NATIONAL RESEARCH UNIVERSITY
HIGHER SCHOOL OF ECONOMICS

Kemal Kıvanç Aköz, Cemal Eren Arbatlı

MANIPULATED VOTERS IN COMPETITIVE ELECTION CAMPAIGNS

BASIC RESEARCH PROGRAM

WORKING PAPERS

SERIES: ECONOMICS
WP BRP 31/EC/2013

Manipulated Voters in Competitive Election Campaigns*

Kemal Kıvanç Aköz† and Cemal Eren Arbatlı ‡

May 29, 2013

Abstract

We provide a game-theoretical model of manipulative election campaigns with two political candidates and a continuum of Bayesian voters. Voters are uncertain about candidate positions, which are exogenously given and lie on a unidimensional policy space. Candidates take unobservable, costly actions to manipulate a campaign signal that would otherwise be fully informative about a candidate's distance from voters relative to the other candidate. We show that if the candidates differ in campaigning efficiency, and voters receive the manipulated signal with an individual, random noise, then the cost-efficient candidate wins the election even if she is more distant from the electorate than her opponent is. In contrast to the existing election campaign models that do not support information manipulation in equilibrium, our paper rationalizes misleading political advertising and suggests that limits on campaign spending may potentially improve the quality of information available to the electorate.

JEL: C72, D72, D82, D84.

Keywords: Hidden actions, election campaigns, manipulation, propaganda, bias.

---

*The authors would like to thank Alessandro Lizzeri, Sergey V. Popov, Alexei Zakharov, Nikolay Arefiev and other seminar participants at HSE for valuable comments. We also thank Pedro Dal Bó, Roberto Serrano and Brian Knight and seminar participants at Brown University; Giri Parameswaran and EconCon conference participants for their helpful comments and suggestions. All remaining errors are our responsibility.

†Aköz: Department of Economics, New York University, 19 W. 4th Street, 6FL, New York, NY 10012 , (e-mail: kka227@nyu.edu).

‡Arbatlı: Department of Economics, National Research University-Higher School of Economics, 26 Shabolovka Street, Building 3, 3116A, Moscow, Russia (e-mail: darbatli@hse.ru).

1 Introduction

While all politicians try to design their political campaigns so as to appeal to the largest possible audience, their policy preferences often do not coincide with that of the electorate. After being elected, most of them stand to gain from reneging on their promises and following policies that serve narrow interests rather than those of the median voter. Therefore, the quality of information supplied in election campaigns is crucial to citizens' ability to choose the best candidate. Politicians typically have more incentive to misrepresent their true policy platforms when voters cannot properly discount misleading campaign messages. This paper offers a model of manipulative election campaigns in which the structure of communication allows the candidates to systematically undermine voters' ability to infer their true positions.

Despite the common wisdom and the evidence that election campaigns can mislead voters (Ramsay et al. 2010, Cassino and Wooley 2011, DellaVigna and Kaplan 2007, Durante and Knight 2009, Puglisi and Snyder 2008), in the literature on electoral competition, political campaigns are modeled as entirely informative activities (Coate 2004a, Gul and Pesendorfer 2009). Although, in most of these models, the informativeness of election campaigns and the uncertainty voters face about candidate qualities are allowed to vary endogenously, the campaign messages that voters receive always remain unbiased. This is because the information structure adopted in these models allows rational voters to discount biased messages to the full extent (e.g., Wittman 1989), leaving no incentive for politicians to provide false information in equilibrium (Banks 1989, Aragones and Neeman 2000, Martinelli 2001). In existing models, even when voters end up holding biased beliefs about the true state of the world, these beliefs are supported in equilibrium only when voters have bounded rationality or when they already had biased beliefs (Heidhues and Lagerlof 2003) prior to the election campaigns.[1] In contrast to existing theoretical approaches to political campaigns, in this paper, we argue that even when voters are perfectly rational and do not hold biased prior beliefs, under a non-transparent social communication environment, deceptive political propaganda can become an equilibrium outcome.

We propose a two-candidate electoral competition model in which the candidates know voters' most-preferred policy, but voters are uncertain about the candidates' positions. Under this informational asymmetry, we characterize the necessary and sufficient conditions for political candidates to successfully bias the beliefs of Bayesian voters about their policy positions and to influence election outcomes. We model information manipulation as directed disinformation on a unidimensional policy space: Candidates try to manipulate voters' beliefs about their relative proximity to voters through public announcements or actions that shift an otherwise unbiased signal. Candidates' objective is to maximize the votes they get, and the outcome of the competition is determined by majority voting among a large set of voters. The campaigning efficiencies of the politicians differ. This is captured by asymmetric costs of manipulative actions.[2] We define *successful manipulation*

[1] For a comprehensive discussion of models with bounded rationality, see Caplan (2006).

[2] Asymmetric costs may stem from the difference between the candidates' personal ability to lie about their true policy positions and get away with it or their unequal access to campaigning resources and technology, which, among other things, may depend on how their policy preferences and personal characteristics appeal to special interest groups or whether they come from an incumbent political party with superior tools of information manipulation at their disposal

as any situation in which the more efficient candidate wins an election that she would not be able to win in the absence of manipulation.

Using this general framework, we analyze when and how information manipulation in election campaigns leads to a disinformed electorate. Our analysis suggests that unrestricted electoral competition in a representative democracy may lead to socially inefficient election outcomes by biasing the electorate's beliefs in favor of the candidate who is more effective at manipulating, but who does not necessarily stand closer to voters. We show that under manipulative election campaigns, even when the cost-efficient candidate is ex-post an inferior choice for voters, this candidate can win the election as long as she does not have too great a positional disadvantage vis-à-vis her opponent.

In addition to the standard assumption of information asymmetry between politicians and the electorate, we argue that there are three key assumptions without which manipulation cannot succeed: (i) asymmetry between the candidates in terms of campaigning efficiency; (ii) an idiosyncratic noise --or bias-- in campaign signals received by the electorate; and (iii) the lack of perfect communication among voters. We show that when these conditions are satisfied, the candidates tailor their manipulative actions in such a way that voters, who cannot perfectly distinguish idiosyncratic noise from systematic bias (the net manipulation) in campaign messages, encounter an inference problem. Although voters hold different beliefs about the candidates due to the heterogeneity in campaign signals, the average bias in their posterior beliefs favors the cost-efficient candidate, who manipulates the electorate to a greater extent than her opponent does. As counterintuitive as it may sound, we show that as long as voters do not communicate perfectly and there is always some noise-induced heterogeneity in campaign signals, the reduction in this heterogeneity makes manipulation more effective in the sense that the cost-efficient candidate gets to win the election under a wider range of relative candidate positions.

We then focus on two main determinants of the bias in average posterior beliefs. First, we analyze how a change in the gap between the candidates' cost efficiency influences the extent of net manipulation in electoral messages. Confirming what Gul and Pesendorfer [15] have found, we demonstrate in our numerical simulations that as the difference in manipulative effectiveness decreases, so does the net bias in campaign messages. Second, we analyze how a change in signal precision affects the candidate strategies and vote shares. Consistent with our theoretical results, numerical simulations demonstrate how higher signal precision renders manipulation more effective and increases the vote share of the cost-efficient candidate over a certain range of relative candidate positions.

Finally, we model the response of a benevolent planner who can limit politicians' campaign spending. We consider a framework in which candidates choose their policy platforms before the elections, and campaign contributions by special-interest groups to the candidates depend on these choices. We prove that a planner who wants to minimize the systematic bias in voters' beliefs will choose to limit the amount of allowable contributions such that candidates will no longer have an incentive to engage in manipulation.

The paper is organized as follows: Section 2 briefly reviews the related literature. Section 3 describes the model and compares the electoral equilibria when manipulation is feasible and when it is not. Section 4 is devoted to a discussion of the necessary and sufficient conditions for manipulative campaigns to succeed in equilibrium. We end this section with an illustrative example demonstrating the nature of voters' inference problem. Section 5 presents numerical simulations of the model with regard to the extent of cost asymmetry between the candidates and the degree of informational heterogeneity among voters. Potential policy implications of these comparative statics analyses are also discussed in this section. Finally, Section 6 concludes.

2 Related Literature

Information manipulation is an outcome of signaling games, where signal receivers' equilibrium beliefs mislead them into strategically inferior actions. This can happen when signal receivers hold biased prior beliefs and the signal they receive does not correct this bias. Cummins and Nyman (2005) study such a case in which better-informed investing firms choose to send signals that confirm the biased priors of less-informed investors. Gentzkow and Shapiro (2006) prove a similar result for media firms. Similarly, Heidhues and Lagerlöf (2003) show that if voters hold biased prior beliefs, political candidates may have an incentive to send biased signals to voters that confirm these priors. None of these papers considers the heterogeneity among signal receivers, and all of them assume that the average prior belief of receivers is biased.

Edmond (2008) and Ekmekci (2009) study political manipulation of information. Edmond (2008) studies the interaction between a manipulative dictator and a continuum of heterogeneous citizens within the framework of global games. Ekmekci (2009) studies the manipulative incentives of an expert to mobilize voters to coordinate around a candidate she favors. They both prove that information manipulation is possible in equilibrium and has a real effect on signal receivers' actions. In both papers, there is one manipulator; signal receivers' incentives are coordinated, and the source of the heterogeneity among agents is an idiosyncratic noise added to the manipulator's messages. Our study differs from them in two aspects. First, in our model, there are two manipulators, and we prove that when one manipulator is more efficient than the other, the advantaged player manipulates voters. When there is no asymmetry between manipulators, their manipulative efforts cancel each other out. Second, we do not assume any strategic complementarity on the side of the voters. Voters care only about the policy outcome and not about the size of the support that the winning candidate gets. These two deviations from the Edmond (2008) model enable us to provide a model that is more readily applicaple to a generic election environment with two dominant candidates.

Our results have different policy implications regarding the regulation of election campaigns from most other studies in the literature. Theoretical explanations for why campaign spending should not be limited commonly assume that campaign messages do not contain misleading information. Coate (2004a) studies such informative campaigns and addresses the question of how campaign contributions of interest groups are related to the election strategies of political parties. He proves that parties randomize between extremist and moderate candidates to cope with the trade-off between attracting the interest groups by supporting an extremist candidate and winning the election by supporting a moderate candidate. Gul and Pesendorfer (2009) propose another approach for informative election campaigns. In their continuous-time model, candidates prefer to provide more information if the information will favor them. They show that when spending is limited, the amount of information revealed is reduced, and, thus, voters become worse-off. Both of these papers suggest that lowering limits on campaign spending is welfare-improving. Our model implies the opposite, as we assume that campaigns are manipulative.

There are a few studies which, like our paper, conclude that limiting campaign contributions can improve welfare (Grossman and Helpman, 2001; Prat, 2002a, 2002b; and Coate, 2004b). However, in existing studies, campaigns do not bias voter beliefs. In contrast, we model political advertising as costly signal jamming, which, in equilibrium, leads to posterior voter beliefs that are biased towards the cost-efficient candidate. As a result, limiting campaign contributions can improve welfare by reducing the ex-ante systematic bias in voters' beliefs. To the best of our knowledge, this type of a justification for campaign spending limits, which relies on the systematic bias in voters' beliefs, has not been formally analyzed before.

Finally, the manipulation-based argument for spending limits, that we develop here, has policy implications for campaign regulation which potentially differ from those implied by the existing models. For instance, in Coate (2004b), replacing "favor financed" campaign spending with a matching tax-financed public subsidy can offset the adverse informational effects of a limit on campaign contributions, further increasing the likelihood that a qualified candidate will come to power. In our model, however, even if we had allowed for informative advertising and campaign contributions from interest groups, changing the source of funding from private to public contributions by itself would not lower the net manipulation in campaign signals. This is because public financing cannot be conditioned on truthful advertising unless regulatory agencies can distinguish manipulative from non-manipulative advertising.

3 Manipulative election campaigns with fully rational voters

In this section, we lay out a model of election campaigns with two political candidates and a large mass of voters. The model features an asymmetric information structure. The candidates are fully informed about voter preferences, while voters are uncertain about candidates' political positions. Voters observe only a noisy signal about candidates' relative proximity to their common bliss point and vote for the candidate that they believe stands closer to them. The key aspect of the model is a "war of disinformation" between the candidates whereby they can distort the signals received by voters via costly campaign messages. One of the candidates is more cost-efficient in manipulation than the other candidate, so that the net distortion in signals favors her whenever there is an incentive for manipulation. Since voters understand the manipulative nature of the campaigns and are fully rational, they correctly anticipate the net direction of the distortion and act accordingly. However, they underestimate the magnitude of the distortion as long as an idiosyncratic noise blurs the communication between voters and the candidates.

3.1 The Model

Consider an electoral competition between two candidates indexed by $j \in P = \{I, C\}$ and a continuum of voters with a unit mass indexed by $i \in [0,1]$.[3] Each candidate holds an exogenously given political position $\omega_j \in \mathbb{R}$.[4] Voters hold a common political position given by $\bar{\omega} \in \mathbb{R}$, which we call voters' bliss point.[5] We define $\theta \in \mathbb{R}$ as an index comparing the political distances of candidates to voters' common bliss point $\bar{\omega}$. That is,

$$\theta = |\omega_C - \bar{\omega}| - |\omega_I - \bar{\omega}| + \frac{1}{2},$$

where $\frac{1}{2}$ is an arbitrary rescaling parameter. Since $\theta > 1/2$ if and only if I stands closer to voters than C, we will refer to θ as the relative proximity of candidate I.

Information structure and manipulative actions

Voters' bliss point $\bar{\omega}$ is common knowledge among all players. While $\omega_I, \omega_C \in \mathbb{R}$ are common knowledge among the candidates, voters do not observe these positions.[6] Therefore, θ is a hidden state variable for voters, whereas candidates I and C are able to condition their actions on the realization of this state variable. In particular, given θ, each candidate $j \in P$ simultaneously chooses a manipulative action $a_j \geq 0$ to signal that she is closer to voters than what is suggested by θ. These manipulative actions can take the form of political statements, advertisements or other political actions with an intention to mislead voters.[7]

Manipulative actions shift the state variable θ in opposite directions, resulting in a public campaign signal $\theta + a_I - a_C$. However, the electorate does not observe this signal directly due to its heterogeneous relationship with the communication environment that interferes with the campaign signals. What they observe, instead, is the public signal that is contaminated with an idiosyncratic noise. Formally, they receive a private signal $x_i := \theta + a_I - a_C + \varepsilon_i$, where $\varepsilon_i \in \mathbb{R}$ is the random noise. $\{\varepsilon_i\}_{i \in [0,1]}$ are assumed to be i.i.d across the voters and have a normal distribution with zero mean and precision $\alpha > 0$ ---i.e., variance α^{-1}.

Voters have uninformative priors about θ.[8] Therefore x_i is the only information that voter i

[3]Given the setup of the strategic interactions that we consider in the model, the assumption of a continuum of voters is without loss of generality. Even if we assumed that cardinality of voters is finite, without changing the other assumptions, our arguments about the effectiveness of manipulation would continue to hold. To keep our model tractable, we abstract from such considerations.

[4]Endogenizing political platforms would lead to policy convergence given that the candidates care only about the vote share and would want to economize on the costs of manipulation. However, it turns out that in equilibrium, the net bias in political messages would still be positive because the candidates would have an incentive to misrepresent each other's position to increase their vote share.

[5]Heterogeneity among voters will be introduced when we discuss the information structure characterizing the election campaign.

[6]This assumption is essential to capture the idea that voters and candidates are asymmetrically informed. Without such asymmetry, there cannot be any incentive to manipulate voters' beliefs.

[7]Note that this broad conceptualization of information manipulation includes actions that disparage the opponent, such as attack ads, as well as any statement that paints an overly optimistic image of the manipulator. However, we do not make any distinction between these two forms of manipulation.

[8]Technically, voters' priors follow an improper uniform distribution, which is the limiting distribution of a sequence of symmetric uniform distributions, whose bounds diverge to infinity. We abstain from imposing specific

can use to make a decision. This information may be biased due to both the net manipulation $\bar{a} := a_I - a_C$ (or the systematic bias) and the idiosyncratic noise ε_i, which can also be interpreted as a non-systematic individual bias. In the presence of idiosyncratic noise, voters cannot identify the location of their private signals in the population distribution unless there is perfect communication among them. Hence, we assume, without loss of generality, that voters do not communicate with each other.[9]

Preferences and campaign costs

We denote the political position of the winning candidate as $\omega_W \in \{\omega_I, \omega_C\}$, and each voter receives a payoff $-|\bar{\omega} - \omega_W|$ after the elections. Under this specification, each voter supports the candidate who is expected to hold a political position closest to the common bliss point $\bar{\omega}$. More formally, voter $i \in [0,1]$ supports C if and only if her posterior belief conditional on her private signal x_i satisfies $Pr\left(\theta \leq \frac{1}{2}|x_i\right) \geq \frac{1}{2}$. This voting decision is described by a binary voting function $s_i(\cdot) : \mathbb{R} \rightarrow \{0,1\}$ such that $s_i(x_i) = 1$ if and only if i supports candidate C. After all voters simultaneously vote, the masses of supporters for both parties are determined. We denote the mass of supporters (vote share) of candidate C by $S(a_I, a_C; \theta)$, which is given by

$$S(a_I, a_C; \theta) \equiv \int_0^1 s_i(x_i)di. \tag{1}$$

For a candidate or political party that expects to run in an election only once, the most natural goal would be to win the election. However, many times, maximizing vote share would be a more realistic goal, especially when a larger vote share leads to higher legislative power in the parliament or when it helps the party raise more resources from interest groups. Candidates try to maximize their vote share subject to costs of manipulation.[10] Candidates' payoffs are given by

$$u_C(a_C, a_I; \theta) = 2S(a_I, a_C; \theta) - 1 - H_C(a_C)$$
$$u_I(a_I, a_C; \theta) = 1 - 2S(a_I, a_C; \theta) - H_I(a_I),$$

where $H_j : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, for $j \in P$ govern the costs associated with manipulative actions.[11] We assume the following three properties about the cost functions:

(A1) For each $j \in P$, $H_j(\cdot)$ is a twice continuously differentiable function such that $\forall a > 0$ $H_j'(a) > 0$ and $H_j''(a) > 0$.

priors that are biased towards one or the other candidate to focus on how information manipulation affects voting behavior of voters who were neutral towards both candidates prior to the election campaign.

[9] Our analysis would not be qualitatively affected if we instead assumed imperfect communication. To locate their private signals precisely, they would need to know the mean of $\{x_i\}_{i\in[0,1]}$. Even if each voter asks around to learn about the signals others have received, she cannot estimate the population mean without noise unless each voter can perfectly communicate with all the rest of the electorate.

[10] If, instead, we assumed that candidates care only about winning the election, the model would be less interesting, while our theoretical results would not be affected qualitatively. In that case, the manipulation game between the candidates would lead to corner outcomes where a candidate would choose zero manipulation if she is bound to lose the election, while the other candidate would engage in manipulation but still would not receive all the votes.

[11] Since the paper's focus is on the role of information manipulation, the level of campaigning is completely parameterized by manipulative actions a_j.

(A2) For each $j \in P$, $H'_j(0) = 0$, and $\forall a > 0$, $H''_j(a) = k_j > \bar{\varphi}(\alpha)$ where $\bar{\varphi}(\alpha) = \max_{x \in \mathbb{R}-\{0\}} [2\alpha\phi'(\sqrt{\alpha}x)]$ $\alpha > 0$ and $k_C > k_I$.[12]

(A3) $H'_I(a) < H'_C(a)\ \forall a > 0$.

The first property states that campaigning costs are strictly increasing and convex in manipulative actions. As the level of campaigning increases, it becomes harder for candidates to send additional messages to the public. The second property serves as a sufficient condition for the existence and uniqueness of the equilibrium that will be defined below.[13] This property implies that manipulation involves very low costs at the beginning, but later on, the marginal cost of manipulation rises faster than its marginal effect on voter share. The third property on the cost functions formalizes the manipulative advantage of candidate I. For each additional message, I has to pay a lower cost than C. Due to her cost advantage, we will refer to candidate I also as the incumbent and to candidate C as the challenger, although this labeling does not imply any other difference between them. Cost functions are also common knowledge, but we assume that the costs associated with manipulative actions of political candidates are not observable by voters.[14]

The sequence of events

Figure 1 describes the sequence of events in the election game. After nature determines the relative proximity θ of candidate I to voters vis-à-vis candidate C, each candidate strategically chooses her manipulative action. In its isolated form, the first stage of the election game is a simultaneous-move, two-person, normal-form game with strategic substitutes.

Figure 1: Timing of events in the election game

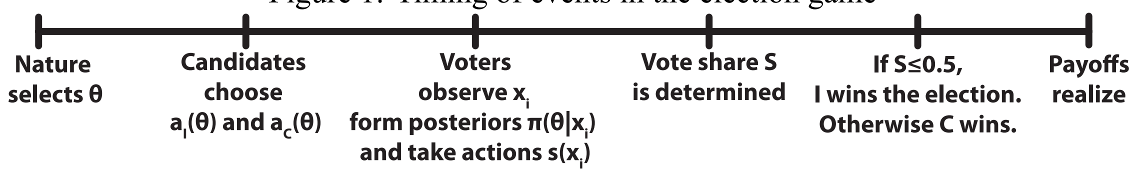

In the second stage, voters form their beliefs about θ conditional on their private signal x_i, which is jointly determined by the hidden state variable θ, manipulation actions of the first stage game and the random draws of individual bias ε_i. Given their beliefs, they vote for the candidate who they

[12] Note that the maximum value of the expression in brackets is finite for a given α, as the function ϕ' converges exponentially to 0 as its argument diverges to infinity. Imposing the lower-bound $\bar{\varphi}(\alpha)$ on the second derivatives of the cost functions guarantees the uniqueness of the equilibrium. This lower bound has to depend on α because the function in the brackets has no finite maximum over (α, x) space. If we relax the assumption of a lower bound on k_j for $j \in P$, there may be more than one equilibrium, but the existence result in Theorem 1 remains intact, and the conclusions in Theorem 2 would apply equally to all equilibria.

[13] As will be seen in the proof of Theorem 1, the equilibrium is characterized by an equation that is a linear combination of first derivatives of cost functions and the probability density function $\phi(\cdot)$. Non-invertibility of $\phi(\cdot)$ on its full support allows for multiple solutions to the equation, which implies multiplicity of equilibria. The condition that we put on second derivatives of the cost function guarantees that the equilibrium is unique.

[14] This is a simplifying assumption that is consistent with the idea that voters may not know how campaign budgets are allocated between activities aiming to deceive voters about candidate qualities, on the one hand, and truly informative campaign activities ---from which our model abstracts on the other. Qualitatively, our results would go through as long as we assume that manipulative campaign spending by candidates is only imperfectly observed.

think is more likely to stand closer to their bliss point $\bar{\omega}$. Resulting vote shares determine the winner of the election and the associated payoffs for all players. Candidate I wins the election if and only if $S(a_I, a_C; \theta) \leq 1/2$, and C wins otherwise. In the next section, we define the equilibrium of the electoral game and present the main theoretical results of the paper.

3.2 Electoral Equilibrium

We employ a perfect Bayesian equilibrium (PBE) concept, which requires that each candidate develops a separate strategy in response to each possible voting strategy by the electorate and chooses the best action given her opponent's action. On the voters' side, the PBE rules out any belief about candidates' actions that contradict the information that voters have. Therefore, in equilibrium, voter beliefs and actions should be consistent with the fact that the incumbent will choose a higher level of manipulation than her opponent since she is the candidate with access to a superior manipulation technology. The following definition lays out the conditions characterizing the electoral equilibrium:

Definition 1 (Electoral Equilibrium) *An electoral equilibrium consists of voter posterior beliefs* $\pi(\cdot)$*, voting behavior* $s(\cdot)$*, mass of supporters of the challenger* $S(\cdot, \cdot; \cdot)$ *and hidden actions* $a_I(\cdot)$*,* $a_C(\cdot)$ *such that* $\forall i \in [0, 1]$*,* $\theta, x_i \in \mathbb{R}$*.*

$$\pi(\theta|x_i) = \frac{f(x_i|\theta, a_I(\theta), a_C(\theta))}{\int_{-\infty}^{\infty} f(x_i|\theta, a_I(\theta), a_C(\theta))d\theta}$$

$$s(x_i) = \begin{cases} 1, & \textit{if } \int_{-\infty}^{1/2} \pi(\theta|x_i)d\theta \geq \frac{1}{2} \\ 0, & \textit{if otherwise} \end{cases}$$

$$S(a_I(\theta), a_C(\theta); \theta) = \int_{-\infty}^{\infty} s(x_i) f(x_i|\theta, a_I(\theta), a_C(\theta))dx_i$$

$$a_I(\theta) \in arg \max_{a \geq 0} \{1 - 2S(a, a_C(\theta); \theta) - H_I(a)\}$$

$$a_C(\theta) \in arg \max_{a \geq 0} \{2S(a_I(\theta), a; \theta) - 1 - H_C(a))\}$$

In the first expression above, $f(x_i|\theta, a_I(\theta), a_C(\theta))$ is the p.d.f. of private signals x_i conditional on θ and the manipulative actions of the candidates in response to θ. Posterior voter beliefs $\pi(\theta|x_i)$ about θ have been written under the assumption that voters form their posterior beliefs starting with an uninformative prior that can be represented as an improper prior, such as a uniform distribution over an infinite interval.

Because of their simple structure and intuitive appeal, we will concentrate on equilibria in pure strategies where voters use monotonic voting rules. Under a monotonic strategy, a voter will support the incumbent candidate if and only if her private signal is above a threshold.[15] Formally, the

[15] Our focus on the class of monotonic equilibria is appealing because they are conceptually closer to simple rules of thumb that voters employ than more complicated non-monotonic strategies are. Since our model does not involve any strategic complementarity among voters' behavior, we cannot rule out mixed-strategy or non-monotonic pure-strategy equilibria. We believe that an analysis of such counter-intuitive equilibria will not contribute to the conceptual

monotonic strategy entails a threshold $\hat{x} \in \mathbb{R}$ such that $s(x_i) = 1 \Leftrightarrow x_i > \hat{x}$.

To highlight the impact of manipulative campaigns on aggregate voting behavior, we first analyze the electoral equilibrium without manipulation as a benchmark case. The following proposition shows that when the candidates cannot engage in manipulation ---i.e., $\forall i \in [0, 1]$, $x_i = \theta + \varepsilon_i$--- the candidate who stands closer to voters always wins the election.

Proposition 1 (Equilibrium without manipulation) *If manipulation is not feasible ---i.e., $x_i = \theta + \varepsilon_i$--- there is a unique electoral equilibrium. The voter behavior in this equilibrium is such that for any voter $i \in [0, 1]$, i votes for C if and only if her private signal is greater than or equal to $1/2$ ---i.e., $s(x_i) = 0 \Leftrightarrow x_i \geq 1/2$. Moreover, I wins the election if and only if she is closer to voters than C is ---i.e., $\theta \geq 1/2$.*

Proof Since manipulation is not feasible, $a_I(\theta) = a_C(\theta) = 0$ for all θ, and this is common knowledge. So the p.d.f. of private signals $f(.)$ does not depend on a_I and a_C. In particular, $f(x_i|\theta) = \sqrt{\alpha}\phi(\sqrt{\alpha}(x_i - \theta))$. Note that voter i with signal x_i supports C if and only if

$$\begin{aligned}
\int_{-\infty}^{1/2} \pi(\theta|x_i)d\theta > \int_{1/2}^{\infty} \pi(\theta|x_i)d\theta \Leftrightarrow \\
\int_{-\infty}^{1/2} f(x_i|\theta)d\theta > \int_{1/2}^{\infty} f(x_i|\theta)d\theta \Leftrightarrow \\
\Phi(\sqrt{\alpha}(x_i - 1/2)) - 1 < -\Phi(\sqrt{\alpha}(x_i - 1/2)) \Leftrightarrow \\
\Phi(\sqrt{\alpha}(x_i - 1/2)) < \frac{1}{2},
\end{aligned}$$

where $\Phi(\cdot)$ is the c.d.f. of the standard normal distribution.

Since $\Phi(\cdot)$ is a strictly increasing and continuous function, by intermediate-value theorem, there is a unique x^* such that $\Phi(\sqrt{\alpha}(x_i - 1/2)) \leq 1/2$ if and only if $x_i \leq x^*$. Since $\Phi(0) = 1/2$, $x^* = 1/2$. Then, the vote share of C is $S(\theta) = \int_{-\infty}^{1/2} f(x_i|\theta)dx_i = \int_{-\infty}^{1/2} \sqrt{\alpha}\phi(\sqrt{\alpha}(x_i - \theta))dx_i = \Phi(\sqrt{\alpha}(1/2 - \theta))$. Then, $S(\theta) > 1/2$ if and only if $\theta < 1/2$, as required. ■

In what follows, we demonstrate that the conclusion in Proposition 1 cannot be generalized to election campaigns in which candidates can engage in manipulative campaigning. Before doing so, we establish that the election model with manipulation has a unique monotonic equilibrium. In the equilibrium, both the behavior of voters and the outcome of the election are determined by unique thresholds. Voters vote for the incumbent if and only if they observe a signal above a threshold x^*, and the incumbent wins the election if and only if the relative proximity parameter θ is above a threshold θ^*.

Theorem 1 (Existence and Uniqueness of Equilibrium with Manipulation) *Suppose that (A1)--(A3) hold. Given the state variable θ, there exists a unique monotonic electoral equilibrium in pure strategies. It is characterized by two thresholds x^* and θ^*, such that voter i supports candidate I if and only if her private signal is above x^*,i.e., $s(x_i) = 0 \Leftrightarrow x_i \geq x^*$, and I wins the election if and*

development (or understanding) of manipulation in election campaigns anyway.

only if $\theta \geq \theta^*$ ---i.e., I's political proximity to voters is sufficiently high relative to the proximity of candidate C. In this unique equilibrium, manipulative actions are strictly positive, and candidate I always manipulates to a higher extent than candidate C ---i.e., $a_I(\theta) > a_C(\theta) > 0$ for any $\theta \in \mathbb{R}$--- implying that net manipulation $\bar{a}(\theta)$ is strictly positive.

Proof See the Appendix.

Theorem 1 establishes that there is a unique equilibrium in which (i) voters devise a monotonic strategy that takes into account the fact that the "war of disinformation" between candidates would distort the campaign signals in favor of the incumbent, and (ii) the candidates' best response to this common voting strategy and to their opponents' manipulative actions results in a net bias in campaign signals that favors the cost-efficient manipulator ---i.e., the incumbent. That both candidates are compelled to do positive manipulation follows from the fact that for a given monotonic voting strategy, candidates can steal votes from each other by shifting the mean of the campaign signals. That the net manipulation $\bar{a}$ is strictly positive stems from I's cost advantage, which enables her to distort the signals by more than what her opponent finds optimal. Aware of I's cost advantage, voters use a voting rule that tries to filter out the bias in campaign signals. This reinforces the necessity of manipulation from the candidates' point of view.

However, Theorem 1 by itself does not necessarily imply that voters are successfully manipulated. The natural question is whether the incumbent, through manipulation, can secure a victory that she would not be able to obtain in the absence of manipulation ---i.e., a victory when $\theta < 1/2$. The theorem suggests that the answer depends on the magnitude of a unique threshold, which, if sufficiently low ---i.e., lower than $\frac{1}{2}$--- implies that the incumbent can get an undeserved victory whenever relative proximity of I satisfies $\theta \in \left[\theta^*, \frac{1}{2}\right)$. As will be evident from the proof of the theorem, there is a closed form solution to neither the net manipulation function nor the equilibrium thresholds x^* and θ^*. Therefore, one cannot address this important question analytically under a set of finite parameters. Yet, it is possible to analyze the limiting behavior of x^* and θ^* as the precision of campaign signals α goes to infinity. The next theorem demonstrates that, in the limit, the incumbent wins the election under any realization of θ.

Theorem 2 (Signal Precision and the Effectiveness of Manipulation) *As the precision of campaign signals, α, goes to infinity, both x^* and θ^*, which are defined in Theorem 1, diverge to $-\infty$.*[16]

Proof See the Appendix.

Theorem 2 implies that for any realization of θ, there must be a high enough signal precision α at which manipulation is capable of changing the identity of the winning candidate by resulting in a value of θ^* that is below $\frac{1}{2}$. Moreover, as precision goes to infinity, candidate I wins the election under almost any realization of θ.

[16]Here, we assume a condition similar to a transversality condition. We restrict the limiting behavior of the equilibrium system indexed by α with an auxiliary assumption, which is stated in the Appendix.

Numerical solutions to the system of equations, which characterizes the equilibrium defined in Definition 1, show that manipulative campaigns can tilt the election results in favor of the incumbent candidate. Figure 2 plots the equilibrium level of net manipulation $\bar{a}(\theta)$ as a function of relative proximity of candidate positions to voters' bliss point.[17]

Figure 2: Equilibrium net manipulation ($\alpha = 1$, $k_I = 1$ and $k_C = 5$)

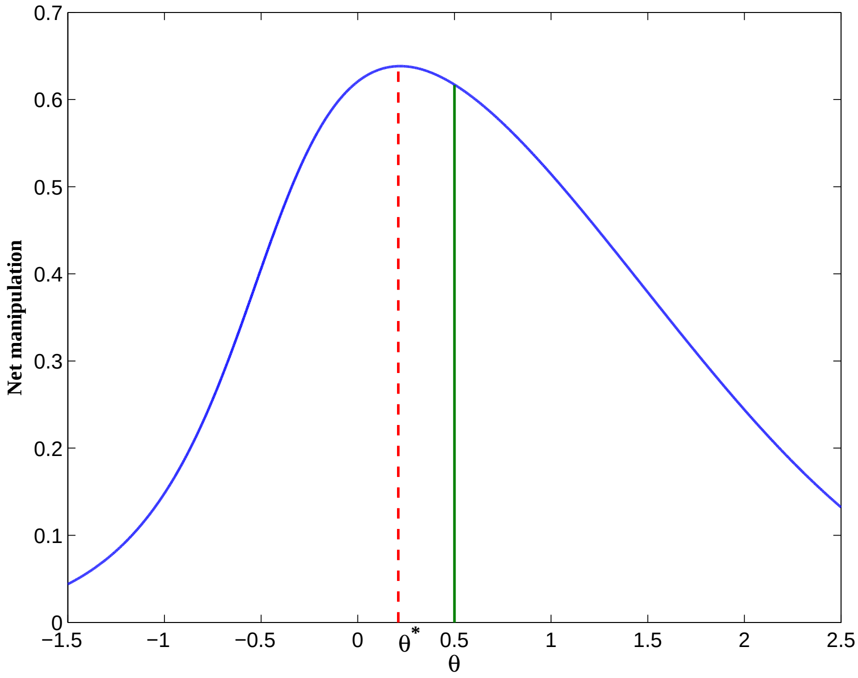

The value of relative candidate-voter distance θ under which each candidate receives exactly half of the votes is denoted by θ^* in the figure. The figure illustrates that $\theta^* < \frac{1}{2}$ under the given parameter values for α and the second derivatives of the cost functions, denoted by k_I and k_C. Hence, the figure suggests that when the relative proximity of I to voters falls between θ^* and $1/2$, candidate I wins an election, that would have been won by C in the absence of manipulation (see Proposition 1). When, for example, $\theta = 0.5$, so that the candidates are equally distant from voters, the latter would be indifferent between supporting I or C, had they been able to observe θ. Nonetheless, since $\theta^* < 1/2$, by Theorem 1, I receives a higher share of votes and wins the election. Only when θ is sufficiently low --- i.e., below θ^*--- does candidate C achieve a deserved victory in the election.

The contrast between the following two figures illustrates the impact of manipulation on election outcomes. Figure 3 plots the vote share of the cost-advantaged candidate as a function of her relative proximity to voters at various levels of signal precision α, under the assumption that manipulation is not feasible. As expected, I's vote share increases with θ and exactly equals $1/2$ when she and her opponent are equally close to voters' bliss point, regardless of the value of α. Hence, as already proved in Proposition 1, whichever candidate is closer to voters gets to win the election.

[17] In all numerical solutions, we assume that voters and the candidates form their strategies as if the relative proximity of candidate I θ has a bounded support and can take a finite number of equally spaced values on the interval $[-1.5, 2.5]$ with a mid-point of $1/2$. We also assume that the cost functions take the form $H_j(a) = \frac{1}{2}k_j a^2$ for $j = C, I$

Figure 3: Vote share of candidate I when manipulation is not feasible

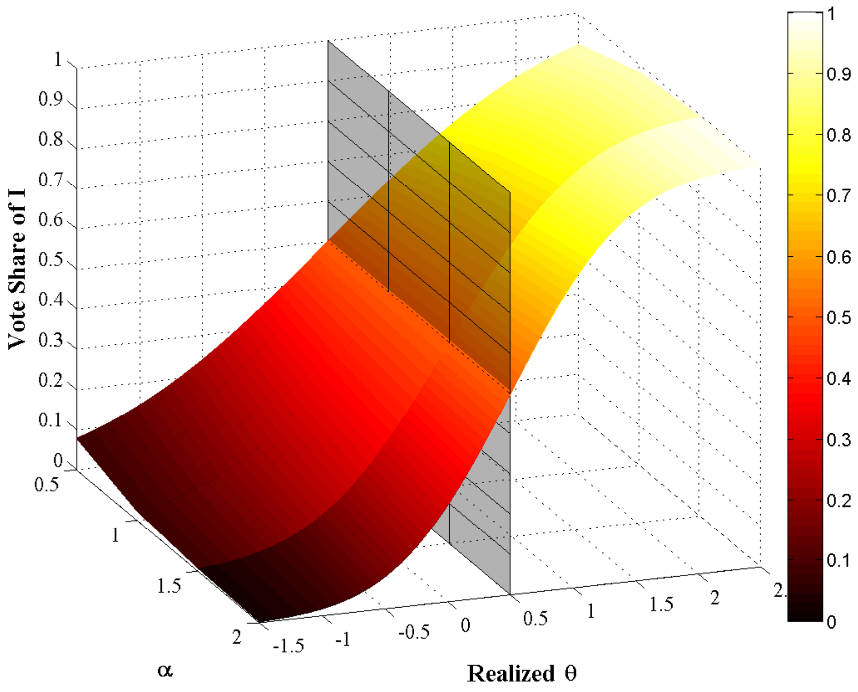

In Figure 4, we numerically solve and plot the equilibrium vote share of I when manipulation is feasible. In contrast to the case without manipulation, the incumbent candidate gets more than half of the votes when $\theta = 1/2$, and this electoral advantage increases with α. The net bias (manipulation) in signals has an adverse overall effect on voter choices.

At this stage one might wonder why voters do not ignore the manipulative signals altogether. The answer lies in the fact that, in our model, manipulative campaign signals are the only sources of information for voters, and although these signals are somewhat misleading, they reveal some useful information about the hidden state variable θ. To clarify the point, consider a slightly modified model in which θ is priorly distributed according to normal distribution with mean $1/2$. Then, if a voter ignores the signals, she will be left with her unbiased prior, which leaves her with a success rate of $1/2$. However, if she takes the signal into account and forms her beliefs in the optimal way, then she expects to choose the wrong candidate only when θ falls into the interval $(\theta^*, 1/2)$, of which the probability of occurrence is clearly smaller than $1/2$.

4 Why does manipulation work? Necessary conditions for successful manipulation

Comparing the electoral equilibria with and without manipulation, it is compelling to conclude that the only reason for the equilibrium bias in voters' beliefs is the presence of manipulative actions in an environment characterized by informational asymmetry between the candidates and the electorate. However, a closer look at the model reveals that (i) the cost asymmetry between the candidates and (ii) an idiosyncratic noise in private signals under the lack of perfect communication

Figure 4: Vote share of candidate I when manipulation is feasible ($k_I = 1$, $k_C = 5$)

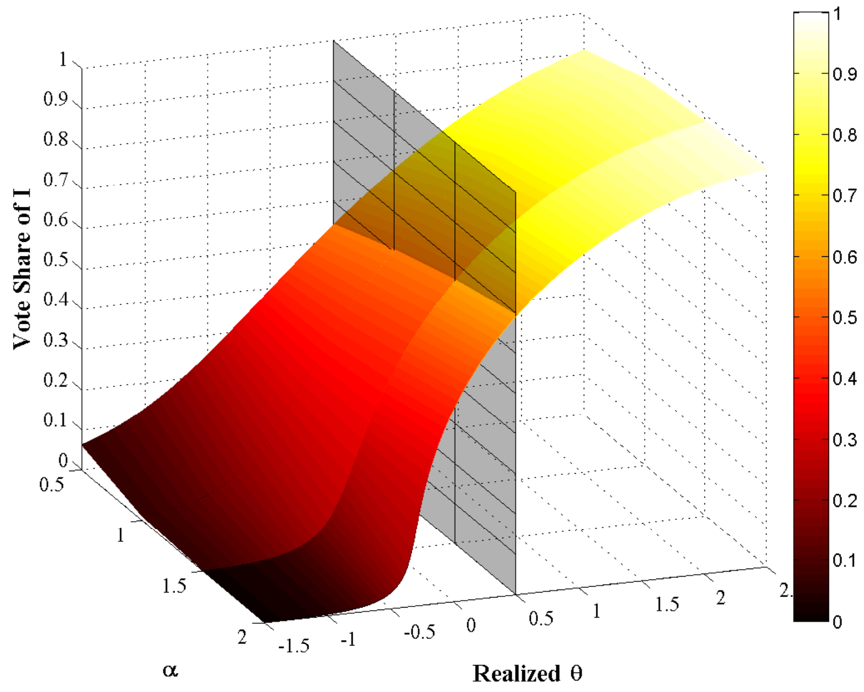

among voters are equally necessary for manipulation to be effective in equilibrium.

4.1 Asymmetric manipulation costs

If there is no cost asymmetry between I and C, it is straightforward to show by going through the proof of Theorem 1 that both candidates would choose the same level of manipulation, and the net manipulation will always be zero. Figure 5 demonstrates that as the manipulation technology of candidate I, captured by $H_I''(a) = k_I$, converges to that of candidate C, $k_C = 5$, the level of net manipulation approaches zero, while the equilibrium threshold θ^* --at which $\bar{a}(.)$ assumes a maximum-- approaches $1/2$, suggesting a move towards the socially efficient election outcome associated with the equilibrium without manipulation. When $\bar{a} = 0$, the median voter with $\varepsilon = 0$ will receive an unbiased signal, and by Proposition 1, the majority will choose the right candidate, the one whose position is closer to voters' bliss point.

4.2 Noisy campaign signals and imperfect communication

Without the uncertainty voters face about relative candidate proximity θ, there would clearly be no incentive for manipulative campaigns. What is less straightforward is the essential role of idiosyncratic noise in campaign signals in rendering manipulation effective and, hence, providing incentives for costly manipulation. In fact, voters' uncertainty about θ, by itself, would not be sufficient for the low-cost candidate to bias the posterior beliefs of voters in equilibrium. This is so because the consistency requirement of a perfect Bayesian equilibrium would rule out any such bias in the absence of some heterogeneity in private signals. To formally demonstrate this

Figure 5: Equilibrium net manipulation as candidate I becomes more cost-efficient ($\alpha = 1$ and $k_C = 5$)

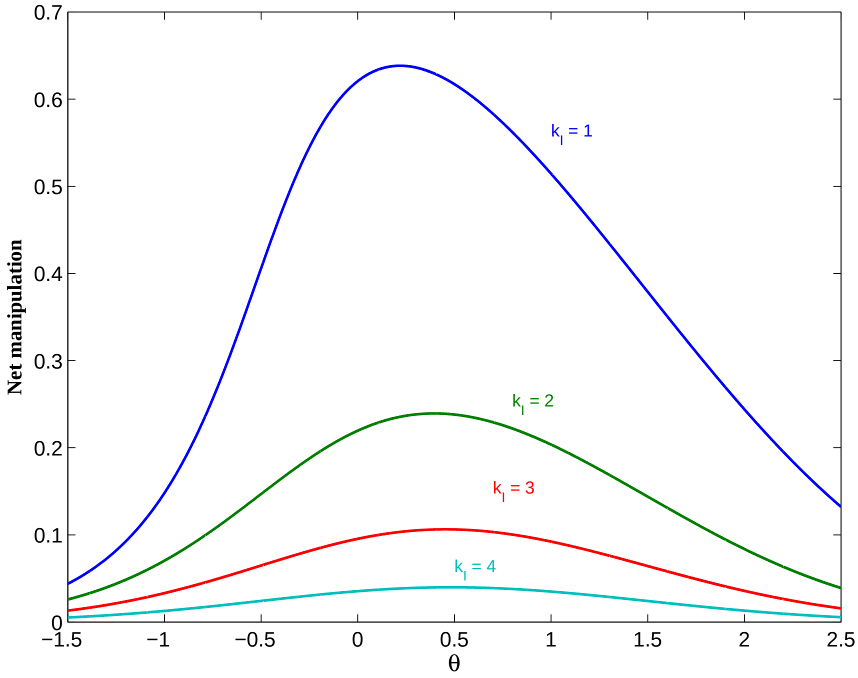

point, we next study the case of purely public signals, where each voter receives a common signal $x = \theta + a_I - a_C$ without any individual noise. Clearly, x becomes common knowledge after its realization.

The following proposition states that there is no equilibrium in which the cost-efficient candidate ---i.e., candidate I--- has an ex-ante advantage in the elections when voters receive a common public signal without any individual noise (or bias).

Proposition 2 *In the public-signal case, where $\varepsilon_i = 0$ for all $i \in [0, 1]$, in any equilibrium, the cost-efficient candidate is more likely to win the election if and only if she is closer to voters than the other candidate ---i.e., $Pr(C \text{ wins the election}) > 1/2 \Leftrightarrow \theta < 1/2$.*

Proof See the Appendix.

Proposition 2 demonstrates that Theorem 1 and the limit result in Theorem 2 hinge on the assumption that voters observe manipulated signals with an idiosyncratic noise. This noise can be broadly interpreted as a byproduct of a social communication environment in which information received by candidates exhibits diversity due to different social habits of information collection and processing, or due to differential access to media channels and social communication networks. Studying such a social structure is far beyond the scope of this paper. It is sufficient for our purpose to represent this structure as an exogenous random process that converts social actions by influential players into noisy signals to be received by the electorate. ε_i can also be interpreted as the level of ideological bias in voters' beliefs about the candidates, relative to that of the median voter. This relative bias is unknown to each voter, but its distribution is common knowledge.

How does the individual noise in campaign signals enable the candidates to manipulate voters? When campaign messages contain an idiosyncratic noise, voters become uncertain about the population mean of the signals assuming that they do not communicate perfectly. Therefore, when evaluating their private signal x_i, they cannot perfectly differentiate between the systematic bias $\bar{a}$ and the individual bias ε_i that it contains. This informational problem gives the candidates an incentive to tailor their campaign messages in a particular way so as to bias voters' inference. In the next section, we discuss in more detail exactly how the cost-efficient candidate conceals part of the systematic bias $\bar{a}(\theta)$ in campaign signals behind a cloud of idiosyncratic noise and distorts voters' posterior beliefs in her favor.

4.3 Voters' inference problem: How does manipulation work?

As proved in Theorem 2 and illustrated in Figure 2, when election campaigns are manipulative, there exists an interval of relative candidate positions $\left[\theta^*, \frac{1}{2}\right]$ over which the cost-efficient candidate ends up winning the election, despite being an inferior candidate ex-post. In this section, we address the question of how θ^* could be smaller than 1/2.

First, note that voters are uncertain about candidates' positions, while the candidates are able to observe θ and adjust their manipulation level based on its particular realizations. When the communication environment that intermediates campaign messages adds an individual noise on manipulated public signals, voters cannot perfectly distinguish between the bias due to manipulative actions and the bias due to the individual noise that they cannot observe. Candidates exploit this additional uncertainty by tailoring their manipulative actions in such a way that voters' posterior beliefs remain biased in equilibrium.[18] In particular, the shape of the equilibrium net manipulation function $\bar{a}(\theta)$ is such that voters end up overestimating the likelihood that θ is above 1/2 when θ lies in a region where gains to manipulative actions for the incumbent are highest.

Figure 6 offers a simplified illustration of how this happens. The curve that lies above the 45-degree line sketches how public signals $\theta + \bar{a}(\theta)$ change as a function of the hidden state variable θ. This public-signal curve lies above the 45-degree line because the competition between incumbent and challenger candidates leads to a positive net bias $\bar{a}(\theta) = a_I(\theta) + a_C(\theta)$. Although voters do not observe θ, since their beliefs cannot contradict what they know about the model, in equilibrium, they form their voting strategies based on the full knowledge of this function. As re-

[18]As previously stated in footnote 3, assuming a finite number of voters --instead of a continuum-- would not change this result qualitatively. If we keep the assumption that voters' behavior is not strategic in the sense that voters form their beliefs exclusively conditional on their private signals --instead of taking other voters' behavior also into account-- then voters will not be able to differentiate between individual bias and systematic bias. This in turn leads to effective manipulation. However, in the case of finite voters, strategic considerations might arise due to costly voting. In this case, voters would choose to abstain from voting unless they think they are pivotal in the election. Those voters, who end up casting a vote, would condition their voting strategy on the event that they are pivotal, i.e., on those realizations of campaign signals that would make them pivotal. However, being pivotal would reveal only the number of private signals that are greater than the threshold that voters condition their behavior on, not the values of the signals themselves, as long as the noise in signals has an infinite support. Thus, even though voters' posterior beliefs would be different from what we find in this model, there would still remain some uncertainty on the voters' side, which would prevent them from differentiating individual bias and systematic bias. We chose to work with continuum voters to crystallize our analysis on the emergence of manipulation in election campaigns by renouncing such strategic considerations and employ the Law of Large Numbers.

flected in Figure 2, the shape of $\bar{a}(\theta)$ resembles a distorted normal curve. Therefore, the difference between public signals and the true value of the state variable θ increases as the latter approaches θ^* ---i.e., the lowest level of θ under which the incumbent wins the election. Hence, the difference between the level of public signals and θ assumes its maximum value at θ^*. Beyond this point, $\bar{a}$ starts declining and so does the difference between public signals and θ.

Figure 6: How does manipulation work?

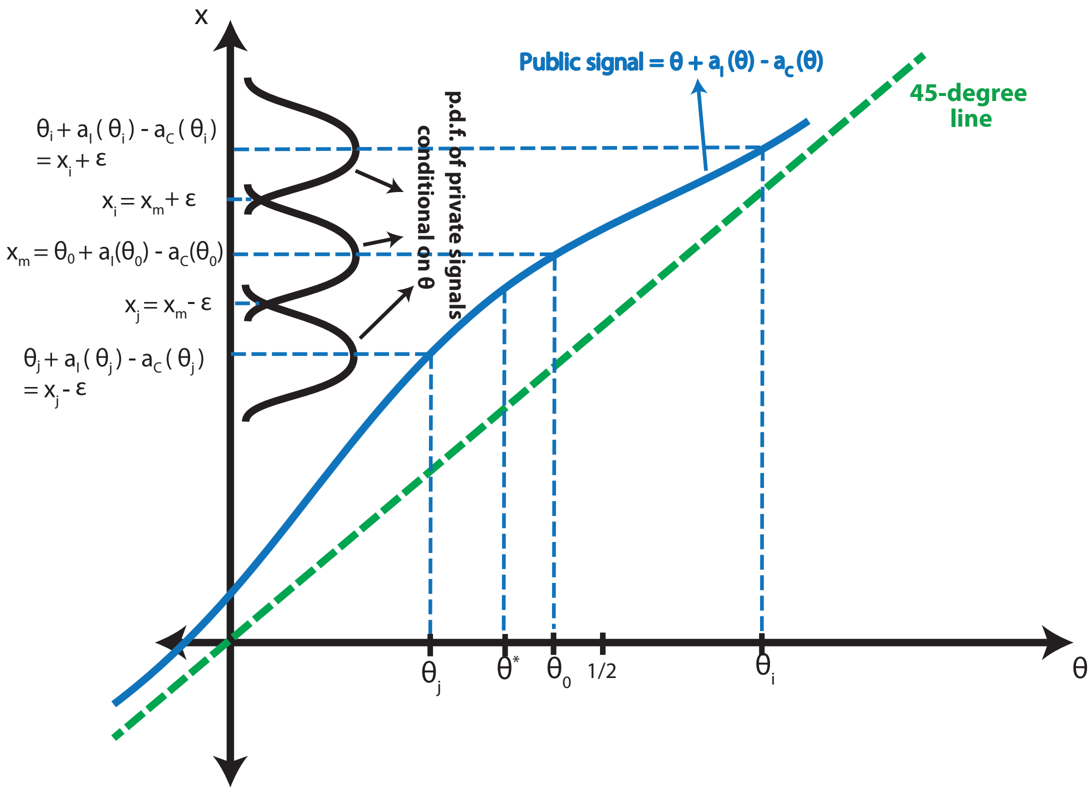

If, as in Figure 6, the true state is $\theta_0 < 1/2$, then voters would be better off ex-post if the challenger had won the election. However, since $\theta_0 \geq \theta^*$, the incumbent ends up winning the election. To see why, consider two voters $i, j \in [0, 1]$. Under θ_0, the equilibrium level of public signal will be given by x_m. In the model, individual biases are normally distributed around a zero mean, and the normal curves in the figure represent the p.d.f.s of private signals x under different possible realizations of θ. For the sake of argument, suppose, for a moment, that the noise in signals can take one of two possible values, ε or $-\varepsilon$, with equal probability. Voter i, who draws an individual bias ε, would only observe a signal $x_i = x_m + \varepsilon$. Note that her priors are such that we can imagine that she assigns equal likelihood to all realizations of θ.[19] Knowing how equilibrium net manipulation behaves as a function of θ, she deduces that, conditional on observing x_i, the probability that she is in state θ_0 *and* has received a positive shock ε should be equal to the probability that the state is $\theta_i > \theta_0$ *and* has received a negative shock $-\varepsilon$.This equivalence is illustrated in Figure 6. Then, formally, we have $Pr(\theta = \theta_0|x_i) = Pr(\theta = \theta_i|x_i)$. Similarly, consider voter j, who receives $\varepsilon_j = -\varepsilon$. By the same logic, for her it must be that

[19]Remember that in the model, voters were assumed to have uniform priors, whose support approaches to infinity, and which become uninformative at the limit.

$Pr(\theta = \theta_0|x_j) = Pr(\theta = \theta_j|x_j)$. Therefore, for voter i [voter j], the event that $\theta = \theta_i$ [$\theta = \theta_j$] and $\theta = \theta_0$ will be assigned equal weight conditional on signal x_i.

Although voters i and j receive signals that are equally distant from x_m, due to the particular shape of the net manipulation function $\bar{a}(\theta)$, we have that $\theta_i - \theta_0 > \theta_0 - \theta_j$. Hence, if these voters were to form an expectation of θ conditional on their signals, voter i would overestimate θ by more than voter j would underestimate it. Obviously, in the model, a continuum of pairs of voters i and j will consider a continuum of possible combinations (θ, ε_i) and (θ, ε_j) that would result in signals x_i and x_j, respectively, when forming their expectations of the true state of the world. Yet, this would not change the conclusion that a larger mass of voters would end up believing that $\theta > 0.5$ and, therefore, voting for the incumbent when, in reality, $\theta = \theta_0 < 0.5$.

5 Comparative statics and policy implications

In this section, we first numerically analyze how the equilibrium bias (net manipulation) in campaign messages, total election spending and the outcome of elections change in response to changes in the candidates' relative campaigning efficiency, as captured by cost asymmetry in manipulation. To complement this analysis, we present a small extension of our main electoral model to demonstrate why limiting campaign contributions to political candidates may be desirable when costs of manipulation are asymmetric, and candidates must cater to non-median voter interests to access campaign funds. Then, we demonstrate through numerical simulations how a change in the degree of heterogeneity in private campaign signals (signal variance) affects the ability of the more-efficient candidate to manipulate the electorate over different regions of θ.

5.1 The degree of asymmetry in campaigning efficiencies

Gul and Peserdorfer (2009) show that, as the difference between candidates' costs of campaigning vanishes, voters' welfare increases. Theorem 1 had a similar implication. As the difference in marginal costs of campaigning diminishes, so does the level of net manipulation --as we have already seen in Figure 5--, and the systematic bias in voters' posterior beliefs disappears.

In most of the real settings, however, politicians and political parties are not equally efficient in spending their campaign resources. In our model, this fact has been captured by asymmetric marginal costs of manipulation between the candidates.This asymmetry can best be interpreted as a difference in campaigning efficiency between the candidates: Each additional dollar spent by the incumbent on manipulative actions leads to greater manipulation than a dollar spent by the challenger.

In the information age, rapid advancements in social communication technologies offer more-efficient campaigning tools, and candidates do not always catch up with these new innovations at the same rate. This may lead to a larger gap in campaigning efficiency. The 2008 U.S. presidential elections is a good example. In their book *The Obama Victory: How Media, Money, and Message Shaped the 2008 Election* Kate Kenski, Bruce W. Hardy, and Kathleen Hall Jamieson argue that among the main factors that gave the "Obama model" the upper hand against the McCain campaign were the former's relative effectiveness in the use of digital technology and the Internet to inform and mobilize the electorate and an innovative strategy of interpersonal, cable and radio microtargeting to sway citizens that were more susceptible to persuasion.[20] In our model, such a widening of the efficiency gap between the candidates induces more manipulation by the advantaged candidate at any level of θ. Whether this has a direct negative impact on the identity of the election's winner -and, hence, on voters' welfare-- depends on the true level of θ. If we denote the equilibrium winning thresholds before and after the increase in candidate I's manipulative efficiency by θ_0^* and θ_1^* respectively, then we have $\theta_1^* < \theta_0^* < 1/2$ because the net manipulation and the equilibrium bias increase in favor of candidate I following a reduction in k_I. Therefore voters will be adversely affected by an efficiency gain of candidate I if and only if $\theta \in [\theta_1^*, \theta_0^*]$. In other words, unless it is the the more-distant candidate who experiences the efficiency gain, and unless the decline in her marginal costs is sufficiently strong, the shift in voter shares will not be enough to make voters worse off.

However, this does not mean that voters are immune from indirect welfare costs associated with higher campaign spending. As evident from the left panel of Figure 7, the campaign costs of I rise over the entire range of θ as the cost parameter k_I decreases. This is because the convexity of the cost functions ensures that the increase in I's manipulative actions a_I more than offsets the reduction in costs due to higher efficiency.

The less-efficient candidate, however, does not experience any change in her manipulative efficiency since k_C is assumed to stay constant. Therefore, the change in campaign costs of C is solely due to an adjustment of a_C as a response to the increase in a_I. As seen in the right panel of Figure 7, whether campaign costs of C rise or fall depends on the range of θ. When, I has some, but not too much, positional advantage over C in terms of the distance from voters, marginal gain for C from manipulative actions decreases as a_I goes up. Then, C chooses to lower a_C to economize on campaign costs. When, on the other hand, C has some positional advantage over I - i.e., when θ is sufficiently, but not too, low - the stakes for C in terms of potential votes she may lose to I is high. Therefore, a rise in her opponents' manipulation actions raises the marginal value of manipulation for C. At the extremes - i.e., when one candidate has substantial positional advantage over the other - qualitative conclusions are reversed. Since net manipulation is very low at these values of θ, only a very small mass of voters - those at the tail of the normal distribution - receive signals that are near the voting threshold x^*. Therefore, the shift in voters' strategy x^*, triggered by the reduction in k_I, does not affect the marginal gain in votes from moving the mean of the campaign signals through manipulation as much as it does at intermediate values of θ.

Therefore, when she has substantial positional advantage, candidate C starts to spend less on

[20]The authors quote David Plouffe, the manager of Barack Obama's 2008 presidential campaign, recalling how, at some point, the campaign's e-mail list had reached 13 million people and how they directly reached about 20 percent of the votes they would need to win the elections by creating their own television network.The book also offers many interesting campaign-related statistics. For example, they cite a 2009 report by Edelman Research titled "The Social Pulpit: Obama's Social Media Toolkit," according to which Internet users spent over 14 million hours on YouTube watching more than 1,800 Obama campaign-related videos. The authors also report that during a typical general-election week, among the respondents to the National Annenberg Election Survey (NAES), those who reported having received campaign information from Internet sources in the previous week were 1.033 percentages more likely to vote for Obama.

Figure 7: Campaign costs of the candidates as cost advantage of I increases ($\alpha = 1$, $k_C = 5$)

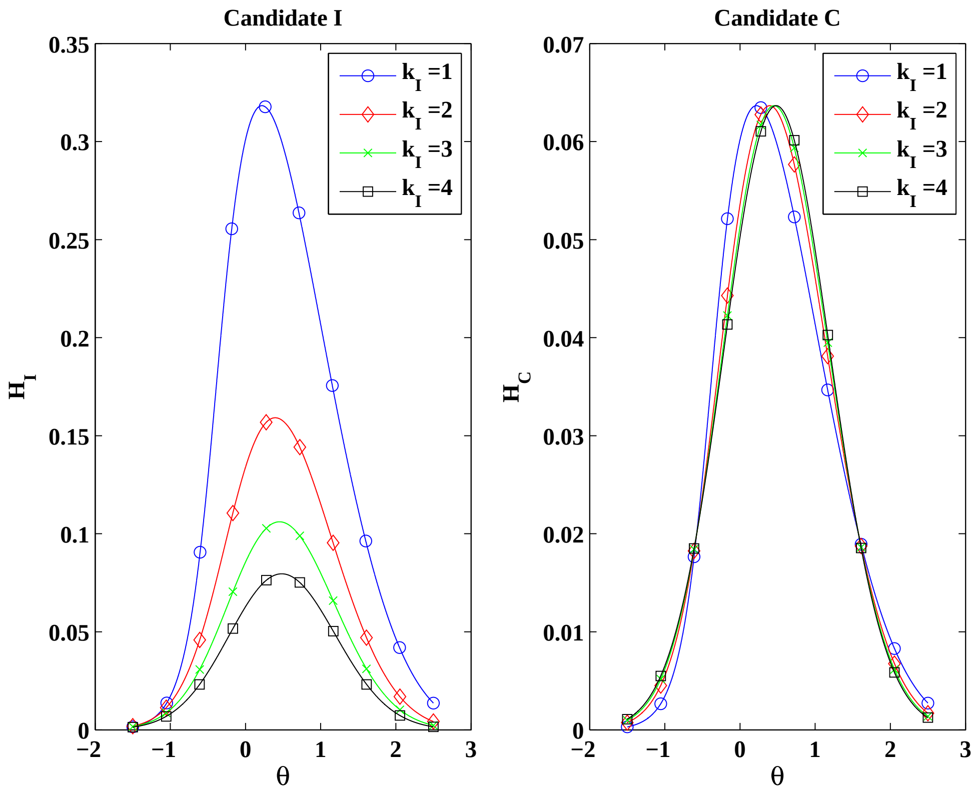

manipulation because the increase in a_I is so minimal that reducing a_C leads to a fall in her costs that more than offsets her vote loss. However, when candidate I is significantly close to voters than C is, she raises a_I by enough for candidate C to feel compelled to respond with higher a_C.

Although candidate C's costs do not increase uniformly, Figure 8 suggests that total manipulation and total campaign spending by the two candidates unambiguously increase despite the fall in k_I. These numerical simulations suggest that a more pronounced incumbent advantage in the form of higher efficiency of campaign spending can be detrimental to social welfare both because wasteful campaign spendings increase and because, in some circumstances, the incumbent may end up winning the election even when the other candidate was closer to voters' bliss point. The potential policy implication is that reducing the cost asymmetry between the candidates can improve welfare not only by limiting disinformation caused by manipulation, but also by lowering the amount of campaign spending.

5.2 The role of informational heterogeneity

The assumptions that manipulated signals arrive with a noise and that voters do not communicate with each other perfectly are crucial for the conclusions in Theorems 1 and 2 to go through. Indeed, as shown in Proposition 2, without any informational heterogeneity among voters, manipulative campaigns do not lead to a systematic bias in voter beliefs. This section analyzes how a change in the degree of informational heterogeneity, as captured by the variance in private campaign signals, affects the electoral equilibrium.

Figure 8: Total campaign costs and manipulative actions as cost advantage of I increases ($\alpha = 1$, $k_C = 5$)

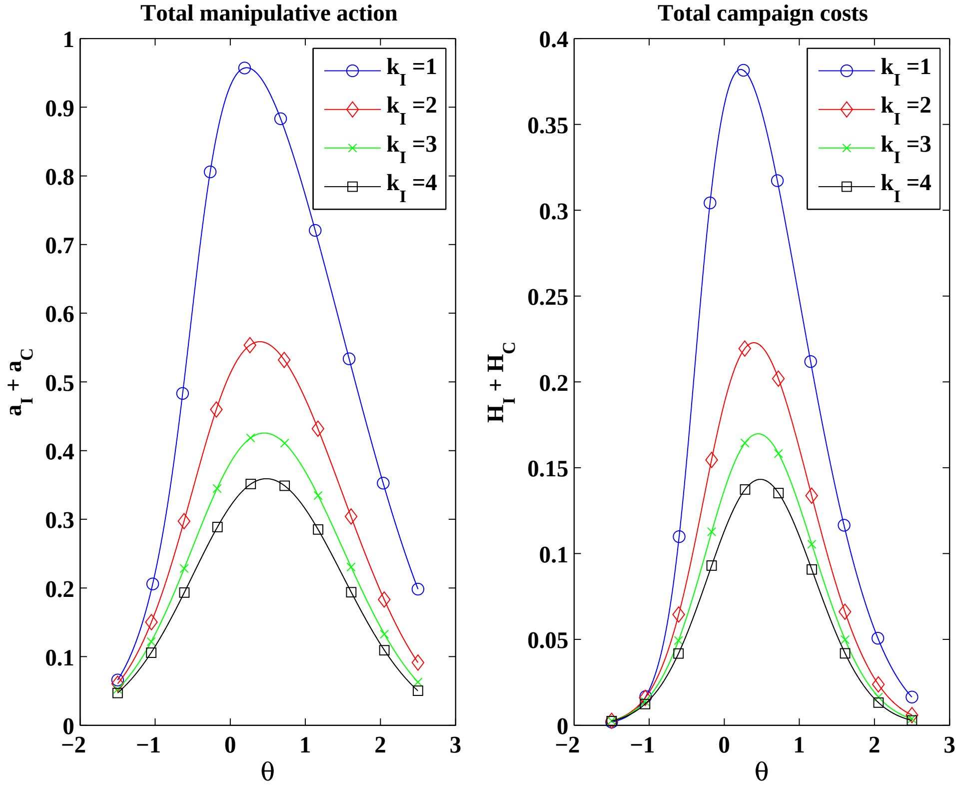

A decrease in the variance of the idiosyncratic noise means that the signals become more precise (higher α). One might be tempted to think that higher precision would lead to better information for voters by lowering the uncertainty about the true nature of the manipulation. However, this is usually not the case. As we have already seen, the net manipulation declines as θ moves away from θ^* towards the extremes. One implication of this is that the mean campaign signal is closer to the voting threshold $x^* = \theta^* + \bar{a}(\theta^*)$ when θ is near θ^*. When campaign signals become more precise, relatively more voters will tend to receive signals that are closer to the mean signal $\bar{x}(\theta) = \theta + \bar{a}(\theta)$, which, if θ is near θ^*, means that a greater mass of signals will accumulate around the indifference threshold x^*. Therefore, a given shift in the mean campaign signal through higher manipulation allows the candidates to steal a larger fraction of votes than before. As θ moves towards the extremes, at some point, the mean signal will be sufficiently distant from x^* so that only those voters at one of the far tails of the signal distribution will be indifferent between voting for I and C. Moreover, their size will shrink further with greater signal precision, which, in turn, will lower the marginal gain in vote share from manipulation.

To illustrate the point, Figure 9 plots the net manipulation function at various levels of signal precision. As expected, in middle ranges of θ, where $\bar{x}(\theta)$ is sufficiently close to x^*, more-precise signals lead to greater net manipulation. At the tails, the opposite is true. Overall, this simulation suggests that campaign signals that are more spread out across the population --perhaps due to a more decentralized media-- tend to discourage manipulation of voters unless candidates do not hold wildly different political views relative to each other.

Figure 9: Net manipulation function as signal precision α increases ($k_I = 1$, $k_C = 5$)

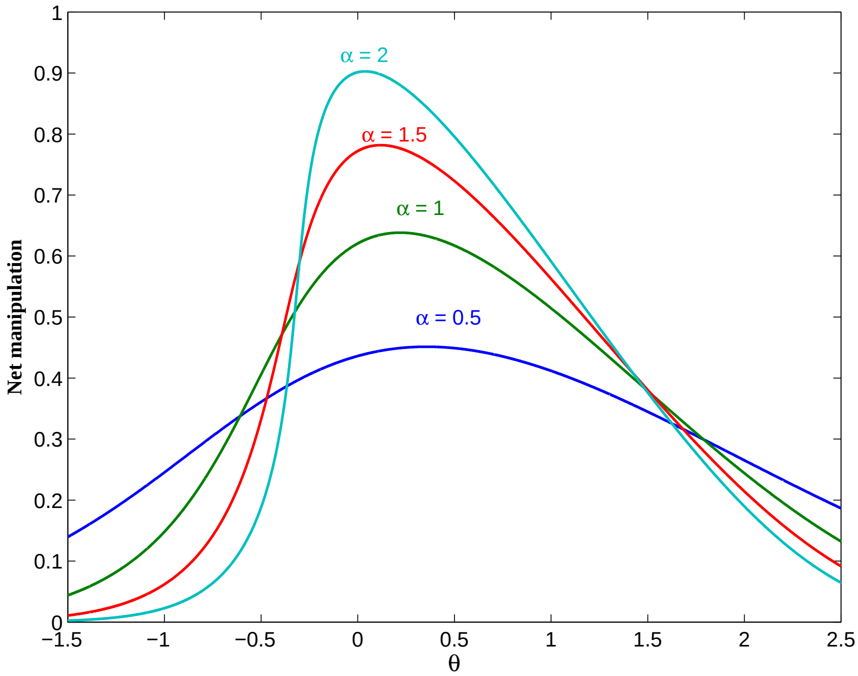

From a policy perspective, campaign signals would be more precise when candidates are forced to communicate with voters from common platforms, where they have relatively equal access to same set of voters. Intuitively, this would reduce the variance in the content of messages the candidates send. One example would be presidential debates, in which each candidate, using the same medium, tries to convince voters that she is a better choice for them. The opposite example would be a segregated communication environment in which candidates have their own TV channels to support them, and voters are more likely to watch one over the other.

We conclude this section by summarizing the effects of cost asymmetry in campaigning and informational heterogeneity on equilibrium vote shares. Figure 10 shows the share of votes the less-efficient candidate (candidate C) receives under different combinations of k_I and α values, when both candidates are situated in equal proximity to voters' blisspoint - i.e., when $\theta = 0.5$. As before, we let k_I take four different values, while α increases from 0.5 to 2 in three equal increments. As seen in the figure, as long as k_I remains below k_C, the incumbent receives a strictly higher share of the votes, although candidates are equally distant from the common voter position. Challenger's vote share falls further when the cost-asymmetry increases. At the current value of θ, a rise in precision leads to higher net manipulation and lower vote share for C. The negative impact of a rise in signal precision gets lower when k_I approaches to k_C. Similarly, the negative impact of a decline in k_I is more pronounced when α is higher.

5.3 Limits on campaign budgets

A typical model of informative political advertising finds that increasing costs of election campaigns harm voters since campaigns provide some information about candidates' characteristics or

Figure 10: Vote share of C in equilibrium when $\theta = 0.5$ ($k_C = 5$)

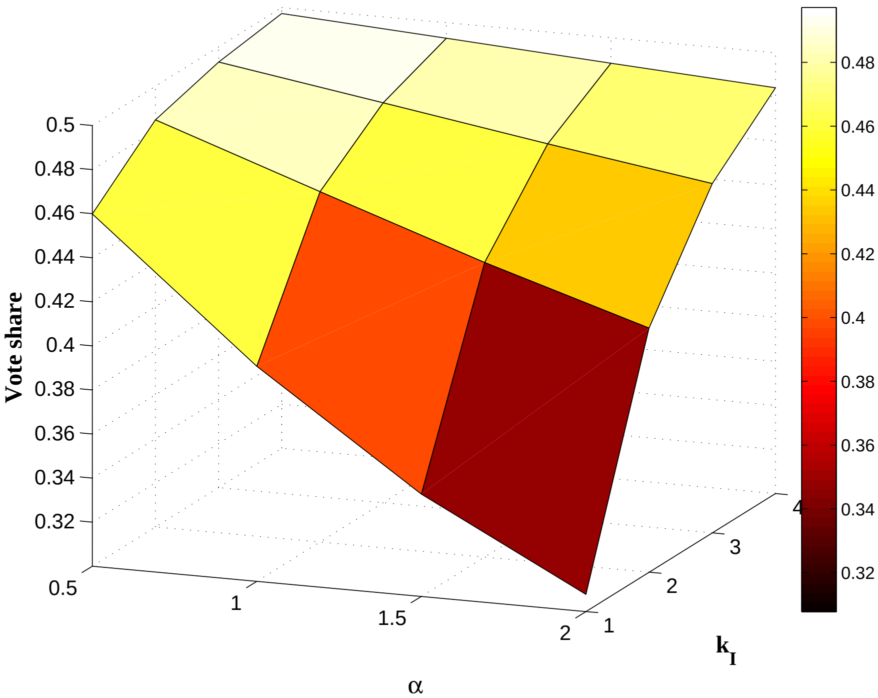

the state of the economy. Parallel to this simple intuition, Gul and Pesendorfer (2009) finds that if costs of campaigning increase for both candidates, voters' welfare decreases. They do not provide a welfare analysis for the case of distorted campaign signals, where candidates' actions may not provide any information. Similarly, Coate (2004a) shows that legal bounds on campaign spending transfer welfare from the median voter to interest groups because campaigns provide information to median voters. Prat (2002a, 2002b) study models of campaign finance with rational voters, politicians and lobbies, where political advertisements are only non-directly informative - i.e., campaign contributions by the lobbies serve as credible signals to voters about politicians' valence. In contrast to Coate (2004a) he concludes that a ban on campaign contributions may enhance the median voter's welfare under reasonable parameter values.

Our model departs from the previous studies in that we allow an electoral race between the candidates in which political advertising can be misinformative. Considering the comparative statics results in Figure 8 about the impact of cost asymmetry on net manipulation, it is not very surprising that, under our setting, limiting campaign spending can increase voter welfare. In practice, a direct cap on manipulative actions would be infeasible, but a uniform upper bound on campaign spending, if binding for the incumbent candidate, would reduce the level of net manipulation and improve welfare. This is because a mandatory reduction in manipulation would be more pronounced for I than for C, and, therefore, net manipulation would decrease.[21]

To demonstrate that limiting campaign expenditures can be welfare-improving also by promot-

[21]Our model abstracts from potential benefits of election campaigns where some of the information provided to voters by the candidates is useful. Since regulators cannot perfectly distinguish between useful and manipulative information, there is inevitably a social trade-off between restricting manipulation and increasing the transmission of useful information.

ing policy convergence towards voters' bliss point, we offer a simple extension of our original model in which access to campaign contributions are conditioned on the policy platform that candidates select. Suppose that voters are positioned at $\bar{\omega} = 0$, and each candidate $j \in P$ has to make a binary choice of political positions from $\{0, \omega_j\}$. We endogenize candidates' platforms in a rather reduced-form way. Assume that if any candidate j chooses ω_j instead of $\bar{\omega} = 0$, then she earns access to a budget of a fixed amount G, perhaps from special interest groups that are willing to fund the candidate if she positions herself at $\omega_j \neq 0$. We assume that G can be used for manipulation or be saved for later use, and this amount is given in the same units as the net benefit derived from vote shares. Suppose that $\omega_C < 0 < \omega_I$ are such that

$$\Phi(\sqrt{\alpha}(x^* - \theta_I - a_I(\theta_I))) < 1/2 < \Phi(\sqrt{\alpha}(x^* - \theta_C + a_C(\theta_C))), \tag{2}$$

where $\theta_I = 1/2 - |\omega_I|$, $\theta_C = 1/2 + |\omega_C|$ and x^* is the equilibrium voting strategy that voters would employ in the original model.[22] The first term θ_I is equal to the relative political position of candidates vis-à-vis voters when candidate C positions herself at $\omega_C = \bar{\omega} = 0$, while candidate I chooses $\omega_I > 0$. If, however, candidate C chooses the platform $\omega_C < 0$, while, this time, I positions herself at $\omega_I = \bar{\omega} = 0$, then $\theta = \theta_C$ instead. Equation (2) simply states that the political positions desired by special interest groups from each candidate --in exchange for access to campaign resources-- are such that the candidates can raise their vote shares to above $1/2$ if they unilaterally deviate from voters' bliss point and, at the same time, engage in manipulation.

Suppose that there is a state agency that is entitled to limit the level of budget, G, to maximize voters' ex-post information about candidates' choice of political positions. In this case, the state agency will choose an upper-bound for G to guarantee that there is no manipulative action, assuming that voters will form their voting strategy as in the benchmark case - i.e., as if the manipulation game the candidates play is the same as in benchmark model. Formally:

Proposition 3 *On any equilibrium of the extended model, the state agency will set the amount of allowable campaign contributions G so that it satisfies*

$$G \leq -1 + 2\Phi(\sqrt{\alpha}(x^* - \theta_I - a_I(\theta_I))) + H_I(a_I(\theta_I))$$
$$G \leq 1 - 2\Phi(\sqrt{\alpha}(x^* - \theta_C + a_C(\theta_C))) + H_C(a_C(\theta_C)),$$

where x^*, $\{a_j(\theta_j)\}_{j \in P}$ are as defined in Theorem 1. Under this restriction, the candidates do not engage in manipulation.

Proof If either of the candidates $j \in P$ chooses ω_j, while the other chooses 0, which is the voters' position, then the condition in equation 2 guarantees that j wins the election. This lowers the voters' ex-post payoff compared to the case where both candidates choose 0. However, if the state agency

[22] The assumption that voters ignore the state agency's behavior is not realistic. However, in our framework of imperfect social communication, even if voters take the state agency into account, they would not get perfectly precise signals about its possible behavior. Voters would use the signal about the state agency to further update their beliefs, but they would behave similarly anyway.

chooses G such that the hypothesis holds, then no candidate would want to deviate from the case where both choose 0. Since bounding G as in the hypothesis is the only way to force candidates to coordinate on the no-manipulation equilibrium, the agency chooses an upper-bound on G as in the hypothesis. ■

When equation (2) holds, it means that if some candidate j deviates by choosing ω_j, then she will definitely win the election, although the other candidate stands closer to voters. Thus, the state agency wants to prevent such a deviation. If equation (2) does not hold, the agency will not care to satisfy the constraint on G given in Proposition 3. In other words, if, to access the resources, the candidates need to be so extreme that they would lose the election despite manipulative advertising, then there is no manipulation in the equilibrium as long as the winning candidate does not deviate from 0, too.

6 Conclusion

This paper has proposed a theoretical model of electoral competition between two candidates who have incentives to send costly manipulative signals to bias public opinion and increase their voter shares. Due to the information asymmetry between the electorate and candidates, voters have to base their decisions on campaign signals about candidate positions. These positions are both systematically manipulated and contain a private noise due to the heterogeneous relationship of voters with the social communication structure that intermediates political messages. In other words, voters face two uncertainties: They are uncertain about both the candidates' real political platforms and the location of their private signal in population distribution of signals. These two factors, combined, make voters unable to filter out the bias correctly and create an incentive for politicians to send manipulated signals. Given such an incentive, the difference in the candidates' cost structures (campaigning efficiency) results in a positive net bias in voters' posterior beliefs in favor of the more-efficient candidate. Under such a setting, we prove that monotonic equilibria involve positive manipulation by both candidates, and the candidate with a cost advantage is capable of biasing the electorate's posterior beliefs in her favor.

We have taken a simplistic approach when modeling political propaganda by assuming that all of the candidates' announcements and actions, which, in reality, refer to multiple political dimensions, are aggregated into signals that can be interpreted as claims of political closeness to voters. This allowed us to model the competition between the candidates as a race to shift the value of a single indicator. We have also abstracted from messages that do not aim at manipulating the voters' opinion, but contains truthful messages. Although we have discussed the welfare implications of our model, to address questions about optimal campaign regulations in the context of a formal welfare analysis, it is necessary to explicitly model social costs of campaign spendings, as well as the benefits and losses associated with election outcomes. An interesting extension left for future work is to combine informative and manipulative political campaigns in a single model.

Another possible extension is to fully endogenize candidate platforms to conditional support from interest groups and lobbyists, so that politicians face a trade-off between approaching the median voter's bliss point to economize on manipulative campaign spending and moving towards the bliss point of special interest groups to finance their election campaigns.

7 Appendix

Proof of Theorem 1 The proof consists of three stages. In the first stage, we show that, given monotonic strategies of voters, the net manipulation function is necessarily a positive, bounded and continuously differentiable function of θ. Then, given any positive, bounded and continuously differentiable net manipulation function, voters uniformly use a monotonic strategy as a function of private signals. In the third stage, we will see that what we have in the first two stages requires a unique equilibrium.

Stage 1

Given the voting threshold $\hat{x}$, the set of ε_i's that will induce voter i to vote for candidate C is

$$\{\varepsilon_i | \varepsilon_i < \hat{x} - \theta - a_I + a_C\}.$$

Then, the size of supporters for C will be

$$S(a_I, a_C; \theta) = \Phi(\sqrt{\alpha}(\hat{x} - \theta - a_I + a_C)). \tag{3}$$

Once we have the functional form of the supporters' size $S(\cdot)$, we can move on to electoral candidates' actions. Given a_C, $u_I(a_C, a; \omega_C, \omega_I) = 1 - 2\Phi(\sqrt{\alpha}(\hat{x} - \theta - a + a_C)) - H_I(a)$. Since $\forall l \in P$, $H_j'(0) = 0$, there is no corner solution to the optimization problem of I. Then, the condition that characterizes the interior solution is

$$2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x} - \theta - a_I + a_C)) = H_I'(a_I),$$

and, similarly,

$$2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x} - \theta - a_I + a_C)) = H_C'(a_C).$$

Note that at any equilibrium, these two conditions are satisfied simultaneously, so we have

$$H_I'(a_I) = 2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x} - \theta - a_I + a_C)) = H_C'(a_C). \tag{4}$$

It follows from (4) that $a_C = {H_C'}^{-1}(H_I'(a_I))$, which is a well-defined continuously differentiable, strictly increasing function, since H_C and H_I are strictly increasing and convex. Then, (4) becomes

$$H_I'(a_I) = 2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x} - \theta - a_I + a_C)). \tag{5}$$

This equation has a unique solution due to the assumptions on the second derivatives of the cost functions. Those conditions imply that $H_I'(\cdot)$ can intersect with the right-hand side only once.

Note that $H_I'(a_I) = H_C'(a_C) \Rightarrow \forall\theta\ a_I(\theta) > a_C(\theta)$. So, the net manipulation function $\bar{a}(\cdot)$ of θ, defined as $\bar{a}(\theta) = a_I(\theta) - a_C(\theta)$, is always a positive function. By a rearrangement of equation

(5), we get

$$\begin{aligned}\bar{a}(\theta) = \\ H_I'^{-1}\left(2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x}-\theta-\bar{a}(\theta)))\right) \\ -H_C'^{-1}\left(2\sqrt{\alpha}\phi(\sqrt{\alpha}(\hat{x}-\theta-\bar{a}(\theta)))\right).\end{aligned} \tag{6}$$

This clearly implies that the maximum value that $\bar{a}$ can take is

$$\hat{a} = H_I'^{-1}\left(\sqrt{\frac{2\alpha}{\pi}}\right) - H_C'^{-1}\left(\sqrt{\frac{2\alpha}{\pi}}\right), \tag{7}$$

which is independent of any endogenous variables. From equations (6) and (7), it is clear that $\bar{a}$ is a continuously differentiable, and bounded function of θ. Using equation (6), we can calculate the derivative of $\bar{a}(\theta)$:

$$\bar{a}'(\theta) = \frac{1}{2}\left[\frac{2\alpha\phi'\left(\sqrt{\alpha}(\hat{x}-\theta-\bar{a}(\theta))\right)\left(\frac{H_I''-H_C''}{H_C''H_I''}\right)}{1-2\alpha\phi'\left(\sqrt{\alpha}(\hat{x}-\theta-\bar{a}(\theta))\right)\left(\frac{H_I''-H_C''}{H_C''H_I''}\right)}\right]$$

It is straightforward to show that there is a $\rho \in (0,1)$ such that $\bar{a}'(\theta) \in (-\rho, \rho)$, by employing the assumptions on the second derivatives of the cost functions H_I, H_C.

Stage 2

Given what we have in the first stage, we can assume without loss of generality that voters take a continuously differentiable, positive and bounded net manipulation function $\bar{a}$ such that $\bar{a}'(\theta) \in (-\rho, \rho)$ for some $\rho \in (0,1)$. Then, a voter with signal x_i assigns probability $P(x_i) := Pr(\theta \geq \frac{1}{2}|x_i)$ so that I is better than C, where

$$Pr(\theta \geq\ 1/2|x_i) = \frac{\int_{1/2}^{\infty}\sqrt{\alpha}\phi(\sqrt{\alpha}(x_i-\theta-\bar{a}))d\theta}{\int_{-\infty}^{\infty}\sqrt{\alpha}\phi(\sqrt{\alpha}(x_i-\theta-\bar{a}))d\theta}.$$

Differentiating this gives that $P'(x_i) > 0$ is equivalent to

$$\begin{aligned}&\int_{-\infty}^{\frac{1}{2}}\sqrt{\alpha}(\theta+\bar{a}-x_i)\phi(\sqrt{\alpha}((\theta+\bar{a}-x_i)))d(\theta)\int_{\frac{1}{2}}^{\infty}\sqrt{\alpha}\phi(\sqrt{\alpha}(\theta+\bar{a}-x_i))d(\theta) \\ &< \int_{\frac{1}{2}}^{\infty}\sqrt{\alpha}(\theta+\bar{a}-x_i)\phi(\sqrt{\alpha}((\theta+\bar{a}-x_i)))d(\theta)\int_{-\infty}^{\frac{1}{2}}\sqrt{\alpha}\phi(\sqrt{\alpha}((\theta+\bar{a}-x_i)))d(\theta).\end{aligned} \tag{8}$$

If $u = \sqrt{\alpha}(\theta+\bar{a}-x_i)$, we have

$$\int_{-\infty}^{\frac{1}{2}} \sqrt{\alpha}(\theta+\bar{a}-x_i)\phi(\sqrt{\alpha}((\theta+\bar{a}-x_i)))d(\theta)$$

$$= \int_{-\infty}^{\sqrt{\alpha}((\theta+\bar{a})(\frac{1}{2})-x_i)} u\phi(u)\frac{du}{(\theta+\bar{a})'((\theta+\bar{a})^{-1}(-\frac{u}{\sqrt{\alpha}}+x_i))},$$

which is in $[-\frac{1}{\rho\sqrt{\alpha}}\phi(\sqrt{\alpha}(\theta+\bar{a}-x_i)), -\frac{1}{2\sqrt{\alpha}}\ \phi(\sqrt{\alpha}(\theta+\bar{a}-x_i))]$.

Thus, the condition in equation (8) is implied by the following:

$$-\frac{1}{2\sqrt{\alpha}}\phi(\sqrt{\alpha}((\theta+\bar{a})(\frac{1}{2})-x_i))A < \frac{1}{2\sqrt{\alpha}}\phi(\sqrt{\alpha}((\theta+\bar{a})(\frac{1}{2})-x_i))B,$$

where $A := \int_{\frac{1}{2}}^{\infty} \sqrt{\alpha}\phi(\sqrt{\alpha}((\theta+\bar{a}-x_i)))d(\theta)$, and $B := \int_{-\infty}^{\frac{1}{2}} \sqrt{\alpha}\phi(\sqrt{\alpha}(\theta-\bar{a}-x_i))d(\theta)$. Note that the inequality above is always true since the right-hand side is always positive and the left-hand side always negative. It is easy to check that $\lim_{x_i\to-\infty} P(x_i) = 0$ and $\lim_{x_i\to\infty} P(x_i) = 1$, which implies that there is a unique $\hat{x}$ that equates $P(x_i)$ to $1/2$.

Stage 3

Existence and uniqueness directly follow by combining what we have in the first two stages. We can also find a threshold θ^* as follows: In this unique equilibrium, the mass of supporters of C, $S(\theta) = \Phi(\sqrt{\alpha}(x^* - \theta - \bar{a}(\theta)))$, is a strictly decreasing function of θ, which implies that there is a unique θ^* such that $S(\theta^*) = 1/2$. ■.

Proof of Theorem 2 We reproduce the equilibrium conditions below for convenience. Given the net manipulation function $\bar{a}(\cdot)$, the equilibrium threshold x^* for signals is determined by the following condition on posterior beliefs:

$$\frac{1}{2} = \frac{\int_{1/2}^{\infty} \sqrt{\alpha}\phi(\sqrt{\alpha}(x^*-\theta-\bar{a}(\theta)))d\theta}{\int_{-\infty}^{\infty} \sqrt{\alpha}\phi(\sqrt{\alpha}(x^*-\theta-\bar{a}(\theta)))d\theta}. \tag{9}$$

Given x^*, the net manipulation function is determined as follows:

$$\bar{a}(\theta) = \left[(H_I')^{-1} - (H_C')^{-1}\right](\sqrt{\alpha}\phi(\sqrt{\alpha}(x^*-\theta-\bar{a}(\theta)))). \tag{10}$$

A double use of the implicit function theorem on these two equations shows the existence of equilibrium. Uniqueness comes from the assumptions on the cost functions.[23] Thus, for each α, we have a unique net manipulation function and signal threshold. Given these two, the threshold θ^* that determines the winner of the election is given by the following equation:

$$\Phi(\sqrt{\alpha}(x^*-\theta^*-\bar{a}(\theta^*))) = \frac{1}{2}, \tag{11}$$

where $\Phi(\cdot)$ is the c.d.f. of the normal distribution. Note that equation (11) is satisfied if and only if

[23] To maintain the uniqueness result at the limit, we adopt the limiting counterpart of assumption (A2) that $\lim_{\alpha\to\infty} H_j''(a) \geq \lim_{\alpha\to\infty} \bar{\varphi}(\alpha)$ for $j \in P$.

$$x^* - \theta^* - \bar{a}(\theta^*) = 0. \tag{12}$$

First consider equation (10). The net manipulation function assumes its maximum value when $x^* - \theta - \bar{a}(\theta) = 0$ holds. However, a straightforward inspection of the first derivative of the net manipulation function shows that this condition holds if and only if equation (12) is satisfied. The maximum value $\bar{a}$ can take is given by equation (7), which clearly diverges to ∞ as $\alpha \to \infty$. However, we face an indeterminacy regarding the limiting behavior of the net manipulation function at other values. We solve this by imposing the following restriction:

$$\text{For any } \theta \neq \theta^*, \quad \bar{a}(\theta) \to 0 \quad \text{as} \quad \alpha \to \infty. \tag{13}$$

Now, since $\sqrt{\alpha}\phi((y)) \to 0$ as α and y diverges to infinity with the same speed, the convergence of the net manipulation function is characterized by the convergence of $\phi(\sqrt{\alpha}(x^* - \theta - \bar{a}(\theta)))$. Thus, the restriction in equation (13) requires point-wise convergence of $\phi(\sqrt{\alpha}(x^* - \theta - \bar{a}(\theta)))$ to zero. However, this does not guarantee its convergence in integrals we see in equation (9). Indeed, we face a second indeterminacy and solve this by requiring that the denominator in equation (9) does not converge to zero.

By equation (10), and restriction (13) for any $\theta \neq \theta^*$, x^* does not converge to θ. If θ^* converges to some finite number, and x^* converges to θ^*, then $x^* - \theta^*$ would converge to zero, but that's impossible as this expression equals $\bar{a}(\theta^*)$, which converges to infinity. Then, we know that x^* diverges either to positive or negative infinity.

By stage 1 of the proof of Theorem 1, $\bar{a}(\theta) + \theta$ is a strictly increasing function. This implies

$$x^* = \bar{a}(\theta^*) + \theta^* > \bar{a}(\theta) + \theta \quad \forall \theta < \theta^*.$$

This implies that $\lim_{\alpha \to \infty} x^* \geq \lim_{\alpha \to \infty} \theta^*$. On the other hand, for any $\theta > \theta^*$:

$$x^* - \bar{a}(\theta^*) - \theta^* > x^* - \bar{a}(\theta) - \theta.$$

Taking the limit on both sides by selecting appropriate subsequences over α reveals that $\lim_{\alpha \to \infty} x^* \leq \lim_{\alpha \to \infty} \theta^*$. Thus, the limit values of x^* and θ^* are the same.

Now, if x^* converges to a finite number, then from Equation (10), $\bar{a}(\theta) > 0$ for some θ that is not equal to the limit value of θ^*, which contradicts our transversality condition. Then, the limit value of both x^* and θ^* is either positive or negative infinity.

Suppose that x^* diverges to positive infinity; then, so does θ^*. This implies that for large α, $\bar{a}(\theta)$ is very close to 0 for every θ smaller than $1/2$. If the posterior probability that candidate C stands closer to voters is greater than $1/2$ - i.e., if

$$\frac{1}{2} < \frac{\int_{-\infty}^{1/2} \sqrt{\alpha}\phi(\sqrt{\alpha}(x^* - \theta - \bar{a}(\theta)))d\theta}{\int_{-\infty}^{\infty} \sqrt{\alpha}\phi(\sqrt{\alpha}(x^* - \theta - \bar{a}(\theta)))d\theta}, \tag{14}$$

then for any x that is less than 0 but whose absolute value is large, we will see that the denominator of (14) is very close to 0. However, the denominator is not zero since there is a θ that is larger than $1/2$ for which $x - \theta - \bar{a}(\theta) = 0$. This implies that, irrespective of the value of α, there is a region with a nonempty interior on which the integrand of the denominator assumes a positive value. However, by the definition of the equilibrium, the posterior probability should be greater than $1/2$. This contradiction shows that x^* diverges to $-\infty$, and so does θ^*. ■

The following lemma shows that no pure-strategy profile of the candidates can be supported as an equilibrium in the public-signals case.

Lemma 1 *In the public-signal case - i.e., when $\varepsilon_i = 0$ for all $i \in [0, 1]$ - there are no pure strategy electoral equilibria.*

Proof Suppose that there is a finite equilibrium threshold $\hat{x}$ for all voters. Then, $\forall \theta \in [\hat{x} - (H_I'^{-1}(2) - H_C'^{-1}(2)), \hat{x}]$, $a_C(\theta) = 0$, and $a_I(\theta) = H_C'^{-1}(2) + \theta - \hat{x}$. Therefore, for such θ, I wins the election, although voters believe that I is worse than C. Such an inconsistency can be solved by a change in voters' threshold $\hat{x}$ to $\hat{x} - (H_I'^{-1}(2) - H_C'^{-1}(2))$. But applying the same argument, this is not an equilibrium either, and so on. If the equilibrium threshold is in $\{-\infty, \infty\}$, as candidates will find it worthless to manipulate the signal, again there will exist some θ such that voters' beliefs are inconsistent with strategies of candidates. Thus, there is no pure-strategy electoral equilibrium. ■

Proof of Proposition 2 $\forall j \in P$, let ρ_j be the mixed strategy that j plays, where $\rho_j \in \Delta(\mathbb{R}_+)$ given the aggregate voter response function $\beta(\cdot)$. Now, given θ and $\rho_C(\cdot)$, the expected utility of playing ρ_I to I is

$$\int_{-\infty}^{\infty} \int_{0}^{\infty} (1 - 2\beta(x) - H_I(a_I))\rho_I(a_I)\rho_C(\theta + a_I - x)da_I dx,$$

and the expected utility of playing a degenerate strategy a_I is

$$\int_{-\infty}^{\infty} (1 - 2\beta(x) - H_I(a_I))\rho_C(\theta + a_I - x)dx.$$

Then, $\rho_C(\cdot)$ should be such that $\forall a, a'$ in the support of ρ_I.

$$\int_{-\infty}^{\infty} (1 - 2\beta(x) - H_I(a))\rho_C(\theta + a - x)dx = \int_{-\infty}^{\infty} (1 - 2\beta(x) - H_I(a'))\rho_C(\theta + a' - x)dx,$$

and, similarly, $\rho_I(\cdot)$ should be such that $\forall a, a'$ in the support of ρ_C.

$$\int_{-\infty}^{\infty} (2\beta(x) - 1 - H_C(a))\rho_I(x - \theta + a)dx = \int_{-\infty}^{\infty} (2\beta(x) - 1 - H_C(a'))\rho_C(x - \theta + a')dx.$$

Along with the equations above, the following defines ρ_C and ρ_I

$$\int_{0}^{\infty} \rho_I(a_I)da_I = \int_{0}^{\infty} \rho_C(a_C)da_C.$$

Now, given ρ_I and ρ_C, we will calculate $\beta(x)$. Note that given voters' expectation $E(\theta|x)$ of θ conditional on x, $\beta(x) = 1$ if and only if $E(\theta|x) < 1/2$. Now, the p.d.f. of the public signal x conditional on θ is

$$f(x|\theta) = \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I.$$

So voters' posterior belief is

$$\pi(\theta|x) = \frac{\int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I}{\int_{-\infty}^\infty \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}.$$

Note that the support of π is the same as that of f.

Now, let's calculate $E(\theta|x)$:

$$E(\theta|x) = \int_{-\infty}^\infty \theta\pi(\theta|x)d\theta$$

$$= \frac{\int_{-\infty}^\infty \theta \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}{\int_{-\infty}^\infty \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}$$

$$= \frac{\int_{-\infty}^\infty \int_0^\infty (x - a_I + a_C)\rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}{\int_{-\infty}^\infty \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}$$
$$= \frac{x}{\int_{-\infty}^\infty \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta} -$$
$$\frac{\int_{-\infty}^\infty \int_0^\infty (a_I - a_C)\rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}{\int_{-\infty}^\infty \int_0^\infty \rho_I(a_I)\rho_C(\theta + a_I - x)da_I d\theta}$$

The second term in the last equation is finite since manipulation is costly, so manipulations are uniformly bounded due to decreasing marginal returns to manipulation. This way, we can use the Fubini theorem to separate the double integral.

A change-of-variables argument shows that $E(\theta|x) = x - E(a_I) + E(a_C)$. Then, $E(\theta|x)$ is a linear function of x, so there is a unique $\hat{x}$ such that $E(\theta|\hat{x}) = 1/2$ and $x \geq 1/2 \Leftrightarrow E(\theta|x) \geq 1/2$.

As both I and C commonly observe θ, they choose their mixed strategies simultaneously. Thus, candidates will have ex-ante equilibrium expectations over the possible realizations of a_I and a_C, which will determine the ex-ante expectations of candidates over voters' expectations which is a function of the ex-ante uncertain public signal. Note that equilibrium expectations about the manipulation outcome x are accurate in the sense that both candidates use the right distribution over x, which is the convolution of ρ_I and ρ_C. Then, for any $\hat{\theta} \in \mathbb{R}$,

$$E_{I,C}(x|\hat{\theta}) = \int_0^\infty \int_0^\infty (\hat{\theta} + a_I - a_C)\rho_I(a_I)\rho_C(a_C)da_I da_C =$$
$$= \hat{\theta} + E(a_I) - E(a_C).$$

So, $E_{I,C}(E(\theta|x(\hat{\theta}))|\hat{\theta}) = \hat{\theta}$ since we have found above that $E(\theta|x) = x - E(a_I) + E(a_C)$ for any θ. This implies that

$$E_{I,C}(\beta(x)) = Pr(\hat{\theta} < 1/2) = 1 \quad \Leftrightarrow \hat{\theta} < 1/2,$$

as required. ■

References

Aragones, E., and Neeman, Z. Strategic Ambiguity in Electoral Competition. *Journal of Theoretical Politics 12*, 2 (2000).

Banks, J. S. A Model of Electoral Competition with Imperfect Information. *Journal of Economic Theory 50* (1989), 309--325.

Caplan, B. *The Myth of Rational Voter: Why Democracies Choose Bad Policies*. Princeton University Press, New Jersey, 2006.

Cassino, D., and Wooley, P. Some News Leaves People Knowing Less, November 21, 2011.

Coate, S. Political Competition with Campaign Contributions and Informative Advertising. *Journal of the European Economic Association 2*, 5 (2004a), 772--804.

Coate, S. Pareto-improving campaign finance policy. *The American Economic Review 94*, 3 (2004b), 628--655.

Cummins, J. G., and Nyman, I. The Dark Side of Competitive Pressure. *The RAND Journal of Economics 36*, 2 (2005).

DellaVigna, S., and Kaplan, E. The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics 122*, 3 (2007), 1187--1234.

Durante, R., and Knight, B. Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy. *NBER Working Paper No. 14762* (2009).

Edelman, R. The Social Pulpit: Obama's Social Media Toolkit. *Research report* (2009).

Edmond, C. Information Manipulation, Coordination, and Regime Change. mimeo, March 2008.

Ekmekçi, M. Manipulation Through Political Endorsements. *Journal of Economic Theory 144* (2009), 1227--1248.

Gentzkow, M., and Shapiro, J. M. Media Bias and Reputation. *Journal of Political Economy 114*, 2 (2006).

Grossman, G., and Helpman, E. *Special interest politics*. MIT Press, Cambridge, MA, 2001.

Gul, F., and Pesendorfer, W. War of Information. mimeo, August 2009.

Heidhues, P., and Lagerlof, J. Hiding Information in Electoral Competition. *Games and Economic Behavior 42* (2003), 48--74.

Kenski, K., Hardy, B. W., and Jamieson, K. H. *The Obama Victory: How Media, Money and Message Shaped the 2008 Election*. Oxford University Press, Oxford, England, 2010.

Martinelli, C. Elections with Privately Informed Parties and Voters. *Public Choice 108*, 1,2 (2001), 147--167.

Prat, A. Campaign Spending with Office-Seeking Politicians, Rational Voters, and Multiple Lobbies. *Journal of Economic Theory*, 103 (2002a), 162--189.

Prat, A. Campaign Advertising and Voter Welfare. *Review of Economic Studies*, 69 (2002b), 999--1017.

Puglisi, R., and James M. Snyder, J. Media Coverage of Political Scandals. *NBER Working Paper No. 14598* (2008).

Ramsay, C., Kull, S., Lewis, E., and Subias, S. Misinformation and the 2010 Election: A Study of the US Electorate, December 10, 2010.

Wittman, D. Why Democracies Produce Efficient Results. *Journal of Political Economy 97*, 6 (1989), 1395--1424.

Kemal Kıvanç Aköz
Department of Economics,
New York University.

Cemal Eren Arbatlı
Department of Theoretical Economics,
National Research University, Higher School of Economics.
Арбатли Джемаль Эрэн
Департамент теоретической экономики
+74957729590*026174 , +79252288945
darbatli@hse.ru