



HAL
open science

Confirmation bias and signaling in Downsian elections

Antony Millner, H el ene Ollivier, Leo Simon

► **To cite this version:**

Antony Millner, H el ene Ollivier, Leo Simon. Confirmation bias and signaling in Downsian elections. 2017. halshs-01631494

HAL Id: halshs-01631494

<https://shs.hal.science/halshs-01631494v1>

Preprint submitted on 9 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destin ee au d ep ot et  a la diffusion de documents scientifiques de niveau recherche, publi es ou non,  emanant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv es.



PARIS SCHOOL OF ECONOMICS
ÉCOLE D'ÉCONOMIE DE PARIS

WORKING PAPER N° 2017 – 49

Confirmation bias and signaling in Downsian elections

**Antony Millner
Hélène Ollivier
Leo Simon**

JEL Codes: D72, D91

Keywords: Confirmation bias, electoral competition, pandering, signaling



PARIS-JOURDAN SCIENCES ÉCONOMIQUES

48, Bd JOURDAN – E.N.S. – 75014 PARIS
TÉL. : 33(0) 1 43 13 63 00 – FAX : 33 (0) 1 43 13 63 10
www.pse.ens.fr

Confirmation bias and signaling in Downsian elections

Antony Millner^{*1}, H el ene Ollivier^{†2}, and Leo Simon^{‡3}

¹London School of Economics and Political Science

²Paris School of Economics – CNRS

³UC Berkeley

August 25, 2017

Abstract

How do voters’ behavioural biases affect political outcomes? We study this question in a model of Downsian electoral competition in which office-motivated candidates have private information about the benefits of policies, and voters may infer candidates’ information from their electoral platforms. If voters are Bayesian, candidates have strategic incentives to ‘anti-pander’ – they choose platforms that are more extreme than is justified by their private beliefs. However, anti-pandering incentives are ameliorated if voters’ inferences are subject to confirmation bias. Voter confirmation bias can thus counteract distortions due to the strategic interaction between candidates, potentially leading to welfare improvements. Indeed, we show that all observers, whether biased or Bayesian, would like the representative voter in our model to exhibit *more* confirmation bias than they do themselves.

Keywords: Confirmation bias, electoral competition, pandering, signaling.

JEL codes: D72, D91

1 Introduction

Since the 1970s the ‘heuristics and biases’ research programme, initiated by Kahneman and Tversky, has documented a litany of deviations from rationality in human judgement and decision-making. Behavioural anomalies such as the availability heuristic, anchoring, base-rate neglect, and confirmation bias, are now firmly established empirical facts, and have been shown to affect doctors’ treatment decisions (Bornstein & Emler, 2001), consumer behaviour

^{*}Millner: a.millner@lse.ac.uk

[†]Ollivier: helene.ollivier@psemail.eu

[‡]Simon: leosimon@berkeley.edu

(Thaler, 2016), and voters’ political preferences (Healy et al., 2010), to name but a few applications. In the face of such pervasive irrationality one could be forgiven for despairing of humanity’s ability to make good decisions. But are behavioural biases always ‘bad’? Does individual irrationality necessarily translate into bad equilibria in strategic situations, such as those that occur in the political arena? Or could behavioural effects perhaps counteract other distortions due to the strategic interactions between political actors? We investigate this question in a simple Downsian model of electoral competition in which office-motivated candidates signal their private information about the consequences of policy choices through their electoral platforms, and voters suffer from confirmation bias when interpreting candidates’ information. We show that, in this context at least, Bayesians can find biased behaviour beneficial.

Of all the well known behavioural biases, confirmation bias is surely one of the most consequential for political decision-making. As Barack Obama stated in his farewell speech on January 11, 2017, we have a tendency to ‘accept only information, whether true or not, that fits our opinions, instead of basing our opinions on the evidence that’s out there’. Psychologists have identified several specific behaviors that the term ‘confirmation bias’ subsumes (Kahneman & Tversky, 1973; Nickerson, 1998). These include a tendency to overweight evidence that confirms our prior beliefs, limit the range of sources we consult to those that are consistent with our priors, and more easily recall evidence that favours our priors. Taber & Lodge (2006) show experimentally that citizens exhibit these behaviours when forming opinions on politically contentious issues such as gun control and affirmative action. In this paper we focus on analyzing how confirmation bias that distorts voters’ interpretation of information might affect the outcome of competitive elections.

More specifically, we construct a simple model that allows us to investigate the interaction between confirmation bias and electoral pandering. The model is a Downsian electoral signaling game in which office-motivated candidates possess private information about policy consequences, and voters may infer candidates’ information from their electoral platforms. In this context a candidate is said to pander to the median voter¹ if his electoral platform is distorted towards the voter’s prior beliefs, and away from his private, more informed, beliefs. If pandering is an equilibrium behaviour, it is because candidates have strategic incentives to underweight their information about policy consequences, and shade their platforms towards the voter’s prior (Heidhues & Lagerlof, 2003).² Similarly, candidates who anti-pander have strategic incentives to choose platforms that are more extreme than is justified by their private beliefs, i.e they push their platforms away from the voter’s prior, and from what they privately believe to be optimal. Kartik et al. (2015) (henceforth KST) have shown that

¹Since we work in the Downsian framework only the median voter’s preferences are relevant for determining equilibrium behaviour. We thus refer to ‘the voter’ from now on.

²Other studies of pandering in a Downsian framework include Laslier & Van der Straeten (2004) and Gratton (2014).

anti-pandering equilibria arise in Downsian electoral signaling games in which candidates can choose platforms from a continuous interval of the real line. Our model marries a version of the KST model with a model of confirmation bias due to Epstein (2006), allowing us to illustrate how voter confirmation bias affects candidates’ behaviour. We show that when the voter is subject to confirmation bias, candidates’ strategic incentives to anti-pander are ameliorated – confirmation bias and strategic electoral incentives partially offset one another. If the degree of confirmation bias is not too extreme, voter confirmation bias is welfare improving from the perspective of a rational Bayesian observer. Moreover, from the perspective of *any* observer, whether biased or Bayesian, the optimal level of the voter’s confirmation bias is *higher* than the observer’s own level.

Our paper contributes to an emerging literature on the potential positive impacts of behavioural biases on political equilibria.³ Using a model of electoral accountability, Ashworth & Bueno De Mesquita (2014) show that voter bias can be beneficial when the strategic responses of politicians are taken into account. For instance, voters who ‘fail to filter’, in the sense that their voting behaviour is influenced by events such as natural disasters that are beyond political control, may induce politicians to make more moderate policy choices in equilibrium, raising welfare. Levy & Razin (2015) find that correlation neglect (i.e. erroneously treating correlated signals as if they were independent) can improve outcomes for voters.⁴ In their setting voting does not aggregate information efficiently because voters underweight their information relative to their ideological preferences when deciding how to vote. However, when voters receive correlated signals they overweight their information because of correlation neglect, counteracting the inefficiency of voting. Our model differs from these approaches in a number of ways. Candidates (not voters) are privately informed in our model, and voters’ beliefs are determined in equilibrium, rather than exogenously specified. In addition, unlike both these papers, candidates and voters have no ideological policy biases in our model. The distortions in our model are entirely due to the strategic interaction between privately informed candidates, which arise even though candidates’ objectives are perfectly aligned with the voter’s.

Our paper is perhaps closest in spirit to Lockwood (forthcoming), who introduces voter confirmation bias into the political agency model of Maskin & Tirole (2004). Although a limited amount of confirmation bias raises welfare in both our models, the comparative statics of pandering with respect to the degree of confirmation bias has opposite signs in Lockwood’s model and ours. This difference highlights the different roles that pandering and confirmation bias play in models involving agency (Lockwood) versus electoral competition (us). In agency models candidates pander to influence the voter’s inference about their *type*, i.e. whether their

³Of course, behavioural biases do not always lead to welfare improvements in political contexts. Biases that lead to negative welfare consequences include context-dependent voting (Callander & Wilson, 2008), and time inconsistency (Bisin et al., 2015; Lizzeri & Yariv, forthcoming).

⁴Levy & Razin (2015) also study a version of their model in which voters exhibit confirmation bias, but find no welfare change relative to rational voters.

policy preferences are aligned with the voter’s or not.⁵ Confirmation bias reduces the benefits of pandering in this framework, since the voter may misinterpret signals, weakening the case for ‘lying’ in order to try to pass for a type who shares the voter’s preferences. In our electoral competition model however, candidates anti-pander as a strategic response to the voter’s inferences about their *information* about the consequences of policy. Confirmation bias weakens the response of the voter’s beliefs to any information she may infer from candidates’ platforms, thus reducing the strategic incentive to anti-pander. The mechanism, and direction of the effects, in these two models is thus quite different.

2 The model

We study a model of Downsian electoral competition (Downs, 1957) in which candidates can signal their private information about the benefits of policy choices through their electoral platforms. Voters are assumed to have preferences over policies $y \in [-1, 1]$ that depend on the value of an unknown state of the world, $\theta \in \Theta \equiv \{-1, 0, 1\}$. Since we work in the Downsian framework, voters’ preferences can be represented by a single median voter, whose preferences are taken to be

$$u(y, \theta) = -E(y - \theta)^2, \tag{1}$$

where E denotes the expectation operator. θ may be thought of as the optimal policy choice given perfect knowledge of the state of the world. The voter believes that the cost of a non-optimal choice is quadratic in the distance between the chosen policy and the optimal *ex-post* choice. It is clear from (1) that in general the voter’s optimal policy choice is

$$\hat{y} = E\theta. \tag{2}$$

Since preferences are single peaked, the voter will always prefer platforms that are closer to \hat{y} to those farther away. The probability distribution from which θ is drawn is denoted by the 3-vector

$$\mathbf{p} = (\Pr(\theta = -1), \Pr(\theta = 0), \Pr(\theta = 1)) \equiv (p_{-1}, p_0, p_1).$$

This probability distribution is itself uncertain, and beliefs about \mathbf{p} are subject to revision if new information becomes available. The vector \mathbf{p} , and not θ itself, is the primitive uncertain quantity in the model, as we explain below.

We assume that there are two office-motivated candidates, A and B . At the beginning of the game the voter and the candidates have common prior beliefs about the possible values of

⁵Candidates have two possible types in the Maskin & Tirole (2004) model: Consonant or Dissonant. Consonants share the voters’ policy preferences, while Dissonants’ preferences are diametrically opposed to the voter’s.

\mathbf{p} . We assume that prior beliefs over \mathbf{p} are given by a Dirichlet distribution with probability density

$$g(\mathbf{p}; \boldsymbol{\alpha}) = \frac{\Gamma(\sum_{k \in \Theta} \alpha_k)}{\prod_{k \in \Theta} \Gamma(\alpha_k)} \prod_{k \in \Theta} p_k^{\alpha_k - 1}, \quad (3)$$

where

$$\boldsymbol{\alpha} = (\alpha_{-1}, \alpha_0, \alpha_1)$$

is a vector of hyper-parameters, and $\Gamma(\cdot)$ is the Gamma function. The Dirichlet distribution is a multidimensional generalization of the Beta distribution, and is a natural prior on the set of discrete probability distributions. The support of the prior distribution in (3) is the 2-dimensional simplex of three element probability distributions. The parameter vector $\boldsymbol{\alpha}$ controls the shape and concentration of prior beliefs on this simplex.

Given these assumptions we can write the voters' ex-ante bliss point as

$$\hat{y}_0 = E\theta = \int \left(\sum_{\theta \in \Theta} p_\theta \theta \right) g(\mathbf{p}; \boldsymbol{\alpha}) d\mathbf{p}. \quad (4)$$

Some simple calculations show that

$$\int p_\theta g(\mathbf{p}; \boldsymbol{\alpha}) d\mathbf{p} = \frac{\alpha_\theta}{\sum_{k \in \Theta} \alpha_k}$$

and thus we have

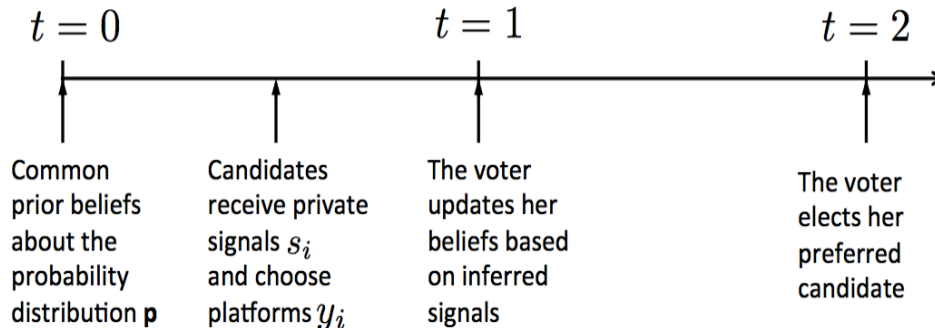
$$\hat{y}_0 = \frac{\alpha_{-1} - \alpha_1}{\sum_{k \in \Theta} \alpha_k}.$$

We assume a symmetric model,⁶ in which prior beliefs place equal weight on all pairs of distributions of the form $\mathbf{p} = (a, b, c)$ and $\mathbf{p}' = (c, b, a)$. Thus $\theta = -1$ and $\theta = 1$ are treated as equally probable ex-ante. This requires $\alpha_{-1} = \alpha_1$, implying that $\hat{y}_0 = 0$. In addition, we assume that prior beliefs are concentrated on probability distributions \mathbf{p} in which $\theta = 0$ is the most likely realized value of θ . This requires $\alpha_0 > \alpha_1$. $\theta = -1, 1$ are thus seen as 'outliers' *a priori*, relative to the modal value of $\theta = 0$.

The timing of events is as follows. In period 0 each candidate $i \in \{A, B\}$ receives a private signal $s_i \in \{-1, 0, 1\}$ about the uncertain state θ . After observing their signals, candidates simultaneously choose their platforms, denoted y_A and y_B , in order to maximize their probability of winning the election. In period 1 the voter observes candidates' platforms, perhaps infers something about the signals the candidates received, and updates her beliefs about the probability distribution \mathbf{p} accordingly. Finally in period 2 an election takes place,

⁶The symmetry assumption allows us to describe the degree of confirmation bias in our model with a single parameter q , defined below (see eq. (9)). An asymmetric model would require more parameters, but lead to qualitatively similar results.

Figure 1: Timing of events in the electoral game



and the candidate whose platform is closest to the median voter’s updated bliss point (2) is elected. If the voter is indifferent between two platforms she randomizes with equal probability between the two candidates. Figure 1 summarizes the sequence of events in the electoral game.

2.1 Updating with confirmation bias

The source of novelty in our approach is the behavioural model we adopt for how the voter’s beliefs change when she infers candidates’ information. We adopt a model of non-Bayesian updating due to Epstein (2006) to represent confirmation bias.⁷ In Epstein’s model agents act as if they are subject to self control problems: they are tempted to deviate from rational behaviour after information has been received, and sustain a cost for retaining a measure of self control. The tension between pursuing ‘rational’ objectives and not paying too high a self-control cost gives rise to a model that allows us to parameterize the degree of confirmation bias the voter exhibits.

To get an intuitive understanding of Epstein’s model it is best to work directly with his representation result.⁸ Let $F(S)$ be a menu of acts (maps between states and outcomes), conditional on receiving a signal S about the state of the world σ . Denote a generic act by $f = f(\sigma) \in F(S)$. Let $P_B(\sigma|S)$ be the rational Bayesian update of the agent’s prior beliefs about σ conditional on receiving signal S . We can think of the Bayesian rule as the updating procedure the agent considers ‘correct’ before she has received the signal S . In addition, let $P_T(\sigma|S)$ be an alternative ‘non-rational’ updating rule which the agent is *tempted* to use once S has been received. Epstein shows that agents who obey the behavioural axioms he

⁷Epstein’s model is in turn based on Gul & Pesendorfer’s (2001) work on temptation and self control in intertemporal choice.

⁸Our presentation of this result is simplified for expositional purposes. In particular, we do not emphasize Epstein’s axioms on state-contingent menus of acts, but rather work directly with his representation result for choices conditional on a given menu. We refer the reader to the original paper for further details.

stipulates will choose the act f that solves:

$$\max_{f \in F(S)} \left\{ \int U(f(\sigma))P_B(\sigma|S)d\sigma + \gamma(S) \left[\int U(f(\sigma))P_T(\sigma|S)d\sigma - \max_{f' \in F(S)} \int U(f'(\sigma))P_T(\sigma|S)d\sigma \right] \right\} \quad (5)$$

The interpretation of the components of this expression is as follows. The first term represents the agent's rational Bayesian expected utility if she were able to commit herself to actions ex-ante, i.e. before receiving the signal S . However the agent suffers from self-control problems: once S has been received, she is tempted to update her beliefs using the non-rational rule $P_T(\sigma|S)$. If the agent manages to partially avoid the temptation to use the non-rational rule, she pays a self-control cost, represented by the terms in square brackets. The cost of self-control increases with the distance between her actual choice and the choice she would have made if she had given in to temptation. The agent's choices are thus a compromise between pursuing her ex-ante rational objectives, and avoiding large self-control costs ex-post. The cost of exhibiting self-control may be signal dependent, and is captured by the weight function $\gamma(S)$.

To apply this preference representation in our model we need to specialize it further. In a fully revealing equilibrium of our electoral signaling game the voter will be able to infer the private signals of both candidates from their equilibrium platform choices. Thus, the signals the voter infers are of the following form:

$$S_{ij} = (s_i, s_j)$$

where $s_i, s_j \in \{-1, 0, 1\}$ are the possible realizations of the candidates' private information. Each observed signal pair S_{ij} allows the voter to update her beliefs about the probability vector \mathbf{p} , which determines the likely values of the state of the world θ . To model confirmation bias, we assume that after receiving a signal S_{ij} , the voter is tempted to deploy an updating rule of the following form:

$$P_T(\mathbf{p}|S_{ij}) = (1 - \lambda(S_{ij}))P_B(\mathbf{p}|S_{ij}) + \lambda(S_{ij})g(\mathbf{p}; \boldsymbol{\alpha}) \quad (6)$$

where $g(\mathbf{p}; \boldsymbol{\alpha})$ is the voter's prior beliefs about \mathbf{p} , given in (3), and $\lambda(S_{ij})$ is a signal-dependent weight, defined below.⁹ By assumption, the agent's prior beliefs place more weight on the central value $\theta = 0$ than the extreme values $\theta = \pm 1$ (i.e. $\alpha_0 > \alpha_1$). Thus, if she infers a signal (0, 0), this confirms her prior expectations. Any other signal, however, diverges from her prior beliefs. A voter who exhibits confirmation bias will treat signals that confirm her prior beliefs differently from those that do not. We adopt the simplest possible choice of the

⁹An axiomatic justification for this choice for $P_T(\mathbf{p}|S_{ij})$ is provided by Epstein (2006).

weight function $\lambda(S_{ij})$ that achieves this:

$$\lambda(S_{ij}) = \begin{cases} 0 & \text{if } S_{ij} = (0, 0) \\ \bar{\lambda} & \text{otherwise} \end{cases} \quad (7)$$

where $\bar{\lambda} \in (0, 1)$ is a parameter. Substituting these choices into the representation (5), and assuming that $\gamma(S)$ is a constant $\bar{\gamma}$, it is easily seen that the voter will act as if she updates her beliefs according to the compromise rule:

$$P^*(\mathbf{p}|S_{ij}) = \left(1 - \frac{\bar{\gamma}\lambda(S_{ij})}{1 + \bar{\gamma}}\right) P_B(\mathbf{p}|S_{ij}) + \frac{\bar{\gamma}\lambda(S_{ij})}{1 + \bar{\gamma}} f(\mathbf{p}). \quad (8)$$

Defining a new parameter $q \equiv \frac{\bar{\gamma}\bar{\lambda}}{1 + \bar{\gamma}} \in [0, 1)$, this updating rule can be written simply as:

$$P_q^*(\mathbf{p}|S_{ij}) = \begin{cases} P_B(\mathbf{p}|S_{ij}) & \text{if } S_{ij} = (0, 0) \\ (1 - q)P_B(\mathbf{p}|S_{ij}) + qg(\mathbf{p}) & \text{otherwise.} \end{cases} \quad (9)$$

Thus, if the voter infers that the candidates' signals are $(0, 0)$ her updating rule is unbiased. In all other cases, her updating rule is a convex combination of the Bayesian rule and her prior beliefs, causing her to under-react to information, and biasing her beliefs towards her prior. The closer q is to 0, the closer the voter's updating rule is to Bayesian rationality; q is thus a measure of the magnitude of the voter's confirmation bias.

Since the voter's bliss point (2) is a sufficient statistic for her preferences between candidate's platforms, we will be interested in computing her posterior expected value of θ conditional on receiving a signal S_{ij} , denoted by $E_q[\theta|S_{ij}]$. Some calculations with the Dirichlet distribution (3) and the updating rule (9) show that:

$$E_q[\theta|S_{ij}] = \int \left(\sum_{\theta \in \Theta} p_\theta \theta \right) P_q^*(\mathbf{p}|S_{ij}) d\mathbf{p} \Rightarrow \begin{cases} E_q[\theta|(1, 0)] = \frac{1-q}{2\alpha_1 + \alpha_0 + 2} = -E_q[\theta|(-1, 0)] \\ E_q[\theta|(0, 0)] = E_q[\theta|(-1, 1)] = E_q[\theta|(1, -1)] = 0 \\ E_q[\theta|(1, 1)] = \frac{2(1-q)}{2\alpha_1 + \alpha_0 + 2} = -E_q[\theta|(-1, -1)]. \end{cases} \quad (10)$$

2.2 Equilibria

A pure strategy for candidate i is a function $y_i(\cdot) : \Theta \rightarrow [-1, 1]$, where $y_i(s_i)$ is the platform chosen by i when his signal is s_i . We focus throughout on *symmetric, fully revealing* strategy profiles, in which $y_i(\cdot) = y(\cdot)$, for $i = A, B$, with the property that if $s_i \neq s'_i$ then $y(s_i) \neq y(s'_i)$. Given such a profile, and a vector of platforms $\zeta = (\zeta_i, \zeta_j) \in y(\Theta) \times y(\Theta)$, denote by $S_{ij}(\zeta)$ the uniquely identified signal pair $(y^{-1}(\zeta_i), y^{-1}(\zeta_j))$ that the voter can infer from observing ζ . Next, given a bias level q , let $\omega_i^y(\zeta; q)$ denote the probability that the voter selects i if the candidates' symmetric strategy is y and the announced platform vector is ζ , with the

restriction:

$$\text{if } \zeta \in y(\Theta)^2 \text{ then } \omega_i^y(\zeta; q) = \begin{cases} 1 & \text{if } |\zeta_i - E_q[\theta|S_{ij}(\zeta)]| < |\zeta_j - E_q[\theta|S_{ij}(\zeta)]| \\ 0 & \text{if } |\zeta_i - E_q[\theta|S_{ij}(\zeta)]| > |\zeta_j - E_q[\theta|S_{ij}(\zeta)]| \\ 0.5 & \text{otherwise} \end{cases} \quad (11)$$

Restriction (11) is the analog of the restriction in sender-receiver games that if a sender makes an “on the equilibrium path” deviation from his declared strategy, then the receiver must use Bayes’ Law to update her beliefs about the sender’s private information. When $\zeta \notin y(\Theta)^2$, we impose no *a priori* restriction on $\omega_i^y(\zeta; q)$; this is analogous to the assumption underlying Perfect Bayesian equilibria that after observing an “off the equilibrium path” deviation, the receiver’s beliefs are unrestricted. (When we support an equilibrium with specific beliefs, we will further restrict ω_i^y .)

We study the *Symmetric Perfect Confirmation-Biased Equilibria (SPCBE)* of the game between candidates and the voter. A strategy profile $y(\cdot)$ is an SPCBE if: A) the voter updates her beliefs in accordance with the non-rational updating rule (9) for some q , then chooses between the candidates so as to maximize her expected utility (1), conditional on these beliefs; and B) for each signal s_i received by candidate i , the platform $y(s_i)$ maximizes his posterior probability of being elected given the voter’s behavior rule A). In symbols, let $\rho(s_j|s_i) = \frac{\alpha_j + I_{j=i}}{\sum_{\theta} \alpha_{\theta} + 1}$ denote candidate i ’s posterior predictive probability that candidate j ’s signal is s_j , given that i ’s signal is s_i . If i observes signal s_i and chooses platform ζ against j ’s strategy $y(\cdot)$, i ’s expected probability of being elected is $\sum_{\Theta} \rho(s_j|s_i) \omega_i(\zeta, y(s_j); q)$. Then $y(\cdot)$ is a SPCBE if for $i = A, B$,

$$\forall s_i \in \Theta, \forall \zeta \in [-1, 1], \quad \sum_{\Theta} \rho(s_j|s_i) \left(\omega_i^y(\zeta, y(s_j); q) - \omega_i^y(y(s_i), y(s_j); q) \right) \leq 0 \quad (12)$$

Our first result shows that there is a unique fully revealing, symmetric equilibrium (all proofs are contained in the appendix):

Proposition 1. *For each $q \in [0, 1)$, there is a unique fully revealing, SPCBE in pure strategies $y_q(\cdot)$, given by:*

$$y_q(-1) = -\bar{y}_q; y_q(0) = 0; y_q(1) = \bar{y}_q, \text{ where } \bar{y}_q = \frac{2(1-q)}{2\alpha_1 + \alpha_0 + 2}. \quad (13)$$

Each candidate is elected with probability 1/2 regardless of his signal realization.

To understand the important features of the equilibrium in Proposition 1, begin by considering the case of a Bayesian voter, i.e. $q = 0$. In this case candidates’ equilibrium platforms exhibit anti-pandering: their platforms are *more extreme* than their private beliefs about the

voters' posterior bliss point. To see this, note that

$$E[\theta|1] = \frac{1}{2\alpha_1 + \alpha_0 + 1} < \bar{y}_0 \quad \text{and} \quad E[\theta|-1] = \frac{-1}{2\alpha_1 + \alpha_0 + 1} > -\bar{y}_0, \quad (14)$$

where expectation operators without subscripts are Bayesian. Thus, for example, a candidate who receives a signal 1 believes that the voter's most preferred platform will be at $E[\theta|1]$, however in equilibrium this candidate plays $\bar{y}_0 > E[\theta|1]$. Candidates thus overreact to their private information, choosing equilibrium platforms that are more extreme than they believe they should be.

To understand why candidates anti-pander it is helpful to see why a strategy profile in which each candidate plays his true posterior expected value of θ is *not* an equilibrium. Suppose that candidates play their true posterior expectations, i.e. on receiving signals $(-1, 0, 1)$ they play $(-y^*, 0, y^*)$ where $y^* = E[\theta|1]$. To make the discussion concrete, suppose that $\alpha_0 = 1 + \delta$ and $\alpha_1 = 1 - \frac{\delta}{2}$ for some $\delta > 0$. In this case we have $y^* = E[\theta|1] = \frac{1}{4}$. Now consider what happens if candidate A plays 0, while B plays y^* . In this case the voter would infer the signal pair $(0, 1)$, and from (10), her posterior expectation of θ would be $E[\theta|(0, 1)] = \frac{1}{5}$. Since $\frac{1}{5}$ is closer to $y_B = \frac{1}{4}$ than to $y_A = 0$, candidate A will lose the election in this case. Thus, A 's strategy admits a profitable deviation: if A plays y^* when he observes a signal 0, rather than playing 0, he can tie, rather than lose, the election. The proposed 'truthful revelation' strategy thus cannot be an equilibrium. To eliminate the possibility of a profitable deviation for A , we must increase the value of $y(1)$ from y^* to the point where $y(1)$ and 0 are equidistant from the voter's posterior, conditional on her inferring the signal pair $(0, 1)$. This point is given by \bar{y}_0 in (13). This equilibrium exhibits *anti-pandering*: each candidate's platform after observing signals ± 1 is greater in absolute value than their private expectations.

The ultimate reason for anti-pandering when voters act as Bayesians is that the voter's posterior beliefs after inferring both candidates' signals are more extreme than the average of each candidate's private expectations, which are based on the observation of just one signal. Indeed, some calculations using the expressions for $E[\theta|(s_i, s_j)]$ in (10), and $E[\theta|s_i]$ in (14), show that

$$E[\theta|(s_i, s_j)] = \left[2 \times \frac{2\alpha_1 + \alpha_0 + 1}{2\alpha_1 + \alpha_0 + 2} \right] \left[\frac{1}{2} (E[\theta|s_i] + E[\theta|s_j]) \right] \quad (15)$$

$$\Rightarrow |E[\theta|(s_i, s_j)]| \geq \left| \frac{1}{2} (E[\theta|s_i] + E[\theta|s_j]) \right|, \quad (16)$$

where the inequality follows since $2 \times \frac{2\alpha_1 + \alpha_0 + 1}{2\alpha_1 + \alpha_0 + 2} \in (1, 2)$, and is strict whenever $|s_i| \neq |s_j|$. Thus, if the candidates truthfully reported their private expectations, then whenever the voter infers signals that are non-identical in absolute value, her posterior beliefs will be closer to the private expectations of the candidate with the more extreme signal. This gives rise to

incentives for candidates to overreact to their private information, choosing platforms that are more extreme than they believe to be optimal. Indeed, using equations (15), (14), and (13) we see that

$$y_0(1) = E[\theta|(1, 1)] = -y_0(-1) = -E[\theta|(-1, -1)]. \quad (17)$$

Thus, candidates choose their platforms at the voter's bliss point, assuming that the other candidate received the same signal as them. These results are a version of those in Kartik et al. (2015), who show that such anti-pandering equilibria arise in a large class of electoral signaling games in which the set of possible platforms is continuous.

Now consider how confirmation bias modifies the equilibrium of the electoral game. When the voter is subject to confirmation bias, her posterior beliefs are *less* responsive to her inferences about candidates' signals than those of a Bayesian voter. Equation (10) shows that her posterior expected value of θ after inferring signal pair S_{ij} is now

$$E_q[\theta|S_{ij}] = (1 - q)E[\theta|S_{ij}] \Rightarrow |E_q[\theta|S_{ij}]| \leq |E[\theta|S_{ij}]|$$

where the inequality is strict as long as $S_{ij} \neq (0, 0)$. Confirmation bias thus draws the voter's posterior expected value of θ towards her prior (note that $E_q[\theta] = 0$ *a priori*). Rational candidates will adapt their strategies to the biased voter's beliefs. Since confirmation bias makes the voter's posterior beliefs *less* extreme, this counteracts the strategic incentive to overreact to information we observed in the Bayesian case. Indeed, we see that platforms in the equilibrium with confirmation bias satisfy

$$y_q(1) = E_q[\theta|(1, 1)] < E[\theta|(1, 1)] = y_0(1).$$

2.3 Welfare

Because confirmation bias counteracts the overreaction effects we observed when the voter is a Bayesian, a Bayesian observer may view voter confirmation bias as welfare improving. To see the intuition for this consider the effect of a very small amount of confirmation bias on candidates' platforms. Equation (17) shows that when $q = 0$ candidates' equilibrium platforms are ex-post optimal from the perspective of a Bayesian observer whenever candidates' signal pairs are $(0, 0)$, $(1, 1)$, $(-1, -1)$, or $(-1, 1)$. Thus, increasing q to a small positive value $\epsilon \ll 1$ has only a second order negative effect on ex-post Bayesian welfare if candidate's receive these signal pairs. However when the inferred signal pair is $(0, 1)$ or $(0, -1)$ candidates' platforms are *not* ex-post optimal according to the Bayesian observer. When $q = \epsilon$ the Bayesian observer would calculate ex-post welfare if signals are $(0, 1)$ as

$$\frac{E[(y_\epsilon(0) - \theta)^2|(0, 1)]}{2} + \frac{E[(y_\epsilon(1) - \theta)^2|(0, 1)]}{2} = \frac{(y_\epsilon(1) - E[\theta|(0, 1)])^2}{2} + \text{terms independent of } \epsilon$$

where we have used the fact that $E(y - \theta)^2 = (y - E\theta)^2 + \text{Var}(\theta)$, and $y_\epsilon(0) = 0$. Since

$$E[\theta|(0, 1)] < y_\epsilon(1) < y_0(1) = E[\theta|(1, 1)]$$

if ϵ is small, increasing confirmation bias from 0 to ϵ has a first order positive effect on ex-post welfare in this case. Similar reasoning holds when the inferred signal pair is $(0, -1)$. Since the positive effects of a small increase in confirmation bias are first order, but the negative effects are second order, a small increase in a Bayesian voter's confirmation bias *increases* ex-ante Bayesian welfare.

We can make this more precise by defining ex-ante Bayesian welfare as follows:

$$W = - \sum_{s_i \in \Theta} \rho(s_i) \sum_{s_j \in \Theta} \rho(s_j|s_i) \left(\frac{1}{2} \sum_{t \in \{s_i, s_j\}} E[(y_q(t) - \theta)^2 | (s_i, s_j)] \right), \quad (18)$$

where

$$\rho(s_i) = \int \text{Prob}(s_i | \mathbf{p}) g(\mathbf{p}; \boldsymbol{\alpha}) d\mathbf{p} = \frac{\alpha_i}{\sum_{k \in \Theta} \alpha_k}$$

is the prior probability of receiving signal s_i , and we have used the fact that the voter randomizes between candidates' platforms in equilibrium. Straightforward, if lengthy, computations¹⁰ show that

$$W \propto -[(\alpha_0 + 2\alpha_1 + 1)q^2 - (\alpha_0 + 2\alpha_1)q] + \text{terms independent of } q.$$

W is thus a concave function of q , which achieves its maximum at some $q^* > 0$. Calculations with the full analytic expression for W yield the following result:

Proposition 2. *A Bayesian observer believes that confirmation bias improves equilibrium welfare (relative to the case of a Bayesian voter) if*

$$q < q_C \equiv \frac{\alpha_0 + 2\alpha_1}{\alpha_0 + 2\alpha_1 + 1}.$$

The value of q that maximizes ex-ante Bayesian welfare is $q^ = q_C/2$.*

Figure 2 illustrates this result, and the dependence of ex-ante Bayesian welfare on the confirmation bias parameter q .

To interpret the optimal value of voter confirmation bias q^* further, notice from (13), (14),

¹⁰These computations are best performed with a computer algebra package. As the proofs of the results in this subsection are simple, if lengthy, algebraic manipulations, we do not provide details here. The Matlab code used to generate the results is available at <https://tinyurl.com/ybfkc7c2>.




Figure 2: Dependence of ex-ante Bayesian welfare on confirmation bias. $\alpha_0 = 3, \alpha_1 = 1$ in this example.

and q^* that

$$y_{q^*}(1) = \frac{2(1 - q^*)}{2\alpha_1 + \alpha_0 + 2} = 2 \frac{1 - \frac{1}{2} \frac{\alpha_0 + 2\alpha_1}{\alpha_0 + 2\alpha_1 + 1}}{2\alpha_1 + \alpha_0 + 2} = \frac{1}{\alpha_0 + 2\alpha_1 + 1} = E[\theta|1].$$

Thus, at the optimal level of voter confirmation bias, candidates play their unbiased private best guess of the value of θ . They thus act as if they were in a non-strategic situation, effectively forgetting that the voter will in fact infer *two* signals in equilibrium. At first this may seem surprising, given that welfare depends on the ex-post expected value of θ conditional on *both* signals being observed. To understand this further, note that the optimal value of $y(1)$ (i.e. candidates' platforms when they receive signal 1) for an ex-ante Bayesian observer must satisfy

$$\max_{y(1)} \left\{ - \sum_{s_j \in \Theta} \rho(s_j|1) E[(y(1) - \theta)^2 | (s_j, 1)] \right\}.$$

Denoting the solution of this optimization problem by $y^*(1)$, we have

$$\begin{aligned}
y^*(1) &= \sum_{s_j \in \Theta} \rho(s_j|1) E[\theta|(s_j, 1)] \\
&= \sum_{s_j \in \Theta} \rho(s_j|1) \int \left(\sum_{\theta \in \Theta} p_\theta \theta \right) \text{Prob}(\mathbf{p}|(s_j, 1)) d\mathbf{p} \\
&= \int \left(\sum_{\theta \in \Theta} p_\theta \theta \right) \left[\sum_{s_j \in \Theta} \text{Prob}(\mathbf{p}|(s_j, 1)) \rho(s_j|1) \right] d\mathbf{p} \\
&= \int \left(\sum_{\theta \in \Theta} p_\theta \theta \right) \text{Prob}(\mathbf{p}|1) d\mathbf{p} \\
&= E[\theta|1].
\end{aligned}$$

Since the voter's payoff function is quadratic, the platform that maximizes ex-ante Bayesian welfare when one of the candidates receives a signal 1 is just the expected value of θ , conditional on 1 being observed. This expectation runs over the values of the other candidate's signal s_j . However, by the Law of Iterated Expectations, the expected value of θ conditional on the signal pair $(s_j, 1)$ being observed, averaged across all values of s_j , is equivalent to the expected value of θ if just the single signal 1 were observed. Unbiased strategies are thus ex-ante optimal from the perspective of a Bayesian observer.

The Bayesian welfare measure we have used thus far provides a natural normative benchmark for welfare judgements. However, it is also interesting to ask how a *confirmation-biased* observer would evaluate the equilibrium outcome. An observer who has a confirmation bias parameter q_w would compute ex-ante welfare as follows:

$$W_{q_w} = - \sum_{s_i \in \Theta} \rho(s_i) \sum_{s_j \in \Theta} \rho_{q_w}(s_j|s_i) \left(\frac{1}{2} \sum_{t \in \{s_i, s_j\}} E_{q_w} [(y_q(t) - \theta)^2 | (s_i, s_j)] \right), \quad (19)$$

where the expectation E_{q_w} is now taken with respect to the observer's biased posterior beliefs, given by (9) with $q = q_w$. Notice that two occurrences of confirmation bias enter this welfare calculation: the observer's bias q_w and the voter's bias q , and only the latter influences the candidates' strategies. Aside from modifying the observer's views on the 'correct' platform conditional on candidate's signals S_{ij} , confirmation bias also affects the observer's views on the posterior predictive probability $\rho_{q_w}(s_j|s_i)$. Using an updating rule analogous to (9) but

for one signal,¹¹ and the properties of the Dirichlet distribution, one can show that

$$\rho_{q_w}(s_j | s_i) = \begin{cases} \frac{\alpha_j + I_{s_j=0}}{\sum_k \alpha_k + 1} & s_i = 0 \\ (1 - q_w) \frac{\alpha_j + I_{s_i=s_j}}{\sum_k \alpha_k + 1} + q_w \frac{\alpha_j}{\sum_k \alpha_k} & s_i \neq 0. \end{cases} \quad (20)$$

The prior probability of receiving a signal s_i , $\rho(s_i)$, is unaffected by confirmation bias, as the observer is only tempted to deviate from the Bayesian updating rule after a signal has been received (see the discussion in Section 2.1).

Calculations similar to those used in Proposition 2 now yield the following:

Proposition 3. *For all $q_w \in [0, 1)$, the biased ex-ante welfare measure W_{q_w} is a concave function of the voter's confirmation bias parameter q . It attains its maximum at $q = q^*(q_w) \in (0, 1)$, where the function $q^*(q_w)$ is increasing, and satisfies $q^*(q_w) > q_w$ for all $q_w \in [0, 1)$.*

The upshot of this proposition is that all observers, whether biased or Bayesian, would prefer the voter to be more biased than themselves. Figure 3 illustrates this result. The intuition is very similar to that in the Bayesian case. Consider the case where the voter's confirmation bias parameter q is just above the observer's, i.e. $q = q_w + \epsilon$. A voter with $q = q_w$ induces candidates to choose platforms that are ex-post optimal from the perspective of an observer with bias q_w when the inferred signals are $(-1, 1)$, $(0, 0)$, $(1, 1)$ or $(-1, -1)$. A small increase in q gives rise to second order losses in these cases. However, platforms are non-optimal when inferred signals are $(0, 1)$ or $(0, -1)$, and a small increase in q gives rise to first order gains in these cases since

$$E_{q_w}[\theta | (0, 1)] < y_{q_w + \epsilon}(1) < y_{q_w}(1) = E_{q_w}[\theta | (1, 1)].$$

Thus, increasing q above q_w is always welfare improving for the biased observer. Similar logic shows that a small decrease in q below q_w reduces welfare, since in addition to the familiar second order losses, this also gives rise to first order losses when inferred signals are $(0, 1)$ or $(0, -1)$. Thus all observers, whether biased or Bayesian, would like the voter to be more biased than themselves.

3 Concluding remarks

We have presented a simple model of the effect of voter confirmation bias on electoral competition, demonstrating that confirmation bias can counteract strategic distortions to candidates'

¹¹Note that, unlike Bayesian updating, updating with confirmation bias will in general be sensitive to the order in which signals are received (see also Rabin & Schrag (1999)). Thus, posterior beliefs after observing two signals in sequence will differ from posterior beliefs after observing the same two signals simultaneously. Since signals are always inferred simultaneously in our model this poses no difficulties for us, and we use the same updating rule for individual signals s_i and simultaneous signal pairs S_{ij} .




Figure 3: Optimal level of voter confirmation bias $q^*(q_w)$ according to a biased observer with confirmation bias q_w . $\alpha_0 = 3, \alpha_1 = 1$ in this example.

electoral platforms, and hence increase equilibrium welfare. Indeed, we have shown that all observers in our model, whether biased or Bayesian, would prefer the representative voter to exhibit more confirmation bias than they do themselves.

Our work contributes to an emerging group of papers suggesting that while deviations from rationality may be unambiguously undesirable for individuals acting in isolation, the implications of behavioural biases for society at large are more complex. These observations illustrate the theory of the second-best (Lipsey & Lancaster, 1956) in political contexts. Since politics invariably leads to non-optimal policy choices, increasing the ‘rationality’ of political actors does not necessarily improve outcomes, and may sometimes be counterproductive from a purely instrumental perspective.¹²

¹²Some readers may object that truth has intrinsic value, independent of its instrumental consequences. We concur, and do not view our work as a justification for biased interpretations of evidence. Rather, our findings illustrate that rational information processing is neither necessary, nor sufficient, for good collective decision-making.

References

- S. Ashworth & E. Bueno De Mesquita (2014). ‘Is voter competence good for voters?: Information, rationality, and democratic performance’. *American Political Science Review* **108**(3):565–587.
- A. Bisin, et al. (2015). ‘Government policy with time inconsistent voters’. *The American Economic Review* **105**(6):1711–1737.
- B. H. Bornstein & A. C. Emler (2001). ‘Rationality in medical decision making: a review of the literature on doctors’ decision-making biases’. *Journal of Evaluation in Clinical Practice* **7**(2):97–107.
- S. Callander & C. H. Wilson (2008). ‘Context-dependent voting and political ambiguity’. *Journal of Public Economics* **92**(3):565–581.
- A. Downs (1957). ‘An economic theory of political action in a democracy’. *Journal of Political Economy* **65**(2):135–150.
- L. G. Epstein (2006). ‘An axiomatic model of non-Bayesian updating’. *The Review of Economic Studies* **73**(2):413–436.
- G. Gratton (2014). ‘Pandering and electoral competition’. *Games and Economic Behavior* **84**:163–179.
- F. Gul & W. Pesendorfer (2001). ‘Temptation and self-control’. *Econometrica* **69**(6):1403–1435.
- A. J. Healy, et al. (2010). ‘Irrelevant events affect voters’ evaluations of government performance’. *Proceedings of the National Academy of Sciences* **107**(29):12804–12809.
- P. Heidhues & J. Lagerlof (2003). ‘Hiding Information in Electoral Competition’. *Games and Economic Behavior* **42**(1):48–74.
- D. Kahneman & A. Tversky (1973). ‘On the psychology of prediction’. *Psychological Review* **80**(237–251).
- N. Kartik, et al. (2015). ‘Information Revelation and Pandering in Elections’. *Working Paper*.
- J.-F. Laslier & K. Van der Straeten (2004). ‘Electoral competition under imperfect information’. *Economic Theory* **24**(2):419–446.
- G. Levy & R. Razin (2015). ‘Correlation neglect, voting behavior, and information aggregation’. *The American Economic Review* **105**(4):1634–1645.

- R. G. Lipsey & K. Lancaster (1956). 'The General Theory of Second Best'. *The Review of Economic Studies* **24**(1):11–32.
- A. Lizzeri & L. Yariv (forthcoming). 'Collective Self-Control'. *American Economic Journal: Microeconomics*. .
- B. Lockwood (forthcoming). 'Confirmation Bias and Electoral Accountability'. *Quarterly Journal of Political Science* .
- E. Maskin & J. Tirole (2004). 'The Politician and the Judge: Accountability in Government'. *American Economic Review* **94**(4):1034–1054.
- R. S. Nickerson (1998). 'Confirmation Bias: a ubiquitous phenomenon in many guises'. *Review of General Psychology* **2**(2):175–220.
- M. Rabin & J. L. Schrag (1999). 'First impressions matter: A model of confirmatory bias'. *The Quarterly Journal of Economics* **114**(1):37–82.
- C. S. Taber & M. Lodge (2006). 'Motivated Skepticism in the Evaluation of Political Beliefs'. *American Journal of Political Science* **50**(3):755–769.
- R. H. Thaler (2016). *Misbehaving: The Making of Behavioral Economics*. W. W. Norton & Company, United States of America.

Appendix

A Proof of Proposition 1 (existence)

For $\delta \in [0, \bar{y}]$, the following belief system supports (13) as a SPCBE: for $i \in \{A, B\}$

$$\text{if } i \text{ plays } \begin{cases} y' \leq -\bar{y} \\ -\bar{y} + \delta, \\ \bar{y} - \delta, \\ y' > \bar{y} \end{cases} \quad \text{voter assign beliefs } \begin{cases} (1, 0, 0) \\ (1 - \frac{\delta}{\bar{y}}, \frac{\delta}{\bar{y}}, 0) \\ (0, \frac{\delta}{\bar{y}}, 1 - \frac{\delta}{\bar{y}}) \\ (0, 0, 1) \end{cases} \quad \text{to } i\text{'s signal vector } (-1, 0, 1). \text{ Using}$$

these beliefs we can extend (in the obvious way) the specification (13) of ω_i^y to cover all unilateral deviations from $y(\cdot)$.

To verify that (13) is a best response to itself, it is straightforward to check, from the definition of \bar{y} and the expressions in (10) for the voter's posterior expectations of θ , that for any signal pair, the two platforms defined by \bar{y} are equi-distant from the voter's posterior bliss point conditional on that pair.

We now show that player i cannot do better by deviating to any platform in \mathfrak{R} . If:

- 1) i plays $y' < -\bar{y}$; the voter believes that with probability 1, i 's signal was -1. For each possible realization of s_j , it follows immediately from (10) that y' is further than $y_j(s_j)$ from the voter's ideal point, $E[\theta|(-1, s_j)]$; thus i loses the election, for any signal realization by candidate j .
- 2) i plays $-\bar{y} + \delta$, for $\delta \in (0, \bar{y})$; the voter assigns beliefs $(1 - \frac{\delta}{\bar{y}}, \frac{\delta}{\bar{y}}, 0)$ to the possibilities that $s_i = (-1, 0, 1)$. We will show that with these beliefs, the voter is indifferent between i and j , for any signal s_j received by j .
 - a) suppose $s_j = -1$. The voter has to consider 2 possibilities: if $s_i = -1$, then $E[\theta|(s_i, s_j)] = -\bar{y}$, so she receives utility $-\delta^2$ if she votes for i and zero if she votes for j ; if $s_i = 0$, then $E[\theta|(s_i, s_j)] = -\bar{y}/2$, so she receives utility $-(\bar{y}/2 - \delta)^2$ if she votes for i and $-\bar{y}^2/4$ if she votes for j . Her expected utility difference between voting for i and j is thus

$$\left(1 - \frac{\delta}{\bar{y}}\right)\delta^2 + \frac{\delta}{\bar{y}}\left(\left(\frac{\bar{y}}{2} - \delta\right)^2 - \bar{y}^2/4\right) = \delta^2 - \delta^3/\bar{y} + \frac{\delta}{\bar{y}}\left(\delta^2 - \delta\bar{y}\right) = 0;$$

That is, the voter is indifferent between the two candidates.

- b) suppose $s_j = 0$. If $s_i = -1$, then $E[\theta|(s_i, s_j)] = -\bar{y}/2$, so the voter receives utility $-(\bar{y}/2 - \delta)^2$ if she votes for i and $-\bar{y}^2/4$ if she votes for j ; if $s_i = 0$, then $E[\theta|(s_i, s_j)] = 0$, so she receives utility $-(\bar{y} - \delta)^2$ if she votes for i and 0 if she votes for j . The expected utility difference between voting for i and j is thus

$$\left(1 - \frac{\delta}{\bar{y}}\right)\left(\left(\frac{\bar{y}}{2} - \delta\right)^2 - \bar{y}^2/4\right) + \frac{\delta}{\bar{y}}(\bar{y} - \delta)^2 = \left(1 - \frac{\delta}{\bar{y}}\right)(\delta^2 - \bar{y}\delta) + \frac{\delta}{\bar{y}}(\bar{y} - \delta)^2 = \delta/\bar{y}\left((\bar{y} - \delta)(\delta - \bar{y}) + (\bar{y} - \delta)^2\right) = 0$$

Again, the voter is indifferent between the two candidates.

- c) suppose $s_j = +1$. If $s_i = -1$, then $E[\theta|(s_i, s_j)] = 0$; so the voter receives utility $-(\bar{y}-\delta)^2$ if she votes for i and $-\bar{y}^2$ if she votes for j ; if $s_i = 0$, then $E[\theta|(s_i, s_j)] = +\bar{y}/2$, so she receives utility $(3\bar{y}/2 - \delta)^2$ if she votes for i and $-\bar{y}^2/4$ if she votes for j . The expected utility difference between voting for i and j is thus

$$\left(1 - \frac{\delta}{\bar{y}}\right) \left((\bar{y} - \delta)^2 - \bar{y}^2 \right) + \frac{\delta}{\bar{y}} \left((3\bar{y}/2 - \delta)^2 - \bar{y}^2/4 \right) = \delta/\bar{y} \left((\bar{y} - \delta)(\delta - 2\bar{y}) + 2\bar{y}^2 - 3\bar{y}\delta + \delta^2 \right) = 0$$

Once again, the voter is indifferent between the two candidates.

- 3) i plays $\bar{y} - \delta$, for $\delta \in (0, \bar{y})$; the voter assigns beliefs $(0, \frac{\delta}{\bar{y}}, 1 - \frac{\delta}{\bar{y}})$ to the possibilities that $s_i = (-1, 0, 1)$. The proof that the voter is indifferent between i and j , for any signal s_j received by j , is the mirror image of proof step 2).
- 4) i plays $y' > \bar{y}$; the voter believes that with probability 1, i 's signal was +1. The proof that i loses the election, for any signal realization by j , is the mirror image of proof step 1).

B Proof of Proposition 1 (uniqueness)

Let $y(\cdot)$ be a fully revealing SPCBE.

- 1) All type-pairings result in a tied election: Clearly (A) if $s_i = s_j$, then the election result is a tie, because i and j play identical strategies. Second, observe that (B) if candidate j of type 0 is paired with a candidate of type $s_i \neq 0$ then the result must be a tie. To see this, suppose (w.l.o.g) that $s_i = -1$.

- a) *if j wins against i* : Then condition (12) fails for $s_i = -1$ and $\zeta = y_i(0)$, since:

$$\sum_{\Theta} \rho(s_j|-1) \omega_i^y(y(-1), y(s_j); q) \leq \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left((\alpha_{-1} + 1)/2 + \alpha_{+1} \right) = \frac{3\alpha_{-1} + 1}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

since $y(-1)$ ties against $y(-1)$, loses against $y(0)$ and, at best, wins against $y(1)$, while

$$\sum_{\Theta} \rho(s_j|-1) \omega_i^y(y(0), y(s_j); q) \leq \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left((\alpha_{-1} + 1) + \alpha_{0/2} \right) = \frac{2\alpha_{-1} + \alpha_0 + 2}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

since $y(0)$ wins against $y(-1)$, ties against $y(0)$ and, at worst, loses against $y(1)$. Since $\alpha_0 > \alpha_{-1}$, this shows $y(0)$ is a profitable deviation for i from $y(-1)$ conditional on $s_i = -1$.

- b) *if j loses against i* : Then condition (12) fails for $s_j = 0$ and $\zeta = y(-1)$, since:

$$\sum_{\Theta} \rho(s_i|0) \omega_j^y(y(0), y(s_i); q) \leq \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left((\alpha_0 + 1)/2 + \alpha_{+1} \right) = \frac{\alpha_0 + 1 + 2\alpha_{-1}}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

since $y(0)$ loses against $y(-1)$, ties against $y(0)$ and, at best, wins against $y(1)$, while

$$\sum_{\Theta} \rho(s_i|0) \omega_j^y(y(-1), y(s_i); q) \leq \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left((\alpha_{-1} + 1)/2 + \alpha_0 \right) = \frac{\alpha_{-1} + 1 + 2\alpha_0}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

since $y(-1)$ ties against $y(-1)$ wins against $y(0)$, and, at worst, loses against $y(1)$. This shows $y(-1)$ is a profitable deviation for j from $y(0)$ conditional on $s_j=0$.

It now follows from (B) that if candidate i of type -1 is paired with candidate j of type $+1$, then the election result must be a tie. To see this, suppose w.l.o.g. that to the contrary

c) *i loses against j*: Then condition (12) fails for $s_i = -1$ and $\zeta = y(+1)$, since:

$$\sum_{\Theta} \rho(s_j|-1) \omega_i^y(y(-1), y(s_j); q) = \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left((\alpha_{-1} + 1)/2 + \alpha_0/2 \right) = \frac{\alpha_{-1} + \alpha_0 + 1}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

since $y(-1)$ ties against $y(-1)$ and $y(0)$, and loses against $y(+1)$, while

$$\sum_{\Theta} \rho(s_j|-1) \omega_i^y(y(+1), y(s_j); q) = \frac{1}{1 + \sum_{\theta} \alpha_{\theta}} \left(3(\alpha_{-1} + 1)/2 + \alpha_0/2 \right) = \frac{3\alpha_{-1} + \alpha_0 + 1}{2(1 + \sum_{\theta} \alpha_{\theta})}$$

This shows $y(+1)$ is a profitable deviation for i from $y(-1)$ conditional on $s_i=-1$.

2) The unique vector of platforms s.t. all type pairings result in ties is (13):

Assume that $y(\cdot)$ satisfies this property. The following properties are all implications of the calculations (10).

- i) $E[\theta|(-1, +1)] = 0$ implies that $y(-1) + y(+1) = 0$.
- ii) property i), together with $|E[\theta|(-1, 0)]| = |E[\theta|(0, +1)]|$, implies $y(0) = 0$.
- iii) $y(0) = 0$ together with $E[\theta|(-1, 0)] = \frac{-(1-q)}{2\alpha_1 + \alpha_0 + 2}$ implies $y(-1) = \frac{-2(1-q)}{2\alpha_1 + \alpha_0 + 2}$.
- iv) $y(0) = 0$ together with $E[\theta|(0, +1)] = \frac{1-q}{2\alpha_1 + \alpha_0 + 2}$ implies $y(+1) = \frac{2(1-q)}{2\alpha_1 + \alpha_0 + 2}$.