

Voter Confirmation Bias and Electoral Accountability*

Ben Lockwood[†]
University of Warwick

First version: 8 February 2015

This version: 7 June 2015

Abstract

This paper considers the implications of an important cognitive bias in information processing, confirmation bias, in a political agency setting. In the baseline two-period case where only the politician's actions are observable before the election, we show that when voters have this bias, it decreases pandering by the incumbent, and can raise voter welfare as a consequence. This result is robust in several directions, including to the case where the voter can also observe payoffs with some probability before the election (Maskin and Tirole's "feedback" case). In the three-period case, with two elections, the dynamic evolution of confirmation bias can lead to more pandering before the first election. Finally, we show that when confirmation bias is present, other things equal, the case for decision-making by an elected rather than an appointed official is greater.

KEYWORDS: confirmation bias, selective exposure, voting, pandering, elections

JEL CLASSIFICATION: D72, D83

Address for correspondence; Department of Economics, Warwick University, Coventry, CV4 7 AL, United Kingdom. E-mail B.Lockwood@warwick.ac.uk.

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

[†]Address for correspondence: Department of Economics, University of Warwick, Coventry CV4 7AL, UK. Email: b.lockwood@warwick.ac.uk

1 Introduction

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 Rabin 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 bias¹; 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, Hutchison, and Frank (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. Biased processing of exogenously presented information, is discussed by Rabin and Schrag (1999) amongst others². Notable examples include experiments where subjects were initially questioned in a salient policy issue (Lord, Ross, and Lepper (1979), capital punishment, Plous (1991), 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. In the classic experimental selective-exposure research paradigm, participants work on a binary decision problem and come to a preliminary conclusion (such as choosing one of two investment strategies). Participants are then given the opportunity to search for additional information, which is typically received in the form of short statements indicating the perspectives of newspaper articles, experts, or former participants. In the design, half of available statements will be positive, and half will be negative about each choice, and the participants are asked to indicate those pieces that they would like to read in more detail later on. 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.

There is also non-experimental evidence, mainly in the context of political campaigns, showing that voters are more likely to access media outlets that confirm their prior beliefs, (for example, Chaffee and

¹So salient is confirmation bias that it has been noted long before modern psychology came into being: "*The human understanding when it has once adopted an opinion draws all things else to support and agree with it. And though there be a greater number and weight of instances to be found on the other side, yet these it either neglects and despises, or else by some distinction sets aside and rejects, in order that by this great and pernicious predetermination the authority of its former conclusion may remain inviolate.*". Francis Bacon, 1620, quoted in Rabin and Schrag (1999).

²Rabin and Schrag do not discuss, or model, selective exposure.

Miyo (1983), Stroud (2010), Iyengar and Hahn (2009), Jerit and Barabas (2012)), or talk to friends who share one's political views (Huckfeldt and Sprague (1988)).

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 confirmation bias, because as professionals, with access to expert advice, politicians and bureaucrats are perhaps less likely to suffer from this bias. We focus on the first form of bias, preferential treatment of given evidence; selective exposure requires also the modelling of the supply of information by the media and government, and is beyond the scope of this paper.

In this paper, we introduce voter confirmation bias into a fairly general political agency model. 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 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). To model confirmatory 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. 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.

In the observable action case, the signal of incumbent quality observed by the voter is the binary action taken by the incumbent. So, in this setting, confirmation bias means that the voter "mis-reads" the action with some probability³. Confirmation bias is modeled in a similar way in the observable payoff case; that is, confirmation bias means that the voter "mis-reads" the payoff with some probability.

The key feature of the Maskin-Tirole model is that it can explain political pandering i.e. choice of policy to follow popular opinion even when this conflicts with what the benevolent politician knows is best. Our baseline finding is that in the observable action case, voter confirmation bias *reduces* pandering, as it lowers the electoral "reward" for this behavior by reducing the increase in the probability of being elected from pandering. As pandering generally has an ambiguous effect on voter welfare, it is possible that an increase in confirmation bias (parametrized by the probability of misreading the signal) increases voter welfare.

There are two kinds of intuition for this result. The first is a general second-best one. That is, there is initially a distortion present, pandering behavior, so introducing another distortion, confirmation bias, need not make the voter worse off. The second is more specific; as argued by Prat (2005), when voters can observe the actions of the politicians, this is the "wrong" kind of transparency; when the voters (the principal) can observe the actions, this can induce pooling or pandering by the agent (the voter). So,

³Note that in our model, the relationship between incumbent quality and the signal is thus determined endogenously in political equilibrium, as opposed to Rabin and Schrag(1999), where the relationship is exogenously specified.

from this general intuition, it is reasonable that confirmation bias, being formally a kind of "garbling" of the action signal, can improve the welfare of the principal. However, confirmation bias is not equivalent to just garbling the signal of the type, because the voter *does not understand that he is making these errors*. Specifically, to capture this bounded rationality, following Rabin and Schrag (1999), we assume that the voter ignores his own bias when performing Bayesian updating.

We then consider the case where voters observe their payoffs before the election, as in Besley (2006). Here, we show that the equilibrium structure is completely different; the good incumbent behaves non-strategically, and the bad incumbent imitates ("pools" with) the good incumbent with some probability; this probability is decreasing in confirmation bias. In this case, confirmation bias is no longer welfare-improving.

We also consider the robustness of our baseline results by allowing the voter to observe the action of the incumbent, *and* also the payoff, with some probability (Maskin and Tirole (2004) call this the "feedback" case). If this probability is less than one-half, there continues to be a pandering equilibrium, and again, the amount of pandering is decreasing in confirmation bias.

In another extension, we consider the dynamic evolution of confirmation bias in a three-period model with two elections. At the beginning of the first period, the voter is initially neutral i.e. believes the incumbent is equally likely to be good or bad. If the incumbent is retained, the voter must then revise his beliefs upwards, inducing confirmation bias. In this case, it can be shown that confirmation bias increases the probability of pandering in the first period. Nevertheless, voter welfare can still increase.

Finally, we revisit the choice between a politician and an unelected official, the focus of Maskin and Tirole's original paper. For some parameter values, the politician always dominates the bureaucrat, in the sense that voter welfare is higher with the former. When parameter values are such that this is *not* the case, i.e. when the choice between an elected and appointed official is not trivial, confirmation bias always works in favour of choosing the elected official; this is because bias reduces pandering. So, in policy areas where 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, technical decisions, such as those concerning monetary policy or utility regulation, are usually taken by appointed officials.

The remainder of the paper is structured as follows. Section 2 discusses related literature, Section 3 describes the model, and Section 4 derives baseline results. Sections 5 and 6 consider the case of observable payoffs and multiple elections respectively. Finally, Section 7 revisits the choice between a politician and an unelected official, the focus of Maskin and Tirole's original paper, in the light of previous results.

2 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. 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"⁴. This refers to the stylized fact that voters vote for or against the incumbent partly in response to "events (like natural disasters, economic shocks, or football losses) that the voters should know are outside of the politicians' control". They capture this by the assumption that "the voter's welfare.... is now a function of observable events not affected by the politician in office", specifically modelled by a random shock to voter payoffs. Failing to filter can make the voter better off⁵.

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. Experimental evidence suggests that failure to filter is probably driven by an affective, rather than cognitive, process, namely a well-being spillover, where a shock that increases well-being makes the voter better disposed to the incumbent. For example, Healy, Malhotra, and Mo (2010) find that when respondents are "treated" by being told the outcome of an important football match for a team that they supported before being asked for their rating of President Obama's job performance, each additional adjusted win experienced by respondents significantly increased approval, with the effect size being 2.3 percentage points. They conclude, on the basis of these results, plus other evidence, that "changes in well-being induced by the surprise component of sporting events affect people's evaluations of the incumbent".

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 interests of the voters. 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 i.e. pool with, the good incumbent or not. In turn, it is argued in the political science literature that pandering is in practice, one of the most important forms of political inefficiency (Canes-Wrone, Herron, and Shotts (2001)).

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 incentives for the bad incumbent (the extremist in their model) to imitate the good one (the moderate) are stronger under a fail-to-filter voter than under a rational voter, and so failing to filter buys the voter better discipline of

⁴They also show that if the voter can choose a re-election rule that is more stringent than the sequentially rational one, voter welfare can be improved. However, there is no particular support for such a behavioural bias in the psychology literature.

⁵Spiegler (2013) studies a related phenomenon in a political economy setting, where salient, recent events are intuitively perceived to be causes of an observed outcome.

the incumbent at the cost of worse selection. In contrast, as described above, we find that confirmation bias implies worse discipline but better selection. So, overall, the results of this paper are complementary to thiers.

Our paper is also close in spirit to Levy and Razin (2014), 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, due to 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 i.e. a referendum on two alternatives, and correlation neglect causes individuals base their vote more on their information rather than on their preferences⁶.

There are also a number of papers in both the economics and political science literature showing that additional information may not be to the benefit of voters, because it may induce a strategic response in the incumbent politician. As already remarked, Prat (2005) makes the point in a general agency model; 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 can make the principal worse off, as it induces the bad agent to pool with the good one, and thus worsens selection. Besley (2006) shows in a rather different setting, that allowing voters some direct information about the incumbent can reduce the discipline effect of elections to the point where the voter is worse off. Ashworth and Bueno de Mesquita (2014) make the same point as Prat(2005), but in an explicit political economy setting. However, unlike what is attempted here, these contributions assume full rationality of voters.

Finally, there is a growing literature introducing other behavioral biases to political economy models. One early contribution is Callander and Wilson (2006), (2008) who introduce a theory of context-dependent voting, where for example, for a left wing voter, the attractiveness of a left wing candidate is greater the more right wing is the opposing candidate, and apply it to the puzzle of why candidates are so frequently ambiguous in their policy. Another is Ghirardato and Katz (2006), who show that if voters are ambiguity-averse, they might strictly prefer abstaining to voting, even if voting is costless⁷. Other recent contributions include Ellis (2012), who extends the arguments of Ghirardato and Katz (2006) to investigate information aggregation in large elections, Passarelli and Tabellini (2013), who introduce features based on loss-aversion into a model of redistribution with a benevolent government, and Alesina and Passarelli (2015), who study the effect of loss aversion on majority voting.

⁶Ortoleva and Snowberg (2015), in a related paper, show theoretically that overconfidence and ideological extremeness are connected; empirically they find, using a large US election study, that overconfidence is the most reliable predictor of ideological extremeness and an important predictor of voter turnout.

⁷The paper is motivated by the empirical observation that voters who arguably face no cost of voting might still abstain, as in the case of case of multiple elections on one ballot.

3 Set-Up

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 payoff to the action depends on a state of the world $s_t \in \{A, B\}$. Prior to choosing x_t , the incumbent observes the state.

The voter gets utility v_t in period t , which is 1 if the incumbent's action in period t matches the state, and 0 otherwise. 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, as described below. Both voter and the incumbent discount payoffs by δ .

Politicians are of two types, congruent, denoted C , and non-congruent, denoted D . 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. They also have an ego-rent E from office. Dissonant politicians when in office, get u_t if $s_t \neq x_t$, 0 if $s_t = x_t$, and they also have an ego-rent E from office. Finally, we assume without loss of generality that $E[u_t] = 1$, and we assume $\bar{u} > \delta(1 + E)$.

The reason why we assume that politicians' payoffs from their most preferred outcome are determined by random draw from a continuous distribution (rather than being fixed at 1, as in Maskin and Tirole) is to ensure 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 applied⁸.

3.1 Order of Events and Information Structure

All agents i.e. incumbent, voter, and challenger have a prior $p > 0.5$ that state A will occur. The incumbent is consonant with probability π , and faces a randomly drawn challenger at the end of period 1, who is consonant with the same probability π . Initially, we assume, following Maskin and Tirole(2004), that the voter observes the action x_1 before election, but not the payoff v_1 generated by x_1 : so in equilibrium, x_1 is a signal for the voter - in fact, the *only* signal - about the politician type. In Section 5 below, we instead assume that the incumbent observes his payoff v_1 , or both v_1, x_1 , rather than x_1 ; this case is similar to Besley (2006). In Section 5, we also consider the hybrid case where both v_1, x_1 are observed. Whether payoffs or actions are observed makes a big difference to the effect of confirmation bias.

⁸A 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.

3.2 Modelling Confirmation Bias

We follow Rabin and Schrag in modelling confirmation bias. 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)

In our setting, the voter cares about the incumbent type, so voter confirmation bias can be modelled as follows. Say that the voter is *optimistic* (resp. *pessimistic*) about the quality of the incumbent if $\pi > 0.5$ (resp. $\pi < 0.5$). In our simple setting, if the agent is optimistic, the agent misreads $x_t = B$ as $x_t = A$ with probability q . If the agent is pessimistic, the agent misreads $x_t = A$ as $x_t = B$ with probability q . Note that a key difference between Rabin and Schrag (1999)'s set-up and this one is that in theirs, the signal is exogenously generated, whereas in our set-up, voter signals are actions generated by equilibrium -play of the game between incumbent and voter⁹.

Let the action that the voter *thinks* he observes be \tilde{x}_t ; following Rabin and Schrag, we assume that the voter performs ordinary Bayesian updating, conditional on \tilde{x}_t being the "true" action - i.e. he does *not* adjust for the fact that \tilde{x}_t is a noisy signal of x_t . Finally, we assume that when the voter is neutral i.e. $\pi = 0.5$, he has no confirmation bias i.e. $q = 0$.

Finally, when it comes to politician behavior, 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)).

3.3 Equilibrium Selection

We focus on pure-strategy perfect Bayesian equilibrium in what follows. As we show below, any equilibrium in pure strategies must have the cutoff property that the incumbent panders in the first period (i.e. chooses an action that is not short-run optimal) if and only if the return to the short-run action, u_1 , is below some cutoff \hat{u} . However, as remarked by Maskin and Tirole(2004), without any further restrictions on beliefs, there are two types of pandering equilibrium. The first is characterized in Proposition 1 below, where the voter re-elects the incumbent only if he thinks he observes action A . The second is a "perverse" pandering equilibrium, where the voter re-elects the incumbent only if he thinks he observes action

⁹The formal definition of Rabin and Schrag (1999) is the following. Assume a binary state of the world, $s = \{A, B\}$, and sequence $t = 1, ..T$ of informative signals $\sigma_t \in \{A, B\}$ about the state, where $\Pr(\sigma_t = K | s = K) = \theta > 0.5$. If π_t is the decision-maker's prior 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 .

B ¹⁰. Maskin and Tirole argue that this equilibrium is implausible, and show that in their set-up, where $u_t = 1$ with probability 1, it is not robust to the introduction of a small probability that the incumbent is a "mechanical" type who always matches.

We follow them by focussing on the first type of pandering equilibrium¹¹. It is convenient to do so by imposing a monotonicity assumption on beliefs that the probability that the incumbent is congruent, conditional on $x_1 = A$ is at least as great as the probability that the incumbent is congruent, conditional on $x_1 = B$. To avoid repeated statement of these qualifications, we just refer to any pure-strategy equilibrium that is consistent with the monotonicity assumption on voter beliefs just as a *political equilibrium*.

4 The Baseline Case

4.1 Political Equilibrium

In the second period, consonant (dissonant) politicians match the action to the state according to their preferences i.e. 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 $\delta(E[u] + E) = \delta(1 + E) \equiv V$. Moreover, clearly, the voter prefers to re-elect the incumbent if and only if he is consonant. So, in what follows, we focus on the first period, and so we can drop time subscripts without ambiguity.

We will show that there is a unique political equilibrium, which has the following structure. Let the policy that the voter *thinks* he observes be denoted \tilde{x} . First, voters re-elect the incumbent if and only if $\tilde{x} = A$. Second, both incumbent types follow a cutoff rule i.e. a politician of type $i = C, D$ "panders" if and only if his payoff u in period 1 from his short-run optimal action is less than some critical value. Here, pandering means choosing action A even when the short-run optimal action is B .

Assume first that $\pi > 0.5$, so that the voter is an optimist. We consider first the consonant type. A conflict between short-run and long-run payoffs arises when $s = B$. Here, the payoffs are as follows. If the incumbent chooses A , he gets a short-run payoff of zero, but will be re-elected with probability 1. If he chooses B , he gets a short-run payoff of u , but will be re-elected with probability q , as the voter mis-reads $x = B$ as $x = A$ with probability q . (Recall that we are assuming that the politician understands that the voter has confirmation bias). Finally, the payoff to being re-elected is V . Then, the consonant incumbent

¹⁰To see this, note that in this equilibrium, pandering is defined as choosing B when it is short-run optimal to choose A . Then an argument as in Section 4.1 shows that it is optimal to pander whenever $u \leq \delta(1 + E)(1 - q)$. But, for this to be an equilibrium, it must be that $\Pr(i = C | x = B) \geq \pi > \Pr(i = C | x = A)$. Now, in this equilibrium, the re-election probabilities for C, D types are $r_C = \lambda + (1 - \lambda)(1 - p)$, $r_D = \lambda + (1 - \lambda)p$, and formulae (5) still hold. Then, as $r_C < r_D$, $\Pr(i = C | x = B) \geq \pi > \Pr(i = C | x = A)$ is verified.

¹¹Prat (2005) also addresses this issue, and simply rules out the perverse equilibrium as not being of interest.

prefers to pander i.e. take action $x = A$ even when $s = B$ if

$$\underbrace{V(1-q)}_{\text{period 2 gain}} \geq \underbrace{u}_{\text{period 1 loss}} \quad (1)$$

So, the cutoff for the consonant type is $\hat{u} = V(1-q)$.

Now consider the dissonant type. A conflict between short-run and long-run payoffs now arises when $s = A$. Here, the payoffs are as follows. If the incumbent chooses A , he gets a short-run payoff of zero, but will be re-elected with probability 1. If he chooses B , he gets a short-run payoff of u , but will be re-elected with probability q . Finally, his payoff to being re-elected is also V . So, the model is symmetric, in the sense that the dissonant type has the same cutoff \hat{u} .

The model also has another important symmetry: the equilibrium cutoff, \hat{u} , is the same for both incumbent types also if the voter is a pessimist. To see this, note that in this case, if the consonant incumbent chooses A , he gets a short-run payoff of zero, but will be re-elected with probability $1-q$, as the voter mis-reads $x = A$ as $x = B$ with probability q . If he chooses B , he gets a short-run payoff of u , but will be re-elected with probability 0. So, the consonant incumbent again prefers to pander when (1) holds.

So, for all values of π , we can define the probability that either type of incumbent panders as

$$\lambda = F(V(1-q)) \quad (2)$$

Note also that by the assumption $\bar{u} > V$, λ is always strictly less than 1.

To complete the description of equilibrium, we need to show that it is ex post rational for the voter to re-elect the incumbent if and only if he observes $\tilde{x} = A$. Recall that Rabin and Schrag (1999) assume that an agent with confirmation bias engages in Bayesian updating that *would be fully rational given his environment if he were not misreading evidence*. In this context, this means that the voter is willing to re-elect the incumbent if and only if $\tilde{x} = A$ whenever the following holds:

$$\Pr(i = C | x = A) \geq \pi > \Pr(i = C | x = B) \quad (3)$$

where x is the actual action of the incumbent, because the voter believes that $x = \tilde{x}$. Note that the probabilities that the consonant and dissonant type choose $x = A$ are

$$r_C = \lambda + (1-\lambda)p, \quad r_D = \lambda + (1-\lambda)(1-p) \quad (4)$$

respectively. So, from Bayes' rule, we have

$$\begin{aligned}\Pr(i = C | x = A) &= \frac{r_C \pi}{r_C \pi + r_D (1 - \pi)} \\ \Pr(i = C | x = B) &= \frac{(1 - r_C) \pi}{(1 - r_C) \pi + (1 - r_D) (1 - \pi)}\end{aligned}\tag{5}$$

But from (4) and (5), we see that $r_C > r_D$ as long as $\lambda < 1$, and so (3) holds. But, by assumption that $\bar{u} > V$, λ is always strictly less than 1. This establishes that voters will in fact behave as claimed in equilibrium. We can summarize as follows;

Proposition 1. *Whether voters are optimists or pessimists, there is a unique political equilibrium where both types pander with probability $\lambda = F(V(1 - q))$. So, pandering λ is decreasing in voter confirmation bias.*

The proof that this equilibrium is unique is in the Appendix.

4.2 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 confirmation bias or not?¹² Rabin and Schrag(1999) argue that the voter is, by definition, unaware of confirmation bias. This certainly must be true in the ex post sense; if the voter knows that he has mis-classified the signal, *after* the signal is received, an otherwise rational voter would correct the mis-classification. However, the voter may or may not be aware, before getting the signal x , that he will mis-classify it with probability q . If he is not aware of this possibility, we call the voter *naive*. However, the voter is aware ex ante that he will mis-classify the signal, we could call the voter *sophisticated*¹³, and in that case, the welfare calculation should allow for the consequences of politician selection based on noisy signals.

We begin with the naive case. Here, it is easily calculated that the naive voter expects a payoff

$$W = \lambda(p + \delta\pi) + (1 - \lambda)(\pi + \delta(r\pi_A + (1 - r)\pi))\tag{6}$$

where

$$r = p\pi + (1 - \pi)(1 - p)\tag{7}$$

is the unconditional probability that incumbent is re-elected if he does not pander, and $\pi_A = p\pi/r$ is the probability that the incumbent is consonant, conditional on him not pandering and being re-elected.

This is explained as follows. With probability λ , the incumbent panders, and thus generates expected

¹² See Bernheim and Rangel(2007) for a more general discussion of welfare evaluations when agents have behavioral biases.

¹³ This terminology is based on O'Donoghue and Rabin (1999)'s study of the decision-makers with time-inconsistent preferences.

payoff p for the voter in the first period, as he takes the right action for the voter only if the state is A . Following this, he is re-elected with probability 1, and then in the second period, will generate a payoff of 1 for the voter only if he is consonant. Note that both of these are calculated *ignoring* confirmation bias, as the voter is naive.

With probability $1 - \lambda$, the incumbent takes the short-run optimal action. In what follows, for brevity, we refer to choice of the short-run action as *matching*, because the action is matched to the state¹⁴. Clearly, matching only generates a payoff of 1 for the voter if the incumbent is consonant, and so the expected payoff to the voter from matching is just equal to the probability that the incumbent is consonant. In the first period, this is just π . Following this, the incumbent is re-elected with probability r . If he is re-elected, he retains office, and, ignoring his own confirmation bias, the voter calculates that the incumbent is consonant with probability π_A via application of Bayes' rule. If he is not re-elected, a challenger sets policy in the second period, generating payoff π .

Combining (6), (7), and π_A , after some re-arrangement, we get:

$$W = \pi(1 + \delta) + \lambda(p - \pi) - (1 - \lambda)\delta\pi(1 - \pi)(2p - 1) \quad (8)$$

This makes clear that following an increase in pandering, λ , there is welfare loss through decreased selection of consonant incumbents via the election (the term $\pi(1 - \pi)(2p - 1)$) but a change in first-period conduct of the incumbent, which Besley (2006) calls the discipline effect of the election. This discipline effect on welfare is positive and only if pandering is better than short-run optimization i.e. $p > \pi$.

Now consider the effect of an increase in confirmation bias q on (8). This works through λ , so we have;

$$\frac{\partial W}{\partial q} = -\frac{\partial W}{\partial \lambda} f(\hat{u}) = f(\hat{u}) \times \left\{ \underbrace{\delta(\pi(1 - \pi)(2p - 1))}_{\text{selection gain (+)}} - \underbrace{(p - \pi)}_{\text{discipline loss (?)}} \right. \quad (9)$$

That is, increased confirmation bias reduces pandering, and this always increases the quality of selection of politicians, as the dissonant incumbent is more likely to be fired. The effect on discipline is ambiguous; if $p < \pi$, increased confirmation bias can actually improve discipline. In fact, it is clear from (9) that $\frac{\partial W}{\partial q} \geq 0$ when

$$\delta \geq \hat{\delta} \equiv \frac{(p - \pi)}{\pi(1 - \pi)(2p - 1)}$$

We can summarize:

Proposition 2. *Assume only incumbent actions are observable before the election. Then, an increase in confirmation bias q always makes the voter weakly better off if $\delta \geq \hat{\delta}$, and worse off otherwise. In particular, if pandering worsens discipline, i.e. $p \leq \pi$, an increase in confirmation bias q always makes*

¹⁴Of course, in the case of a dissonant incumbent, the matching is the reverse to that valued by the voter.

the voter strictly better off.

Now consider the case where the voter is sophisticated. If the incumbent panders, welfare is as before. Note that when $\pi < 0.5$, there is a probability q that $x = A$ is misread as $x = B$, and the incumbent is fired; however, the replacement will have the same expected quality and thus generate the same second-period payoff of π for the voter as retaining the incumbent. So, welfare conditional on pandering is still $p + \delta\pi$.

If the incumbent matches, welfare is

$$W_M = \begin{cases} \pi + \delta[r^+\pi_A^+ + (1 - r^+)\pi], & \pi > 0.5 \\ \pi + \delta[r^-\pi_A + (1 - r^-)\pi] & \pi < 0.5 \end{cases} \quad (10)$$

Here,

$$r^+ = r + (1 - r)q, \quad r^- = r(1 - q)$$

are the probabilities of retaining the incumbent, taking into account positive or negative confirmation bias. With positive confirmation bias, the voter rationally anticipates that with probability q , the incumbent will be retained even if the "true" signal is $x = B$. Hence, conditional on retention, by Bayes' rule, the probability that the incumbent is consonant is

$$\pi_A^+ = \frac{(p + (1 - p)q)\pi}{r^+}$$

With negative confirmation bias, if the incumbent is retained, he has surely chosen $x = A$, and so π_A replaces π_A^+ . It is easily checked that both formulae in (10) reduce to

$$W_M = \pi + \delta[\pi + r(1 - q)(\pi_A - \pi)] \quad (11)$$

So, then overall welfare is

$$W = \lambda(p + \delta\pi) + (1 - \lambda)W_M = \pi(1 + \delta) + \lambda(p - \pi) + (1 - \lambda)\delta\pi(1 - \pi)(2p - 1)(1 - q) \quad (12)$$

Differentiating (12), and recalling that λ depends on q via (2), we see that

$$\frac{\partial W}{\partial q} = f(\hat{u}) [\delta(\pi(1 - \pi)(2p - 1) - (p - \pi))] - (1 - \lambda)\delta\pi(1 - \pi)(2p - 1) \quad (13)$$

Comparing (13) to (9), we see that there is a new third term $(1 - \lambda)\delta\pi(1 - \pi)(2p - 1)$ which measures the efficiency loss due to less "accurate" selection via the election of the consonant type, due to the errors induced by confirmation bias. So, confirmation bias is less likely to be welfare-improving if the welfare calculation is sophisticated. Nevertheless, it can still be positive, for example if $\delta > \hat{\delta}$, and the density $f(\hat{u})$ is large enough.

4.3 Extensions

As long as the assumption that only actions are observable is retained, these results are generalizable in several directions. Here, we just summarize further results retaining this assumption; detailed proofs are available on request. So far, we have assumed two periods. As pointed out by Besley (2006), this has the very artificial feature that if the challenger is elected, he will also be a lame duck in period 2. We can relax this assumption by supposing that the voter is infinitely lived, but the politicians are term-limited. For tractability, we restrict attention to a two-term limit.

Because of the two-term limit, the incumbent faces the same decision problem as in the two-period case. But now, for a pandering equilibrium to exist, the voter must be willing to retain a "lame duck" for a second term if his choice of x in his first term of office is A , rather than replace him with a new incumbent who will pander in the first period. It is intuitive that this requires $\pi \geq p$. So, subject to this parameter restriction, Proposition 1 continues to hold. This of course, means that pandering involves a loss of discipline. In turn, this implies that confirmation bias will *always* increase welfare, a stronger result than in our model.

Second, the assumption that politicians differ in benevolence can be replaced by the assumption that they differ in competence. Specifically, we could adopt the framework of Canes-Wrone, Herron, and Shotts (2001), who suppose that all politicians are benevolent, but that they differ in competence, with better politicians receiving more accurate signals about the state of the world. In this case, if only actions are observable, a pandering equilibrium similar to the type described in Proposition 1 exists, but with a slightly more complex structure; the more able type is more likely to pander, because he has a higher continuation payoff from office. But, as in Proposition 1, an increase in q reduces pandering for both types.

Third, while the evidence here more limited, there are some findings that suggest that legislators and administrators suffer bias, even though they are "professional" decision-makers (Moynihan and Lavertu (2012)). Because beliefs about the state of the world are optimistic by assumption i.e. $p > 0.5$, incumbent confirmation bias can be modelled by supposing that the incumbent misreads $s_t = B$ as $s_t = A$ with probability $1 > q > 0$. To understand the effect of politician confirmation bias, we need to make a distinction between naive and sophisticated politicians, as we did with the voter. Say that the politician is *naive* if he is unaware of his own behavioral bias, and *sophisticated* if he is aware¹⁵. In the baseline model, we have implicitly assumed that politicians are sophisticated, because they understand the effect of voter confirmation bias on their re-election probabilities, so we continue with this assumption. If the politician is sophisticated, an increase in confirmation bias q lowers the incumbent's continuation payoff

¹⁵By "aware", we mean that the politician knows that in the future he will misinterpret $s_t = B$ as $s_t = A$; obviously, he cannot be aware *ex post* of a misinterpretation, for then, by definition, he would no longer suffer from confirmation bias.

V , because he rationally anticipates that he will make mistakes. So, he values re-election less, and so is less willing to pander. So, politician confirmation bias has the same qualitative effect on the equilibrium as voter confirmation bias, but via a different mechanism.

5 Observable Payoffs

The assumption that the voter only observes actions before voting is a strong one; it may be appropriate for choice of e.g. an infrastructure project, but not for policies that directly impact voters, such as changes in tax rates. Now assume that instead, the voter observes the payoff $v \in \{0,1\}$ from the policy choice before the election, as in Besley (2006). Initially, we assume that the action cannot be observed.

First, we need to specify confirmation bias in this case. We conjecture that $v = 1$ (resp. 0) is a signal that the incumbent is consonant; we will verify this in equilibrium. So, it is natural to define positive confirmation bias in this case as a mis-classification of $v = 0$ as $v = 1$ with probability q , and negative confirmation bias as a mis-classification of $v = 1$ as $v = 0$ with probability q . Let this noisy observation of v be denoted \tilde{v} .

In this case, the equilibrium has an entirely different structure. In fact, we can show:

Proposition 3. *Assume that the voter only observes \tilde{v} . Then, whether voters are optimists or pessimists, there is a unique political equilibrium where the consonant type matches i.e. chooses $x = s$, $s = A, B$, and the dissonant type imitates him with probability $\lambda_D = F(V(1 - q))$. So, imitation λ is decreasing in voter confirmation bias.*

In this case, the consonant incumbent generates welfare $1 + \delta$ for the voter, because he chooses optimally in both periods, and the dissonant one generates a payoff of 1 if he pools and is re-elected, and zero otherwise. But, in the second case, he is replaced by the challenger, who generates π . So, overall, welfare is

$$W = \pi(1 + \delta) + (1 - \pi)(\lambda_D + (1 - \lambda_D)\delta\pi) \quad (14)$$

It is clear from (14) that welfare is increasing in the pooling probability λ_D , so now increased confirmation bias q is unambiguously *bad* for welfare.

This result raises the question of whether the baseline results for the Maskin-Tirole version of the model are very special, or whether they are robust to some observability of payoffs. To investigate this, we assume that the voter always observes the action, as in the baseline case, but that with probability f , he also observes the payoff as well. Maskin and Tirole (2004) call this the case with feedback¹⁶.

¹⁶Canes-Wrone et al (2001) in their study of pandering, also consider this case.

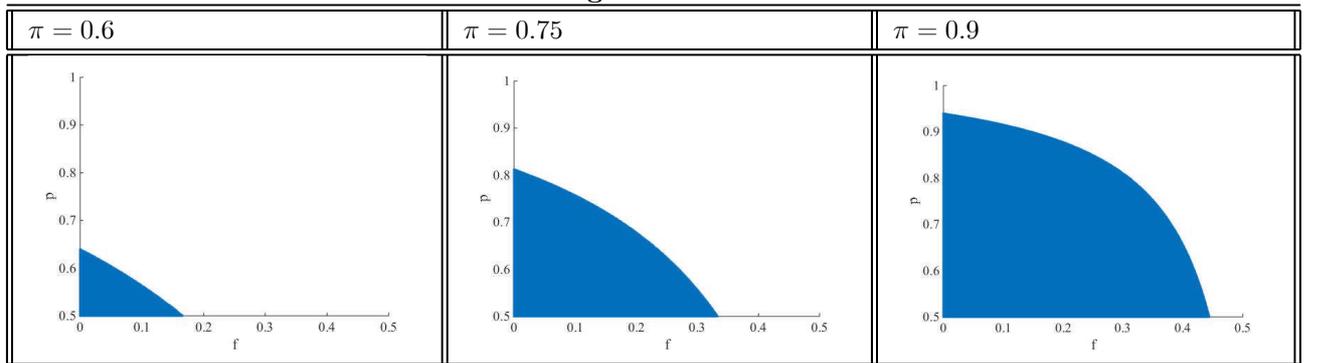
When both action and payoff are observed, the voter has a two-dimensional signal (\tilde{x}, \tilde{v}) of incumbent type, where the mapping from (x, v) to (\tilde{x}, \tilde{v}) is already specified i.e. an optimistic voter mis-classifies $x = B$ or $v = 0$ as $x = A$ or $v = 1$ respectively with probability q , and similarly a pessimistic incumbent mis-classifies $x = A$ or $v = 1$ as $x = B$ or $v = 0$ respectively with probability q .

We can then show;

Proposition 4. *Whether voters are optimists or pessimists, then for $0 \leq f \leq 0.5$, there is a unique political equilibrium with the following structure. First, the voter re-elects the incumbent either when (a) \tilde{v} is not observed, if $\tilde{x} = A$; (b) \tilde{v} is observed, if $\tilde{v} = 1$. Second, the dissonant type panders with probability $\lambda_D = F(V(1-q))$, and the consonant type panders with probability $\lambda_C = F(V(1-2f)(1-q))$. So, pandering by either type is decreasing in voter confirmation bias.*

This demonstrates that Proposition 1 is not just a knife-edge result. Moreover, for f close to 0, welfare can be increasing in q , confirmation bias. To see this, we can compute welfare directly as in Section 4.2, and then calculate the derivative analogous to (13). The details are more complex and are reported in the Appendix. To do this, we consider only the case of the naive voter, and also assume that F is uniform. As a result of a uniform F , W is linear in q , so voter welfare is higher with confirmation bias than without if and only if $\partial W / \partial q > 0$. The key parameters that we vary are f , the probability of observing the payoff, and p , the degree of initial bias in favour of $x = A$. Other parameters are fixed as explained in the Appendix. The Figure below shows the area (colored in blue/dark grey) in the space $(f, p) \in (0.5, 0.1] \times [0, 0.5]$ where $\frac{\partial W}{\partial q}$ is positive.

Figure 1



We can see that as expected, conditional on p , $\frac{\partial W}{\partial q}$ is more likely to be positive when f is small. Also, conditional on f , $\frac{\partial W}{\partial q}$ is more likely to be positive when p is small. This is because when $p < \pi$, the discipline loss due to confirmation bias actually becomes negative. Finally, as is clear from Figure 1, $\frac{\partial W}{\partial q}$ is more likely to be positive when π is large. When $\pi > p$, it is clear that $\frac{\partial W}{\partial q}$ can be positive even when f is a long way from zero.

6 Multiple Elections

So far, we have assumed that the voter starts with a bias already in place. This is a restrictive assumption. It is worth asking how bias evolves over time as the voter updates beliefs. To do this, we assume an additional period 0 to the baseline model, so there are now periods $t = 0, 1, 2$. Moreover, we assume that the voter is initially *neutral*- that is, his initial prior about the quality of the incumbent is $\pi = 0.5$. Then, at the first election, the voter has no confirmation bias. However, if he re-elects the incumbent, this will create confirmation bias $q > 0$ because the posterior belief of the voter after the first election is now greater than 0.5, and so he is now an optimist.

It is then possible to show that there is a unique equilibrium in cutoffs, $\hat{u}_{C,0}, \hat{u}_{D,0}, \hat{u}_1$ where as before, both types of incumbent pander at $t = 1$ iff $u_1 \leq \hat{u}_1 = V(1 - q)$, and types C, D pander at $t = 0$ if $u_0 \leq \hat{u}_{C,0}, u_0 \leq \hat{u}_{D,0}$. The formulae for $\hat{u}_{C,0}, \hat{u}_{D,0}$ are more complex, and is derived as part of the proof of Proposition 5 below. The corresponding pandering probabilities are $\lambda_{C,0} = F(\hat{u}_{C,0}), \lambda_{D,0} = F(\hat{u}_{D,0}), \lambda_1 = F(\hat{u}_1)$. Moreover, cutoffs $\hat{u}_{C,0}, \hat{u}_{D,0}, \hat{u}_1$ depend on q , and so do the probabilities, so we write $\lambda_{C,0}(q), \lambda_{D,0}(q), \lambda_1(q)$.

Then, we can show:

Proposition 5. *Assume that $\pi = 0.5$, and $0.5 < p < \bar{p}$. With confirmation bias q , $\lambda_{T,0}(q) > \lambda_{T,0}(0) > \lambda_1(0) > \lambda_1(q) > 0$, for $T = C, D$. That is, pandering is higher in the first period, and lower in the second, than without confirmation bias, for both types.*

The intuition for lower pandering in the second period has already been discussed. The intuition for higher pandering in the first period is simple. First, the higher q , the less likely is the incumbent to lose the second election. So, the higher q , the higher the continuation value of office if the incumbent wins the first election, and this in turn increases the incentives for pandering in the first election. Note now, however, that the incentive to pander is higher for the consonant type.

7 The Politician and the Judge Revisited

One of the main purposes of the original Maskin and Tirole (2004) paper was to provide a framework for thinking about when a politician (an elected official) would be more attractive than a judge (an appointed one). In our simple framework, the voter payoff from an elected official has already been calculated in (6). The voter welfare from an appointed official is simply $\pi + \delta\pi$, as the official will match the action to the state in both periods if and only if he is benevolent. So, the welfare gain to an elected official is easily computed to be

$$\Delta W = \lambda(p - \pi) + (1 - \lambda)\delta(2p - 1)(1 - \pi)\pi \quad (15)$$

We can see that this gain depends on the probability of pandering by the elected official i.e. the higher λ , the higher is the discipline effect of election, $p - \pi$, and the lower is the selection effect, $\delta(2p - 1)(1 - \pi)\pi$.

From (16), the effect of a change in effective confirmation bias on the relative advantage of an elected official is this

$$\frac{\partial \Delta W}{\partial \hat{q}} = -f(\hat{u}) \frac{\partial \Delta W}{\partial \lambda} = f(\hat{u})[\delta(2p - 1)(1 - \pi)\pi - (p - \pi)] \quad (16)$$

So, everything turns on the sign of $p - \pi$. On issues where voters have a strong prior about what is the "right" policy ($p > 0$), from (15), a politician always dominates a judge, and but this advantage may increase or decrease with confirmation bias. Second, if $p < \pi$, it is ambiguous whether a politician dominates a judge, but the relative advantage of the politician is strictly increasing with confirmation bias. It is helpful to express this in the following table:

Table 1

	Politician better than Judge?	Effect of CB on ΔW
$p > \pi$	Always	positive only if $\delta \geq \hat{\delta}$
$p < \pi$	only if $\frac{\lambda}{1-\lambda} < \frac{\pi-p}{\delta(2p-1)(1-\pi)\pi}$	positive

So, we see from Table 1 that when the choice between an elected and appointed official is not trivial, confirmation bias always works in favour of choosing the elected official. This is because bias reduces pandering.

So, there is a complementarity, from the voter's point of view, between a more confirmation bias and elected officials. If confirmation bias is strong, any advantage of elected officials is higher i.e. it is more likely that condition $\frac{\lambda}{1-\lambda} < \frac{\pi-p}{\delta(2p-1)(1-\pi)\pi}$ holds, and so that elected officials dominate appointed ones.

8 Conclusions

This paper considers the implications of voter confirmation bias in a political economy setting. In the baseline case, we show that confirmation bias reduces pandering, as it lowers the electoral "reward" for this behavior by reducing the increase in the probability of being elected from pandering. As pandering generally has an ambiguous effect on voter welfare, it is possible that an increase in confirmation bias (parametrized by the probability of misreading policy choice, the signal about the benevolence of the politician) increases voter welfare. This baseline result on pandering is robust to a number of extensions of the model, but does not extend to the case where the incumbent can have three (or more) terms of office, or when payoffs are observable with a high probability. Finally, we have shown that when

confirmation bias is present, other things equal, the case for decision-making by an elected rather than an appointed official is greater.

9 References

- Alesina, A. and F. Passarelli, (2015), Loss Aversion in Politics. NBER Working Paper 21077
- Ashworth, S., and Bueno de Mesquita, E (2014). Is Voter Competence Good for Voters?: Information, Rationality, and Democratic Performance. *American Political Science Review*, 108(03), 565-587.
- Besley, T. (2006), *Principled Agents? : The Political Economy of Good Government*, Oxford University Press
- Bernheim, B. D., & Rangel, A. (2007). Toward choice-theoretic foundations for behavioral welfare economics. *The American Economic Review*, 464-470.
- Besley, T., & Prat, A. (2006). Handcuffs for the Grabbing Hand? Media Capture and Government Accountability. *American Economic Review*, 96(3), 720-736.
- Callander, S., & Wilson, C. H. (2006). Context-dependent voting. *Quarterly Journal of Political Science*, 1(3), 227-254.
- Callander, S., & Wilson, C. H. (2008). Context-dependent voting and political ambiguity. *Journal of Public Economics*, 92(3), 565-581.
- Chaffee, Steven H. and Y. Miyo. (1983) Selective exposure and the reinforcement hypothesis: An intergenerational panel study of the 1980 Presidential campaign. *Communication Research*, 10(1), 3-36.
- Canes-Wrone, B., Herron, M. C., & Shotts, K. W. (2001). Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science*, 532-550.
- Doll, B. B., Hutchison, K. E., & Frank, M. J. (2011). Dopaminergic genes predict individual differences in susceptibility to confirmation bias. *The Journal of Neuroscience*, 31(16), 6188-6198.
- Ellis, A. (2012). Condorcet Meets Ellsberg. Mimeo, Boston University, Boston (MA).
- Ghirardato, P., & Katz, J. N. (2006). Indecision Theory: Weight of Evidence and Voting Behavior. *Journal of Public Economic Theory*, 8(3), 379-399.
- Gibson, R. K., & Römmele, A. (2009). Measuring the professionalization of political campaigning. *Party Politics*, 15(3), 265-293.
- Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4), 555.
- Huckfeldt, R., and J. Sprague. (1988) Choice, social structure, and political information: The information coercion of minorities. *American Journal of Political Science*, 32(2), 467-482.

- Iyengar, S., & Hahn, K. S. (2009). Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1), 19-39.
- Jerit, J., & Barabas, J. (2012). Partisan perceptual bias and the information environment. *The Journal of Politics*, 74(03), 672-684
- Levy, G. and R. Razin. (2014) Correlation Neglect, Voting Behaviour and Information Aggregation. *American Economic Review*, forthcoming
- 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, 2098–2109
- Maskin, E., & Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 1034-1054.
- Moynihan, D. P., & Lavertu, S. (2012). Cognitive Biases in Governing: Technology Preferences in Election Administration. *Public Administration Review*, 72(1), 68-77.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175.
- O'Donoghue, T., & Rabin, M. (1999). Doing it now or later. *American Economic Review*, 103-124.
- Passarelli, F. and G. Tabellini, (2014), Emotions and Political Unrest. Unpublished paper, Bocconi
- Petrova, M. (2008). Inequality and media capture. *Journal of Public Economics*, 92(1), 183-212.
- Plous, S., (1991). Biases in the Assimilation of Technological Breakdowns: Do Accidents Make Us Safer?. *Journal of Applied Social Psychology*, 21, 1058–1082
- Prat, A. (2005). The wrong kind of transparency. *The American Economic Review*, 95(3), 862-877.
- Rabin, M., & Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 37-82
- Spiegler, R. (2013). Placebo reforms. *The American Economic Review*, 103(4), 1490-1506.
- Stroud, N. J. (2010). Polarization and partisan selective exposure. *Journal of Communication*, 60(3), 556-576.

A Appendix

Proof of Proposition 1. (i) Assume that voters only re-elect if they observe $\tilde{x} = A$. Then, the only other possible equilibrium is where one (or both) 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 voters re-elect iff they observe $\tilde{x} = B$. But this is ruled out by the monotonicity assumption.

(iii) The third and fourth possibilities are that voters always or never re-elect the incumbent, whatever \tilde{x} . But in this case, both types will choose their short-run optimal actions, whatever u , so that $\Pr(i = C | x = B) < \pi < \Pr(i = C | x = A)$. So, neither voting strategy can be sequentially rational. \square

Proof of Proposition 3. Assume first that the voter is an optimist. Then, the payoffs to imitation and short-run optimal actions in either state are Vq , $u + V$ respectively. 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. The same argument applies if the voter is a pessimist. To show uniqueness, we can apply the same argument as in the proof of Proposition 1. \square

Proof of Proposition 4. (i) Assume that consonant and dissonant incumbents pander - i.e. always choose A - with probabilities $\lambda_C, \lambda_D \in [0, 1]$ respectively. First, recall that the consonant incumbent will produce payoff $v = 1$ with probability 1 if he matches, and probability p if he panders, whereas the dissonant incumbent will produce payoff $v = 1$ with probability 0 if he matches, and probability p if he panders. So, by straightforward application of Bayes' rule;

$$\begin{aligned} \Pr(C|B, 1) &= 1 \\ \Pr(C|A, 1) &= \frac{(\lambda_C p + 1 - \lambda_C)\pi}{(\lambda_C p + 1 - \lambda_C)\pi + \lambda_D p(1 - \pi)} \\ \Pr(C|A, 0) &= \frac{\lambda_C(1 - p)\pi}{\lambda_C(1 - p)\pi + (1 - p)(1 - \pi)} \\ \Pr(C|B, 0) &= 0 \end{aligned} \tag{17}$$

Now, note that $\lambda_C p + 1 - \lambda_C \geq \max\{\lambda_C p, \lambda_D p\} \geq \lambda_D p$, so $\Pr(C|A, 1) \geq \pi$.

(ii) Assume that $\lambda_C < 1$, so that $\Pr(C|A, 0) < \pi$, implying that by sequential rationality, the voter will re-elect the incumbent iff he observes $\tilde{v} = 1$. Now consider the incentives for pandering for the consonant and dissonant type respectively. Assume w.l.o.g. that the voter has positive confirmation bias. If $s = A$, the best choice for the consonant incumbent is unambiguously $x = A$. If $s = B$, the consonant type gets

$$[1 - f + fq]V, u + [(1 - f)q + f]V \tag{18}$$

from pandering ($x = A$) and matching ($x = B$) respectively. The first payoff is calculated as follows. With probability $1 - f$, the voter observes only $x = A$ and will re-elect the incumbent. With probability $1 - f$, the voter will observe $v = 0$ but will mis-classify it as $v = 1$ with probability q . Similarly, the second payoff is calculated as follows. With probability $1 - f$, the voter observes B but mis-classifies it as A , and re-elects the incumbent. With probability f , the incumbent produces $v = 1$ for the incumbent, and is re-elected.

So, from (18), the consonant type will pander if

$$u \leq \hat{u}_C \equiv V(1 - 2f)(1 - q)$$

Now consider the dissonant type. If $s = A$, the payoffs from choosing $x = A, B$ are respectively

$$V, u + qV \tag{19}$$

The explanation is as follows. With probability $1 - f$, the voter observes only $x = A$ and will re-elect the incumbent. With probability $1 - f$, the voter will observe $v = 1$ and will also re-elect him. The second payoff is calculated as follows. With probability $1 - f$, the voter observes B but mis-classifies it as A , with probability q , and re-elects the incumbent. With probability f , the incumbent produces $v = 0$ but this will be mis-classified as $v = 1$ with probability q . In either case, the incumbent is re-elected with probability q .

If $s = B$, then the the choices from $x = A, B$ respectively are;

$$u + (1 - f + fq)V, ((1 - f)q + f)V \tag{20}$$

The explanation is as follows. With probability $1 - f$, the voter observes only $x = A$ and will re-elect the incumbent. With probability f , the voter will observe $v = 0$ and will re-elect him with probability q . The second payoff is calculated as follows. With probability $1 - f$, the voter observes B but mis-classifies it as A , with probability q , and re-elects the incumbent. With probability f , the incumbent produces $v = 1$ and the incumbent re-elects him. So, the dissonant type will pander i.e. always choose $x = A$, if

$$u \leq V(1 - q), u + (1 - f + fq)V \geq ((1 - f)q + f)V \tag{21}$$

A sufficient condition for the second inequality in (21) to hold is $f \leq 0.5$. Then, both (19),(20) require simply $u \leq V(1 - q)$, and so the dissonant type will pander if $u \leq V(1 - q) \equiv \hat{u}_D$. So, $\lambda_C = F(\hat{u}_C)$ if $f < 0.5$, and $\lambda_D = F(\hat{u}_D)$ as required. Finally, note that $\lambda_C < \lambda_D$ in equilibrium, confirming the assumption that $\Pr(C|A, 0) < \pi$, and thus confirming the voter re-election rule.

(iii) It remains to rule out the equilibrium where $\Pr(C|A, 0) \geq \pi$. From (17), this can only $\lambda_C = 1$.

But, by assumption, $\bar{u} > \delta(1 + E) = V$, so $\lambda_C < 1$. \square

Proof of Proposition 5. (i) We solve the model backwards, starting with the election at the end of period 1. Suppose that the probability that pandering occurred in period 0 was $\lambda_{C,0}$ $\lambda_{D,0}$ for the consonant and dissonant types respectively. So, let the voter's posterior at the beginning of period 1, having observed $x_0 = A$, be

$$\pi_1 = \frac{(\lambda_{C,0} + (1 - \lambda_{C,0})p)\pi}{(\lambda_{C,0} + (1 - \lambda_{C,0})p)\pi + (\lambda_{D,0} + (1 - \lambda_{D,0})(1 - p))(1 - \pi)} \quad (22)$$

We need to show that if the voter then observes $x_1 = B$, he will *not* wish to re-elect the incumbent (otherwise, the incumbent will always match in period 1, as there is no electoral cost to doing so). This requires

$$\frac{(1 - \lambda_1)(1 - p)\pi_1}{(1 - \lambda_1)(1 - p)\pi_1 + (1 - \lambda_1)p(1 - \pi_1)} < \pi \quad (23)$$

So, substituting (22) onto (23), and rearranging, we need:

$$\frac{(\lambda_{D,0} + (1 - \lambda_{D,0})(1 - p))}{(\lambda_{C,0} + (1 - \lambda_{C,0})p)} > \frac{1 - p}{p} \quad (24)$$

Then, given (24), the incumbent at $t = 1$ faces a choice between pandering and being re-elected, or matching and being re-elected with probability p , so will pander with probability $\lambda_1 = F(V(1 - q))$.

(ii) Now consider the decision at $t = 0$. The continuation payoff to being re-elected for the consonant type is

$$V_C(q) = p(1 + V) + (1 - p)E[\max\{V, u + qV\}] \quad (25)$$

This is explained as follows. With probability p , the state is $s = A$, so pandering is also short-run optimal, i.e. the incumbent behaves the same whether pandering or matching, giving a payoff of μ plus certain re-election at the second election. With probability $1 - p$, pandering is not short-run optimal, so the incumbent chooses the better of the two options.

For the dissonant type, using the same reasoning, we have,

$$V_D(q) = (1 - p)(1 + V) + pE[\max\{V, u + qV\}] \quad (26)$$

So, in the first period, the pandering probabilities are

$$\lambda_{C,0} = F(\delta V_C(q)), \lambda_{D,0} = F(\delta V_D(q)) \quad (27)$$

Note from (25),(26) that $V_C(q) \geq V_D(q)$, so $\lambda_{C,0} \geq \lambda_{D,0}$ as claimed. Also, clearly, from (25), (26), $V_C(q), V_D(q)$ are increasing in q , as claimed.

(iii) Next, we need to show that pandering is declining in the benchmark without confirmation bias.

Note from (27) that $\lambda_{0,C}(0) > \lambda_1(0)$ if

$$p(1 + V) + (1 - p) \left[\int_{\hat{u}}^u u f du + \lambda_1 V \right] > V$$

which requires

$$p + (1 - p) \int_{\hat{u}}^u u f du > V(1 - p)(1 - \lambda_1)$$

But

$$p + (1 - p) \int_{\hat{u}}^u u f du = p + (1 - p)E[u | u \geq V](1 - \lambda_1) > p + (1 - p)V(1 - \lambda_1)$$

as required. So, $\lambda_{0,C}(0) > \lambda_1(0)$. As similar argument implies that $\lambda_{0,D}(0) > \lambda_1(0)$.

(iv) Finally, we need to verify condition (24). From (25),(26), as $p \rightarrow 0.5$, $V_C(q) \rightarrow V_D(q)$, so $\lambda_{D,0}, \lambda_{C,0} \rightarrow \lambda_0$, so the LHS of (24) tends to

$$\frac{(\lambda_0 + (1 - \lambda_0)(1 - p))}{(\lambda_0 + (1 - \lambda_0)p)}$$

which is clearly greater than $\frac{1-p}{p}$, so this holds. \square

Simulations for the Welfare Effect of Confirmation Bias with Observable Payoffs. We compute this for the case of the naive voter. Define the re-election probabilities for the two types, depending on whether they pander "P", or do not pander, "N", as

$$\begin{aligned} r_{CP} &= f + (1 - f)p, r_{DP} = f + (1 - f)(1 - p) \\ r_{CN} &= fp + 1 - f, r_{DN} = f(1 - p) \end{aligned} \quad (28)$$

Then, voter welfare is

$$\begin{aligned} W &= \pi \lambda_C(p + \delta(r_{CP} + (1 - r_{CP})\pi)) + (1 - \pi)\lambda_D(p + \delta(1 - r_{DP})\pi) \\ &\quad + \pi(1 - \lambda_C)(1 + \delta(r_{CN} + (1 - r_{CN})\pi)) + (1 - \pi)(1 - \lambda_D)(\delta(1 - r_{DN})\pi) \end{aligned} \quad (29)$$

For the simulations, we assume $\delta = 1$, $E = 1$, $\sigma \geq 2$. Moreover, we assume that u is uniformly distributed on $[0, \sigma]$, $\sigma \geq 2$ so from Proposition , the pandering probabilities are $\pi = 0.75$,

$$\lambda_D = \frac{2(1 - q)}{\sigma}, \quad \lambda_C = \frac{2(1 - 2f)(1 - q)}{\sigma} \quad (30)$$

Clearly, from (28)-(30), with F uniform, W is linear in q , so $W(q) > W(0)$, $q > 0$ if and only if $\partial W / \partial q >$

0. From (29),(30), we have;

$$\begin{aligned} \frac{\partial W}{\partial q} &= \pi \frac{\partial \lambda_C}{\partial q} (p + r_{CP} + (1 - r_{CP})\pi - 1 - r_{CN} - (1 - r_{CN})\pi) \\ &\quad + (1 - \pi) \frac{\partial \lambda_D}{\partial q} (p + r_{DP} + (1 - r_{DP})\pi - (1 - r_{DN})\pi) \end{aligned} \quad (31)$$

We then compute $\frac{\partial W}{\partial q}$ from (31) for the values $f \in [0, 0.5]$, $p \in (0.5, 1]$, holding other values fixed as specified above. \square