous ideological preferences in favor of one of the two candidates. A voter's preference over candidates then consists of both a common component (quality) and a private component (ideology). It is thus possible for a voter to prefer the ideologically less-favored candidate if she believes that the relative quality of this candidate is sufficiently high. Voters can vote for either candidate or abstain, but voting is costly and the costs are private information. We assume the election is symmetric in the sense that *a priori* (i) the qualities of both candidates are identical, and (ii) voters' ideological preferences do not systematically favor one candidate over the other. Therefore, from the ex-ante perspective, voter welfare is higher if the quality-superior candidate is elected.

We assume that voters cannot directly get access to candidates' qualities. These are, however, (precisely) observed by mass media outlet(s). Voters, then, can infer about candidates' relative qualities from candidate endorsements released by these media outlet(s). Importantly, the media outlet(s) may have different preferences over candidates than voters and can publish endorsements in an attempt to influence the election outcome in favor of the candidate they prefer. An *unbiased* media outlet prefers the candidate with the higher quality, and her preferences are therefore in line with ex-ante voter welfare maximization. A *biased* media, however, may prefer a candidate even if its quality is inferior. Throughout the analysis, we assume that voters know exactly each media's bias.

As an example, assume that the New York Times was biased towards candidate Hillary Clinton in the 2016 presidential elections, while Fox News was biased towards Donald Trump. Our model assumes 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 beliefs about candidates' quality levels, compared to the uninformative prior. Through affecting voters' posterior beliefs, media –even if voters precisely know their biases– can affect the election outcomes and voter turnout.

We assume that an unbiased media always endorses the candidate with a higher realized quality whereas a biased media outlet may endorse its ex-ante favored candidate even if its quality is inferior.[4] Hence, an endorsement for candidate A from a media outlet biased towards A is less convincing than the same endorsement from an unbiased media. A biased media, however, may endorse the 'other' candidate B if the candidate A towards which it is biased has a realized quality that is too low. The larger the bias towards A, the lower is the likelihood of endorsing B. Therefore, as the media bias in favor of candidate A increases, an endorsement for A becomes less credible, and less likely to signal a large quality difference between the candidates. In contrast, an endorsement of candidate B becomes more credible and signals a larger quality difference between candidates as media bias (in favor of candidate A) increases. As an example of the latter, imagine that Fox News had endorsed Clinton in the 2016 elections. This would have been very informative about the quality difference between Trump and Clinton. In this way, media bias in our model affects voters' decisions (only) through its influence on voters' posterior beliefs of the quality difference.

We study the impact of media bias (say, towards candidate A) on voter behavior and election outcomes in the presence of a single media outlet. On the one hand, with a higher bias the outlet is more likely to endorse candidate A; this effect *per se* increases A's winning chances. On the other hand, this outlet's endorsement for candidate A becomes less credible, whereas its endorsement for the other candidate B becomes more credible, as media bias increases. These effects *per se* decreases candidate A's winning chances. The aggregate effect, therefore, depends on which effect is strongest. With a small degree of bias, the first effect dominates and an increase in media bias raises candidate A's winning probability. With a large degree of bias, the opposite applies. Therefore, from the ex-ante perspective, an increase in media bias affects candidate A's winning probability in a non-monotonic way.

The impact of media bias on voter turnout hinges on what an endorsement reveals about the closeness of the election. With a strong media bias, the outlet's endorsement for candidate A (e.g., a Republican endorsement from the Fox News) is incredible and signals, if any, a small expected quality difference between candidates. This implies that the election is more likely to end up in a close race, because the vote shares for both candidates are expected to be close. In this situation, voter turnout will be high due to the "competition effect" (Levine and Palfrey, 2007).[5] In contrast, this outlet's endorsement for candidate B (e.g.,

[4]In SSS18 we show, in a more general setup where media communicate to voters via cheap talk (Crawford and Sobel, 1982), that this simple binary-endorsement strategy holds robustly in equilibrium.

[5]The competition effect predicts voter turnout to be high if the election is expected to be close. This is because voting is mostly likely to be decisive in close elections.

a Democratic endorsement from the Fox News) is very credible in signaling the superiority of candidate B and implies a sharp quality difference between candidates. This in turn implies that the election is most likely to end up in a landslide victory of candidate B, who is expected to gain a substantially larger share of votes than A. As a result, voters have little incentive to cast costly votes and the turnout rate will be low. Therefore, the expected voter turnout increases (decreases) in media bias conditional on an endorsement for candidate A(B). Ex-ante, however, the expected voter turnout increases unambiguously with media bias, and it achieves its peak in an environment with no information at all.

Finally, we explore how increasing the number of media outlets affects voting behavior and the election outcome. To do so, we analyze the comparative statics of introducing a second media outlet on top of an existing one. We demonstrate that, depending on the biases of both outlets, having an extra outlet can systematically drive the election outcome and voter turnout in either direction.

To test these theoretical predictions, we implemented a laboratory experiment. Aside from enabling a direct test of our theory, laboratory control has important advantages for addressing our research questions.[6] First, it avoids the self-selection into treatments that hinders the casual inference from observational data. For example, media exposure outside the laboratory may be highly correlated with a voter's ideological preference (Gentzkow and Shapiro, 2011; Durante and Knight, 2012). Second, it allows us to directly induce (and, therefore, measure) crucial elements of the model such as quality levels, voter preferences, and the extent of media bias. Finally, it enables direct measurement of beliefs, which may play an important role in voters' decisions.

Our experimental results show that the behavior of aggregate vote share and candidates' winning chances are well predicted. Voter turnout, however, is –contrary to the theoretical prediction– irresponsive to media bias. Overall, standard BNE predictions explain 61.6% to 74.5% of the variations we observe in a candidate's winning chances and vote share, but only 13% of the variation in observed voter turnout. This difference between the model's predictive power for candidate choice (winning probabilities and vote shares) and for turnout warrants further investigation.

We discuss three alternative explanations: (i) quantal response equilibria (QRE) with distinct noise parameters for candidate choice and turnout decisions; (ii) non-Bayesian belief

[6] See Falk and Heckman (2009) for a general discussion of these advantages and a comparison to empirical analyses based on observational field data. One possible disadvantage of testing the model in the laboratory is the small size of the electorate that this allows for. This disadvantage is mitigated, however, when laboratory data are used to test a theory (Schram, 2005). Studies comparing various electorates support the size effects predicted by theory (Levine and Palfrey, 2007). This suggests that the theory itself can be used to extrapolate from laboratory results.

updating; (iii) competition neglect (i.e., voters' imperfect ability to infer the closeness of elections from messages announced by media outlets). We show that a hybrid structural model combining all three features can, to a substantial extent, rationalize subjects' observed behavior in the experiment.

The remainder of this paper is organized as follows. Section 2 summarizes the related literature and our contributions. Section 3 introduces the theoretical model, which is subsequently analyzed in Section 4. Section 5 presents our experimental design and Section 6 discusses the results. In Section 7 we explore possible explanations for observed deviations in subjects' behavior in the experiment. Section 8 concludes.

2 Related literature

Our work relates to a large and growing literature that empirically documents media's influence on voters' behavior (including party vote shares and voter turnout) using observational field data. One strand of this literature examines media influence in a particular election (Strömberg, 2004; DellaVigna and Kaplan, 2007; Gerber et al., 2009; Chiang and Knight, 2011; Enikolopov et al., 2011; Adena et al., 2015). The typical finding is that media exposure can shift party vote shares in favor of the media endorsed candidate, and may either increase or decrease voter turnout. The other strand of this literature investigates the long run impacts of media entry/exit (Gentzkow, 2006; Gentzkow et al., 2011; Drago et al., 2014; Gavazza et al., 2015; Cagé, 2017; Ellingsen et al., 2018). These studies provide mixed evidence regarding how media entry affect party vote shares and voter turnout. Nevertheless, these field studies provide little direct evidence on how biased media affect the election outcome, because these studies rely on data from only parts of the electorate. We contribute to these field studies in three ways. First, we develop a theoretical model that generate testable predictions which may explain these styled empirical findings in a unified framework.[7] Second, we test these predictions in a well controlled laboratory experiment. Third, to the best of our knowledge, we provide the first direct experimental evidence regarding both the interim and ex-ante influence of biased media on election outcomes.

A second strand of literature that relates to our work empirically identify the persuasion effect of media outlets. Once again, the findings are mixed. There is evidence that voters understand that information provided by the media may be biased. Chiang and Knight (2011) investigate the influence of media bias on voting intention by exploiting the exact timing of newspaper endorsements in the context of 2004 US presidential election, and find that only

[7] Piolatto and Schuett (2015) provide an alternative theoretical explanation for the mixed empirical evidence on the impacts of media entry on political participation, based on an ethical voter framework.

highly credible newspaper endorsements affect voters. Durante and Knight (2012) show evidence that viewers in Italy increase their propensity to watch left-leaning television channels to offset the movement of public slant to the right under Berlusconi's administration. Both findings suggest that voters are sufficiently sophisticated to filter out the impact of media bias. However, Cain et al. (2005) find evidence in a controlled laboratory experiment that decision makers fail to sufficiently account for the interest misalignment between themselves and an information provider. DellaVigna and Gentzkow (2010) provide a comprehensive review of the empirical evidence on persuasion. Most of these works exploit observational or survey data from the field. In contrast, we elicit voters' beliefs in a controlled laboratory experiment, which allows us to directly investigate how these beliefs are affected by media bias.

Finally, our paper relates to a large strand of literature studying voting games using both experiments and field data. Most relevant for our analysis of the effects of media bias on turnout is the "competition effect", which predicts higher voter turnout in closer elections. This result has gained empirical support from both field (Blais, 2000) and laboratory experiments (Levine and Palfrey, 2007; Grosser and Schram, 2010).

3 The Election Game

In this section, we present a two-candidate election game with media, which is implemented in the laboratory experiment.

Candidates. There are two candidates, A and B, each characterized by their own qualities. Without loss of generality, we normalize the quality of candidate B to zero and let k be the relative quality of candidate A. Therefore, if $k > (<)0$, A has a higher (lower) quality than B. We assume that k is drawn from a uniform distribution $F(\cdot)$ on the discrete set $\{-1, -0.9, \cdots, -0.1, 0.1, \cdots, 0.9, 1\}$.[8] Unless explicitly stated otherwise, k is assumed to be unobservable to voters but precisely observable to media.

Electorate. The electorate consists of $\mathcal{N} = 25$ eligible voters. Each voter can choose an action from $\{A, B, O\}$, representing voting for candidate A, candidate B, and abstaining, respectively. Voter i is characterized by a 2-dimensional private type, (v_i, c_i). v_i represents voter i's ideological preference, and c_i represents her voting costs, which are executed if and only if she casts a vote. We assume that both v_i and c_i are independently and identically distributed across voters. Each v_i is drawn from a uniform distribution $G(\cdot)$ on the discrete set $\{1, -0.5, -0.2, 0.2, 0.5, 1\}$. Note that a voter with $v_i > (<)0$ has an ideological preference leaning towards A(B). For each voter i, c_i is drawn from a uniform distribution $\rho(\cdot)$ on

[8] Note that the value $k = 0$ cannot occur, so we do not consider the case of equal quality.

the discrete set $\{0.01, 0.02, \cdots, 0.15\}$. We normalize voter i's utility to 0 if candidate B is elected, and let her payoff if A is elected be $k + v_i$. Hence, voter i prefers candidate A(B) if $k > (<) - v_i$, and her individual stake at the election is reflected by $|k + v_i|$. Formally, given k, v_i and c_i, voter i's utility function is given by

$$u(k, v_i, c_i, a_i, \omega) = (k + v_i) \cdot \mathbb{1}_{\omega=A} - c_i \cdot \mathbb{1}_{a_i=O} \tag{1}$$

where $\omega \in \{A, B\}$ indicates the winning candidate, $a_i \in \{A, B, O\}$ indicates voter i's action, and $\mathbb{1}_E$ is an indicator function that equals one if the event E is true.

Media outlets. There is a set of media outlets, M. In this paper, we we consider only cases with one or two media outlets, i.e., $|M| \in \{1, 2\}$. Each outlet $m \in M$ precisely observes k and can send a binary public message $s_m \in \{A, B\}$ to all voters, which are interpreted as "media endorsement". Each outlet $m \in M$ is characterized by a commonly known parameter $\chi_m \in [-1, 1]$, and sends public endorsement using the following cutoff strategy:[9]

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

Namely, outlet $m \in M$ sends endorsement $A(B)$ if and only if $k > (\leq) - \chi_m$. Figure 1 illustrates this endorsement strategy for the case where $\chi_m > 0$.

Figure 1: An illustration of endorsement strategy (2)

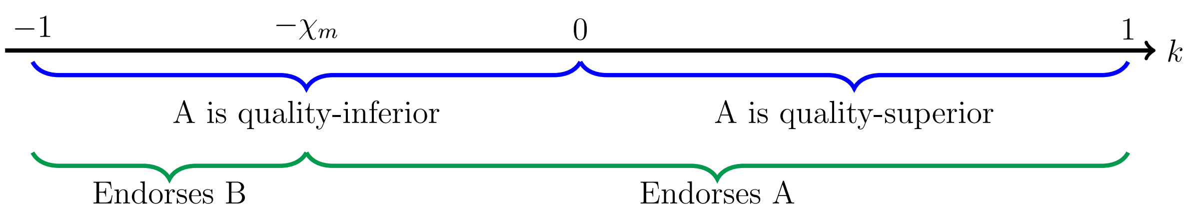

Note: The black line denotes candidate A's relative quality advantage. $\chi_m > 0$ indicates that the media is biased towards A. It endorses A even in some cases where B has the higher quality.

If $\chi_m = 0$, m is said to be *unbiased* since it always endorses the candidate with a higher quality. If $\chi_m > (<) 0$, the outlet is said to be *A(B)-biased* as it may endorse candidate A(B) even if her quality is inferior. For this reason, we refer χ_m as the media bias.

Timing. The timing of the game is as follows:

1. Nature draws k and voter type profile $\{(v_i, c_i)\}_{i=1}^{25}$.

[9] In SSS18, we show that the cutoff strategy (2) endogenously arises in equilibrium in an environment where each media outlet $m \in M$ has utility function $V(k, \chi_m, \omega) = (k + \chi_m) \cdot \mathbb{1}_{\omega=A}$ and communicate to voters via cheap talk (Crawford and Sobel, 1982).

2. Observing k, outlet $m \in M$ sends public message $s_m \in \{A, B\}$ using strategy (2).

3. Observing the message profile $\{s_m\}_{m \in M}$, voters simultaneously decide on their votes (or abstain).

4. The winning candidate is determined by simple majority rule, with ties broken by a fair coin toss. All payoffs then realize.

4 Equilibrium Analyses

4.1 Bayesian Nash Equilibrium

Because this is a game of incomplete information, we first derive the Bayesian Nash Equilibrium (BNE). We focus on type-symmetric BNE, where voters with the same type (v_i, c_i) adopt the same voting strategy in equilibrium.[10] Under incomplete information, voters need to use the available information to: (i) infer which candidate is better, and (ii) decide whether or not to cast a vote for the better candidate. With voting being costly, casting a vote for the un-preferred candidate is strictly dominated by abstaining. Were k common knowledge, (i) is be straightforward; candidate A(B) is strictly better for voter i whenever $k + v_i > (<)0$. The decision whether or not to cast a vote (ii) is much more difficult; the voter has to weigh the expected benefits against the costs of voting.

The turnout decision is critically driven by the closeness of elections. If the two candidates are of equal qualities (i.e., $k = 0$), then the election is expected to end up in a close race because both candidates are expected to receive equal shares of votes, as ideological preferences are symmetrically distributed. In this case, the chances of being pivotal are high and voters have strongest incentives to cast costly votes, leading to a high voter turnout. If instead the quality difference is large (i.e., $|k| \gg 0$), then the election is likely to end up in a landslide victory of the quality-superior candidate, who is expected to get substantially higher shares of votes than her opponent. Consequently, the pivotal chances are low and voters abstain to economize on voting costs. In this way, the quality difference $|k|$ is negatively associated with the closeness of elections. Voter turnout is thus a decreasing function of $|k|$, consistent with the "competition effect" (Levine and Palfrey, 2007).

We formalize this reasoning in SSS18 in a set of propositions. This is based on a model with population uncertainty (Myerson, 2000). Our numerical calculations confirm that the

[10] We ignore asymmetric equilibria (if any) because they are implausible without an omniscient mediator coordinating voters' behavior, in an environment with strategic uncertainty (e.g., voters' v_i and c_i). See, for example, Palfrey and Rosenthal (1985) and Levine and Palfrey (2007). We numerically calculate the BNE for all k concerned in this paper. The details can be found in our online Appendix A.

predictions derived from those propositions carry over to our election game with $\mathcal{N} = 25$ voters. Here, we present those theoretical predictions. Let $VS_A(k)$ denote the expected vote share for candidate A in equilibrium, i.e., the fraction of votes for A among all votes cast if k is common knowledge and the electorate size is $\mathcal{N} = 25$. Define $\pi_A(k)$ as the probability that candidate A wins the election and $T(k)$ as the expected voter turnout rate in equilibrium, both conditional on k and $\mathcal{N} = 25$. Proposition 1 summarizes the comparative statics properties of $VS_A(k)$, $\pi_A(k)$ and $T(k)$ when k is commonly known to all voters.

Proposition 1 *For all $k, k' \in [-1, 1]$, it holds that*

1. $VS_A(k) > VS_A(k')$ *if and only if* $k > k'$; $VS_A(0) = \frac{1}{2}$.
2. $\pi(k) > \pi(k')$ *if and only if* $k > k'$; $\pi(0) = \frac{1}{2}$.
3. $T(k) > T(k')$ *if and only if* $|k| < |k'|$.

Figure 2: BNE Predictions

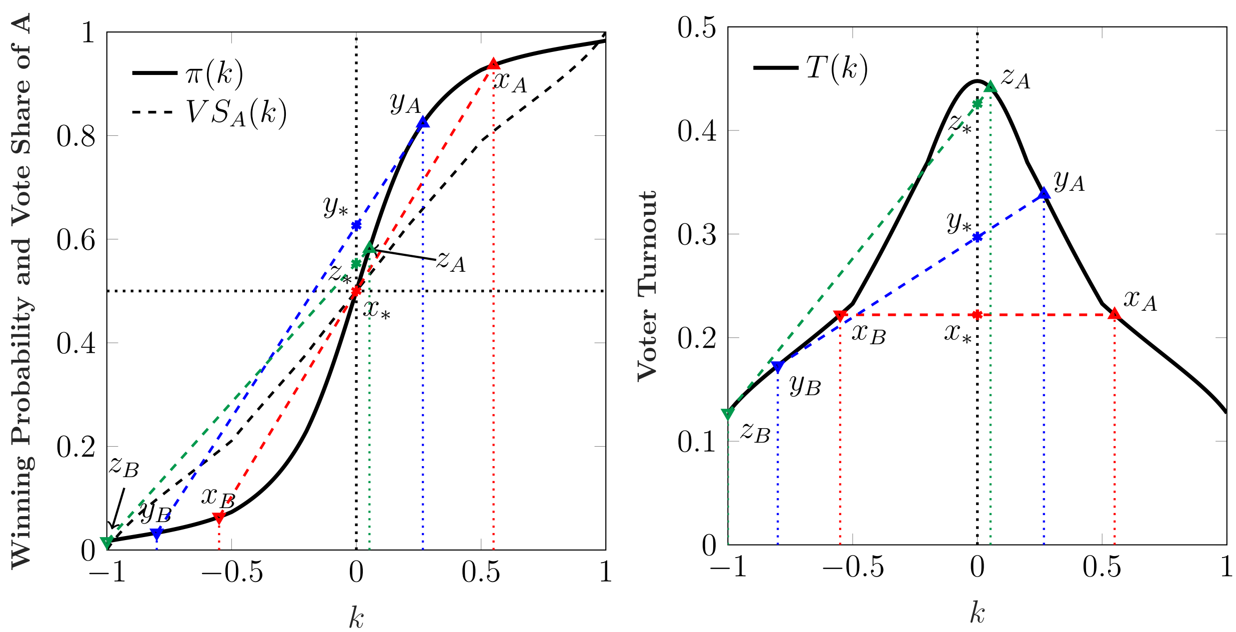

Note: The left panel shows the equilibrium functions $VS_A(k)$ and $\pi_A(k)$, and the right panel shows the equilibrium function $T(k)$. Media bias increases from red ($\chi = 0$) to blue ($\chi = 0.55$) to green ($\chi = 0.95$). w_A, w_B and w_*, $w \in \{x, y, z\}$, represent candidate A's winning chances 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 $w \in \{x, y, z\}$.

To illustrate Proposition 1, Figure 2 shows the functions $VS_A(k)$ and $\pi_A(k)$ (left panel), and $T(k)$ (right panel) for our election game. Though the equilibrium vote share increases almost linearly in k, the winning probability of candidate A is an S-shape function of k; it increases in k convexly (concavely) for $k < (>) 0$. Intuitively, This is because an increase in the vote share that follows from a shift in k is very unlikely to affect the election outcome if either party has a large expected vote share to start with. Only if the support for the two parties is balanced, can a minor change in support have a strong effect on the probability of winning. As for turnout, note that the competition effect is clearly evident in the right panel of Figure 2. Turnout is much higher when quality is expected to be equal ($k = 0$) than when the quality difference is expected to be very large. Note also that the expected voter turnout predicted by BNE remains below 50% for all k's.

Proposition 1 allows us to derive equilibrium predictions for the case where k is observable to media only. Essentially, any information obtained from the media, denoted by I, allows voters to rationally update their beliefs about k. Since voters' utility function is linear in k (eq. (1)), information affects voters' behavior only through its impact on $E[k|I]$, the (common) posterior mean of k conditional on I.[11] Consequently, voters behave as if $E[k|I]$ is common knowledge and equilibrium behavior can be derived from Proposition 1 by replacing k with $E[k|I]$. In section 4.1.1, we study the electoral influence of media bias in the scenario with a single media outlet ($|M| = 1$). In section 4.1.2, we vary $|M|$ from 1 to 2 study the electoral consequences of introducing a second media outlet.

4.1.1 The Influence of Media Bias

Let $|M| = 1$ and χ be the bias of this single media outlet. Without loss of generality assume $\chi \geq 0$, that is, the outlet is either unbiased or A-biased. In our experiment we consider three distinct $\chi \in \{0, 0.55, 0.95\}$, i.e., the media is either *unbiased* ($\chi = 0$), *weakly biased* towards A ($\chi = 0.55$) or *strongly biased* towards A ($\chi = 0.95$). Voters' (rational) posterior expectations about k, depend on both χ and the message $s \in \{A, B\}$ sent by the media. Denote these posterior expectations by $k_s(\chi) \equiv E[k|\chi, s]$ for $s \in \{A, B\}$. As can be inferred from the cutoff strategy (eq. (2); see Figure 1), ex-ante a media outlet with bias χ sends message A with probability $Pr(A|\chi) \equiv 1 - F(-\chi)$ and sends message B with probability $Pr(B|\chi) \equiv F(-\chi)$. Therefore, *a priori* the likelihood of sending a message in favor of A increases with the bias χ. Table 1 summarizes these endorsement probabilities and the Bayesian posterior expectations $k_s(\chi)$ for our experimental parameters.

As is evident from Table 1, an increase in media bias χ has three effects:

Effect I $Pr(A|\chi) = 1 - F(-\chi)$ decreases in χ.

[11] Since all voters are equally informed and have common priors, the posteriors must be common across the electorate.

Table 1: Likelihood of Endorsements and the Rational Posterior Expectations

	$\chi = 0$		$\chi = 0.55$		$\chi = 0.95$	
	$s = A$	$s = B$	$s = A$	$s = B$	$s = A$	$s = B$
$Pr(s\|\chi)$	0.5	0.5	0.75	0.25	0.95	0.05
$k_s(\chi)$	$\frac{11}{20}$	$-\frac{11}{20}$	$\frac{4}{15}$	$-\frac{4}{5}$	$\frac{1}{19}$	-1

Note: $Pr(s|\chi)$ denotes the probability that signal s is sent, conditional on bias χ (cf. Figure 1). $k_s(\chi)$ denotes voters' rationally updated beliefs about χ after receiving signal s. $F(-\chi)$ is assumed to be discrete-uniform on $[-1, 1]$ with steps of 0.1, excluding the value $k = 0$. The values for $k_s(\chi)$ use the assumption that a media that is indifferent between A and B sends signal B (cf. equation (2)).

Effect II Both $k_A(\chi)$ and $|k_A(\chi)|$ decrease in χ.

Effect III $k_B(\chi)$ decreases whereas $|k_B(\chi)|$ increases in χ.

In words, Effect I suggests that a more A-biased media is *a priori* more likely to send message A. Effect II implies that, with a higher biased towards candidate A, the media's message A becomes less credible in signaling the superiority of candidate A and implies a lower expected quality difference between candidates. Effect III suggests that message B becomes more credible in signaling the inferiority of candidate A and implies a larger expected quality difference between candidates, as the media becomes increasingly A-biased. From the ex-ante perspective, however, media bias cannot systematically affect rational voters' posterior expectations. This is because, by law of iterate expectations, voters' posterior means $\sum_{s\in\{A,B\}} Pr(s|\chi) \cdot k_s(\chi)$ must average back to the prior mean 0 (as evident from Table 1). Intuitively, even though message A (in favor of candidate A) is more likely to be sent ex-ante, message A becomes less weaker whereas the other message B becomes stronger in influencing voters' beliefs. On average, these effects balance out, leaving voters' posterior beliefs unaffected. Let $E[\xi(k)|\chi] \equiv [1 - F(-\chi)]\xi(k_A(\chi)) + F(-\chi)\xi(k_B(\chi))$ denote the expectation of outcome $\xi(k)$, with $\xi \in \{VS_A, \pi_A, T\}$, from the ex-ante perspective. Proposition 2 summarizes the interim and ex-ante impacts of an increase in media bias on party vote shares and the election outcome.

Proposition 2 *Suppose $|M| = 1$, the following properties hold*

1. $\forall\chi, \chi' \in \{0, 0.55, 0.95\}$ *such that* $\chi' > \chi$ *and* $s \in \{A, B\}$*, there are*
    (a) $VS_A(k_s(\chi)) > VS_A(k_s(\chi'))$*, and* $VS_A(0) = \frac{1}{2}$.
    (b) $\pi_A(k_s(\chi)) > \pi_A(k_s(\chi'))$*, and* $\pi(0) = \frac{1}{2}$.
2. $E[\pi(k)|0.55] > E[\pi(k)|0.95] > E[\pi(k)|0] = \frac{1}{2}$.

Proposition 2.1 claims that both the expected vote share and winning probability of candidate A decrease in media bias χ, conditional on either message $s \in \{A, B\}$. This is because both $k_A(\chi)$ and $k_B(\chi)$ decrease in χ (cf. Effect II and III). Proposition 2.1 then holds because both $VS_A(k)$ and $\pi_A(k)$ are increasing functions of k (cf. left panel of Figure 2). Intuitively, as media bias towards candidate A increases, both messages signal lower relative quality of candidate A, which causes the expected vote share and winning probability of candidate A to decrease. Since message A (B) is always "good" ("bad") news for candidate A regardless of bias χ[12], it always pushes the expected vote share and winning prospects in favor of candidate A (B) compared to the outcomes under the uninformative prior mean $k = 0$.

Although candidate A's winning probability decreases in bias χ conditional on both messages A and B (due to Effect II and III, respectively), Effect I provides a force in the opposite direction; with a higher χ the media is *a priori* more likely to send message A and induces posterior expectation $k_A(\chi)$. Because $k_A(\chi) > k_B(\chi)$ for all χ, Effect I *per se* increases candidate A's winning probability. The net impact of an increased media bias on candidate A's winning probability thus depends on which of these opposite forces dominate. For a small degree of bias (χ close to 0), Effect I dominates and a marginal increase in bias increases A's winning probability. For a large degree of bias (e.g., $\chi = 0.9$), the opposite applies. For this reason, the relationship between media bias and candidate A's ex-ante winning probability is non-monotonic, as predicted by Proposition 2.2.[13]

The interim and ex-ante impacts of media bias on the election outcome can also be straightforwardly illustrated through a geometric approach demonstrated in the left panel of Figure 2. First suppose $\chi = 0$, that is, the media is unbiased. Then $k_A(0) = 0.55$ and $k_B(0) = -0.55$ (cf. Table 1). As a result, candidate A's winning probability conditional on message A (B) equals $\pi_A(0.55)$ ($\pi_A(-0.55)$), represented by the red node x_A (x_B). Ex-ante, candidate A's winning probability is a convex combination of $\pi_A(0.55)$ and $\pi_A(-0.55)$, and must lie on the red dashed line segment connecting x_A and x_B. By 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 $x_A x_B$ and the vertical line $k = 0$.

Similar exercises can be done for $\chi = 0.55$, yielding interim winning probabilities represented by nodes y_A and y_B and an ex-ante winning probability represented by y_*. Likewise,

[12]This is because $k_A(\chi) > 0$ and $k_B(\chi) < 0$ for all $\chi \in \{0, 0.55, 0.95\}$, as is evident from Table 1.

[13]Note that we do not predict any effect of media bias on A's ex-ante vote share. This is because the vote share is close to a linear function of k (see the left panel of Figure 2). Nevertheless, precisely equilibrium calculations do predict a non-monotonicity between the expected vote share and media bias (see Figure 3, below), albeit the impacts of media bias seem small in magnitude.

for $\chi = 0.95$, the interim winning probabilities conditional on message A, B and the ex-ante winning probability are represented by nodes z_A, z_B and z_*, respectively. The interim impacts of an increase in media bias can thus be straightforwardly illustrated by the movement from x_s to y_s to z_s, for $s \in \{A, B\}$. This confirms Proposition 2.1, which predicts that candidate A's winning probability decreases as media bias increase. The ex-ante impact of increasing media bias is presented by the movement from x_* to y_* to z_*. With $y_* > z_* > x^*$, A's winning probability is highest when the media is weakly biased ($\chi = 0.55$), consistent with the rank predicted in Proposition 2.2.

Proposition 3 summarizes the interim and ex-ante impacts of media bias on voter turnout. Proposition 3.1 suggests that the interim impact depends on the message sent by the media. Conditional on message A, voter turnout increases in bias χ. This is driven by Effect II; $|k_A(\chi)|$ decreases in χ and voter turnout $T(k)$ is a decreasing function of $|k|$ (cf. right panel of Figure 2). Intuitively, although message A always indicates the superiority of candidate A, it becomes weaker and signals a smaller expected quality difference, as media bias increases. As explained above, a smaller expected quality difference drives the expected vote share (for either candidate) closer to 50%. As a result, conditional on message A, electoral competition becomes fiercer and expected voter turnout is higher, with increased bias χ. Based on a similar argument (cf. Effect III), voter turnout decreases in media bias conditional on message B. Intuitively, message B reveals a larger expected quality difference between candidate, as the media becomes increasingly A-biased. Consequently, the expected vote share will be driven away from 50%, leaving the electoral competition less fierce. Voter turnout will then be lower because voting is less likely to be pivotal.

Proposition 3 *Suppose* $|M| = 1$, $\forall \chi, \chi' \in \{0, 0.55, 0.95\}$ *such that* $\chi' > \chi$ *it holds that*

1. $T(k_A(\chi')) > T(k_A(\chi))$ *and* $T(k_B(\chi')) < T(k_B(\chi))$.

2. $E[T(k)|\chi'] > E[T(k)|\chi]$.

3. *Ex-ante, the expected voter turnout is highest without information.*

Although the interim impact of media bias on voter turnout is message-dependent, the ex-ante impact is not; Proposition 3.2 suggests that *a priori* the expected voter turnout increases unambiguously with media bias. This is because with a higher bias χ the media is *a priori* more likely to send message A (Effect I), which indicates a smaller expected quality difference (Effect II) and hence boosts voter turnout. Although message B suggests a larger expected quality difference and leads to lower turnout as χ increases, it is never strong enough to cause an ex-ante reduction in expected voter turnout. Finally, when voters are uninformed, they behave completely based on their (common) prior belief, which indicates no quality difference between candidates. As a consequence, the expected vote share will precisely equal to 50% and the electoral competition is fiercest. Voter turnout will then be highest because voting is most likely to be pivotal and voters have strongest incentives to cast costly votes.

The previously adopted geometric approach can also be applied to derive Proposition 3 (cf. right panel of Figure 2). Similar to the previous analyses with election outcomes, the interim impact conditional on message A and B, as well as the ex-ante impact of media bias on voter turnout, are geometrically represented by w_A, w_B and w_*, respectively, for $w \in \{x, y, z\}$ (again, corresponding to $\chi \in \{0, 0.55, 0.95\}$). The interim impact conditional on message A (B) again follows from the movement from x_A to y_A to z_A (x_B to y_B to z_B). The ex-ante impact is reflected by the movement from x_* to y_* to z_*. Proposition 3 follows evidently from this geometrical exercise.

4.1.2 The Influence of Introducing a Second Media Outlet

We investigate the ex-ante impacts of introducing a second media outlet on election outcomes and voter turnout. Without loss of generality, let the biases of the existing media (outlet 1) and second media (outlet 2) be χ_1 and χ_2. In our experiment, $\chi_1 \neq \chi_2$ and both χ_1 and χ_2 take values from $\{-0.55, 0, 0.55, 0.95\}$. Under the cutoff endorsement strategy (2) (cf. Figure 1), outlet 1 partitions the state space[14] into two subsets, depending on whether the realized k lies above or below $-\chi_1$. Having a second media outlet refines the information partition; it partitions the state space into three intervals, depending on whether $k > \max\{-\chi_1, \chi_2\}$, $k \leq \min\{-\chi_1, -\chi_2\}$ or k lies in between. Essentially, introducing a second media affects the election outcomes and voter turnout entirely through its influence on the distribution of voters' posterior beliefs. These predictions are summarized in Propositions 4 and 5 below.

Proposition 4 *Suppose* $|M| = 2$ *and* $\chi_1, \chi_2 \in \{-0.55, 0, 0.55, 0.95\}$, *then* $E[\pi_A(k)|\chi_1, \chi_2] > (<) E[\pi_A(k)|\chi_1]$ *if* $\chi_2 > (<) \chi_1$.

Proposition 5 *Suppose* $|M| = 2$ *and* $\chi_1, \chi_2 \in \{-0.55, 0, 0.55, 0.95\}$, *it holds that*

1. *If* $(\chi_1, \chi_2) = (0, 0.55)$, *then* $E[T(k)|\chi_1, \chi_2] > E[\pi(k)|\chi_1]$.

2. *If* $(\chi_1, \chi_2) = (0.55, -0.55)$, *then* $E[T(k)|\chi_1, \chi_2] > E[\pi(k)|\chi_1]$.

3. *If* $(\chi_1, \chi_2) = (0.95, 0)$, *then* $E[T(k)|\chi_1, \chi_2] < E[\pi(k)|\chi_1]$.

[14] The state space refers to the set of all possible values of k, which equals $\{-1, -0.9, \cdots, -0.1, 0.1, \cdots, 0.9, 1\}$ (cf. Section 3).

In words, Proposition 4 claims that introducing media outlet 2 can increase (decrease) candidate A's ex-ante winning probability if and only if outlet 2 is more (less) biased towards candidate A compared to the existing outlet 1. Proposition 5 suggests that how introducing a second media outlet on ex-ante voter turnout depends on the biases of both media outlets in more subtle ways. If the existing media is unbiased ($\chi_1 = 0$), introducing a biased outlet unambiguously increases the ex-ante voter turnout (Proposition 5.1). If instead the existing media is weakly biased, introducing a second outlet that is equally biased but in favor of the opposite candidate can systematically increase the expected voter turnout (Proposition 5.2). Finally, if the existing media is strongly biased, then introducing an unbiased outlet can systematically decrease voter turnout (Proposition 5.3).[15]

4.2 Quantal Response Equilibrium

Various studies on voter behavior have shown that Quantal Response Equilibrium (QRE; see McKelvey and Palfrey (1995)) better describes voter behavior than BNE (Goeree and Holt, 2005; Grosser and Schram, 2006, 2010; Tyszler and Schram, 2016). To derive QRE we first denote a voter's expected utility of voting for candidate A, B or abstaining (O), conditional on model parameters (k, v, c), by $EU^A(k, v, c)$, $EU^B(k, v, c)$ and $EU^O(k, v, c)$, respectively. It is without loss of generality to normalize $EU^O(k, v, c)$ to 0. QRE allows voters to make errors in their voting decisions. This can be modeled by adding a stochastic error term to expected utilities, which yields:

$$\begin{aligned} &EU^A(k, v, c) + \frac{\varepsilon_A}{\lambda} \\ &EU^B(k, v, c) + \frac{\varepsilon_B}{\lambda} \\ &EU^O(k, v, c) + \frac{\varepsilon_O}{\lambda} = \frac{\varepsilon_O}{\lambda} \end{aligned} \tag{3}$$

In (3), ε_d for $d \in \{A, B, O\}$ denote i.i.d. noise and $\lambda > 0$ is a noise parameter that captures the relative weight of noise and utility. As $\lambda \to 0$, noise dominates and as $\lambda \to +\infty$, noise disappears. This parameterization is general enough to capture different sources of noise, such as distractions, perception biases, miscalculations or limited computational capability (Goeree and Holt, 2005). Note that we assume the same noise parameter applies for all the three different actions; voting for A, voting for B, or abstaining. We will return to this assumption when discussing our results.

In QRE, a voter will still choose the action that yields the highest expected utility,

[15] The precise mechanisms driving Proposition 4 and 5 are related to the curvatures of functions $\pi_A(k)$ and $T(k)$, respectively. Interested readers can refer to our companion paper SSS18 for more detailed explanations.

but this is now a stochastic event. For example, she will vote for A if $EU^A(k,v,c) + \frac{\varepsilon_A}{\lambda} > EU^B(k,v,c) + \frac{\varepsilon_B}{\lambda}$ and $EU^A(k,v,c) + \frac{\varepsilon_A}{\lambda} > \frac{\varepsilon_O}{\lambda}$ holds simultaneously. Or equivalently, $\varepsilon_B - \varepsilon_A < \lambda \cdot \big(EU^A(k,v,c) - EU^B(k,v,c)\big)$ and $\varepsilon_O - \varepsilon_A < \lambda \cdot EU^A(k,v,c)$.[16] 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. Assuming that the ε_d follows the extreme value type 1 distribution for all $d \in \{A, B, O\}$, the (multinomial) probabilities are given by:

$$\begin{aligned}
p^A(k,v,c|\lambda) &= \frac{e^{\lambda EU^A(k,v,c)}}{e^{\lambda EU^A(k,v,c)} + e^{\lambda EU^B(k,v,c)} + 1} \\
p^B(k,v,c|\lambda) &= \frac{e^{\lambda EU^B(k,v,c)}}{e^{\lambda EU^A(k,v,c)} + e^{\lambda EU^B(k,v,c)} + 1} \\
p^O(k,v,c|\lambda) &= \frac{1}{e^{\lambda EU^A(k,v,c)} + e^{\lambda EU^B(k,v,c)} + 1}
\end{aligned} \tag{4}$$

The expected utilities on the right hand side of (4) depend on the probabilities of being pivotal, which in turn depends on the probabilities on the left hand side of the equations. For any λ, the set of equations (4) can then be solved to obtain the QRE by finding the fixed point. For more details, see online Appendix B.

To derive the QRE predictions, we need to specify λ. We obtained out-of-sample maximum-likelihood estimates for this parameter using data from a previous experiment with an electorate of size 11 (as opposed to the 25 used here) conducted in the summer of 2016.[17] Figure 3 shows the BNE and QRE predictions for the ex-ante election outcome and voter turnout with $|M| = 1$; these are point predictions conditional on the realizations of all the model parameters in the 2017 experiment, which is introduced in Section 5.

Note that regarding candidate A's expected vote share and winning probability, the BNE and QRE predictions are quite similar. The only exception is the comparison between the scenario with $\chi = 0.95$ and the scenario with no information; BNE predicts a decrease in candidate A's expected vote share and winning probability, where QRE predicts the opposite. The non-monotonic relationship predicted in Proposition 2.2 is also predicted by QRE. Differences are larger for (ex ante) predicted turnout. The comparative statics predictions regarding the expected voter turnout (Proposition 3.2 and 3.2) are also valid under QRE, though turnout levels predicted by QRE are substantially higher than the levels predicted by BNE. In short, all but one of the comparative statics predictions under BNE are still valid under QRE.

[16] Here, we assume that the decision to vote or abstain and the party chosen in the former case are made simultaneously. Our online Appendix B discusses an alternative model in which these decisions are made sequentially.

[17] The estimated λ's approximately equal 19.45 for $|M| = 1$ and 17.97 for $|M| = 2$. Further details about our previous experiments and the QRE predictions are available from the authors upon request.

Figure 3: The Comparisons between QRE and BNE Predictions

(a) Vote Share of A (b) Winning Probability of A (c) Voter Turnout

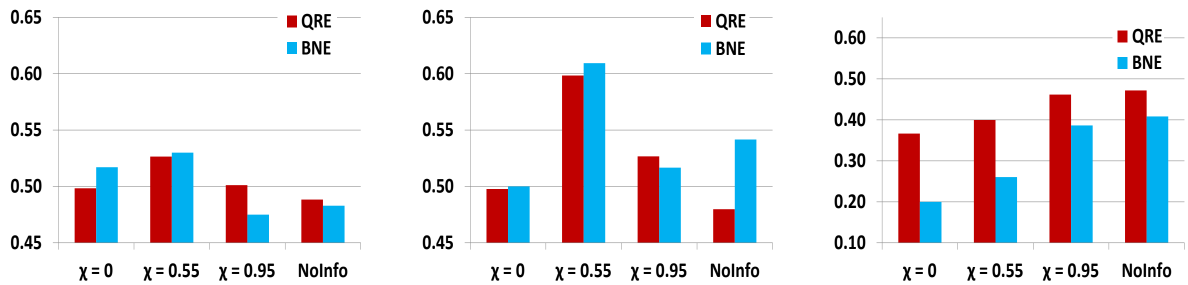

Note: These panels show the BNE (blue bars) and QRE (with noise parameter λ estimated out-of-sample; red bars) predictions for the ex-ante vote share of A (left), winning probability of A (central), and voter turnout (right). In the horizontal axis, "NoInfo" denotes the scenario where voters are making decisions without any information. All predictions are generated based on the realized model parameters drawn for the real experiment.

5 Experimental Design and Hypotheses

5.1 Frame

To test our key theoretical predictions, we designed a laboratory experiment that implements the two candidate, 25-voter election game presented in Section 3. The candidates are represented by vases labeled A and B. Each voter is independently and equally likely to be assigned to either team A or team B. If the elected vase matches a voter's team (e.g, vase A is elected and the voter belongs to team A), the voter receives a private bonus of either 20, 50 or 100 "tokens".[18] The bonus is determined randomly with equal probability and independently across voters and elections. If the elected vase does not match the voter's team, the voter receives no private bonus. The "partisan preference" that this induces is thus either *centralist* (20 tokens), *moderate* (50 tokens) or *extreme* (100 tokens).

To induce the candidate qualities, we allow voters to receive additional earnings from so called "diamonds", depending on which vase is elected. At the beginning of each election, one of the two vases is randomly chosen (with equal probability) and filled with k "diamonds", where k is a random draw from a discrete-uniform distribution on multiples of ten between 10 and 100. Δk, the difference between the numbers of diamonds in vase A and B is then discrete-uniformly distributed on multiple of tens between -100 and 100 (except 0) and can be interpreted as the relative quality of vase A. The other vase has no diamonds. If the vase filled with diamonds wins the election, each voter irrespective of their partisan preference earns

[18]Note that we convert the fractions in the model to whole numbers by multiplying all numbers by 100. In the experiment each 100 tokens are exchanged to 2 euros or, equivalently, each token is worth 2 eurocents.

k tokens, on top of any private bonus they might have earned. Consistent with our model, voters are not directly informed of diamond allocations. However, voters may in some cases receive (imperfect) information about diamond allocations from one or two media outlets, who are framed as "robots". After observing the public message(s) voters simultaneously make their voting decisions. The election outcome is determined by simple majority rule, with ties broken by a coin toss. Finally, for each voter, casting a vote induces a private cost, independently drawn from a discrete-uniform distribution on integers between 1 and 15 tokens.

5.2 Treatments

We distinguish between seven treatments. In a benchmark ($NoInfo$), voters are completely uninformed about the diamond allocations. In the other six treatments, there are either one ($|M| = 1$) or two media outlets ($|M| = 2$) providing information about diamond allocations to voters. These outlets ("robots") are computerized and communicate to voters by sending binary public messages $s \in \{A, B\}$ to all voters with the following pre-specified strategy. In line with our theoretical model (cf. Figure 1), a robot is characterized by a parameter χ, sends message A if and only if $\Delta k \geq -\chi$, and sends message B otherwise. In treatments with only one robot ($|M| = 1$), we distinguish between an unbiased robot (UB, $\chi = 0$), a weakly biased robot (WB_A, $\chi = 55$) and a strongly biased robot (SB_A, $\chi = 95$). In the instructions, we label these robots as "ALPHA", "BETA" and "GAMMA", respectively.

In treatments with two media outlets ($|M| = 2$), we let the combination of media biases be (UB, WB_A), (WB_A, WB_B) or (SB_A, UB), where a subscript denotes the candidate to which the media is biased. We label the additional robot as "DELTA", with $\chi = -55$. The public messages sent by the robot(s) give voters coarse information about the diamond allocations. Table 2 shows the table format used to inform subjects about the robots' strategies.

We implement a 2×4 factorial design where we use between-subject comparisons in the dimension of the number of media outlets ($|M| = 1, 2$), and within-subject comparisons in the dimension of media bias ($NoInfo$, UB, WB_A, SB_A, or $NoInfo$, (UB, WB_A), (WB_A, WB_B), (SB_A, UB)). Within each electorate, the 25 subjects (voters) played 24 rounds of elections for each of the four media bias treatment conditions for the $|M|$ concerned (for a total of 96 elections per session). To facilitate direct comparisons we used the same realized parameter draws for the 24 elections (including realizations of diamonds allocations, voters' team assignment, private bonus and voting costs) across all treatments and sessions.

To avoid order effects, we randomly divided the 24 elections for each treatment into three distinct blocks, each containing eight consecutive elections. This gives 12 blocks of elections

Table 2: Robot Strategies

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

Note: This table was presented to subjects in the experimental instructions (cf. online Appendix D). They were also given a copy in print to have available throughout the experiment. Shaded cells indicate cases where the robot sends message **A**; unshaded cells indicate cases where the robot sends message **B**. Robots ALPHA, BETA, and GAMMA represent, respectively, an unbiased, weakly biased, and strongly biased media. DELTA is a media outlet weakly biased in favor of B and was only included in the sessions with $|M| = 2$.

for each electorate, three for each media bias treatment. We then randomized the order of these blocks independently for each electorate in the following manner: each set of 4 blocks (block $1-4$, block $5-8$, and block $9-12$) included one block from each of the 4 distinct treatments. The orders of these blocks were completely randomized. Aside from controlling for order effect, this allows us to investigate learning effect by comparing subsamples of elections in later blocks to their early counterparts. In summary, a subject always faces the same $|M|$. She starts with eight elections facing one randomly chosen (pair of) media bias. This is followed by eight elections with different media bias. She faces all four bias treatments before returning to a previously faced one for another eight elections.

5.3 Belief Elicitation

In the last election of each block with robot(s) (UB, WB_A, SB_A) or ((UB, WB_A), (WB_A, WB_B), (SB_A, UB)), we elicit subjects' posterior beliefs about candidates' relative qualities conditional on each possible message (combination) being sent. As explained below, this allows us to test whether subjects' beliefs are systematically affected by media bias. More specifically, we elicit subjects' beliefs about the probability that they will receive strictly higher earnings if vase A is elected than if vase B is elected, conditional on the message(s) $s \in \{A, B\}$ sent by the robot(s). Note that a Bayesian voters' average posterior means do not vary with media bias.[19] This allows us to directly use reported beliefs to test for persuasion bias (which we report in Section 7.2).

To elicit these beliefs, we ask: "How likely do you think that vase A, if elected, pays you more than vase B, if the robot X sends message s?".[20] Depending on the treatment, X can be ALPHA, BETA or GAMMA. We adopt the strategy method for this task (subjects report their beliefs for each of the two possible signals) and ask subjects to complete the task before they start the election game.

To incentivize subjects to report their beliefs truthfully, we use the choice list approach (Andersen et al., 2006) by asking subjects to fill out choice lists as shown for $|M| = 1$ in Figure 4. In each choice list, the left option (Option 1) is a lottery that pays ten euros if and only if vase A pays strictly more than vase B. The right option (Option 2) is a lottery that pays the subject ten euros with probabilities varying from 0% to 100% (with increments of

[19] A pays more than B if and only if $k + v_i > 0$. The elicited posterior belief is thus interpreted as $Pr[k + v_i > 0|s, \chi]$. If voters are Bayesian, then $Pr[k + v_i > 0|\chi] = Pr[k + v_i > 0|A, \chi] \cdot Pr[s = A|\chi] + Pr[k + v_i > 0|B, \chi] \cdot Pr[s = B|\chi] \equiv Pr[k + v_i > 0]$.

[20] This is the text we use under $|M| = 1$. Under $|M| = 2$, the text depends on whether the two robots send the same message. If so, the text reads "How likely do you think that vase A pays you more than vase B, if both robot X and robot Y send message s?" If they send different messages, the text reads "How likely do you think that vase A pays you more than vase B, if robot X sends message B and robot Y sends message A?".

5%-point).

Figure 4: The sample choice list for the estimation task

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

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

OK

Note: If robot ALPHA sends message **A** in this election, then one of the 21 choices from this list will be randomly selected to determine the relevant lottery for you. All these choices are equally likely to be selected.

Note: In the final election of each block with robots (for a total of nine times), each participant filled out two choice lists (one for message A, one for message B) like the one depicted in this screen shot.

In each belief elicitation round, subjects are asked to fill out multiple such lists, one for each possible message (combination) sent by the robot(s). We impose choice consistency to avoid multiple switching points and reduce subjects' burden by allowing them to complete the choice list by simply making one click.[21] After subjects have made their decisions, the relevant list (that matches the actual message(s) 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 played out. This mechanism is incentive compatible; it guarantees that subjects have the highest chance of winning the ten-euro prize if they report their beliefs truthfully (Andersen et al., 2006). At the end of the experiment one out of the (nine) tasks is chosen (for each subject separately) to determine whether she earns the ten-euro prize.

[21] If a subject ticks Option 1 for choice X (a number between 1 to 21), then the computer automatically ticks Option 1 for all choices $Y < X$ and Option 2 for all choices $Z > X$. If the subject ticks Option 2 for choice X, then the computer automatically ticks Option 1 for all choices $Y < X$ and Option 2 for all choices $Z > X$.

5.4 Procedure

The experiment was programmed in PHP-MYSQL and run at the CREED laboratory of the University of Amsterdam in the summer of 2017. A total of 300 subjects were recruited through the CREED recruitment database. In total, we have data for 12 electorates, six for each of the treatments $|M| = 1$ and $|M| = 2$. Each session consisted of exactly one electorate. Random assignment of subjects across treatments was successful; descriptive statistics (Table C.1 in online Appendix C) show no evidence that subjects' background characteristics differ between the two treatment groups ($|M| = 1, 2$). Sessions took on average 150 minutes when there was one media and 200 minutes when there were two. Average payment was about 35.7 euros (with a minimum of 18.9 euros and a maximum of 53 euros).[22]

At the beginning of each session, subjects read instructions (also distributed in a handout) on screen and have to answer a set of control questions correctly before they can start with the experiment. At the end of the experiment, they are invited to fill out a post-experimental questionnaire that includes questions about their decision making. The full texts of instructions, control questions and post-experiment questionnaire can be found in online Appendix D. Subjects' payments composed of 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.

5.5 Hypotheses

Propositions 2 and 3 directly yield three (composite) hypotheses regarding the influence of media bias that can be tested with our experimental data.

Hypotheses 1 *Influence of media bias on the election outcome (Proposition 2):*

(a) Conditional on message A, both candidate A' expected vote share and winning probability are highest in UB and lowest in SB_A.

(b) Conditional on message B, both candidate A' expected vote share and winning probability are highest in UB and lowest in SB_A.

(c) Ex-ante, candidate A's winning probability is highest in WB_A and lowest in UB.

Hypotheses 2 *Influence of media bias on voter turnout (Proposition 3):*

[22] This includes a show-up fee of ten euros. Because the sessions with two media were expected to take longer (which, in fact, they did), we gave subjects an ex-post (and unannounced) additional ten euros.

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

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

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

Regarding the influence of introducing a second media outlet, Proposition 4 and 5 yield two additional sets of hypotheses to be tested.

Hypotheses 3 *Influence of introducing a second media outlet on the ex-ante election outcome (Proposition 4):*

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

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

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

Hypotheses 4 *Influence of introducing a second media outlet on the ex-ante voter turnout (Proposition 5):*

(a) *Ex-ante, voter turnout is higher in treatment (UB, WB_A) than in UB.*

(b) *Ex-ante, voter turnout is higher in treatment (WB_A, WB_B) than in WB_A.*

(c) *Ex-ante, voter turnout is higher in treatment (SB_A, UB) than in SB_A.*

We will test Hypotheses 1 to 4 non-parametrically in the following section.

6 Results

We present the results in three subsections. First, we present an aggregate overview of subjects' voting behavior in our experiments. This is followed by statistical tests for the hypotheses presented in the previous section. Finally, we investigate the extent to which behavior in our experiments can be explained by our theoretical predictions. We use data from the full sample of 12 blocks (of eight elections each). All tests used are two-sided even though many of our hypotheses are directional. The results may therefore be seen as a conservative approach to testing what we predict. Unless explicitly stated otherwise, tests are non-parametric and use the electorate (session) as the unit of independent observation.

6.1 Aggregate Results

We start with the one media treatment. Figure 5 shows the aggregate results for A's vote share (top panels), A's probability of winning (middle panels) and turnout (lower panels). We distinguish between results conditional on message $s = A$ (left panels), message $s = B$ (central panels) and the ex-ante results (right panels; these are the weighted average of the left and central panels based on the ex-ante probabilities of sending each message).

Figure 5: Aggregate Results, $|M| = 1$

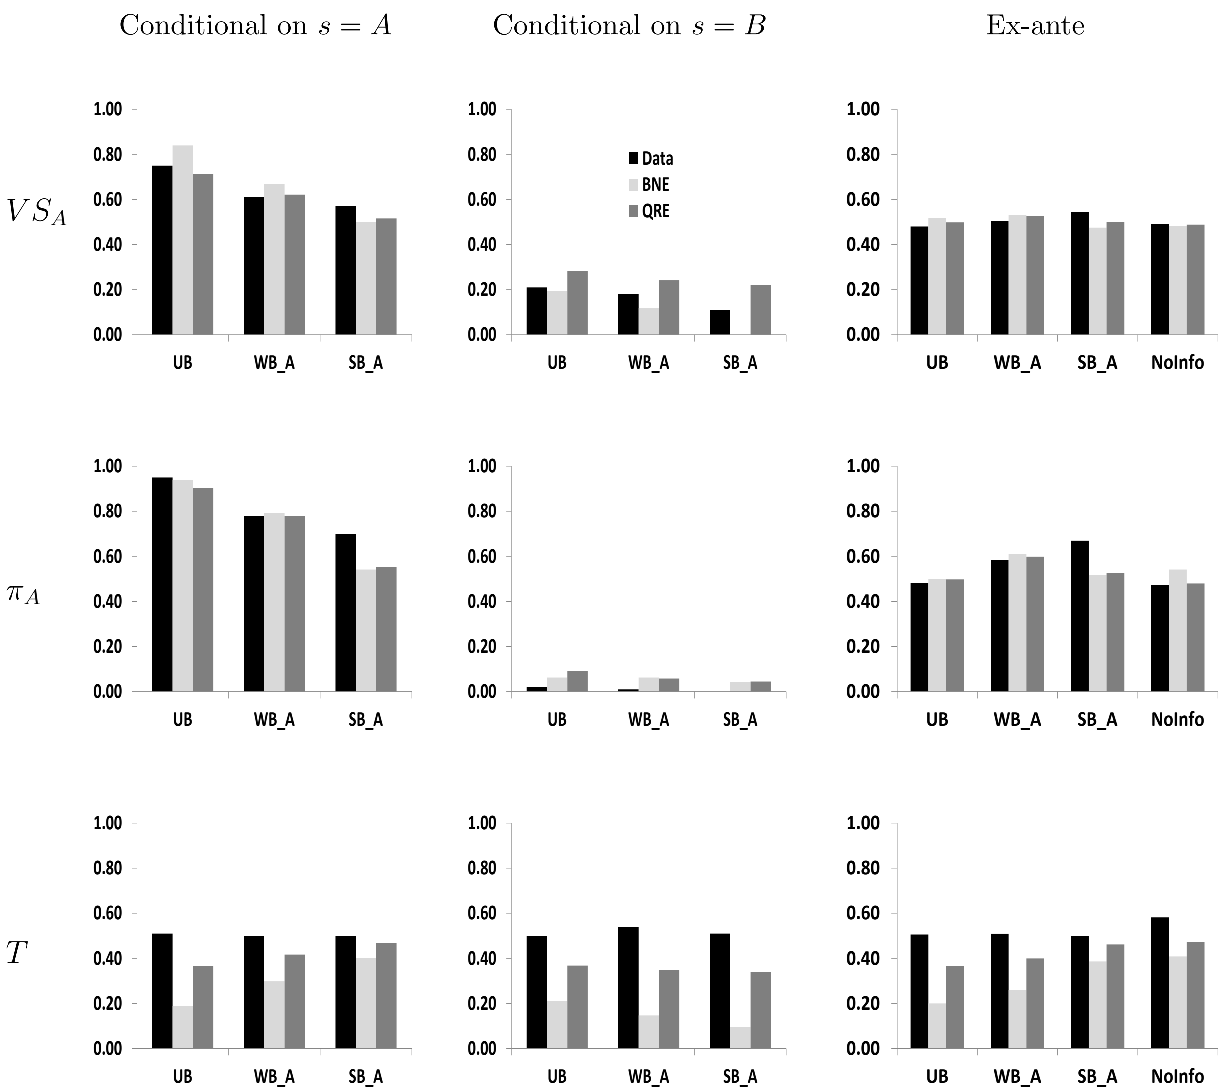

Note: VS_A denotes the vote share of candidate A; π_A denotes the probability that A wins the election (in the data this is measured by the fraction of elections won by A, where ties are computed as a 0.5 score each); T denotes voter turnout. The black bar denotes the actual observed outcomes. The light and dark grey bars denote the corresponding BNE and QRE predictions, respectively.

A bird eye's view of Figure 5 shows that the election outcome (measured by VS_A and

π_A) is better predicted by BNE than turnout is (measured by T). Nonetheless, there are substantial deviations. First, BNE predicts a non-monotonic relationship between media bias and candidate A's winning probability. Instead, our data show a monotonic increase as media bias increases from UB to WB_A, yet the outcomes remain more or less constant from WB_A to SB_A. Second, BNE predicts substantial treatment effects of media bias on turnout (both conditional on the signal and unconditional). The data, however, seem to show that voter turnout responds little to variations in media bias, with only two exceptions. In the bottom-left panel of Figure 5, the observed voter turnout is highest in SB_A, though the difference is small. The bottom-right panel shows that turnout in the treatment without media ($NoInfo$) seems higher than in the media treatments. Formal tests of the relationships between predictions and observations will be presented in the next subsection.

Figure 6: Aggregate Results, $|M| = 2$

$(\chi_1, \chi_2) = (0, 0.55)$ $(\chi_1, \chi_2) = (0.55, -0.55)$ $(\chi_1, \chi_2) = (0.95, 0)$

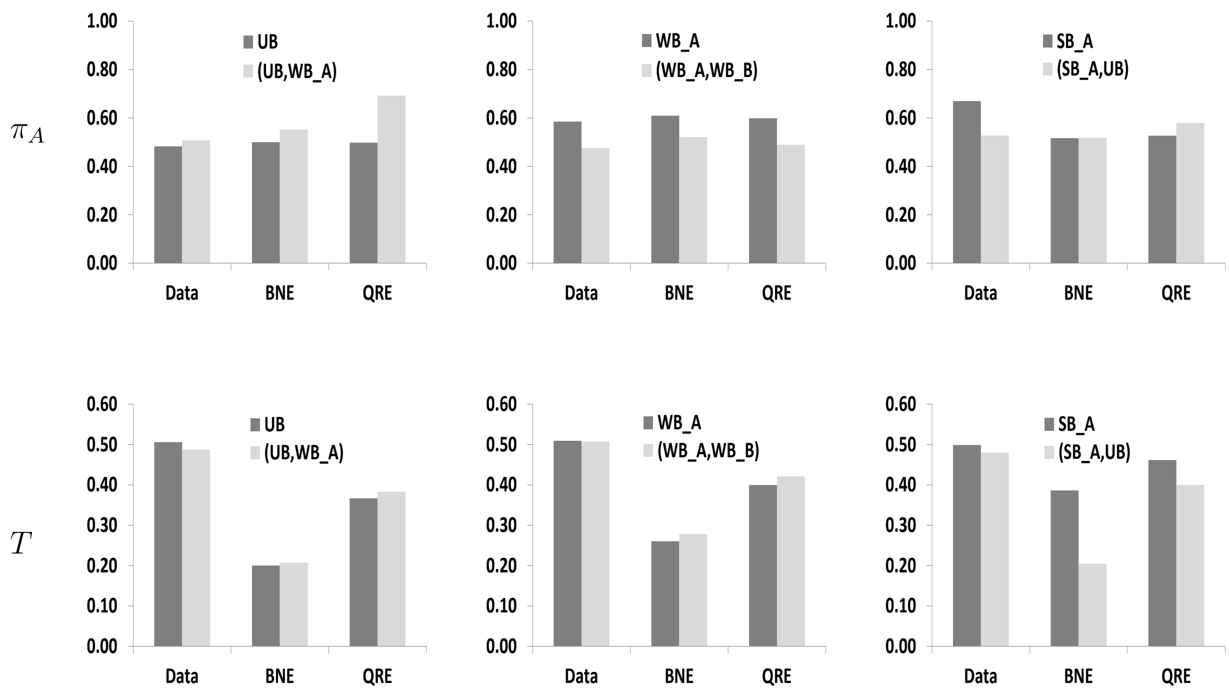

Note: Bars show ex-ante winning probability of candidate, or voter turnout for the treatment concerned and the corresponding predictions from BNE and QRE models. χ_1 and χ_2 represents the biases of the preexisting and entrant media outlet, respectively. χ_1 and χ_2 can take values from 0 (UB), 0.55 (WB_A), 0.95 (SB_A) and -0.55 (WB_B).

When there are two media outlets, our main interest lies in the impacts of introducing a second media outlet on the election outcomes (measured by π_A) and voter turnout (measured by T), from the ex-ante perspective. Figure 6 depicts the predicted and observed election outcome and voter turnout for the three two-media configurations we ran. For ease of comparison, we add the observed outcomes with the corresponding one-media case to which we will compare each configuration.

As is evident from the top panels of of Figure 6, BNE predictions correctly captures, both qualitatively and quantitatively, the comparative statics of the winning probability of candidate A after media entry. the bottom panels of Figure 6 show that voter turnout, however, varies only slightly with media entry, though in the direction predicted by BNE.

6.2 Testing Hypotheses

We start with Hypotheses 1, which concerns the impact of media bias on candidate A's expected vote share and winning probability. These hypotheses predict that both VS_A and π_A increase in media bias conditional on either message A (H1a) or B (H1b). The top-left panel of Figure 5 shows that, conditional on message A, the vote share of candidate A indeed decreases monotonically in media bias, as we predicted. This decrease is statistically significant (Friedman test –henceforth, Ft–, $p = 0.002$, $N = 6$). Similarly, the top-central panel shows that A's expected vote share decreases in media bias conditional on message B. The decrease is also statistically significant (Ft, $p = 0.006$, $N = 6$). For candidate A's winning probability conditional on message A (middle-left panel), we again observe a monotonic decrease as media bias increases, and this is also significant (Ft, $p = 0.004$, $N = 6$). However, candidate A's winning probability conditional on message B seems unresponsive to media bias (middle-central panel), and the effect is indeed statistically insignificant (Ft, $p = 0.607$, $N = 6$). This last result might be attributed to a boundary effect. As both predicted and observed probability of winning are very close to 0 for all levels of media bias, any necessarily small treatment effects are likely to be dominated by noises.

H1c predicts that candidate A's ex-ante winning probability varies non-monotonically in media bias (being highest in WB and lowest in UB). As the middle-right panel of Figure 5 shows, the observed ex-ante winning probability monotonically increases from UB to WB_A (Wilcoxon signed rank –henceforth W–, $p = 0.031$, $N = 6$), and from WB_A to SB_A (W, $p = 0.031$, $N = 6$). Overall, candidate A's ex-ante winning probability does vary significantly with media bias (Ft, $p = 0.002$, $N = 6$). These results partly support H1c, which predicts A's winning probability to increase from UB_A to WB_A but to decrease from WB_A to SB_A. In summary, our data yield mixed support for the three parts of hypotheses 1. This is formally summarized in Result 1.

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

1. *Conditional on message A, candidate A's vote share and winning probability both decrease in media bias.*

2. *Conditional on message B, candidate A's vote share decreases in media bias, but A's probability of winning is irresponsive to such decreases.*

3. *Ex-ante, candidate A's winning probability increases monotonically in media bias.*

Next we examine hypotheses 2, which concerns the impact of media bias on voter turnout. First, H2a (H2b) predicts that conditional on message A (B), voter turnout increases (decreases) with media bias. Observed voter turnout are depicted in the bottom-left and bottom-central panels of Figure 5, respectively. The bottom-left panel shows that, conditional on message A, voter turnout remains almost the same across all levels of media bias; the influence of media bias is indeed statistically insignificant (Ft, $p = 0.957$, $N = 6$). Similarly, conditional on message B (bottom-central panel), voter turnout also seems to vary little with media bias; the differences are indeed statistically insignificant (Ft, $p = 0.200$, $N = 6$). H2c predicts that ex-ante, voter turnout is increasing in media bias and it is highest in the $NoInfo$ treatment. The bottom-right panel of Figure 5 suggests that voter turnout is indeed higher in $NoInfo$ than in other treatments with a single media outlet[23]. However, voter turnout seems to respond little to media bias (Ft, $p = 0.311$, $N = 6$). These are formally summarized in Result 2

Result 2 *Influence of media bias on voter turnout:*

1. *Conditional on message A, voter turnout is irresponsive to media bias.*

2. *Conditional on message B, voter turnout is irresponsive to media bias.*

3. *Ex-ante, voter turnout is highest in the $NoInfo$ treatment and is irresponsive to variations in biases in treatments with a single media outlet.*

Hypotheses 3 concern the ex-ante impacts of introducing a second media outlet on the election outcome (i.e., winning probability of candidate A). H3a predicts that introducing an outlet weakly biased towards candidate A will increase A's ex-ante winning probability, when the preexisting outlet is unbiased. We test this by comparing candidate A's ex-ante winning probability under treatment UB to the corresponding outcome in treatment (UB, WB_A). Note that, contrary to the previous tests, this is a between-subject comparison because we vary the number of media outlets across sessions. We do observe a slight increase in A's

[23]The p values of Wilcoxon signed rank tests for comparisons between $NoInfo$ and UB, WB_A, and SB_A, are all 0.031 ($N = 6$).

winning probability (from 48.3% to 50.7%) in the top-left panel of Figure 6, the difference is statistically significant (Mann Whitney –henceforth, MW–, $p = 0.0.039$, $N = 12$). H3b predicts that, when the existing outlet is weakly biased in favor of A, introducing an opposite-minded outlet with symmetric bias decreases candidate A's winning probability. Indeed, we observe a sizable decrease (from 58.5% to 47.6%) in the top-central panel of Figure 6, and the difference is statistically significant (MW, $p = 0.013$, $N = 12$). Finally, H3c predicts that, if the preexisting outlet is strongly biased towards candidate A, then introducing an unbiased outlet decreases A's winning probability. Again, we observe a sharp decrease (from 66.9% to 52.7%) in the top-central panel of Figure 6, and the difference is statistically significant (MW, $p = 0.002$, $N = 12$). Nevertheless, conditional on the realizations of model parameters for the experiment, both BNE and QRE predictions actually violate H3c; they predict slight treatment effects in the other direction, if any. These results are formally summarized in Result 3.

Result 3 *Influence of introducing a second media outlet on candidate A's ex-ante winning probability:*

1. *A's winning probability is significantly lower in treatment* (UB, WB_A) *than in* UB.
2. *A's winning probability is significantly lower in treatment* (UB, WB_A) *than in* UB.
3. *A's winning probability is significantly lower in treatment* (SB_A, UB) *than in* SB_A.

Hypotheses 4 concern the ex-ante impacts of introducing a second media outlet on voter turnout. H4a predicts that if the preexisting media is unbiased, introducing a biased outlet leads to an increase in voter turnout. Contrary to H3a, our results (bottom-left panel of Figure 6) show a slight decrease in turnout (from 50.6% to 48.8%). This is statistically insignificant, however (MW, $p = 0.699$, $N = 12$). H3b predicts that if the preexisting media outlet is weakly biased (WB_A), the introducing an opposite-minded media with a symmetric bias (WB_A, WB_B) can systematically increase voter turnout. Contrary to H3b, ex-ante voter turnout varies little with (from 44.9% to 49.0%) as shown in the bottom-central panel of Figure 6. The difference is statistically insignificant (MW, $p = 0.937$, $N = 12$). Finally, H3c predicts that if the preexisting media is sufficiently biased (as in SB_A), then introducing an unbiased media (as in (SB_A, UB)) decreases voter turnout. The slight decrease that we indeed observe (from 49.9% to 48.0%, see the bottom-right panel of Figure 6) is again statistically insignificant (MW, $p = 0.461$, $N = 12$). Overall, we find no evidence that introducing a second media outlet affects average voter turnout. This is summarized in Result 4.

Result 4 *Influence of media entry on voter turnout: For all the combinations of media biases we consider, we don't find significant treatment effects of media entry on the ex-ante voter turnout.*

6.3 The Predictive Power of BNE

In this section we evaluate the predictive power of our theory. For this purpose, we regress the observed outcomes on the corresponding theoretical (BNE) predictions.

$$Y_{is} = \beta_0 + \beta_1 BNE_{is} + X_s\gamma + BNE_{is} \times X_s\delta + \varepsilon_{is} \tag{5}$$

In (5), Y_{is} and BNE_{is} are the observed and corresponding BNE 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 Economics students) 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 (i.e., standard errors are clustered at the electorate level). Note that the interaction terms in (5) allow for the impacts of BNE predictions on observed outcomes to be different across distinct compositions of the electorate. Table 3 presents the OLS regression results for (5).

Models (1) and (3) show that variations in the BNE predictions alone can explain 72.4% (76.0%) of the variations in the observed vote share for A in $|M| = 1$ ($|M| = 2$), while (5) and (7) show that they explain 58.5% (64.1%) of the variations in the winning candidate in $|M| = 1$ ($|M| = 2$). (9) and (11), however, show that only 6.5% (19.0%) of the variations in the observed voter turnout is explained for $|M| = 1$ ($|M| = 2$). Adding controls for electoral background variables does not increase the model's explanatory power for the election outcomes, but does increase the model's explanatory power for turnout by 6.2%-points (13.5%-points) when $|M| = 1$ ($|M| = 2$). We conclude that BNE has strong predictive power for the election outcome and vote share, but less predictive power for voter turnout. Similar results are obtained for the predictive power of QRE predictions.[24]

Finally, we check for learning effects by conducting the same regression (5) separately for subsamples in blocks 1-4, 5-9 and 9-12. By and large, we find little evidence of learning except that the explanatory power of BNE for voter turnout increases from 3.9% in blocks 1-4 to 8.7%, both in blocks 5-8 and 9-12 for $|M| = 1$. For $|M| = 2$, these numbers are 11.7%, 23.2%, and 27.2%, respectively. Though this indicates some learning, the percentage

[24]More precisely, the QRE predicts slightly better the variations in election outcomes and vote share than BNE does, but does not do better than BNE in explaining voter turnout. Details for this and the following analyses are available from the authors upon request.

Table 3: The Predictive Power of BNE

Dep. Variable	Vote Share of A				Winning Probability of A				Voter Turnout			
	$\|M\|=1$		$\|M\|=2$		$\|M\|=1$		$\|M\|=2$		$\|M\|=1$		$\|M\|=2$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
BNE	0.667***	0.667***	0.710***	0.710***	0.765***	0.765***	0.806***	0.806***	0.193***	0.193***	0.371***	0.371***
	(0.035)	(0.006)	(0.011)	(0.005)	(0.012)	(0.007)	(0.014)	(0.012)	(0.02)	(0.018)	(0.033)	(0.023)
Age		0.050***		-0.001		-0.004		0.008**		0.036		−0.036***
		(0.008)		(0.005)		(0.004)		(0.003)		(0.018)		(0.009)
$Males$		−0.030***		-0.004		−0.009***		-0.005		-0.008		0.049***
		(0.004)		(0.004)		(0.002)		(0.003)		(0.008)		(0.010)
$Econ$		0.055***		0.010***		−0.012**		0.006*		0.024		−0.051***
		(0.011)		(0.001)		(0.004)		(0.002)		(0.023)		(0.009)
$BNE \times Age$		−0.106***		-0.011		−0.029***		-0.012		0.017		0.018
		(0.010)		(0.007)		(0.005)		(0.014)		(0.028)		(0.029)
$BNE \times Males$		0.053***		0.027**		0.022***		0.017		0.01		-0.018
		(0.004)		(0.007)		(0.003)		(0.018)		(0.012)		(0.030)
$BNE \times Econ$		−0.121***		−0.016*		−0.017**		-0.015		0.003		0.060***
		(0.010)		(0.007)		(0.005)		(0.012)		(0.031)		(0.008)
$Constant$	0.151***	0.151***	0.119***	0.119***	0.093***	0.093***	0.044***	0.044***	0.473***	0.473***	0.411***	0.411***
	(0.018)	(0.007)	(0.006)	(0.004)	(0.007)	(0.003)	(0.004)	(0.002)	(0.017)	(0.014)	(0.024)	(0.007)
$\#.Obs.$	996	996	1,422	1,422	1,008	1,008	1,440	1,440	1,008	1,008	1,440	1,440
R^2	0.724	0.734	0.758	0.76	0.585	0.586	0.641	0.641	0.065	0.127	0.19	0.325

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

of variations in observed turnout that can be explained by BNE predictions still remains low.

7 Explaining our Results

Two findings in our result are markedly different from both BNE and QRE predictions. First, candidate A's ex-ante winning probability does not non-monotonically with media bias. Instead, it increases in media bias monotonically. Second, ex-ante voter turnout varies very little across treatments and is much higher than predicted under either equilibrium concept. In this section we discuss four possible explanations for these observed deviations. These are (i) distinct noise parameters for candidate choice and turnout decisions; (ii) non-Bayesian belief updating by voters; (iii) "competition neglect", that is, voters' imperfect abilities to infer the competitiveness of elections from the messages announced by media outlets; (iv) a hybrid structural model that includes features (i) to (iii). We discuss these alternative explanations for the case with $|M| = 1$. The results are similar to those obtained for $|M| = 2$; details are available upon request.

In searching for these explanations, we apply three important restrictions. First, we derive equilibrium predictions. This is not to say, of course, that non-equilibrium behavior does not matter. It simply reflects the lack of an alternative tractable model that can systematically explain the observed behavior. Second, we assume homogeneous players (aside from the heterogeneity created by design). Though it is intuitive that people differ in many dimensions, it is not *a priori* clear what aspects to include in the model and how to structure this. Third, we assume that all parameters of bounded rationality that we consider are common knowledge.[25] All these assumptions are made to maintain a tractable theory.

7.1 Distinct noise parameters for candidate choice and turnout decisions

When using the data from our previous experiments in 2016 to estimate the noise parameter λ of equations (3), we assumed that this same parameter applies for both turnout and candidate choice (voting for A or B) decisions (cf. equations (4)). Instead, noise might play a stronger role in one of these decisions than in the other. To formally account for this possibility, we construct an alternative "sequential game" in which voters make decisions in two stages. In the first stage (turnout decision), a voter decides between casting a costly

[25] This excludes explanations such as higher-order beliefs, level-k thinking and cognitive hierarchy (Goeree et al., 2016). This is not because we think that these play no role. We simply cannot model this role in a way that helps us explain our data in a tractable and satisfactory way.

vote or abstaining. If the voter chooses to vote, then she comes to the second stage (candidate choice decision) and decides which candidate to vote for, A or B. We allow the voter to possess different noise parameters in the two stages. We present the formal model and estimate the noise parameters for turnout decisions (λ_t) and candidate choices (λ_p) based on our 2016 experiment in online Appendix B.

Intuitively, candidate choices seem cognitively much easier than turnout decisions. To make the former decision, a voter only needs to infer which candidate is better based on her private ideology and the public information. To make the latter decision, however, the voter has to precisely infer the chances of casting pivotal votes and then judge whether the expected benefits from voting dominates its costs. For this reason, we expect voters to behave more rationally in candidate choices than in turnout decisions. The estimation results for the sequential model indeed confirms this intuition; the noise parameter is lower for turnout decisions ($\lambda_t = 16.66$) than for candidate choices ($\lambda_p = 21.83$). This reflects that noises play a more prominent role in turnout decisions than candidate choices. Nevertheless, this alternative sequential model does not outperform the standard QRE model presented in Section 4.2.[26] In addition, predictions of this sequential model cannot explain why we do not observe the predicted non-monotonicity of A's ex-ante winning probability. Therefore, we conclude that incorporating distinct noise parameters for candidate choices and turnout decisions alone cannot satisfactorily explain the observed deviations in our experiment.

The equilibrium analyses in Section 4 reveal that, to make an optimal voting decision, a voter needs not only to form correct posterior beliefs of candidate quality difference k, but also to infer the closeness of election based on their posterior beliefs. Voters may behave sub-optimally if they make mistakes in either of the two aspects. So, could subjects' observed voting behavior be explained by deviations from Bayesian updating, their failure to infer the closeness of elections, or both? We discuss these alternative explanations in the next three subsections.

7.2 Non-Bayesian Belief Updating

We elicited subjects' posterior beliefs about the probability that candidate A generates a strictly higher payoff than B (cf. Section 5.3). Note that our choice-list elicitation procedure yields belief intervals of 5%-points. For the analysis we exclude inconsistent choices (i.e., when vase A was believed to pay less under message A than under message B) and use the

[26] As reported in online Appendix B, predictions of the latter model generates a slightly higher log-likelihood than the predictions from the former. Results from Vuong's closeness tests (Vuong, 1989) cannot reject the null hypothesis that the two models are equally close to the true data generating process.

midpoint of the remaining intervals as an estimate of the reported belief.[27] As argued in section 5.3, if voters are Bayesian, biased media outlets cannot systematically alter the mean of their posterior beliefs (that is, from the ex-ante perspective, voters should still expect no quality difference between candidates). If, however, voters are subject to persuasion bias, then their mean posterior beliefs may vary with media bias.

Table 4 summarizes for each treatment condition the mean reported and rational posterior beliefs that A will give a higher payoff than B. It appears that reported beliefs follow the comparative statics predicted by the Bayesian posteriors reasonably well, though –recalling that the prior probability is 0.5– voters tend to update less than rationally prescribed (with one exception). The differences between reported and Bayesian beliefs are not significant across treatments (Ft, $p = 0.139$, $N = 6$).[28] This suggests that voters' posterior beliefs are not systematically influenced by media bias.

Table 4: Reported versus Bayesian Posterior Beliefs

Message	$s = A$		$s = B$		Ex-ante	
	Reported	Bayesian	Reported	Bayesian	Reported	Bayesian
UB	0.72	0.76	0.26	0.25	0.49	0.51
WB_A	0.60	0.60	0.20	0.13	0.50	0.48
SB_A	0.55	0.52	0.14	0.00	0.53	0.50

Note: Cells report, for sessions with $|M| = 1$, subjects' mean reported and Bayesian posterior beliefs about the probability that candidate A generates strictly better payoffs to them than candidate B. Note that the unconditional (i.e., ex-ante) Bayesian posterior beliefs may differ from the prior 0.5 as a consequence of random draws of model parameters in the experiment.

To investigate in more detail how reported beliefs respond to messages and the Bayesian benchmark, we use the regression framework reported in (6).

$$R_{it} - \eta_{it} = \beta_0 + \beta_1 mA_t + \beta_2(B_{it} - \eta_{it}) + \gamma X_i + \delta X_i \times (B_{it} - \eta_{it}) + \epsilon_{it} \tag{6}$$

where i is an index of the individual subject, t is the index of the election, mA_t is a dummy variable that equals 1 if message A is sent in election t, R_{it} and B_{it} are the reported and Bayesian posterior beliefs, η_{it} is the Bayesian prior belief for subject i (which varies with i's ideological preferences), X_i is a vector of individual characteristics, including age, gender and a dummy variable indicating whether the individual is a student of economics (including

[27] 32% of pairs of elicited beliefs (for $|M| = 1$) had inconsistent choice patterns in this way. These inconsistencies were evenly spread across subjects; nine decisions were made (each involving a reported belief conditional message A and one conditional on message B) and only 7% of the subjects were inconsistent zero or nine times.

[28] Similarly, we observe no significant effects if we test separately per voter type (centralist, moderate, extreme), after correcting for the multiple comparisons problem.

business), and ϵ_{it} represents idiosyncratic noises for each subject i in each election t. In formula (6), $\beta_0 + \beta_1$ (β_0) measures how much a voter's reported posterior deviates from the Bayesian prior conditional on message $A(B)$. β_2 measures how responsive are voters' reported beliefs to variations in the Bayesian posterior. Table 5 reports the estimated coefficients β_0, β_1 and β_2 from regression (6) in each treatment condition under $|M| = 1$.

Table 5: Estimated Coefficients of Regression 6

Treatment	UB	WB_A	SB_A
Constant (β_0)	-0.184^{***}	-0.122^{***}	0.077^{***}
	(0.023)	(0.028)	(0.016)
mA_t (β_1)	0.360^{***}	0.177^{***}	-0.032^{**}
	(0.022)	(0.034)	(0.01)
$B_{it} - \eta_{ij}$ (β_2)	0.169^{*}	0.425^{***}	0.878^{***}
	(0.071)	(0.063)	(0.019)
$\#.Obs.$	596	600	624
R^2	0.401	0.376	0.570

Note: Cells report the estimated coefficients of a subset of independent variables in (6), listed in the first column. Estimates are for $|M| = 1$ and excluding inconsistent reports, as explained in the main text. Robust standard errors (reported in parentheses) are clustered at the electorate level. Complete results are available upon request. ***/**/* indicates statistical significance at the 1%-/5%-/10%-level.

If subjects are perfect Bayesian, we should observe $\beta_0 = \beta_1 = 0$ and $\beta_2 = 1$. If instead voters respond to Bayesian posteriors insufficiently, we have $\beta_2 < 1$. If voters irrationally trust a message at face value to at least some extent, then $\beta_0 + \beta_1 > 0$ and $\beta_0 < 0$. Voters' average posterior beliefs are systematically different from their priors if $(\beta_0 + \beta_1) \cdot Pr[s = A|\tau] + \beta_0 \cdot Pr[s = B|\tau] = \beta_0 + \beta_1 \cdot Pr[s = A|\tau] \neq 0$, where $\tau \in \{UB, WB_A, SB_A\}$ indicates the treatment condition (with $|M| = 1$). Table 6 summarizes the tests for these combinations of coefficients.

Table 6: Reports of Statistical Tests

Treatment	UB	WB_A	SB_A
$\beta_0 = 0$	$(-)^{***}$	$(-)^{***}$	$(+)^{***}$
$\beta_0 + \beta_1 = 0$	$(+)^{***}$	$(+)^{*}$	$(+)^{*}$
$\beta_2 = 1$	$(-)^{***}$	$(-)^{***}$	$(-)^{***}$
$\beta_0 + \beta_1 \cdot Pr[s = A\|\tau] = 0$	$(-)$	$(+)$	$(+)^{*}$

Note: Cells report the test results for the restrictions listed in the first column. In performing these tests, we focus on treatments with $|M| = 1$ and exclude inconsistent reports, as explained in the main text. (–)/(+) indicates a value for the (function of the) coefficient(s) that is below/above the value that is being tested. For example, when testing $\beta_2 = 1$, (–) means that the estimated coefficient is below 1. ***/* indicates significance at the 1%-/10%-level.

In sessions with $|M| = 1$, we observed clear evidence for conservatism in belief updating; $\beta_2 < 1$ for all treatments $\tau \in \{UB, WB_A, SB_A\}$. This suggests that voters' posterior beliefs respond insufficiently to information compared to a Bayesian. In treatments UB and WB_A, voters more or less trust media endorsements at face value; $\beta_0 + \beta_1 > 0$ and $\beta_0 < 0$. They are not systematically persuaded, however, because from the ex-ante perspective their average posterior beliefs do not systematically deviate from their priors.[29] The results are slightly different in treatment SB_A. There, subjects shift their posterior beliefs in favor of candidate A conditional on both message A and B; $\beta_1 + \beta_0 > 0$ and $\beta_0 > 0$. Moreover, the effect is even stronger conditional on message B than on message A ($\beta_0 > \beta_1 + \beta_0 > 0$). There is marginal evidence that voters' posterior beliefs are systematically shifted in favor of candidate A in treatment SB_A.

Because voters' posterior beliefs are relative conservative compared to the Bayesian benchmark (except for SB_A), they expected less quality difference between the candidates than is objectively the case. They might therefore expect a more competitive election. This might partly explain why the observed levels of turnout are systematically higher, and vary less with media bias, than predicted. Nevertheless, this could not explain the irresponsiveness of voter turnout to variations in media bias. The results in Table 4 may, however, provide a clue to why we do not observe the non-monotonic relationship between candidate A' ex-ante winning probability and media bias. Voters' average posterior beliefs are a bit high under SB_A (cf. the last row of Table 4). Recall from the left panel of Figure 2 that candidate A's winning probability is a very steep function at prior $k = 0$. Then, even a small persuasion effect under message A might have a sizable influence on the election outcomes, resulting a higher probability of A winning in treatment SB_A; this could take away the non-monotonicity relationship predicted by both BNE and QRE.

7.3 Competition Neglect

A third possible explanation for the two deviations from the theoretical prediction is that voters do not correctly perceive the implication of media's messages on the competitiveness of elections. More precisely, voters may incorrectly infer the probabilities of casting pivotal votes. This might be aggravated by our experimental design where, after every eight elections, voters change to a new treatment condition with possibly different media biases. This feature may make it difficult for subjects to learn the pivotal probabilities conditional on each specific information environment, which critically drives turnout decisions. As an alternative, we present in online Appendix E a modified model of bounded rationality where

[29]This is because the hypotheses $\beta_0 + \beta_1 \cdot Pr[s = A|\tau] = 0$ are not rejected for $\tau \in \{UB, WB_A\}$, as evident in Table 6.

players are *competition neglect*; they overlook the correlation between the competitiveness of elections and the information environment. In this modified model, voters mistakenly believe pivotal probabilities to be constant across all treatments, and best respond to the average pivotal probabilities.[30] The predictions for voter turnout under this modified model with competition neglect are depicted in Figure 7.

Figure 7: Predicted Voter Turnout under the Modified Model with Competition Neglect

Conditional on $s = A$ | Conditional on $s = B$ | Ex-ante

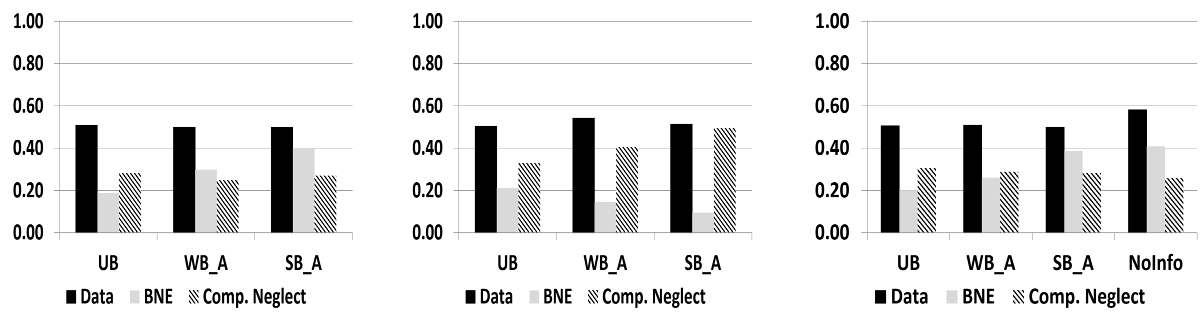

Note: These panels compare the observed voter turnout (black bars) with the predicted turnout from the standard BNE (grey bars) and a modified BNE with competition neglect (strip bars). All predictions are generated based on the realizations for model parameters in the real experiment.

Figure 7 shows that competition neglect fundamentally alters the comparative statics predictions of standard BNE. Conditional on the model parameters realized in the experiment, the modified model with competition neglect predicts voter turnout to remain more or less constant after message $s = A$ (the left panel), to increase in media bias after message $s = B$ (the central panel). Ex-ante, voter turnout is predicted to be slightly *decreasing* in media bias and be lowest in treatment $NoInfo$ (the right panel), if voters are completely competition neglect. This model may partly explain the irresponsiveness of voter turnout with respect to variations in media bias conditional on message $s = A$ or unconditionally. Nevertheless, it cannot explain the irresponsiveness of voter turnout to media biases conditional on message B. In addition, the fact that observed ex-ante voter turnout is *highest* in treatment $NoInfo$ is also inconsistent with the implication of competition neglect. Last but not least, this modified model cannot explain why we do not observe a non-monotonic relationship between media bias and candidate A's winning probability.[31] Therefore, we conclude that a bounded rational model with complete competition neglect (i.e., voters perceiving pivotal probability to be invariant with information) alone cannot provide a satisfactory explanations for the observed deviations in subjects' voting behavior.

[30]We thank David Levine for suggesting this alternative model.

[31]Precise predictions are available from the authors upon request.

7.4 A Hybrid Structural Model

The previous (three) explanations suggest that (i) the irresponsiveness of observe voter turnout to variations in media bias may be partly attributed to voters' limited rationality in turnout decisions, conservative belief formation, and competition neglect, (ii) the failure of non-monotonicity may be attributed to voters non-Bayesian belief updating. Here, we present a hybrid structural model that combines all these features in an attempt to explain both phenomena. The point of departure is the following. Whereas the BNE assumes perfect rationality by voters, the QRE relaxes this by allowing for noisy decision making. QRE assumes, however, that the same decision makers are perfectly rational in belief updating. We now present an alternative model where noise occurs in belief formation. We allow both, for noise in belief updating (as in Section 7.2) and for competition neglect (Section 7.3). Moreover, we allow the noise parameters in response to be different for turnout than for the choice of A or B.

To start, we assume that the electorate forms a common posterior expectation $\widehat{k_\tau}$ for each information environment τ.[32] To allow for partial competition neglect, we consider a simple binary-type model where with probability $\rho \in [0,1]$, a voter is perfectly rational and infers precisely the pivotal probability for any specific information condition and best responds to it; with the remaining probability $1-\rho$ the voter ignores the connection between information environment and pivotal probabilities, and she only best responds to the average pivotal probabilities. Formally, let $PivA_\tau$ and $PivB_\tau$ denote the equilibrium pivotal probabilities of a vote for candidate A and B, respectively. f_τ denotes the ex-ante probability of encountering each information environment τ.[33] Define $\overline{PivA} \equiv \sum_\tau f_\tau \cdot PivA_\tau$ and $\overline{PivB} \equiv \sum_\tau f_\tau \cdot PivB_\tau$ as the (weighted) average pivotal probability of a vote for candidate A and B, respectively, in treatments with $|M|=1$. A voter's perceived pivotal probabilities –denoted by $\widehat{PivA_\tau}$ and $\widehat{PivB_\tau}$– are then given by:

$$\widehat{PivA_\tau} = \rho \cdot PivA_\tau + (1-\rho) \cdot \overline{PivA} \tag{7}$$

$$\widehat{PivB_\tau} = \rho \cdot PivB_\tau + (1-\rho) \cdot \overline{PivB} \tag{8}$$

If $\rho = 0$, the model reduces to the scenario with complete competition neglect as discussed in Section 7.3. 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

[32] Hence, $\tau \in \{NoInfo, UB_A^A, UB_A^A, WB_A^A, WB_A^B, SB_A^A, SB_A^B\}$, where the superscript refers to the message sent by the media (if any). Allowing $\widehat{k_\tau}$ to freely vary across treatment conditions means voters may exhibit different levels of limited rationality in belief formation, depending on the treatments.

[33] In treatments with $|M|=1$, these probabilities are $f_{NoInfo} = \frac{1}{4}$, $f_{UB_A^A} = f_{UB_A^B} = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{2}$, $f_{WB_A^A} = \frac{1}{4} \cdot \frac{3}{4} = \frac{3}{16}$, $f_{WB_A^B} = \frac{1}{4} \cdot \frac{1}{4}$, $f_{SB_A^A} = \frac{1}{4} \cdot \frac{19}{20} = \frac{19}{20}$ and $f_{SB_A^B} = \frac{1}{4} \cdot \frac{1}{20}$.

relationship between pivotal probabilities and the information environment, but insufficiently so compared to a rational agent. For this reason, we say voters are *partial competition neglect* if $\rho \in (0, 1)$; a higher ρ implies a lower degree of competition neglect.

Next, we incorporate the posterior expectation $\widehat{k_\tau}$ and perceived pivotal probabilities in the QRE models formally described in online Appendix B. We jointly estimate the noise parameters (λ_p, λ_t), voters' common posterior beliefs (k_τ for each information environment τ) and the degree of competition neglect (ρ) using data from our 2016 experiment. The results are presented in online Appendix F. The estimation we obtain for ρ equals 0.507. According to (7) and (8), this shows that the actual pivotal probability weighs in at only 50% in the perceived probability. We use the model parameters estimated from our 2016 experiment to predict voters' behavior for the experimental sessions conducted to in 2017. The results are summarized in Figure 8.

Given the restrictions set out at the beginning of this section (an equilibrium model with homogeneous agents and common knowledge of the parameters of bounded rationality), our hybrid model is very flexible; it allows beliefs to vary freely across treatments and includes the possibility of competition neglect. The hybrid model predicts that ex-ante candidate A's winning probability is almost invariant in media bias (cf. middle-right panel of Figure 8). This is neither consistent with the non-monotonic relationship predicted by BNE, nor with the increasing relationship observed in data. As for voter turnout, the hybrid model predicts much smaller treatment effects of media bias than BNE predictions (cf. bottom-right panel of Figure 8). This is in line with the irresponsiveness of voter turnout appeared in our data. At the same time, the observed substantially higher voter turnout under $NoInfo$ compared to other treatments is also inconsistent with the prediction of the hybrid model (cf. bottom-right panel of Figure 8). In sum, even this very flexible model does not explain the non-monotonicity and turnout effects we observe.[34]

8 Conclusion

In this paper, we construct an experimental election game based on a general model in SSS18. We systematically characterizes both the interim and ex-ante impacts of media bias on the election outcome and voter turnout. We show that in an environment with

[34]As a final check, we investigated whether the differences between the 2016 and 2017 data might explain our results. For example, the size of the electorate was 11 in 2016 but 25 in 2017. For this purpose, we randomly split the data from 2017 into two. We used one half to estimate the parameters of the hybrid model and the other to test the predictions. Once again, the substantially higher turnout in $NoInfo$ compared to other treatments with media, and the observed monotonic increase in candidate A's winning probability, could not be explained. More information is available upon request.

Figure 8: Predicted Aggregate Outcomes based on the Hybrid Structural Model

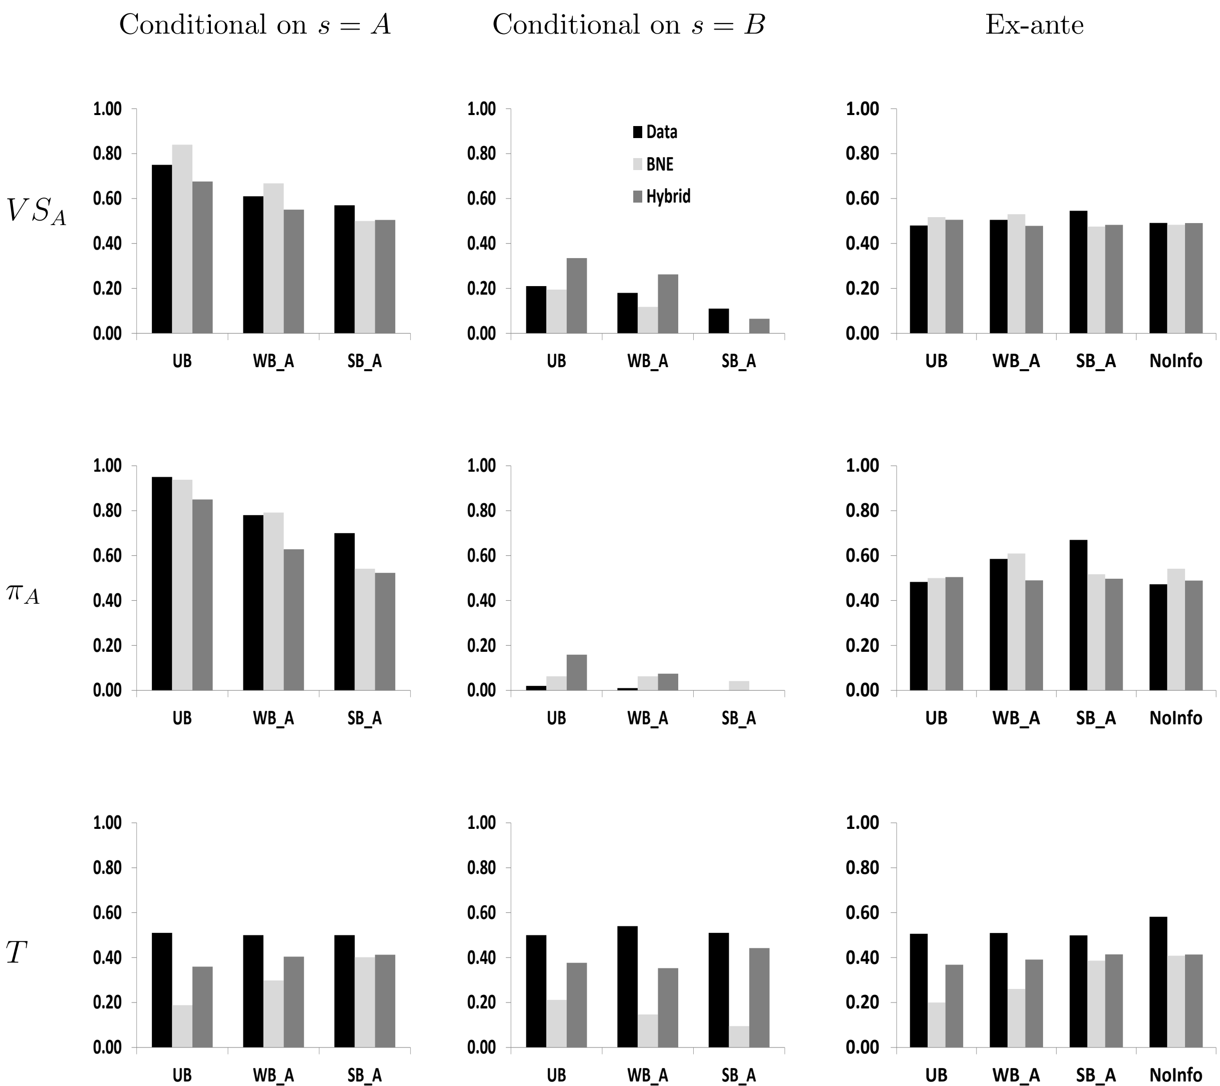

Note: VS_A denotes the vote share of candidate A; π_A denotes the probability that A wins the election (in the data this is measured by the fraction of elections won by A, where ties are computed as a 0.5 score each); T denotes voter turnout. The black and light gray bars denote the observed outcomes and the corresponding BNE predictions, respectively. The dark gray bars denote the predictions based the hybrid structural model with parameters estimated from our 2016 experiment.

a single media outlet, the ex-ante election outcome varies non-monotonically with media bias while voter turnout robustly increases in media bias. We then show that introducing a second media outlet can systematically shift the election outcome and voter turnout in either direction, depending on the biases of both outlets. This result provides a possible rationale for the mixed empirical evidence on the impact of media entry on voter turnout.

We subsequently test our theoretical predictions in a laboratory experiment. We find that party vote shares and the election outcome are well predicted, albeit that predicted non-monotonic relationship is not observed. Observed voter turnout, however, is much less responsive to media bias than predicted. In looking for explanations for our results, we show that voters are able to filter out the influence of media bias (i.e., their posterior beliefs are not systematically affected by media bias), though they do update their beliefs conservatively. We conclude that our data are best explained by a hybrid structural model that combines (i) noisy best response in the tradition of QRE; (ii) noisy belief updating with respect to the qualities of the candidates; and (iii) partial competition neglect, that is, voters' imperfect ability to infer the closeness of elections from the messages announced by media outlets. Even this hybrid model, however, cannot explain some of the patterns in our data.

Our results imply that voters are not fully irrational in their response to the information provide by biased media. 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 some level of competition neglect is also intuitive. In our design (and arguably also in the world outside the laboratory), subjects face an electoral environment that changes regularly. Keeping track of how media information in any particular environment affects the competitiveness of the election is a difficult task.

In conclusion, one might be tempted to conclude that rational choice models do a poor job of predicting voter behavior in an environment with biased media. It is important to stress, however, that only a few features in our data have proven difficult to rationalize. In broad overview, we believe that our results provide a solid ground for the conclusion that voter behavior in an environment with biased media may be noisy, but that it does have a strong core of rational response.

References

Adena, M., Enikolopov, R., Petrova, M., Santarosa, V., and Zhuravskaya, E. (2015). Radio and the Rise of the Nazis in Prewar Germany. *Quarterly Journal of Economics*, 130(4):1885–1939.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–36.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2006). Elicitation using multiple price list formats. *Experimental Economics*, 9(4):383–405.

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

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

Blais, A. (2000). *To vote or not to vote? The merits and limits of rational choice theory.* University of Pittsburgh Pre.

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

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

Chiang, C.-F. and Knight, B. (2011). Media bias and influence: Evidence from newspaper endorsements. *Review of Economic Studies*, 78(3):795–820.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, pages 1431–1451.

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

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

Djankov, S., McLiesh, C., Nenova, T., and Shleifer, A. (2003). Who owns the media? *Journal of Law and Economics*, 46(2):341–382.

Downs, A. (1957). An economic theory of political action in a democracy. *Journal of Political Economy*, 65(2):135–150.

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

Durante, R. and Knight, B. (2012). Partisan control, media bias, and viewer responses: Evidence from Berlusconi's italy. *Journal of the European Economic Association*, 10(3):451–481.

Ellingsen, S., Hernæs, Ø., et al. (2018). The impact of commercial television on turnout and public policy: Evidence from Norwegian local politics. *Journal of Public Economics*, 159(C):1–15.

Ellman, M. and Germano, F. (2009). What do the papers sell? A model of advertising and media bias. *Economic Journal*, 119(537):680–704.

Enikolopov, R., Petrova, M., and Zhuravskaya, E. (2011). Media and political persuasion: Evidence from Russia. *American Economic Review*, 101(7):3253–3285.

Falk, A. and Heckman, J. J. (2009). Lab experiments are a major source of knowledge in the social sciences. *Science*, 326(5952):535–538.

Gavazza, A., Nardotto, M., and Valletti, T. M. (2015). Internet and politics: Evidence from UK local elections and local government policies. (December 2015). CEPR Discussion Paper No.DP10991. Available at SSRN: https://ssrn.com/abstract=3039013.

Gentzkow, M. (2006). Television and voter turnout. *Quarterly Journal of Economics*, 121(3):931–972.

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

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

Gentzkow, M. and Shapiro, J. M. (2011). Ideological segregation online and offline. *Quarterly Journal of Economics*, 126(4):1799–1839.

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

Gentzkow, M., Shapiro, J. M., and Stone, D. F. (2015). Media bias in the marketplace: Theory. In *Handbook of media economics*, volume 1, pages 623–645. Elsevier.

Gerber, A. S., Karlan, D., and Bergan, D. (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.

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

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2016). *Quantal response equilibria*. Springer.

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

Grosser, J. and Schram, A. (2006). Neighborhood information exchange and voter participation: An experimental study. *American Political Science Review*, 100(2):235–248.

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

Levine, D. K. and Palfrey, T. R. (2007). The paradox of voter participation? A laboratory study. *American Political Science Review*, 101(01):143–158.

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

McMillan, J. and Zoido, P. (2004). How to subvert democracy: Montesinos in peru. *Journal of Economic Perspectives*, 18(4):69–92.

Mullainathan, S. and Shleifer, A. (2005). The Market for News. *American Economic Review*.

Myerson, R. B. (2000). Large Poisson games. *Journal of Economic Theory*, 94(1):7–45.

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

Pew Research Center (2008). A Continuum of Condemning the Press. Washington, DC: Pew Research Center. http://www.journalism.org/2008/10/10/a-continuum-of-condemning-the-press/.

Piolatto, A. and Schuett, F. (2015). Media competition and electoral politics. *Journal of Public Economics*, 130:80–93.

Prat, A. and Stromberg, D. (2013). The Political Economy of Mass Media. In *Advances in Economics and Econometrics: Volume 2, Applied Economics: Tenth World Congress*, volume 50, page 135. Cambridge University Press.

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

Puglisi, R. and Snyder, J. M. (2015). The balanced US press. *Journal of the European Economic Association*, 13(2):240–264.

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

Sobbrio, F. (2014). Citizen-editors' endogenous information acquisition and news accuracy. *Journal of Public Economics*, 113:43–53.

Strömberg, D. (2004). Radio's impact on public spending. *Quarterly Journal of Economics*, 119(1):189–221.

Sun, J., Schram, A., and Sloof, R. (2018). A Theory on Media Bias and Elections. Working Paper, University of Amsterdam.

Tyszler, M. and Schram, A. (2016). Information and strategic voting. *Experimental Economics*, 19(2):360–381.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, pages 307–333.