

# Confirmation Bias and Electoral Accountability

Ben Lockwood\*

*University of Warwick, UK; B.Lockwood@warwick.ac.uk*

---

## ABSTRACT

This paper considers the implications of an important cognitive bias in information processing, confirmation bias, in a political agency setting. When voters have this bias and when only the politician's actions are observable before the election, it decreases pandering by the incumbent, and can raise voter welfare as a consequence. This result is driven by the fact that the noise aspect of confirmation bias, which decreases pandering, dominates the bounded rationality aspect, which increases it. The results generalize in several directions, including to the case where the voter can also observe payoffs with some probability before the election. We identify conditions when confirmation bias strengthens the case for decision-making by an elected rather than an appointed official.

---

*Keywords:* Confirmation bias; selective exposure, voting, pandering, elections.

This paper contributes to the growing literature on the effect of voter and politician behavioral biases on the performance of electoral institutions. Our focus here is on a key bias in information-processing, confirmation bias. As Ra-

---

\*I would like to thank Peter Buisseret, Chris Ellis, Gilat Levy, seminar participants at the universities of Oregon, Princeton and Warwick, and two referees and an editor for helpful comments. I also thank Omiros Kouvas and Federico Trombetta for excellent research assistance.

---

Online Appendix available from:

[http://dx.doi.org/10.1561/100.00016037\\_app](http://dx.doi.org/10.1561/100.00016037_app)

MS submitted on 12 April 2016; final version received 10 December 2016

ISSN 1554-0626; DOI 10.1561/100.00016037

© 2017 B. Lockwood

bin and Schrag (1999) put it, “A person suffers from confirmatory bias if he tends to misinterpret ambiguous evidence as confirming his current hypotheses about the world.” This is one of the most pervasive and well-documented forms of cognitive bias; as Nickerson (1998) says, in a recent survey, “If one were to attempt to identify a single problematic aspect of human reasoning that deserves attention above all others, the confirmation bias would have to be among the candidates for consideration.” Indeed, there is even some evidence of a genetic basis for confirmation bias (Doll *et al.*, 2011).

Nickerson (1998) emphasizes two mechanisms underlying confirmation bias; preferential treatment of evidence supporting existing beliefs, and looking only or primarily for positive cases that support initial beliefs. This second mechanism is sometimes called selective exposure. There is considerable evidence for both mechanisms. Evidences for preferential treatment of given evidence include experiments where subjects were initially questioned on a salient policy issue (Lord *et al.*, 1979), capital punishment; (Plous, 1991) and safety of nuclear technology to determine their views, and then presented with the same randomly sampled reading material for and against the issue. After exposure, those initially in favour (against) tended to be more in favour (against), despite having been exposed to the *same* reading material. There is also a large body of experimental evidence that selective exposure occurs.<sup>1</sup>

As confirmation bias is a bias in information processing, it is particularly relevant in political economy settings where decision-makers update their beliefs in response to new information. In particular, voters may be prone to confirmation bias, because as professionals, with access to expert advice, politicians and bureaucrats are perhaps less likely to suffer from this bias.<sup>2</sup> So, we focus on voter confirmation bias in this paper.

In this paper, we introduce voter confirmation bias into a fairly general political agency model. Political agency models are widely used to study the degree to which elections can hold incumbents accountable for their actions while in office.<sup>3</sup> Our model is quite flexible; if the voter only observes the actions of the incumbent before the election, the model is a variant of Maskin and

---

<sup>1</sup>In the classic experimental selective-exposure research paradigm, participants are given the opportunity to search for additional information when faced with a binary choice problem, which is typically in the form of short statements indicating the perspectives of newspaper articles, experts, or former participants. In a meta-analysis of 91 such studies, Hart *et al.* (2009) find significant evidence indicating that participants choose additional information that confirms their initial decisions.

<sup>2</sup>However, there are well-known examples of political leaders ignoring negative evidence about their policies, when they have strong prior beliefs in the efficacy of such policies (Canes-Wrone *et al.*, 2001; Majumdar and Mukand, 2004).

<sup>3</sup>For surveys of the political agency literature, see Besley (2006) and Ashworth (2012).

Tirole's (2004) model of political pandering, and if only payoffs are observed, the model is a variant of that used in Chapter 3 of Besley (2006).

We focus on the first form of bias, preferential treatment of given evidence; a complete study of selective exposure requires also the modelling of the supply of information, for example by the media, and is beyond the scope of this paper. To model confirmation bias, we adopt the approach of Rabin and Schrag (1999), who assume that when the agent gets a signal that is counter to the hypothesis he currently believes is more likely, there is a positive probability that he misreads that signal as supporting his current hypothesis. Moreover, the agent is unaware that he is misreading the signal in this way, and consequently ignores the error when updating his prior.

To understand the effect of confirmation bias in this setting, it is helpful to note that the Rabin–Schrag formulation of confirmation bias has two distinct elements: (1) it introduces (biased) *noise* into the voter's observation of the action of the politician, and (2) the voter is *boundedly rational* in the sense that she fails to take account of the noise when performing Bayesian updating. The question then is how each of these two elements affect the degree of pandering in equilibrium.<sup>4</sup>

Our main results are the following: (1) The noise effect tends to reduce pandering, whereas the bounded rationality effect tends to increase it. (2) The bounded rationality effect is always dominated by the noise effect, so that overall, confirmation bias reduces pandering. (3) The relative contributions of noise and bounded rationality turn on whether the voter *only* cares about the incumbent's quality of decision-making, or whether there is some other dimension of preference over candidates. In the first case, the re-election probability does not depend on exactly how the voter updates, that is the re-election probability is 1 if the voter believes the incumbent is better than the challenger, and 0 otherwise. Then, the effect of confirmation bias on political equilibrium only works via the noise effect, with bounded rationality playing no role.

However, if the voter has non-policy preferences, the re-election probability of the incumbent is a smooth function of the voter posterior belief, and then the bounded rationality element of confirmation bias comes into play. In particular, bounded rationality makes the electoral return to pandering higher, because the voter ignores the noise in the signal of the politician's action, and thus updates as if the signal of the action were perfectly accurate.

We then turn to study voter welfare. As pandering generally has an ambiguous effect on voter welfare, it is possible that an increase in confirmation bias increases voter welfare, and we identify conditions under which this happens. We also consider the robustness of our results in several directions.

---

<sup>4</sup>The effects of changing observability of the incumbent's actions on voter welfare have already been studied in a setting with fully rational voters by Prat (2005) and Fox (2007). Their contributions are discussed later in more detail.

First, we allow the voter to observe not only the action of the incumbent, but also, with some probability, the payoff before the election (Maskin and Tirole (2004) call this the “feedback” case), or indeed just the payoff. Second, we show that our basic argument applies in other leading models of electoral accountability, such as Canes-Wrone *et al.* (2001) and Fox (2007).

Finally, we revisit the choice between a politician and an unelected official, the focus of Maskin and Tirole’s original paper. With voter confirmation bias, when the choice between an elected and appointed official is not trivial, confirmation bias always works in favour of the elected official; this is because bias reduces pandering. So, in policy areas where voter confirmation bias is likely to be strong — perhaps where voters have strong prior beliefs — it is better, other things equal, to have elected officials rather than non-elected officials. This is broadly consistent with the observation that in the public policy arena, decisions concerning, for example taxation are taken by politicians, whereas technical decisions, such as those concerning monetary policy or utility regulation, are usually taken by appointed officials.

## 1 Related Literature

This paper is a contribution to a small but growing literature studying the implications of introducing behavioral and cognitive biases into rational choice models of voting.<sup>5</sup> The most closely related contribution is by Ashworth and Bueno de Mesquita (2014), who are the first to consider deviations from the full rationality of the voter in a political agency setting. In particular, they consider voters who in their words, “fail to filter.” This refers to the stylized fact that voters vote for or against the incumbent partly in response to events such natural disasters or economic shocks, or even changes in personal circumstances, that the voters should know are outside of the politicians’ control. They model this by assuming that in addition to the policy payoff from the incumbent’s action, the voter gets a random shock to his payoff from this exogenous event if he votes for the incumbent.<sup>6</sup>

In their setting, the “good” politician is a non-strategic type that always acts in the interests of the voters, and the “bad” politician is an extremist. Generally, an incumbent extremist chooses a policy which is more moderate than he would like in order to increase his chances of re-election. They then

---

<sup>5</sup>Other contributions include Callander and Wilson (2006, 2008) who introduce a theory of context-dependent voting, Ghirardato and Katz (2006) and Ellis (2016) on ambiguity aversion and voting, Alesina and Passarelli (2015) and Lockwood and Rockey (2016) on the effect of loss aversion on direct democracy and electoral competition, respectively, and Matějka and Tabellini (2016) on rationally inattentive voters.

<sup>6</sup>Analytically, the random shock assumed by Ashworth and Bueno de Mesquita (2014) is similar to the preference parameter  $\theta$  in our model below. It is thus analytically distinct from confirmation bias, which is modelled in our framework by the parameter  $q$ .

show that the random shock to voter preferences can, under some conditions, strengthen the link between policy moderation and re-election, thus inducing more moderation in equilibrium. In turn, this can raise voter welfare.

However, there are a number of differences in our approaches. First, confirmation bias is a distinct type of bias to failing to filter, and the mechanism at work is different. The available evidence suggests that failure to filter is probably driven by an affective, rather than cognitive, process, namely a well-being spillover, where a random shock that increases income of well-being makes the voter better disposed to the incumbent.<sup>7</sup>

Second, Ashworth and Bueno de Mesquita (2014) make the strong assumption that the “good” politician is a non-strategic type that always acts in the best interests of the voters, that is has no re-election motive. This is an important restriction, because it means that they cannot analyze political pandering; rather, a strategic decision is only made by the bad incumbent, who must decide whether to imitate, that is pool with, the good incumbent or not.

Third, for conditions under which failing to filter can improve voter welfare, identified in their Proposition 4, the mechanism at work is the reverse to ours. Specifically, they find that the incentive for the bad incumbent (the extremist in their model) to imitate the good one (the moderate) can be stronger under a fail-to-filter voter than under a rational voter, and so failing to filter buys the voter better discipline of the incumbent at the cost of worse selection. In contrast, as described earlier, we find that confirmation bias implies less pandering but better selection. So, overall, the results of this paper are complementary to theirs.

Our paper is also close in spirit to Levy and Razin (2015), who find that the cognitive bias of correlation neglect can improve outcomes for voters, due to a second-best argument; in their setting, information aggregation via voting is initially inefficient, because voters underweight their information when deciding how to vote. If a voter ignores the fact that two of her signals are correlated, she will “overweight” the signals, and thus put more weight on her information, offsetting the original distortion. However, both the institutions and the mechanism at work are completely different. They consider direct democracy, that is a referendum on two alternatives, and correlation neglect causes individuals base their vote more on their information rather than on their preferences.<sup>8</sup>

---

<sup>7</sup>For example, Bagues and Esteve-Volart (2016) show that random income shocks to Spanish regions due to the national lottery have a positive effect on incumbent vote share. Liberini *et al.* find that a random negative life event (widowhood) can make individuals less willing to support the party of government, using UK panel data.

<sup>8</sup>Ortoleva and Snowberg (2015), in a related paper, show theoretically that correlation neglect, overconfidence and ideological extremeness are connected; empirically they find,

The last related literature is the one that studies the effects of additional voter information on equilibrium outcomes in political agency models when the voter is fully rational. Specifically, there are a number of papers showing that additional information may not be to the benefit of voters, because it may induce a strategic response by the incumbent politician, and in turn, this strategic response may weaken either the selection or discipline effects of elections (Besley, 2006; Fox, 2007; Prat, 2005). For example, Prat (2005) makes this point in a general agency model where the agent varies in competence; starting from a baseline where the principal can only observe the payoff from the action of the agent, allowing the principal to observe the action as well can make the principal worse off, as it induces the bad agent to pool with the good one, and thus worsens selection. However, unlike this paper, all these contributions assume full rationality of voters.

## 2 The Set-Up

Our set-up is a variant of Maskin and Tirole (2004). While this is not the only model of political pandering, it has the advantage of being well-known and relatively simple. We argue in Section 5 that our main insights extend to two other well-known models of political pandering, Canes-Wrone *et al.* (2001), and Fox (2007).

A single voter lives for periods  $t = 1, 2$ . In each of the two periods, a politician chooses a binary policy  $x_t \in \{A, B\}$ . The first-period incumbent faces an election at the end of his first term of office, where the voter can either re-elect the incumbent or elect a challenger. The payoffs of voters and politicians depend on the action and a state of the world  $s_t \in \{A, B\}$ . All agents, that is incumbent, voter, and challenger have a prior belief  $1 > p > 0.5$  that state  $A$  will occur.

### 2.1 Payoffs

The voter gets a policy-related payoff in period  $t$ , which is 1 if the incumbent's action in period  $t$  matches the state, and 0 otherwise. Moreover, if the voter elects the challenger at the end of period 1, the voter gets a non-policy related payoff  $\theta \in \mathfrak{R}$ , which measures, for example attractiveness or the valence of the challenger relative to the incumbent, as in Morelli and Van Weelden (2013).<sup>9</sup> We assume that  $\theta$  is determined by random draw from a continuous distribution

---

using a large US election study, that overconfidence is the most reliable predictor of ideological extremeness and an important predictor of voter turnout.

<sup>9</sup>The problem with the second interpretation is that the valence must be unknown to the incumbent himself until after the action is taken (Morelli and Van Weelden, 2013). Some may find this assumption implausible.

$G$  that is symmetric around 0. The role of  $\theta$  is to smooth the response of the re-election probability to voter beliefs.

Following Maskin and Tirole (2004), we assume that politicians get zero payoff when out of office, and enjoy an exogenous ego-rent  $E$  when in office; they also care about policy choices when in office.

Politicians are of two types, consonant, denoted  $C$ , and dissonant, denoted  $D$ .<sup>10</sup> The incumbent and challenger types are independent draws from the same distribution, where the probability of a consonant type is  $\pi > 0.5$ . All the results of this paper extend to the case where  $\pi < 0.5$ , but allowing for both cases considerably complicates the formal definition of confirmation bias.

Congruent politicians, when in office, get utility  $u_t$  if  $s_t = x_t$ , and 0 if  $s_t \neq x_t$ . Here,  $u_1, u_2$  are i.i.d. random variables with a continuous distribution  $F$  on support  $[0, \bar{u}]$ . So, they share the same basic preferences as voters, but can vary in the extent to which they value an action that matches the state. Dissonant politicians, when in office, get  $u_t$  if  $s_t \neq x_t$ , and 0 if  $s_t = x_t$ . We assume without loss of generality that  $E[u_t] = 1$ , and we assume  $\bar{u} > \delta(1 + E)$ , where  $\delta$  is defined below. This inequality ensures that for some values of  $u_1$ , the incumbent prefers not to pander even at the cost of not being re-elected.

The reason why we assume that politicians' payoffs from their most preferred outcome are determined by random draw (rather than being fixed at 1, as in Maskin and Tirole) is twofold. First, this ensures uniqueness of equilibrium, as explained in Section 2.5. Second, it ensures that in all cases,  $x_t$  is an informative, but not perfect signal of politician type, so that the Rabin–Schrag definition of confirmation bias can be applied.<sup>11</sup>

Finally, both voter and the incumbent discount second-period payoffs by  $\delta$ .

## 2.2 Order of Events and information Structure

The timing of events and the information available to each player at each stage is as follows:

1. In period  $t = 1$ , nature determines the type  $\{C, D\}$  of the incumbent and challenger, the state of the world  $s_1 \in \{A, B\}$ ,  $u_1 \in [0, \bar{u}]$  and the voter's non-policy related payoff  $\theta \in \mathfrak{R}$ . The incumbent observes  $u_1$ , and  $s_1$ , but not  $\theta$ . The voter only observes  $\theta$ .
2. The incumbent chooses  $x_1 \in \{A, B\}$ .
3. The voter observes  $x_1$  and votes to retain the incumbent or to replace him with a challenger.

<sup>10</sup>The terminology “congruent” and “dissonant” is taken from Besley (2006). It is a little more memorable than the Maskin and Tirole (2004) terminology of “congruent” and “non-congruent.”

<sup>11</sup>A problem arises with  $u_t \equiv 1$  because then in the pandering equilibrium,  $x_t$  is not an informative signal of type, as both  $C$  and  $D$  types choose  $x = A$  with probability 1.

4. All players receive their first-period payoffs.
5. In period  $t = 2$ , nature determines the state of the world,  $s_2 \in \{A, B\}$ , and  $u_2 \in [0, \bar{u}]$ . The incumbent observes  $u_2$ , and  $s_2$ , and chooses  $x_2 \in \{A, B\}$ . Then, all players receive their second-period payoffs.

Note that we assume, following Maskin and Tirole (2004), that the voter observes the action  $x_1$  before election, but not the payoff generated by  $x_1$ . In Section 4, we instead assume that the voter observes his payoff  $v_1$ , or both  $v_1, x_1$ , rather than  $x_1$ .

### 2.3 The Second Period

To define confirmation bias in the simplest way, it is helpful to reduce the model to a one-period game between the incumbent and the voter by solving out for the second period. This is of course, consistent with solving the model for the perfect Bayesian equilibrium.

In the second period, consonant (dissonant) politicians match the action to the state according to their preferences, that is consonant politicians choose  $x_2 = s_2$ , and dissonant politicians choose  $x_2 \neq s_2$ . Thus, both types of politicians have an expected continuation payoff from election of  $\delta(E[u_2] + E) = \delta(1 + E) \equiv V$ .

Moreover, as a consonant (resp. dissonant) incumbent generates a payoff of 1 (resp. 0) for the voter in the second period, the voter's expected payoff from re-electing the incumbent is just equal to the posterior probability that he is consonant, and his expected payoff to electing the challenger is  $\pi + \theta$ . Armed with these descriptions of the second-period continuation payoffs of the actors, we can now focus entirely on the first period, and so we can drop time subscripts without ambiguity. So,  $x, s, \dots$  now refer to  $x_1, s_1, \dots$  and so on.

### 2.4 Modelling Confirmation Bias

Rabin and Schrag define confirmation bias in a single-person decision problem, where the decision-maker (agent) gets noisy signals about a payoff-relevant state of the world. They assume that “when the agent gets a signal that is counter to the hypothesis he currently believes is more likely, there is a positive probability that he misreads that signal as supporting his current hypothesis. The agent is unaware that he is misreading evidence in this way and engages in Bayesian updating that would be fully rational given his environment if he were not misreading evidence” (Rabin and Schrag, 1999, p. 48).

To extend this definition, we first need to identify what is the payoff-relevant state of the world, and the signal, for the voter. At the time when the voter acts, that is votes, the payoff-relevant state of the world is the type of the incumbent, because that is persistent by assumption, and thus determines the

expected payoff of the voter in the next period. Also by assumption, the only thing observed by the voter before the election is the action  $x$ , so this is the signal.

The complication here is that the link between the state of the world, thus defined, and the signal, is generated by equilibrium play of the game between incumbent and voter. This is in contrast to Rabin and Schrag (1999) set-up where the link between the signal and the state of the world is exogenous.<sup>12</sup> This, of course, creates a possible problem of circularity — the definition of confirmation bias depends on the equilibrium strategy of the incumbent, which in turn may depend on the definition of confirmation bias.

To deal with this, we will define confirmation bias *conditional on incumbent choices*, and then make confirmation bias part of the definition of equilibrium. Specifically, define  $a_C, a_D$  to be the unconditional probabilities that type  $C, D$  incumbents, respectively, choose action  $A$  in period 1. Note that these probabilities are not conditional on  $u, s$  and are therefore typically between 0 and 1 even though conditional on  $u, s$ , the incumbent is assumed to play a pure strategy, that is to choose  $x \in \{A, B\}$ . Then, following Rabin and Schrag (1999), voter confirmation bias can be defined as follows.<sup>13</sup>

**Definition 1** (Confirmation Bias Conditional on Incumbent Choices). *If  $a_C > a_D$ , the voter misreads  $x = B$  as  $x = A$  with probability  $q > 0$ . If  $a_C < a_D$ , then the voter misreads  $x = A$  as  $x = B$  with probability  $q > 0$ . Moreover, when the voter updates his beliefs, he ignores these errors.*

To interpret this definition, note that as  $\pi > 0.5$ , the voter’s “current hypothesis,” in the sense of Rabin and Schrag, is that the incumbent is consonant and so the voter is biased in favour of mis-reading negative signals as positive signals. The definition says that if choice of action  $A$  is a positive signal of the incumbent being consonant, that is  $a_C > a_D$ , the voter misreads a signal that the incumbent is dissonant ( $x = B$ ) as a signal that the incumbent is consonant ( $x = A$ ) with probability  $q$ . On the other hand, if  $a_C < a_D$ , positive and negative signals are reversed, that is  $x = B$  is now a signal that the incumbent is consonant. Conditional on this reversal, however, the voter still misreads a signal that the incumbent is dissonant as a signal that the incumbent is consonant with probability  $q$ .

---

<sup>12</sup>The formal definition of Rabin and Schrag (1999) is the following. Assume a binary state of the world,  $s = \{A, B\}$ , and sequence  $t = 1, \dots, T$  of informative signals  $\sigma_t \in \{A, B\}$  about the state, where  $\Pr(\sigma_t = K | s = K) = \theta > 0.5$ . If  $\pi_t$  is the decision-maker’s prior that the state is  $A$  at  $t$ , then: (i) if  $\pi_t > 0.5$ , the agent misreads  $\sigma_t = B$  as  $\sigma_t = A$  with probability  $q$ , and (ii) if  $\pi_t < 0.5$ , the agent misreads  $\sigma_t = A$  as  $\sigma_t = B$  with probability  $q$ .

<sup>13</sup>We rule out the borderline case where  $\pi = 0.5$ ; in this case, Rabin and Schrag assume no confirmation bias, that is  $q = 0$ . So, this case, apart from being non-generic, is also uninteresting.

A very helpful way of modelling the misreading of action  $x$  in Definition 1 is to think of the voter as observing a noisy and biased signal of  $x$ . Specifically, we define the signal  $\sigma \in \{A, B\}$  of  $x$  as follows.

**Definition 2** (The Signal  $\sigma$ ). *If  $a_C > a_D$ ,  $\sigma(A) = A$  with probability 1,  $\sigma(B) = A$  with probability  $q$  and  $\sigma(B) = B$  otherwise. If  $a_C < a_D$ , then  $\sigma(B) = B$  with probability 1, and  $\sigma(A) = B$  with probability  $q$ , and  $\sigma(A) = A$  otherwise.*

Then, comparing Definitions 1 and 2, it is clear that confirmation bias is formally equivalent to (i) the voter observing  $\sigma$ , rather than  $x$ , and (ii) believing that he has observed  $x$ , that is ignoring the noise in  $\sigma$  when he updates his beliefs about the quality of the incumbent. We call this last feature *boundedly rational updating*. This characterization will be very useful in what follows.

Finally, when it comes to politician behaviour, we will assume that the politician understands that the voter has confirmation bias, and takes this into account when making his policy choices. This seems a reasonable assumption; in modern politics, political parties conduct extensive research into voter attitudes and behavior (Gibson and Römmele, 2009).

## 2.5 Equilibrium Concept

We focus on perfect Bayesian equilibrium in what follows, which we call just a *political equilibrium*. Write  $\sigma(x, a_C, a_D)$  to emphasize that the signal  $\sigma$  depends on  $x, a_C, a_D$ . Then, a *political equilibrium* is comprised of: (i) a voter decision to elect the incumbent or challenger conditional on  $\sigma(x, a_C, a_D), \theta$ ; (ii) incumbent choices of  $x \in \{A, B\}$ , conditional on incumbent type and state of the world, which maximize incumbent payoffs given voting rule (i); (iii) probabilities  $a_C, a_D$  that are consistent with incumbent choices.

It is understood in this definition that as part of (i), the voter updates in a boundedly rational way. We show here that this equilibrium is unique, given our assumption  $\bar{u} > \delta(1 + E)$ .<sup>14</sup>

We close with a definition of pandering which follows Maskin and Tirole (2004). The incumbent is said to *pander* if he chooses the action matching the state that the voter believes is ex ante more likely, whatever the actual state of the world. As  $p > 0.5$ , pandering in the first period is therefore a choice of

---

<sup>14</sup>In particular, each type of incumbent chooses both actions with positive probability on the equilibrium path; this rules out a “perverse” pandering equilibrium, where the voter re-elects the incumbent only if he thinks he observes action  $B$ . Such an equilibrium can arise in Maskin and Tirole (2004) model, because the incumbent payoff to matching the action to the state is fixed at 1.

$x = A$  for  $s = A, B$ . A central concern of the analysis will be the probability of pandering by the incumbent in equilibrium.

## 2.6 The Rational Voter With Noise

To proceed, it is now helpful to introduce the idea of the rational voter with a noisy signal (the *rational voter with noise* for short, denoted the RN voter). Like the voter with confirmation bias (the CB voter for short), the RN voter observes  $\sigma$ . But, the RN voter updates his prior belief  $\pi$  in a fully rational way, taking into account that  $\sigma$  is defined as in Definition 2, whereas the CB voter updates his prior ignoring this i.e. assumes that  $\sigma = x$ .

The rational voter with noise is just an intermediate construct which allows us to distinguish the noise and bounded rationality effects of confirmation bias on political equilibrium. In what follows, we characterize equilibrium separately for RN and CB voters. This will enable us to decompose the effects of the two components of confirmation bias on the level of pandering in equilibrium. The definition of equilibrium for the RN voter is exactly as defined earlier, except it is now understood that in part (i), given  $\sigma$ , the RN voter updates in a fully rational way.

## 3 Confirmation Bias and Pandering

### 3.1 Political Equilibrium

We give a brief informal description of the structure of the equilibrium before stating our main results. Let  $\pi_{\text{RN}}(\sigma)$ ,  $\pi_{\text{CB}}(\sigma)$  be the posterior belief that the incumbent is consonant for the RN and CB voters respectively, having observed signal  $\sigma \in \{A, B\}$  as defined earlier. The formulae for these are given in the Appendix. Note that because the voter gets a payoff of one in the second period if the politician is consonant and zero otherwise, the voter of type  $k = \text{RN}, \text{CB}$  will re-elect the incumbent, having observed  $\sigma = A, B$  if and only if the difference in perceived quality between the incumbent and challenger,  $\pi_k(\sigma) - \pi$ , exceeds the non-policy preference for the challenger  $\theta$ , that is

$$\pi_k(\sigma) - \pi \geq \theta, \quad k = \text{RN}, \text{CB} \quad (1)$$

This voting rule generates re-election probabilities for the incumbent, conditional on actions, of  $r_k(A), r_k(B)$ ,  $k = \text{RN}, \text{CB}$ . We can show that in equilibrium,  $r_k(A) > r_k(B)$ ; see the Appendix.

So, the incumbent clearly faces a choice of whether to pander, that is always choose  $x = A$ , or to take the short-run optimal action. These two objectives

only conflict when  $A$  is not short-run optimal for the incumbent. In that case, the opportunity cost of pandering is  $u$ , the benefit from the short-run optimal action, whether the incumbent is consonant or dissonant.

The benefit of pandering is the second-period continuation payoff,  $V$ , times the increase in the re-election probability  $r_k$ ,  $\Delta r_k = r_k(A) - r_k(B)$ , from choosing  $x = A$  over  $x = B$ , giving an expected benefit of  $\Delta r_k V$ . So, the incumbent will pander if and only if  $u \leq \Delta r_k V$ , giving a pandering probability of  $\lambda_k = F(\Delta r_k V)$ ,  $k = \text{RN, CB}$ . Note that this probability  $\lambda_k$  is the same for both consonant and dissonant incumbents, as they both have the same continuation payoff. So, our main focus will be on the pandering probability  $\lambda_k$ ,  $k = \text{RN, CB}$ .

We are now ready to state our first result.<sup>15</sup>

**Proposition 1.** *Assume voters have no non-policy preferences ( $\theta \equiv 0$ ). Then, there is a unique political equilibrium where: (a) the voter re-elects the incumbent iff  $\sigma = A$ , whether the voter is a RN or CB type; (b) both incumbent types pander with probability  $\lambda_{\text{RN}} = \lambda_{\text{CB}} = \lambda = F((1 - q)V) < 1$ ; (c) the consonant incumbent is more likely to choose action  $A$ , that is  $a_C > a_D$ .*

There are three notable features of this equilibrium. First, without confirmation bias, the probability of pandering is of course  $\lambda_0 \equiv F(V)$ , so the presence of confirmation bias lowers the probability of pandering from  $F(V)$  to  $F((1 - q)V)$ . Second, the equilibrium is the same whether the voter is a RN type or a CB type; in other words, the bounded rationality aspect of confirmation bias has no effect on the outcome. Third, voter behavior is consistent with  $a_C > a_D$ ; that is, the voter of either type correctly believes that the  $C$  type is more likely to choose  $A$ , and consequently re-elects the incumbent only if  $\sigma = A$ .<sup>16</sup>

The reason why  $\lambda$  depends on  $q$  is the following. In equilibrium, we know  $a_C > a_D$ . So,  $r(A) = 1$ , but  $r(B) = q$ , because if  $x = B$ , the voter will re-elect the incumbent anyway with probability  $q$ , having observed an incorrect signal  $\sigma = A$ . So, the increase in the probability of re-election from choosing  $x = A$  rather than  $x = B$  is just  $1 - q$ , smaller than in the baseline case of a rational voter without any noise; in the latter case, the increase in the probability of re-election is from 0 to 1.

We call this the *mis-classification effect* of confirmation bias, as it arises because the voter mis-classifies the action. So, when voters have no policy preferences, the conclusion is that the presence of confirmation bias lowers the probability of pandering from  $F(V)$  to  $F((1 - q)V)$  due to the mis-classification effect only.

<sup>15</sup>The proofs of all results are in the Appendix.

<sup>16</sup>In fact,  $a_C = \lambda + (1 - \lambda)p$ ,  $a_D = \lambda + (1 - \lambda)(1 - p)$ , because (for example) the  $C$ -type with always choose  $A$  if he panders, and will choose  $A$  with probability  $p$  even if he does not. So, as  $\lambda < 1$ ,  $a_C > a_D$ .

If the voter does have non-policy preferences, this is no longer the case; bounded rationality comes into play. We now assume that the distribution of  $\theta$  has large enough support so that the probability of being re-elected is always strictly between 0 and 1. Assuming that  $\theta$  is distributed in  $[-d, d]$ , it is easily seen from (1) that this requires  $d > \max\{\pi, 1 - \pi\}$ .

Then, the increase in the re-election probability from pandering can be shown to be

$$\Delta r_{RN} = \Delta(q)(1 - q), \quad \Delta r_{CB} = \Delta(0)(1 - q) \tag{2}$$

if the incumbent faces an RN or CB voter, respectively.

Here the new term  $\Delta(q) < 1$  measures the dampening of the voter’s response to observing  $A$  rather than  $B$ , due to the fact that the voter now trades off a non-policy preference for the incumbent versus the challenger against the increase in incumbent quality signalled by observing  $A$ . In full,  $\Delta(q)$  is

$$\begin{aligned} \Delta(q) &= G\left(\frac{\pi(1 - \pi)(1 - h(q))}{\pi + h(q)(1 - \pi)}\right) - G\left(\frac{\pi(1 - \pi)(1 - \tilde{h}(0))}{\pi + \tilde{h}(0)(1 - \pi)}\right) \\ h(q) &= \frac{a_D + (1 - a_D)q}{a_C + (1 - a_C)q}, \quad \tilde{h}(q) = \frac{1 - a_D + a_Dq}{1 - a_C + a_Cq} \end{aligned} \tag{3}$$

where in turn,  $a_C, a_D$  are defined in (A4) of the Appendix. Note also that  $\Delta(q)$  is decreasing in  $q$ , implying that  $\Delta(q) < \Delta(0)$  and thus  $\Delta r_{RN} < \Delta r_{CB}$ . This is because for the RN voter, the sensitivity of the posterior belief  $\pi_{RN}(\sigma)$  to the signal  $\sigma$  is lower, the higher  $q$ , because the RN voter knows that the signal is more noisy and therefore weights it less. For the CB voter, there is no adjustment of updating to  $q$  — because of the boundedly rational updating effect.

We can now state the following proposition.

**Proposition 2.** *Assume the voter has non-policy preferences. If the voter is a RN type, there is a unique political equilibrium where: (i) the incumbent is re-elected iff  $\pi_{RN}(\sigma) - \pi \geq \theta$ ; (ii) both consonant and dissonant incumbents pander with probability  $\lambda_{RN} = F(\Delta(q)(1 - q)V)$ . If the voter is a CB type, there is a unique political equilibrium where (i) the incumbent is re-elected iff  $\pi_{CB}(\sigma) - \pi \geq \theta$ ; (iii) both consonant and dissonant incumbents pander with probability  $\lambda_{CB} = F(\Delta(0)(1 - q)V)$ . Finally, in both cases,  $a_C > a_D$ .*

So, we see that if voters do have non-policy preferences, the picture is more complex. For both the voter types, the mis-classification effect is present, but for the CB voter, the boundedly rational updating effect also comes into play. So now, there are two opposing effects at work; the mis-classification effect *reduces* pandering, but the boundedly rational updating effect *increases* pandering. To see this, note that as  $\Delta(q) < \Delta(0), \lambda_{RN} < \lambda_{CB}$ ; that is, the

Table 1: Pandering probabilities.

	$q = 0.1$	$q = 0.5$	$q = 0.9$
$\lambda_0$	0.606	0.606	0.606
$\lambda_{CB}$	0.552	0.324	0.067
$\lambda_{RN}$	0.547	0.300	0.057

Other parameters:  $\pi = p = 0.75$ ,  $\bar{u} = V = 1$ .

incumbent is more likely to pander if the voter is a CB type and thus boundedly rational.

The next question is which of these two effects dominates. The obvious baseline is the equilibrium level of pandering that occurs when there is no confirmation bias, that is  $q = 0$ , which is  $\lambda_0 = F(\Delta(0)V)$ . So, given that  $\lambda_{CB} = F(\Delta(0)(1-q)V)$ , overall, with confirmation bias, equilibrium pandering decreases, that is  $\lambda_{CB} < \lambda_0$ . To put it another way, the mis-classification effect always dominates the boundedly rational updating effect.

Next, given that the mis-classification effect dominates, it is interesting to know by how much; are they of roughly equal size, or is the mis-classification effect much larger? To investigate this, we report some numerical simulations in Table 1.<sup>17</sup> These simulations take into account the fact that  $\Delta(0)$ ,  $\lambda_{CB}$ ,  $\Delta(q)$ , and  $\lambda_{RN}$  are simultaneously determined. We also assume  $F$  is uniform, that is  $F(u) = u/\bar{u}$ . The equilibrium levels of pandering  $\lambda_{CB}$ ,  $\lambda_{RN}$  are reported, along with the baseline equilibrium level of pandering  $\lambda_0$ .

We can note the following from Table 1. First, as  $q$  rises, the mis-classification effect becomes stronger, and as expected, this reduces the probability of pandering, that is  $\lambda_{CB}$  falls. Also, the bounded rationality effect is small relative to the mis-classification effect; for example, when  $q = 0.9$ , the latter can drastically reduce the pandering probability, down from around 0.6 to around 0.06, but the bounded rationality effect only pushes the probability back up by about 0.01. However, it can be shown that for  $\pi < 0.5$ , the boundedly rational updating effect can almost offset the mis-classification effect.

### 3.2 Welfare

We now turn to consider the effect of changes in confirmation bias  $q$  on welfare. The definition of welfare is not straightforward in this case; should it be calculated taking into account the behavioral bias of the voter or not?

---

<sup>17</sup>The details of these simulations are in the supplementary Appendix.

In the literature on behavioral economics, the focus has been on objective measures of welfare, abstracting from behavioral biases. For example, in their study of decision-making with hyperbolic discounting, O’Donoghue and Rabin (1999) argue that “Since present-biased preferences are often meant to capture self-control problems, where people pursue immediate gratification on a day-to-day basis, we feel the natural perspective in most situations is the long-run perspective,” and consequently their welfare criterion is the decision-maker utility without present bias. A similar assumption is made in Bernheim and Rangel (2004) study of addiction.<sup>18</sup>

So, following these studies, our baseline measure of welfare will be an objective measure, taking into account that the signal  $\sigma$  is noisy. However, there is also some interest in calculating the welfare of the voter who has confirmation bias but is not aware of it. For example, a society that is composed entirely of voters with confirmation bias would make decisions based on this criterion.

Call these two measures *objective* and *subjective* welfare respectively. It will turn out that the only difference between the two is that the subjective welfare calculation overestimates the selection benefit from elections, as it assumes that the decision whether or not to replace the incumbent is based on perfectly accurate information.

To make the comparison as clearly as possible, we focus on the case  $\theta \equiv 0$ , where we know from Proposition 1 that the behaviour of the politician in equilibrium is the same whether the voter is a CB type or an RN type, that is he panders with probability  $\lambda = F(V(1 - q))$ . Then, in the Appendix, we establish the following formula for objective welfare of the voter:

$$W_O = \pi(1 + \delta) + \lambda(p - \pi) + (1 - \lambda)\delta\pi(1 - \pi)(2p - 1)(1 - q) \quad (4)$$

Formula (4) is composed of three parts. The first,  $\pi(1 + \delta)$ , is baseline welfare, which is the present value payoff of the voter if the incumbent did not face an election — in this case, he produces utility of 1 for the voter with probability  $\pi$  in both periods. The second term,  $\lambda(p - \pi)$ , measures the period 1 gain if the incumbent panders rather than chooses his short-run optimal action; with pandering,  $x = A$  is always chosen by the incumbent, and is the correct action with probability  $p$ , whereas if the incumbent does not pander, he chooses the correct action only if he is consonant, which occurs with probability  $\pi$ . Finally, the last term captures the selection benefit of elections.

Now consider the effect of changing  $q$  on objective welfare. First, there is a *direct effect* of  $q$ ; holding the pandering probability  $\lambda$  fixed,  $W_O$  is clearly decreasing in  $q$ . This is because confirmation bias makes selection of a good

---

<sup>18</sup>See Bernheim (2009) for a more general discussion of welfare evaluations when agents have behavioral biases.

candidate at the election more unlikely, in turn because  $\sigma$  becomes a noisier signal of incumbent type as  $q$  rises.

Second, there is also an *indirect effect* on  $W_O$  via the effect of  $q$  on  $\lambda$ . To compute the indirect effect, we use  $\lambda = F((1 - q)V)$  to get  $\partial\lambda/\partial q = -f(\hat{u})V$ , where  $\hat{u} = (1 - q)V$ , and we compute  $\partial W_O/\partial\lambda$  from (4) to get

$$\frac{\partial W_O}{\partial\lambda} \frac{\partial\lambda}{\partial q} = (\pi - p + \xi(1 - q))f(\hat{u})V \tag{5}$$

where  $\xi = \delta\pi(1 - \pi)(2p - 1) > 0$ . So, we see that there are two components to the indirect effect. An increase in  $\lambda$  reduces pandering, which causes a change  $\pi - p$  in welfare. Also, an increase in  $\lambda$  improves the probability that a bad incumbent is detected and replaced by a good incumbent at the election, that is an increase in  $\lambda$  leads to better selection. So, overall, the indirect effect of  $q$  on  $W_O$  is definitely positive if  $\pi \geq p$ .

Moving to subjective welfare, this can be expressed as:

$$W_S = \pi(1 + \delta) + \lambda(p - \pi) + (1 - \lambda)\delta\pi(1 - \pi)(2p - 1) \tag{6}$$

Not that it is identical to (4) except that the voter ignores the effect of the noisy signal on the gain from selection. So, in this case,  $q$  only affects  $W_S$  via the indirect effect on  $\lambda$ , and so the effect of  $q$  on  $W_S$  is positive if  $\pi \geq p$ .

Given this discussion, it is then straightforward to establish sufficient conditions for increased confirmation bias  $q$  to increase welfare.<sup>19</sup>

**Proposition 3.** *An increase in confirmation bias  $q$  raises objective welfare if  $\pi \geq p$  and the elasticity  $\varepsilon = -d\ln(1 - \lambda)/d\ln(1 - q) \geq 1$ . An increase in confirmation bias  $q$  raises subjective welfare if  $\pi \geq p$ ; no elasticity condition is required.*

Note from Proposition 1 that the elasticity condition is required to ensure that the positive indirect effect on selection dominates the negative direct effect. Using  $\lambda = F((1 - q)V)$ , it can be rewritten as a condition

$$\varepsilon = \frac{f((1 - q)V)(1 - q)V}{1 - F((1 - q)V)} \geq 1.$$

For example,  $F$  is uniform, that is  $F(u) = u/\bar{u}$ , this reduces to  $2(1 - q)V \geq \bar{u}$ . We have assumed  $\bar{u} > V$  but these two conditions are consistent as long as  $q \leq 0.5$ .

---

<sup>19</sup>Of course, increased confirmation bias can also reduce welfare, but our focus is on the counter-intuitive effects of bias.

#### 4 Observable Payoffs

The assumption that the voter only observes actions before voting is a strong one. It may be appropriate for choice of, for example an infrastructure project, which may be initiated but not completed before an election. It is less plausible for policies that immediately impact voters' income and well-being, such as changes in tax rates. The question then is whether payoff observability affects our results.

Assume first that the voter *only* observes his first-period payoff  $v_1 \in [0, 1]$  in our model, and not the action, before the election, as in Besley (2006).<sup>20</sup> In this case, our definition of confirmation bias can be straightforwardly extended. To do this, we note that if re-elected, the both consonant and dissonant incumbents have continuation payoffs of  $V = \delta(1 + E)$  as before, and the voter prefers to re-elect a consonant incumbent, as before. So, we now drop time subscripts without ambiguity, as we are only concerned with what happens in the first period.

Let  $v_C, v_D$  be the probabilities that types  $C$  and  $D$ , respectively, generate a payoff of 1 for the voter in the first period. Then we can modify Definition 1 as follows.

**Definition 3** (Confirmation Bias Conditional on Payoffs). *If  $v_C > v_D$ , then the voter misreads  $v = 0$  as  $v = 1$  with probability  $q > 0$ . If  $v_C < v_D$ , then the voter misreads  $v = 1$  as  $v = 0$  with probability  $q > 0$ . Moreover, when the voter updates his beliefs, he ignores these errors.*

The definition says that if the observed payoff  $v$  is a positive signal of the incumbent being consonant, that is  $v_C > v_D$ , the voter misreads a signal that the incumbent is dissonant ( $v = 0$ ) as a signal that the incumbent is consonant ( $v = 1$ ) with probability  $q$ . On the other hand, if  $v_C < v_D$ , positive and negative signals are reversed, that is  $v = 0$  is now a signal that the incumbent is consonant. Conditional on this reversal, however, the voter still misreads a signal that the incumbent is dissonant as a signal that the incumbent is consonant with probability  $q$ .

As in the observable action case, it is very helpful to think of the voter with confirmation bias observing a noisy and biased signal  $\sigma \in \{0, 1\}$  of  $v$ , defined as follows.

**Definition 4** (The Signal  $\sigma$ ). *If  $v_C > v_D$ ,  $\sigma(1) = 1$  with probability 1,  $\sigma(0) = 1$  with probability  $q$  and  $\sigma(0) = 0$  otherwise. If  $v_C < v_D$ , then  $\sigma(0) = 0$  with probability 1, and  $\sigma(1) = 0$  with probability  $q$ , and  $\sigma(1) = 1$  otherwise.*

---

<sup>20</sup>The main difference between our model with only observable payoffs and Besley's (2006) model is that the latter is slightly more parsimonious; the  $u$  for the consonant incumbent is set to 1, and also  $\theta = 0$ .

Moreover, as before, the voter with confirmation bias, having effectively observed  $\sigma$  rather than  $v$ , updates in a boundedly rational way, that is *as if* he observed  $v$ , rather than  $\sigma$ . Write  $\sigma(v, v_C, v_D)$  to emphasize that  $\sigma$  depends on  $v, v_C, v_D$ . Then, a *political equilibrium* is comprised of (i) a voter decision to elect the incumbent or challenger conditional on  $\sigma(v, v_C, v_D), \theta$ ; (ii) incumbent choices of  $x \in \{A, B\}$ , conditional on incumbent type and state of the world, which maximize incumbent payoffs given the voting rule (i); (iii) probabilities  $v_C, v_D$  that are consistent with incumbent choices. It is understood in this definition that as part of (i), the voter updates in a boundedly rational way.

Finally, in this section, we focus on the baseline case where  $\theta \equiv 0$ ; all the results in this section extend in a routine way to allowing for a non-policy preference for the voter.

When the voter *only* observes payoffs, and not actions, then the structure of equilibrium is very different. In particular, in equilibrium, the consonant incumbent always acts in his short-run interest, thus always generating a payoff of 1 for the voter and always being re-elected. The dissonant incumbent will, in equilibrium, imitate, or pool with, the consonant incumbent with some probability. In fact, Proposition 4 follows.

**Proposition 4.** *Assume that the voter only observes his payoff before the election. Then, whatever the type of the voter, there is a unique political equilibrium where (i) the voter re-elects the incumbent iff  $\sigma = 1$ ; (ii) the consonant type chooses  $x = s, s = A, B$  with probability 1, and the dissonant type imitates him with probability  $\lambda_D = F((1 - q)V)$ ; (iii) the consonant incumbent is more likely to generate a payoff of 1 for the voter i.e.  $v_C > v_D$ .*

A first implication of this proposition is that the imitation probability is decreasing in voter confirmation bias, a similar finding to the observable action case. What is the effect of confirmation bias on welfare? Using argument very similar to the derivation of (4), we can show that in this case, the objective welfare of the voter is:

$$W_O = \pi(1 + \delta) + (1 - \pi)\lambda_D + (1 - \lambda_D)\delta\pi(1 - \pi)(1 - q) \tag{7}$$

Formula (7) is composed of three parts. The first,  $\pi(1 + \delta)$ , is baseline welfare, which is the present value payoff of the voter if the incumbent did not face an election. The second term,  $(1 - \pi)\lambda_D$ , measures the first-period benefit to the voter from imitation in the event that the incumbent is dissonant. Finally, the last term captures the selection benefit of elections. Note that for a fixed imitation probability, it is decreasing in  $q$ , as confirmation bias makes selection more noisy.

Using  $\lambda = F(V(1 - q))$  to get  $\frac{\partial \lambda}{\partial q} = -f(\hat{u})V$ , where  $\hat{u} = V(1 - q)$ , and computing  $\partial W_O / \partial \lambda$  from (7), we get

$$\frac{\partial W_O}{\partial q} = -\xi(1 - \lambda_D) + (\xi(1 - q) - (1 - \pi))f(\hat{u})V \tag{8}$$

where now  $\xi = \delta\pi(1 - \pi) > 0$ . Since  $\xi(1 - q) < (1 - \pi)$ , we see that  $\partial W_O / \partial q < 0$ . This is for the following reason. First, increasing  $q$  worsens selection directly. Second, increasing  $q$  lowers the imitation probability  $\lambda_D$ , which worsens discipline but improves selection. However, as  $\delta\pi(1 - q) < 1$ , the discipline effect always dominates.

If we consider subjective welfare  $W_S$ , this is the same as  $W_O$  except without the term  $1 - q$ , so now there is no direct effect of  $q$  on welfare. However, the fact that the discipline effect dominates the welfare effect again implies that  $W_S$  is decreasing in  $q$ .

So, it appears that contrary to the observable action case, confirmation bias appears to always reduce welfare. However, our model is very stylized. In richer models of observable payoffs, imitation can be welfare-reducing and so in those settings, confirmation bias can be good for the voter. For example, Besley and Smart (2007) present such a model.

A simpler and more ad hoc case where this can happen is where there are two periods after the election, rather than one. In this case, the weight  $\delta$  on the last term in (7) becomes  $\delta(1 + \delta)$ , and then we see that after collecting terms in  $f(\hat{u})$ ,  $\partial W_O / \partial q$  can be re-written as:

$$\frac{\partial W_O}{\partial q} = (1 - \pi)(\delta\pi(1 + \delta)(1 - q) - 1)f(\hat{u})V - \xi(1 + \delta)(1 - \lambda_D)$$

Moreover, note that the value of office in the definition of  $\lambda_D$  in Proposition 4 is now  $V = \delta(1 + \delta)(1 + E)$ . Then, it is easy to find a distribution  $F(\cdot)$  and other parameters for which this derivative is positive.<sup>21</sup> So, we see that it is also possible that confirmation bias can increase welfare when only payoffs are observable.

One might ask what happens if both actions and payoffs are observed. To frame this question, suppose that voter always observes the action, as in the baseline case, but that with probability  $\phi$ , he also observes the payoff as well. Maskin and Tirole (2004) call this the case of pandering with feedback. Then, in the supplementary Appendix, we show the following.<sup>22</sup>

If  $\phi > 0.5$ , there is an equilibrium with a similar structure to Proposition 4, where the consonant incumbent always matches the action to the state and the dissonant incumbent imitates him with probability  $\lambda_A = F((1 - q)V)$  if the state is  $A$ , and  $\lambda_B = F((2\phi - 1)(1 - q)V)$  if the state is  $B$ . In this equilibrium, the voting rule depends on whether the payoff is observed or not. If not, the voter re-elects the incumbent if and only if the voter believes he observes  $x = A$ , and if the payoff is observed, voter re-elects the incumbent if and only if he believes he observes a payoff of 1, that is thinks he observes

<sup>21</sup>For example,  $F(u) = u/\bar{u}$ , and if  $\delta = 1$ , then  $1 - \lambda_D$  can be made arbitrarily small by setting  $\bar{u} = 2(1 + E) + \varepsilon$ , and then  $\partial W_O / \partial q > 0$  if  $2\pi(1 - q) > 1$ .

<sup>22</sup>A formal definition of confirmation bias when both actions and payoffs are observed is also given in the supplementary Appendix.

either  $(A, 1)$  or  $(B, 1)$ . So, in this case, the probability of pooling is decreasing in confirmation bias.

On the other hand, if  $\phi < 0.5$ , there can be an equilibrium with pandering. To get a tight characterization of equilibrium, we need to assume that  $u$  is uniformly distributed, that is  $F(u) = u/\bar{u}$ . Then, there is a  $\bar{\phi} \leq 0.5$ , such that as long as  $\phi < \bar{\phi}$ , there is an equilibrium where the dissonant and consonant types pander with probabilities

$$\lambda_D = F((1 - q)V), \quad \lambda_C = F((1 - 2\phi)(1 - q)V) \tag{9}$$

respectively.<sup>23</sup> So, when  $\phi$  is not too high, there is an equilibrium similar to the baseline case in Proposition 1, where both types pander with positive probability. Moreover, the pandering by either type is decreasing in voter confirmation bias, as before.

Note from (9) that when  $\phi > 0$ , the consonant incumbent panders less than the dissonant one. Indeed, when  $\phi \simeq 0.5$ , he hardly panders at all. This is because when payoffs are observed, the consonant incumbent can more accurately signal his type via the payoff he achieves for the voter, rather than the action he chooses.

Moreover, for  $\phi \leq \bar{\phi}$ , welfare can be increasing in  $q$ , confirmation bias. In the Supplementary Appendix, we compute a formula for objective welfare as a function of  $q$ ,  $W_O(q)$ , as in Section 3.2. Figure 1 reports some simulations

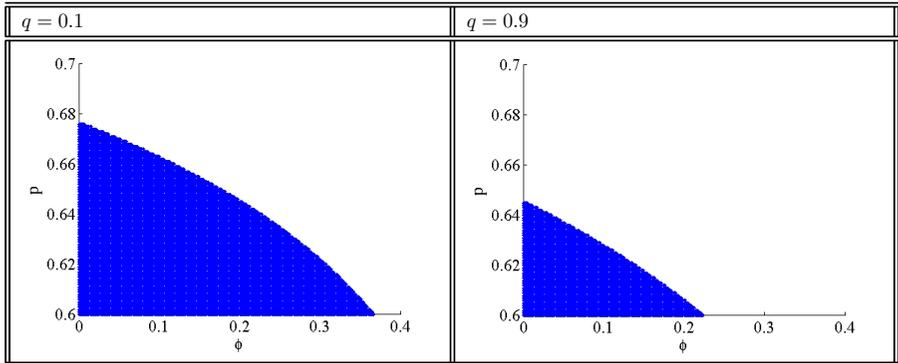


Figure 1: We assume  $F(u) = u/6$ , and  $\delta = E = 1, \pi = 0.75$ . For the parameter values chosen,  $\bar{\phi} = 0.5$ .

<sup>23</sup>It is shown in the Supplementary Appendix that

$$\bar{\phi} \equiv \min \left\{ 0.5, \frac{2p - 1}{2(1 - p)} \left( \frac{\bar{u}}{V(1 - q)} - 1 \right) \right\}.$$

using this formula. The shaded area in Figure 1 shows parameter combinations  $(\phi, p)$  for which welfare is higher with confirmation bias, that is for which  $W_O(q) > W_O(0)$ , for  $q = 0.1, 0.9$ . So, apart from  $q$ , the key parameters that we vary are  $\phi$ , the probability of observing the payoff, and  $p$ , the degree of initial bias in favour of  $x = A$ .

We can see that as expected, conditional on  $p$ , welfare is more likely to be higher with conformation bias when  $\phi$  is small. Also, conditional on  $\phi$ , welfare is more likely to be higher with conformation bias when  $p$  is small. This is because when  $p < \pi$ , the welfare change due to lower pandering is positive. Finally, welfare is more likely to be higher with conformation bias when  $q$  is small, as when  $q$  is very large, the inefficiency due to decreased accuracy of selection via elections plays an important role in decreasing welfare.

## 5 Other Models of Electoral Accountability

As mentioned in Section 1, there are several other leading models of pandering that might have been used as the vehicle for our analysis. Here, we argue that our basic insights are robust to two of these other models. We begin with the model of Fox (2007). The main difference between Fox's model and the one of this paper is that in the former, the "bad" politician no longer wishes to take the action opposite to that preferred to the voter, but *always* prefers action  $B$ , whatever the state.<sup>24</sup> Suppose for convenience that his payoff from doing so is always unity. It could be argued that this is a more realistic kind of non-congruence than in Maskin–Tirole model, because here, the "bad" politician is simply dogmatic, or stubborn, rather than always opposed to the electorate. Call this bad politician a type  $B$  (for biased, or one who always prefers alternative  $B$ ).

It is easily checked that the argument establishing Proposition 1 is virtually unchanged when the bad politician is type  $B$ , rather than type  $D$ . In fact, the only change to voter behaviour is that the policy payoff to re-electing a  $B$  type is now  $1 - p$ , as he will choose the right decision when the state of the world is  $B$ , rather than zero with a  $D$  type. So, the expected payoff to electing a politician who is judged to be good with probability  $\pi$  is  $\pi + (1 - \pi)(1 - p) = 1 - p + \pi p$ , and consequently, the re-election rule (1) becomes

$$(\pi_k(\sigma) - \pi)p \geq \theta \tag{10}$$

So, if  $\theta \equiv 0$ , the behaviour of the voter is the same as in the Maskin–Tirole model, and so Proposition 1 continues to hold.

---

<sup>24</sup>A minor difference is that the "noise" that smooths incumbent behaviour is not a random payoff from choosing the most preferred action in any period, but a random payoff from holding office. However, this just changes minor details in the algebra.

When  $\theta$  is stochastic, (10) implies that  $\Delta(q)$  is modified by multiplying each argument of  $G(\cdot)$  in (3) by  $p$ . So, Proposition 2 continues to hold, with the definition of  $\Delta(q)$  appropriately modified.<sup>25</sup>

The other leading model of pandering is that of Canes-Wrone *et al.* (2001), where the politicians differ in competence, rather than in preferences, with the good incumbent observing a perfect signal of the state, and the bad incumbent only observing a noisy signal of the state. Their model is quite rich, with a number of elements not included here; namely, voters may observe the state directly with some probability, and the ex ante quality of the incumbent may be different to that of the challenger. Here, we study a stripped-down version of their model without these elements, but retaining the features of our baseline model.<sup>26</sup>

Suppose that politicians are now all consonant (i.e. benevolent) but, instead of observing the state directly, get a signal of  $s_t$ , the state of the world. Politicians are of two types,  $H$  and  $L$ . An  $i$ -type politician gets a signal  $\zeta_t^i$ , with accuracy  $\Pr(\zeta_t^i = K | s_t = K) = \mu_i, K = A, B$ . The  $H$ -type has a more accurate signal than the  $L$ -type, that is  $1 \geq \mu_H > \mu_L > 0.5$ . In all other respects, the model is like the baseline model.

The key difference from the baseline model is that now, the continuation payoff of the incumbent depends on his type. In particular, in the second period, the incumbent of type  $i$  can do no better than to match the action to the signal, and so chooses the “right” state with probability  $\mu_i$ . Therefore, he has a continuation payoff  $V_i = \delta(\mu_i E[u_2] + E) = \delta(\mu_i + E)$ . Note that  $V_H > V_L$ . It is then straightforward to show that Proposition 1 continues to hold, but with the pandering probability being type-dependent, that is  $\lambda_H = F((1 - q)V_H), \lambda_L = F((1 - q)V_L)$ . Note that the more competent type is more likely to pander, that is  $\lambda_H > \lambda_L$ . So, in equilibrium, as in the baseline case, the probability that the  $H$ -type chooses action  $A, a_H$ , is greater than the same probability for the  $L$ -type,  $a_L$ .<sup>27</sup>

In the case where  $\theta$  is stochastic, the equilibrium is also much as in the baseline case. The voter re-election rule is still (1), but the details of the

<sup>25</sup>The only difference to the structure of equilibrium is that the  $B$ -type never chooses  $x = A$  unless he panders, so his probability of choosing is  $a_B = \lambda$ . So, the crucial condition for equilibrium,  $a_B < a_C$ , still holds.

<sup>26</sup>Even in the stripped-down version, the pandering equilibrium in the original Canes-Wrone *et al.* model has a complex structure, with both the incompetent politician and the voter randomizing. Our stochastic payoff  $u_t$  implies the existence of a pure-strategy equilibrium with a simpler structure.

<sup>27</sup>It is easy to check that the probability that the  $k$ -type chooses action  $A$  in the first period in equilibrium is:

$$a_k = \lambda_k + (1 - \lambda_k)(p\mu_k + (1 - p)(1 - \mu_k))$$

As  $\lambda_H > \lambda_L, \mu_H > \mu_L, p > 0.5$ , it must be that  $a_H > a_L$ .

computation of  $\pi_k(\sigma)$  are slightly different. However, the basic conclusion that the noise and boundedly rational updating effects move in opposite directions, with the former dominating, are qualitatively unchanged.

### 6 The Politician and the Judge Revisited

Here, we address the issue of how voter confirmation bias affects the choice between elected and appointed officials. This choice has been addressed in Maskin and Tirole (2004), as well as in several subsequent papers (e.g. Iaryczower *et al.* (2013) and Lim (2013)). However, to our knowledge, there is no existing study of how behavioral biases affect this choice.

The appropriate measure of voter welfare, as argued earlier, is objective welfare. Voter welfare from an appointed official is simply  $\pi(1 + \delta)$ , as the official will match the action to the state in both periods if and only if he is consonant. Welfare with an elected official has already been calculated in (4). So, the welfare gain to an elected official is easily seen to be

$$\Delta W = W_O - \pi(1 + \delta) = \lambda(p - \pi) + (1 - \lambda)\xi(1 - q) \tag{11}$$

where  $\xi = \delta(1 - \pi)\pi(2p - 1) > 0$ .

Note that if  $p \geq \pi$ , an elected official is always at least weakly preferred to an appointed one, whatever  $q$ , because the pandering and selection effects on welfare are both positive. So, on issues where voters have a strong prior about what is the “right” policy ( $p > \pi$ ), from (11), a politician always dominates a judge, but this advantage may increase or decrease with confirmation bias, as  $W_O$  can be increasing or decreasing in  $q$ .

The other case is where  $\pi > p$ . In this case, from (11), the effect of a change in confirmation bias on the relative advantage of an elected official is thus

$$\begin{aligned} \frac{\partial \Delta W}{\partial q} &= -(1 - \lambda)\xi + \frac{\partial W_O}{\partial \lambda} \frac{\partial \lambda}{\partial q} \\ &= (1 - \lambda)\xi(\varepsilon - 1) + (\pi - p)f(\hat{u})V \end{aligned} \tag{12}$$

where the second line follows from an argument as in Section 3.2 and  $\varepsilon = -[d \ln(1 - \lambda)]/[d \ln(1 - q)]$ . So, we see that starting from an initial position where a judge might dominate a politician, that is where  $\pi > p$ , confirmation bias always increases the relative attractiveness of a politician when  $\varepsilon > 1$ . Finally, note that if we use subjective welfare as a criterion, the direct effect  $-(1 - \lambda)\xi$  in (12) no longer applies, and so confirmation bias always increases the relative attractiveness of a politician unconditionally.

So, we can summarize as follows.

**Proposition 5.** *If voter welfare is measured by the objective criterion, then if  $\varepsilon = -d\ln(1 - \lambda)/d\ln(1 - q) \geq 1$ , the higher is voter confirmation bias, the more likely it is that an elected official is preferred to an appointed one. If voter welfare is measured by the subjective criterion, then it is unconditionally true that the higher is voter confirmation bias, the more likely it is that an elected official is preferred to an appointed one.*

## 7 Conclusions

This paper considers the implications of voter confirmation bias in a political economy setting. In a baseline model based on Maskin and Tirole (2004), we show that voter confirmation bias reduces pandering, as it lowers the electoral reward for this behaviour by reducing the increase in the probability of being elected from pandering. This result is driven by the fact that the noise aspect of confirmation bias, which decreases pandering, dominates the bounded rationality aspect. Moreover, as pandering generally has an ambiguous effect on voter welfare, it is possible that an increase in confirmation bias increases voter welfare.

These baseline results are robust to a number of extensions and changes to the model, for example when the voter observes his payoff from the election with some probability, or where politicians vary in competence. Finally, we show that voter confirmation bias strengthens the case for the case for decision-making by an elected rather than an appointed official.

Of course, it is possible that politicians, in addition to voters, might suffer from confirmation bias. There are well-known examples of political leaders ignoring the evidence that policies are not working, when they have strong prior beliefs in the efficacy of such policies (Canes-Wrone and Shotts, 2007; Majumdar and Mukand, 2004). As Mukand (2004) remarks, “a striking aspect of the history of policy-making is the apparent unwillingness of leaders to learn from previous experiments. Political leaders are typically reluctant to change course midway, even if the policy is publicly perceived to be failing”.<sup>28</sup> Study of the implications of politician confirmation bias, particularly the interesting case where politicians vary in this kind of bias and the voters try to infer the bias from the incumbent behaviour, is certainly a topic for future work.

---

<sup>28</sup>A prominent example is Margaret Thatcher insistence that a poll tax would be a better method of financing local government in the United Kingdom than a property tax, the face of all the evidence against.

## Appendix

### A.1 Computation of Posteriors

We calculate the posterior probabilities  $\pi_{CB}, \pi_{RN}$  that the incumbent is consonant conditional on  $\sigma$  for the CB and RN voters. The CB voter updates assuming that  $\sigma = x$ . Given that the consonant and dissonant politicians choose  $x = A$  with probabilities  $a_C, a_D$  respectively, Bayes' rule gives

$$\begin{aligned} \pi_{CB}(A) &= \frac{a_C \pi}{a_C \pi + a_D(1 - \pi)} \\ \pi_{CB}(B) &= \frac{(1 - a_C)\pi}{(1 - a_C)\pi + (1 - a_D)(1 - \pi)} \end{aligned} \tag{A1}$$

On the other hand, the RN voter will update taking into account that  $\sigma$  and  $x$  are related as in Definition 2. Then we have:

$$\begin{aligned} \pi_{RN}(A) &= \frac{(a_C + (1 - a_C)q)\pi}{(a_C + (1 - a_C)q)\pi + (a_D + (1 - a_D)q)(1 - \pi)}, \quad a_C > a_D \\ \pi_{RN}(B) &= \frac{(1 - a_C)(1 - q)\pi}{(1 - a_C)(1 - q)\pi + (1 - a_D)(1 - q)(1 - \pi)}, \quad a_C > a_D \\ \pi_{RN}(A) &= \frac{a_C(1 - q)\pi}{a_C(1 - q)\pi + a_D(1 - q)(1 - \pi)}, \quad a_C < a_D \\ \pi_{RN}(B) &= \frac{(1 - a_C + a_Cq)\pi}{(1 - a_C + a_Cq)\pi + (1 - a_D + a_Dq)(1 - \pi)}, \quad a_C < a_D \end{aligned} \tag{A2}$$

### A.2 Proof of Proposition 1

- (a) Assume for the moment that  $0 < a_D < a_C < 1$ . Then, using  $a_D < a_C$ , it is clear from (A1), (A2) that

$$\pi_k(A) > \pi > \pi_k(B), \quad k = RN, CB \tag{A3}$$

That is, the incumbent is judged to be of higher quality than the challenger if he chooses  $A$ , and of lower quality if he chooses  $B$ . So, from (A3), noting  $\theta = 0$ , we see that the voter of type  $k$  re-elects the incumbent iff he observes  $\sigma = A$ , as claimed.

- (b) As explained in Section 3.1, the incumbent of either type will pander iff  $u \leq \Delta r_k V$ . Given the voter re-election rule, it is clear that for both types  $k = RN, CB$ ,  $r_k(A) = 1, r_k(B) = q$ , implying  $\Delta r_k = r_k(A) - r_k(B) = 1 - q$ . So, pandering occurs iff  $u \leq (1 - q)V$ , implying a pandering probability of  $\lambda_k = F((1 - q)V)$ ,  $k = RN, CB$ , as required.

- (c) To complete the construction of equilibrium, we verify that  $1 > a_C > a_D > 0$ . Given that the probability of pandering is  $\lambda$  for both incumbent types, we see that the probabilities that the consonant and dissonant types choose  $x = A$  are

$$a_C = \lambda + (1 - \lambda)p, \quad a_D = \lambda + (1 - \lambda)(1 - p) \tag{A4}$$

Moreover, by the assumption that  $\bar{u} > V, \bar{u} > \Delta rV$  and consequently  $\lambda < 1$ . From A4, we see that  $1 > a_C > a_D > 0$ , confirming our initial assumption.

- (d) Here we prove uniqueness of the equilibrium that we have constructed.
  - (i) Assume that the voter only re-elects the incumbent if he observe  $A$ , but that where one (or both) incumbent types does *not* follow a cutoff rule. But then (say) the  $C$ -type will pander when  $u = u'$ , but not when  $u = u''$ , for some  $u' > u''$ . But then the gain to pandering when  $u = u''$  is  $V(1 - q) - u''$ , which is greater than  $V(1 - q) - u'$ , which is the gain to pandering when  $u = u'$ , a contradiction.
  - (ii) The second possibility is that voter re-elects iff he observes  $B$ . But then an argument as in Section 3.1 shows that it is optimal for the incumbent to pander, that is always choose  $B$  whenever  $u \leq (1 - q)V$ . But then, the probabilities that  $C, D$  choose  $B$  are  $b_C = \lambda + (1 - \lambda)(1 - p), b_D = \lambda + (1 - \lambda)p$ , where  $\lambda = F((1 - q)V) < 1$ . Note that  $b_D > b_C$ , so Bayesian updating implies  $\pi_k(B) < \pi < \pi_k(A), k = CB, RN$ . But then, the voter will not re-elect the incumbent if he observes  $B$ , a contradiction.
  - (iii) As the voter does not randomize, the third and fourth possibilities are that the voter always or never re-elects the incumbent, whatever  $\sigma$ . But in this case, both types will choose their short-run optimal actions, whatever  $u$ , so that  $\pi_k(B) < \pi < \pi_k(A), k = CB, RN$ . So, neither voting strategy can be sequentially rational for the voter.  $\square$

### A.3 Proof of Proposition 2

Assume for the moment that  $0 < a_D < a_C < 1$ . In this case, the re-election probability  $r$  conditional on a type  $k = RN, CB$  incumbent and  $\sigma$  is

$$\hat{r}_k(\sigma) = G(\pi_k(\sigma) - \pi) \tag{A5}$$

where  $G$  is the c.d.f. of  $\theta$ . So, the probability of being re-elected from choosing  $x = A, B, r_k(x)$ , is as follows. As  $a_D < a_C$ , from Definition 1, the voter will interpret  $x = B$  as  $x = A$  with probability  $q$ , so

$$r_k(A) = \hat{r}_k(A), \quad r_k(B) = \hat{r}_k(B)(1 - q) + \hat{r}_k(A)q, \quad k = RN, CB \tag{A6}$$

So, the increase in the re-election probability from choosing  $A$  over  $B$ , that is pandering is

$$\Delta r_k \equiv r_k(A) - r_k(B) = (1 - q)(\hat{r}_k(A) - \hat{r}_k(B)), \quad k = \text{RN, CB} \quad (\text{A7})$$

Defining  $\hat{r}_k(A) - \hat{r}_k(B) \equiv \Delta_k(q)$ , we see from (A5, A7) that for the voter with confirmation bias

$$\begin{aligned} \Delta r_{\text{CB}} &= (1 - q) [G(\pi_{\text{CB}}(A) - \pi) - G(\pi_{\text{CB}}(B) - \pi)] \\ &= (1 - q) \left[ G \left( \frac{\pi(1 - \pi)(1 - h(0))}{\pi + h(0)(1 - \pi)} \right) - G \left( \frac{\pi(1 - \pi)(1 - \tilde{h}(0))}{\pi + h(0)(1 - \pi)} \right) \right] \\ &= (1 - q)\Delta(0) \end{aligned} \quad (\text{A8})$$

where in the second step, we have used (A1) and in the third, the definition of  $\Delta(q)$  from (3). In the same way, for the rational voter with noise, we see after some simplification, using (A2), (A7) that

$$\begin{aligned} \Delta r_{\text{RN}} &= (1 - q) [G(\pi_{\text{RN}}(A) - \pi) - G(\pi_{\text{RN}}(B) - \pi)] \\ &= G \left( \frac{\pi(1 - \pi)(1 - h(q))}{\pi + h(q)(1 - \pi)} \right) - G \left( \frac{\pi(1 - \pi)(1 - \tilde{h}(0))}{\pi + \tilde{h}(0)(1 - \pi)} \right) \\ &= (1 - q)\Delta(q) \end{aligned} \quad (\text{A9})$$

as required. Then, from (A8), (A9) and  $\lambda_k = F(\Delta r_k V)$ , we get  $\lambda_{\text{CB}} = F((1 - q)\Delta(0)V)$ ,  $\lambda_{\text{RN}} = F((1 - q)\Delta(q)V)$  as required. Also, we can verify that  $a_C > a_D$  exactly as in the proof of Proposition 1. Finally, proof of uniqueness is as in Proposition 1.  $\square$

#### A.4 Derivation of Welfare Formula (4)

Suppose first the incumbent panders. As  $\pi > 0.5$ ,  $x = A$  implies  $\sigma = A$ , and so the incumbent certainly retained, and so welfare is  $p + \delta\pi$ . So, welfare conditional on pandering is

$$p + \delta\pi = \pi(1 + \delta) + (p - \pi) \quad (\text{A10})$$

If the incumbent does not pander, voter welfare is computed as follows. First, the objective probability of re-electing the incumbent is the probability that  $\sigma = A$ . This will be  $r(\pi) = r + (1 - r)q$ , where  $r = p\pi + (1 - \pi)(1 - p)$  is the probability of retaining the incumbent if he does not pander and the signal  $\sigma$

is perfectly accurate. Now by Bayes' rule,  $\hat{\pi}$ , the posterior probability that the incumbent is good, given that he does not pander and is re-elected, is

$$\hat{\pi} = \frac{(p + (1 - p)q)\pi}{r(\pi)} \tag{A11}$$

So, we can write welfare with no pandering as;

$$\pi + \delta(r(\pi)\hat{\pi} + (1 - r(\pi))\pi) = \pi(1 + \delta) + \delta r(\pi)(\hat{\pi} - \pi) \tag{A12}$$

Overall welfare is  $\lambda$  times the pandering payoff plus  $1 - \lambda$  times the non-pandering payoff, that is from (A10), (A12):

$$W_O = \lambda(p + \delta\pi) + (1 - \lambda)(\pi + \delta(r(\pi)\hat{\pi} + (1 - r(\pi))\pi)) \tag{A13}$$

Finally, after some computation, using (A11), (A12), it can be shown that

$$r(\pi)(\hat{\pi} - \pi) = \pi(1 - \pi)(2p - 1)(1 - q) \tag{A14}$$

Plugging (A14) into (A13) and simplifying, we get formula (4) as required. □

### A.5 Proof of Proposition 3

This result in the subjective welfare case follows directly from the fact that the indirect effect of  $q$  is positive iff  $\pi \geq p$  from (5). In the objective welfare case, differentiating (4), we get

$$\begin{aligned} \frac{\partial W_O}{\partial q} &= -(1 - \lambda)\delta\pi(1 - \pi)(2p - 1) + \frac{\partial W_O}{\partial \lambda} \frac{\partial \lambda}{\partial q} \\ &= -(1 - \lambda)\xi + (\pi - p + \xi(1 - q))f(\hat{u})V \end{aligned} \tag{A15}$$

where in the second line, we have used (5). So, for  $\partial W_O / \partial q \geq 0$ , if  $\pi \geq p$ , from (A15), we also need

$$(1 - q)f(\hat{u})V \geq 1 - \lambda \tag{A16}$$

But note from  $\lambda = F((1 - q)V)$  that

$$\frac{d(1 - \lambda)}{d(1 - q)} = -\frac{d\lambda}{d(1 - q)} = -f(\hat{u})V \tag{A17}$$

Then, combining (A15), (A17) gives the result. □

### A.6 Proof of Proposition 4

- (i) Assume for the moment that  $v_C > v_D$ . Then it is easy to verify that  $\pi_{CB}(1) > \pi > \pi_{CB}(0)$ , and so the voter will re-elect the incumbent if  $v = 1$ .
- (ii) Now, consider the behaviour of the incumbent. First, given this voting rule, then it is clear that the consonant incumbent does best by setting  $x = s$ . Next, note that whatever  $s$ , if the dissonant incumbent imitates the consonant incumbent, he gets  $V$ , and if he takes his short-run optimal action he gets  $u + qV$ . So, the dissonant incumbent will imitate the consonant one iff  $u \leq V(1 - q)$ , so the probability of imitation is  $\lambda = F(V(1 - q))$ , as required.
- (iii) The next step is to confirm that  $v_C > v_D$ . But given equilibrium behaviour of the incumbents,  $v_C = 1$ ,  $v_D = \lambda < 1$  as required.
- (iv) To show uniqueness, we can apply the same argument as in the proof of Proposition 2; this part of the proof is omitted.  $\square$

## References

- Alesina, A. and F. Passarelli. 2015. "Loss Aversion in Politics". NBER Working Paper 21077.
- Ashworth, S. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work". *Annual Review of Political Science* 15: 183–201.
- Ashworth, S. and E. Bueno de Mesquita. 2014. "Is Voter Competence Good for Voters?: Information, Rationality, and Democratic Performance". *American Political Science Review* 108: 565–87.
- Bagues, M. and B. Esteve-Volart. 2016. "Politicians' Luck of the Draw: Evidence from the Spanish Christmas Lottery". *Journal of Political Economy* 124: 1269–94.
- Bernheim, B. D. 2009. "Behavioral Welfare Economics". *Journal of the European Economic Association* 7: 267–319.
- Bernheim, B. D. and A. Rangel. 2004. "Addiction and Cue-triggered Decision Processes". *The American Economic Review* 94: 1558–90.
- Besley, T. 2006. *Principled Agents?: The Political Economy of Good Government*. Oxford: Oxford University Press.
- Besley, T. and M. Smart. 2007. "Fiscal Restraints and Voter Welfare". *Journal of Public Economics* 91: 755–73.
- Callander, S. and C. H. Wilson. 2006. "Context-Dependent Voting". *Quarterly Journal of Political Science* 1: 227–54.

- Callander, S. and C. H. Wilson. 2008. "Context-Dependent Voting and Political Ambiguity". *Journal of Public Economics* 92: 565–81.
- Canes-Wrone, B., M. C. Herron, and K. W. Shotts. 2001. "Leadership and Pandering: A Theory of Executive Policymaking". *American Journal of Political Science* 45: 532–50.
- Canes-Wrone, B. and K. W. Shotts. 2007. "When Do Elections Encourage Ideological Rigidity?" *American Political Science Review* 101: 273–88.
- Doll, B. B., K. E. Hutchison, and M. J. Frank. 2011. "Dopaminergic Genes Predict Individual Differences in Susceptibility To Confirmation Bias". *The Journal of Neuroscience* 31: 6188–98.
- Ellis, A. 2016. "Condorcet Meets Ellsberg". *Theoretical Economics* 11: 865–95.
- Fox, J. 2007. "Government Transparency and Policymaking". *Public Choice* 131: 23–44.
- Ghirardato, P. and J. N. Katz. 2006. "Indecision Theory: Weight of Evidence and Voting Behavior". *Journal of Public Economic Theory* 8: 379–99.
- Gibson, R. K. and A. Römmele. 2009. "Measuring the Professionalization of Political Campaigning". *Party Politics* 15: 265–93.
- Iaryczower, M., G. Lewis, and M. Shum. 2013. "To Elect or to Appoint? Bias, Information, and Responsiveness of Bureaucrats and Politicians". *Journal of Public Economics* 97: 230–44.
- Levy, G. and R. Razin. 2015. "Correlation Neglect, Voting Behaviour and Information Aggregation". *American Economic Review* 105: 1634–45.
- Liberini, F., M. Redoano, and E. Proto. "Happy Voters". *Journal of Public Economics*, forthcoming.
- Lim, C. S. 2013. "Preferences and Incentives of Appointed and Elected Public Officials: Evidence From State Trial Court Judges". *The American Economic Review* 103: 1360–97.
- Lockwood, B. and J. Rockey. 2016. "Negative Voters? Electoral Competition with Loss-Aversion". Unpublished paper, University of Warwick.
- Lord, C. G., L. Ross, and M. R. Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence". *Journal of Personality and Social Psychology* 37: 2098–109.
- Majumdar, S. and S. W. Mukand. 2004. "Policy Gambles". *The American Economic Review* 94: 1207–22.
- Maskin, E. and J. Tirole. 2004. "The Politician and the Judge: Accountability in Government". *American Economic Review* 94: 1034–54.
- Matějka, F. and G. Tabellini. 2016. "Electoral Competition with Rationally Inattentive Voters". Unpublished paper.
- Morelli, M. and R. Van Weelden. 2013. "Ideology and Information in Policy-making". *Journal of Theoretical Politics* 25: 412–39.
- Nickerson, R. S. 1998. "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises". *Review of General Psychology* 2: 175–220.

- O'Donoghue, T. and M. Rabin. 1999. "Doing it Now or Later". *American Economic Review* 89: 103–24.
- Ortoleva, P. and E. Snowberg. 2015. "Overconfidence in Political Behavior". *The American Economic Review* 105: 504–35.
- Plous, S. 1991. "Biases in the Assimilation of Technological Breakdowns: Do Accidents Make Us Safer?" *Journal of Applied Social Psychology* 21: 1058–82.
- Prat, A. 2005. "The Wrong Kind of Transparency". *The American Economic Review* 95: 862–77.
- Rabin, M. and J. L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias". *Quarterly Journal of Economics* February: 37–82.