Democracy, Populism, and (Un)bounded Rationality *

Johannes Binswanger†

Jens Prüfer‡

April 29, 2009

Abstract

In many instances, both voters and politicians are imperfectly informed about which policies are optimal. We analyze politicians' policy choice in such situations. A distinctive element of our analysis is that we investigate how the strategic sophistication of voters' beliefs about politicians' behavior affects policy choice. This provides a novel approach in political economy that leads to a number of important insights. We show that these beliefs determine the strength of self-serving politicians' incentives to engage in populism. Surprisingly, limited strategic sophistication of voters *weakens* politicians' incentives to pander to public opinion. The reason is that politicians know that such voters *expect* them to choose a policy that is not perfectly pandering to public opinion. Furthermore, when comparing the welfare ranking of different constitutional regimes, we find that limited strategic sophistication of voters makes indirect democracy relatively more attractive compared to the case of full strategic rationality – and often more attractive than alternative constitutional regimes.

Key words: Imperfect information, beliefs, strategic sophistication, democracy, populism, accountability, experts.

JEL classification: D72, D78, D83.

---

*We would like to thank Jan Boone, Reto Föllmi, Alia Gizatulina, Manuel Oechslin, Torsten Persson, Jan Potters, Amrita Ray Chaudhuri, Dominic Rohner, Uwe Walz, Jörgen Weibull, and seminar participants at various universities and participants at the AEA 2009 annual meeting for very helpful comments. All errors are our own.

†*Corresponding author*, Department of Economics, CentER, Netspar, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.binswanger@uvt.nl.

‡Department of Economics, CentER, Tilec, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.prufer@uvt.nl.

1 Introduction

The best argument against democracy is a
five-minute conversation with the average voter.

Winston Churchill

The world is complex and people may not properly understand how the modern society and its economy function. In some cases, voters' beliefs deviate substantially from the view of experts. Caplan (2001) conducted a survey where he compared the opinions of the general public to the opinions of economists with regard to a number of economic issues. For instance, he asked people whether they think that trade agreements between the United States and other countries have helped create more jobs in the U.S. He coded the answer that trade agreements "have cost the U.S. jobs" as 0, that they "haven't made much of a difference" as 1, and that they "helped create jobs in the U.S." as 2. He finds that the mean response among economists is 1.47. In contrast, the mean response among the general audience is only .64. Caplan reports similar discrepancies with respect to other economic issues.

This evidence suggests that voters' are sometimes imperfectly informed about which policies are in their own self-interest. In contrast to voters, politicians have a stronger incentive to acquire scrutinized information about desirable policies. After all, the success or failure of an implemented policy affects their reelection prospects. Second, politicians have simply better access to information than voters as they are endowed with advisor teams.

This notwithstanding, there are good reasons to expect that also politicians may be imperfectly informed about what policies maximize voters' well-being. Their views about optimal policies may be shaped by ideological considerations. Furthermore, they may be shaped by their own professional background and the professional composition of their advisor teams. Different professions, and different groups within a profession, may have diverging views about how voters react to a policy change. To the degree that this is true, there is imperfect knowledge about optimal policies even among politicians, not only among voters.

In this paper, we analyze politicians' policy choice in a setup where both voters and politicians are imperfectly informed about which policy maximizes voters' utility. A main novel aspect of our analysis is that we take into account an important element of bounded rationality: voters may have a limited capacity to anticipate a politician's incentive to higher-order strategic behavior. This gives rise to beliefs of limited strategic sophistication about politicians' behavior that we dub *level-k* beliefs. These beliefs should be understood as a reduced-form version of the *cognitive hierarchy model* of Camerer et al. (2004). We show that the nature of these beliefs is decisive for the type of equilibrium that emerges. Integrating them into a model of public choice leads to interesting insights that have not appeared in the literature yet. Beyond this, we demonstrate how a reduced-form version of the cognitive hierarchy model can be fruitfully used for understanding the impact of strategic thinking on equilibrium outcomes in applied contexts, rather than the stylized games that are analyzed in Camerer et al.'s seminal paper.

So far, there has been little attempt to incorporate elements of behavioral economics and bounded rationality into models of voting. Thus, this becomes an important item on the research agenda in political economy. For instance, Besley (2006) writes "going forward it would be interesting to understand better what the differences are between behavioral models of politics and the postulates of strict rationality" (p. 172). By taking into account limitations on voters' strategic thinking, and by including full rationality as a limit case, we make one step in this direction. For the sake of clarity and parsimony, we assume that voters are rational in any aspect not related to beliefs about a politician's strategic behavior.

In our model, every voter is endowed with a belief about which policy maximizes his expected utility. We refer to this belief as a voter's *opinion*. While opinions may be heterogeneous, we assume that the policy that truly maximizes voters' utility is identical for all voters. In addition, voters are endowed with a belief about the strategic behavior of politicians, that is how they set a certain policy based on their information. It is this belief that may be characterized by bounded rationality.

Politicians in our model receive a signal indicating which policy is optimal. There are two types of politicians dubbed competent and incompetent, respectively. The competent type's signal perfectly reveals the optimal policy (from an ex ante point of view), whereas the incompetent type's signal is noisy. Importantly, politicians do not observe their type. This captures the fact that it may be very difficult to objectively prove whether one policy choice dominates another.

There are two office periods. An incumbent politician selects a policy for the first period. At the end of the first period, an election takes place where the incumbent may get reelected or replaced by a challenger. Then, a policy is chosen for the second period. Voters' aim is to elect a politician for the second period who is perceived as competent *through the lens of their opinions*. This is a rational decision strategy since it maximizes expected utility, given their opinions.

In order to get reelected, politicians have an incentive to pander to the pivotal voter's opinion. However, the inventive to pander is mitigated by voters receiving a (noisy) signal about which policy has been optimal before the election takes place. Voters use this signal to form a "posterior opinion" that they use for judging an incumbent politician's competence. Due to being endowed with better information, a competent politician is more successful anticipating the median voter's posterior opinion and hence faces a higher probability of being reelected.

The idea of level-k beliefs can be understood by means of the following recursive reasoning. Suppose that voters believe that a politician tries his best to maximize voters' objective utility, given the information obtained through his signal. We define such a belief as a level-0 belief. As we will show, this belief induces an incentive for a politician not to maximize voters' objective well-being. Rather, he wants to be perceived as competent by the median voter and, hence, panders to the latter's opinion. If voters' have a belief of strategic sophistication of level-1, they anticipate this first-order strategic incentive to pander to the median voter's opinion. In particular, voters' take this incentive into account when judging the competence of an incumbent politician.

Given that voters hold a level-1 belief, a politician faces a second-order incentive to deviate from this belief. If voters also anticipate this latter strategic incentive, they are said to be endowed with level-2 beliefs. Proceeding with this recursion, we may define a level-k belief, where k may take on any number between zero and infinity. The case of an infinite k corresponds to perfect strategic sophistication, that is full rationality.

In our analysis we take k as given and common across voters. The empirical evidence discussed in Camerer et al. (2004) suggests that realistic values of k may lie between one or two. This stands in sharp contrast to the requirement of unlimited rationality for anticipating strategic reactions for infinitely many orders. A salient example of how difficult it is in practice to anticipate higher-order strategic reactions is provided by the chess game.

Whether k is to be seen as a low number – as suggested by the evidence reported in Camerer et al. (2004) – or rather infinitely high turns out to crucially matter for our results. For finite k, we obtain a separating equilibrium, in which competent and incompetent politicians implement different policies. This equilibrium is characterized by only partial populism in that both the politician's signal and the median voter's opinion have a positive weight in affecting the policy choice. In contrast, for infinite k, we obtain a pooling equilibrium in which both politician types implement identical policies. These are exclusively determined by the median voter's prior opinion about the optimal policy and do not depend on the incumbent politician's signal at all. Thus, the equilibrium is perfectly populist.

Overall, we find that the higher the degree of strategic sophistication k, the lower the influence of a politician's signal on his policy choice. The intuition for this result is that a very sophisticated voter with a high k *expects* a politician to pander to the median voter's beliefs by a high number of orders. In the limit, the median voter's opinion completely determines a politician's policy choice.

Conversely, limited strategic sophistication of voters' beliefs reduces a politician's incentive to pander and increases the importance of his signal for policy making. To the degree that the signal is sufficiently informative, welfare is higher in the case of bounded rationality than in the benchmark case of full rationality. Beyond this, our analysis provides a novel explanation why politician's actions may not be fully populist. This result could also be obtained by assuming that politicians are intrinsically motivated. In contrast, we obtain this result even in the case that politicians are completely motivated by ego-rents from holding office, but where voters are boundedly rational. This is a novel result that has not appeared in the literature yet.

We also consider the implications of the degree of strategic sophistication on how the institution of indirect democracy – the institution that corresponds to our baseline setup – compares to direct democracy and to the case where policies are chosen by non-accountable agents. We show that limited strategic sophistication of voters' beliefs tend to give indirect democracy an edge over direct democracy and delegation to independent agents.

The remainder of this paper is organized as follows. Section 2 briefly discusses the related literature. Section 3 introduces our baseline model of indirect democracy. In Section 4, we solve this baseline model. In Section 5, we compare indirect democracy to the case of direct democracy and to the case of delegation of policy making to non-accountable agents such as judges, experts, or bureaucrats. We conclude in Section 6. Proofs are contained in Appendix A, if not stated otherwise. Several extensions of the model are presented in Appendices B, C, and D.

2 Related Literature

Our analysis is related to a number of existing studies. Maskin and Tirole (2004) analyze the welfare effects of pandering when voters are imperfectly informed about which of two potential policy options is optimal. As in our model, there are two types of politicians. These do not differ in terms of their competence but in terms of whether their preferences are congruent or non-congruent with voters'. In contrast, in our paper, we analyze politicians' policy choice in a setup where both voters and politicians are imperfectly informed about which policy maximizes voters' utility. Unlike Maskin and Tirole, we consider a continuous setup where the policy space corresponds to the entire real line. This allows us to capture the distance between voters' opinions and a politician's signal, on the one hand, and the truly optimal policy, on the other, in a natural way.

The main difference between our analysis and Maskin and Tirole's work arises from our focus on the degree of the strategic sophistication of voters' beliefs about an incumbent politician's behavior. As outlined in the introduction, our results differ fundamentally for the two cases of either perfect or limited strategic sophistication of these beliefs and this insight is novel.

Another related study is Canes-Wrone et al. (2001). They consider a setup where both voters and politicians are imperfectly informed about which of two policy options is optimal. Again, our approach differs in that we analyze a continuous setup and, in particular, focus on beliefs of limited strategic sophistication. Furthermore, Canes-Wrone et al. do not compare the welfare properties of different political institutions. Dixit and Weibull (2006) study the possibility of polarization in voters' beliefs in a setup that is reminiscent of ours.

Alesina and Tabellini (2007, 2008) provide an in-depth analysis of the advantages and disadvantages of accountability. In particular, they compare the performance of a politician who aims to get reelected with the performance of a bureaucrat who is concerned about his career perspective. Pandering and imperfect knowledge about optimal policies do not play a pivotal role in their analysis. Schultz (2008) analyzes the welfare effects of accountability by focusing on the term length of office periods. In our study, we take this term length as given.

Besley and Case (1995) study the effects of diminished accountability of politicians via term-limit rules empirically. Analyzing the behavior of U.S. governors from 1950 to 1986, they find evidence that politicians in fact adjust their behavior if they face a binding term limit, that is if reelection cannot serve as a disciplining device anymore.

Blanes i Vidal and Leaver (2008) go one step further by providing empirical evidence that pandering to voters is not restricted to politicians who face reelection. They show that pandering can also be found in the behavior of public officials – tenured judges in their sample – who do not face the threat of getting ousted from office but of losing decision-making power. This suggests that our analysis of indirect democracy may apply more generally to public decision makers, even if they are not directly accountable.

3 A Model of Indirect Democracy

3.1 Voters

We consider an economy populated by a unit mass of individuals to which we refer as voters. We consider a setup with two periods, indexed by $t = 1, 2$. In each period, voters' utility is determined as

$$V_t = -\left(g_t - x_t^* - \varepsilon_t\right)^2 . \tag{1}$$

The utility function is identical across voters. The variable $g_t \in \mathbf{R}$ denotes a policy action and is set by the office-holding politician. Neglecting ε_t, the utility maximizing level of g_t is given by $x_t^* \in \mathbf{R}$. The crucial assumption in our framework is that x_t^* is unobserved. We assume that x_t^* is drawn at the beginning of each period by nature from a normal distribution with mean Ex_t^* and variance σ_x^2. The mean may vary across periods and is unknown.

The variable ε_t is a normally distributed random variable with an expected value of zero and a variance of σ_ε^2. We assume that ε_t is identically and independently distributed over time and independent of all other random variables in the model. As is the case for x_t^*, ε_t is also unobserved. The distribution of ε_t is common knowledge.

As will be spelled out in more detail in Subsection 3.3, nature first draws x_t^*, before ε_t is realized. The policy action g_t is to be set after x_t^* has been determined but before ε_t is realized. Thus, x_t^* determines the $ex\ ante$ optimal policy in period t. It specifies how, from an ex ante point of view, a choice of g_t translates into voters' utility. In contrast, ε_t represents a short-term shock to x_t^* and determines the $ex\ post$ optimal level of g_t. While x_t^* and ε_t are not observed in isolation, voters do observe the sum $x_t^* + \varepsilon_t$ after g_t has been set. This allows voters to learn, although imperfectly, about x_t^* (see below).

From an ex-ante perspective, voters' utility in period t is given by the expected value of V_t, that is by

$$EV_t = -E\left[(g_t - x_t^* - \varepsilon_t)^2\right]. \tag{2}$$

The loss function specification is chosen for tractability. This utility function should be taken as reflecting indirect utility, meaning that optimal values of all other choices that voters may make are already substituted. There are two essential features of (1) or (2). First, x_t^* determines a unique interior optimum for g_t from an ex ante point of view. Second, there is risk aversion over the realizations of g_t if the latter are uncertain. We assume that x_t^* and, hence, (1) and (2) are common across voters.[1]

To consider an example, suppose that there is a given budget to be spent for combating crime. Suppose that the relevant decision is to determine the share of this budget to be spent on preventive measures (schooling, prevention of youth unemployment, quality of neighborhoods etc.) versus the share to be spent on punishment (e.g. prison infrastructures). In this example, x_t^* refers to the optimal budget share for preventive measures, given the *general* current situation in society. This may refer to the degree of income inequality and ethnic heterogeneity, the degree to which people follow certain norms, the general level of youth unemployment etc. The variable ε_t corresponds to a shock to the "threat of crime" and may originate from a sudden rise in youth unemployment, a sudden increase in immigration or the like.[2]

As already stated, x_t^* is not observed and Ex_t^* is unknown. However, voters have *prior opinions* about x_t^*.[3] Specifically, we make the following assumption.

[1]There are standard examples where heterogeneous voters agree about the optimal level of provision of a public good. For instance, this is the case in the presence of income heterogeneity when the utility function is of Cobb-Douglas type with private consumption and a public good as the arguments and with a linear income tax. This is an important benchmark case (see Atkinson and Stiglitz, p. 302).

[2]We assume that x_t^* is normally distributed because of the high tractability of the normal distribution. For the example of choosing a share of a budget to be spent on preventive measures for combating crime, the policy variable could only take on values between zero and one. This would not be consistent with a normal distribution. However, it is straightforward to find a *transformation* of the domain of admissible policies such that they may take on any real value. Any function that is bijective and maps $[0,1]$ onto the entire real line would achieve this.

[3]Normally, we would speak of a prior *belief* rather than opinion. We prefer to use the latter term in order to prevent confusion between beliefs about optimal policies, on the one hand, and beliefs about a politician's behavior, that is level-k beliefs (see below), on the other.

Assumption 1 *A voter i's prior opinion about x_t^* is given by x_t^i which is a normally distributed random variable with mean μ_t^i and variance σ_x^2. The distribution of prior opinions across voters is common knowledge.*

According to Assumption 1, the prior means of x_t^i may be heterogeneous among voters while, for simplicity, we assume that the variance is common across voters.

3.2 Politicians

The policy action g_t is chosen and implemented by an incumbent politician. An incumbent politician's objective in the first office period is to maximize the probability of getting reelected for a second term.[4] Conditional on being (re)elected for office in the second period, a politician's objective is simply to maximize voters' utility in this period. The latter assumption is to be understood as a shortcut and does not affect our main conclusions in a substantive way.[5]

A politician knows the distribution of voters' prior opinions about x_t^*. However, he does not directly observe x_t^*. Rather, a politician receives a signal ξ_t that is informative about x_t^*. There are two politician types that we dub competent and incompetent, respectively. The prior probability that a politician is competent is denoted by α and is common knowledge. In case of the competent politician, $\xi_t = x_t^*$, that is the signal reveals the truth.[6] An incompetent politician receives a noisy signal. Specifically, in the first period, $\xi_1 = x_1^* + \zeta_1$, where ζ_1 is a random variable with mean zero and variance σ_ζ^2. We dub ζ_t an incompetent politician's $bias$. In the second period, $\xi_2 = x_2^* + \zeta_2$ in the case of a challenger winning the election. We assume that ζ_t is independent of all random variables in the model and that ζ_2 is independent of ζ_1 and identically distributed. Furthermore, the distribution of ζ_t is common knowledge.

For the case of an incumbent politician we make the following assumption.

Assumption 2 *An incompetent incumbent who gets reelected for a second office period keeps his bias ζ_1, that is $\xi_2 = x_2^* + \zeta_1$.*

[4]One interpretation of this is that he derives ego rents from being in office, as in Rogoff (1990). See also Besley (2006).

[5]In particular, we may allow for rent seeking along the lines of a model discussed in Persson and Tabellini (2000, Ch. 4). See footnote 13 below. We exclude rent seeking here in the interest of transparency.

[6]The assumption that the competent politician perfectly observes x_t^* is made for simplicity. The main conclusions from our analysis could also be obtained if the competent type received a more informative, but imperfect, signal than the incompetent type.

We make Assumption 2 because we find it more plausible than assuming that a politician's bias is drawn afresh when he gets elected for a second period. In fact, the analysis would be slightly simpler if we assumed that a politician's bias were determined anew every period.

In principle, it may be natural to allow for $E\zeta_t \neq 0$. One may argue that politicians are drawn from the general population and may thus have systematically biased views about x_t^*. We briefly discuss this case in Appendix C but do not consider it in the main model since it complicates the analysis without leading to substantive additional insights.[7]

We make two further assumptions about a politician's information. First, we follow the literature on career concerns by assuming that a politician does not observe his own type (see Holmström, 1999, or Prat, 2005). This means that he does not observe whether his signal is perfect or noisy. Second, we make the simplifying assumption that a politician treats his signal ξ_t as a best predictor for x_t^*. To state this formally, assume that a politician's belief about the ex ante optimal policy x_t^* is captured by a random variable x_t^p. The superscript is an index for politicians. We then state the mentioned assumptions as follows.

Assumption 3 *(i) A politician does not observe his type. (ii) He believes that* $E\left[x_t^p \,|\xi_t\right] = \xi_t$.

We make the additional simplifying assumption that a politician takes ξ_t as a "point estimate" for x_t^* in the sense of classical statistics and his behavior is only based on this point estimate (rather than on a non-degenerate belief x_t^p). We make this assumption only for simplicity and discuss its relaxation in Appendix D. In practical terms, ξ_t should be interpreted as a policy suggestion that a politician gets from (discussions with) his advisers or his party. Thus, a politician's competence is not only determined by his personal skills but also by the competence of his advisers and party strategists.

3.3 The Political Game

Below we indicate the stages of the political game in a more formal manner. We provide a label for each stage of the game. The letters in the labels refer to the players which have their moves at the respective stages. N denotes nature, P denotes the politician, and V denotes voters. The

[7] Whether politicians should be understood as a "representative sample" drawn from the general population clearly depends on the nature of the political recruitment process and may differ across countries.

first figure after the letter refers to the office period $t = 1, 2$. The second figure indexes moves within an office period for nature, as nature has two moves within one period.

- *Stage N1.1:* Nature draws x_1^*; it determines the type of the incumbent politician and his signal ξ_1.
- *Stage P1:* The incumbent politician chooses g_1.
- *Stage N1.2:* Nature draws ε_1 and sends the signal $x_1^* + \varepsilon_1$ to voters.
- *Stage V1:* Voters decide whether to reelect or oust the incumbent politician.
- *Stage N2.1:* Nature draws x_2^*; if a new politician is in office, nature determines his type and, in case of an incompetent politician, his bias ζ_2; furthermore, nature sends the signal ξ_2.
- *Stage P2:* The politician chooses g_2.
- *Stage N2.2:* Nature draws ε_2.

Politicians have two moves in the above game, since they set g_t in each period. The politician in period 2 may be different from the politician in office in period 1. Voters have only one move in the entire game, that is they decide whether to cast their votes for the incumbent politician or for a challenger.

Our model is comparatively rich. This is due to the fact that it incorporates the feature that both voters and politicians are imperfectly informed about what policy is optimal. Furthermore, the problem under study is only of interest if we allow for the possibility that voters have an opportunity to learn about x_t^*, but imperfectly so. This motivates the inclusion of the random variable ε_t. There is no other more parsimonious setup where we can still analyze the role of imperfect information in politics on both the voters' and the politicians' side in a meaningful way.

4 Analysis of Indirect Democracy

4.1 Overview

In order to understand the logic of our derivation of the equilibrium, it is important to conceptually distinguish between two types of "beliefs" that are part of the solution of the political game. First, there are beliefs about nature's draw of x_1^*[8], the probability that the incumbent is competent, and about the incumbent's potential bias ζ_1. Second, there are beliefs about an incumbent politician's behavior (given his type).[9]

Concerning the first type of beliefs, we assume that voters are perfectly rational and apply Bayes' rule in order to update their beliefs when they receive new information at stage *N1.2* of the game. Concerning the second type of beliefs, we assume that these correspond to the *level-k beliefs* that have been outlined in the introduction. In order to separate the two types of beliefs, we make use of the following terminology. We dub the first type of beliefs *opinions* and will thus speak of a prior or posterior opinion about x_1^*, ζ_1, and α. We reserve the use of the term *belief* for level-k beliefs.

For the derivation of the political equilibrium, it is convenient to first derive voters' posterior opinions at stage *V1* of the game. Due to the fact that an equilibrium is a fixed point, the posterior opinions about ζ_1 and α depend on level-k beliefs. However, we first do not specify the latter explicitly but rather postulate such beliefs in an abstract form. We next discuss a politician's best response to these beliefs, again in an abstract form. It is only then that we specify level-k beliefs explicitly. The reason is that it is only then that we may cut through the fixed-point problem that is associated with the mutual consistency between equilibrium beliefs and equilibrium behavior.

4.2 Voters' Posterior Opinions

In this subsection we characterize voters' sequentially rational opinions at stage V1 of the game: (i) about the ex ante optimal policy level x_1^*; (ii) about the incumbent's bias ζ_1 conditional on the incumbent being incompetent; and (iii) about the probability that the incumbent is compe-

[8]Beliefs about x_2^* are not essential for the solution of the game.

[9]This second type of belief is sometimes called a belief about another player's *behavioral strategy*.

tent. We start with posterior beliefs about x_1^*. These result from observing $x_1^* + \varepsilon_1$ at stage *N1.2* of the game. Although the ex post optimal policy level is given by $x_1^* + \varepsilon_1$, sequential rationality requires voters to be interested in the ex ante optimal level x_1^* since they want to judge a politician's competence and are aware that a politician chooses g_1 before ε_1 is realized.[10]

Concerning notation, we use a *hat* for all variables that are associated with voters' *posterior* opinions (that is opinions at stage *V1*). Variables without a hat refer to prior opinions. The following lemma states a standard result for normally distributed beliefs.

Lemma 1 (Posterior opinions about x_1^*) *Voter i's posterior opinion $\hat{x}_1^i$ about the ex ante optimal policy level x_1^* is normally distributed with mean $\hat{\mu}_1^i = (1-\beta)\,\mu_1^i + \beta\,(x_1^* + \varepsilon_1)$ and variance $\hat{\sigma}_x^2 = \frac{\sigma_x^2 \sigma_\varepsilon^2}{\sigma_x^2 + \sigma_\varepsilon^2}$, where $\beta \equiv \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2}$.*

Lemma 1 characterizes voters' posterior opinion about the ex ante optimal policy level x_1^*. The assumption of normally distributed prior opinions x_1^i implies that the posterior mean $\hat{\mu}_1^i$ is a weighted average of the prior mean and the signal $x_1^* + \varepsilon_1$, observed at stage *N1.2* of the game. The degree of updating of prior opinions depends on the signal-to-noise ratio $\beta = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2}$.

While the result in Lemma 1 is highly standard, it is useful to explain its meaning in the context of the current analysis. Consider again the example of what share of a given budget to spend on preventive measures to combat crime. Ex ante, the optimal share is given by x_1^* and voter i believes that expected welfare is maximized by setting $g_1 = \mu_1^i$. Ex post, voters observe the ex post optimal budget share for preventive measures $x_1^* + \varepsilon_1$. The latter depends on the actual ex post threat of crime according to random short-term factors constituting ε_1. If σ_ε^2 is very low, then observing $x_1^* + \varepsilon_1$ is very informative about x_1^* and voters will put a high weight on the signal. In contrast, if σ_ε^2 is high, voters update their beliefs only in a minor way.

Consider a voter who believes that a high share of the budget to combat crime should be spent on punishment and that the actual share spent on punishment has indeed been high. Suppose that, ex post, the crime rate is high. If σ_ε^2 were low, then the voter would infer that his prior beliefs were probably wrong. But if σ_ε^2 is high, he will conclude that criminal threat must have been unusually high.

We now turn to voters's posterior opinions about the realization of the incumbent politician's bias ζ_1, conditional on the incumbent being incompetent. (For the competent type, the

[10]Remember that we assume that voters are fully rational in any aspect except for beliefs about a politician's strategic behavior.

bias is identical zero.) We denote the posterior opinion about the realization of the bias ζ_1 by $\hat{\zeta}_1^i$. As indicated by the notation, this posterior opinion is heterogeneous since it depends on heterogeneous beliefs x_1^i (or $\hat{x}_1^i$). The reason why a rational voter wants to update his beliefs about ζ_1 is that an incompetent incumbent who gets reelected for the second office period will keep his bias as stated in Assumption 2. Thus, a voter i uses $\hat{\zeta}_1^i$ for determining expected utility in the second period in case of reelection of the incumbent (see Section 4.3 below).

Inferring $\hat{\zeta}_1^i$ requires a belief about how a politician sets g_1 as a function of his signal ξ_1. This is where level-k beliefs come into play. Due to a fixed-point problem, we defer the precise specification of these beliefs to later (see Section 4.5). Until there, we simply postulate that voters believe that $g_1 = G(\xi_1)$, where G is a continuous and strictly increasing function[11] that is $common$ across voters.

Because the function G is continuous and strictly increasing – this will be verified later on – the inverse function of G, denoted by G^{-1}, exists. Thus, voters simply use G^{-1} to infer the unobserved signal ξ_1 from the observed action g_1. We denote this inferred signal by $\hat{\xi}_1$. Since the belief G need not necessarily be correct, due to limited strategic sophistication, it need not be that $\hat{\xi}_1 = \xi_1$, in contrast to the standard case of full rationality in a Bayesian Nash equilibrium. Voter i expects that, conditional on the incumbent politician being incompetent, $\hat{\xi}_1 = \hat{x}_1^i + \zeta_1$. It follows that $\hat{\xi}_1$ serves as a signal that a voter i uses to update his beliefs about the realization of the politician's bias ζ_1, using his $posterior$ opinion about x_1^* (see Lemma 1).[12] We then have the following result.

Lemma 2 (Posterior opinion about $\boldsymbol{\zeta_1}$) *Voter i's posterior opinion $\hat{\zeta}_1^i$ about the realization of the incumbent's bias ζ_1 is normally distributed with mean $E\hat{\zeta}_1^i = \gamma\left(\hat{\xi}_1 - \hat{\mu}_1^i\right)$ and variance $\hat{\sigma}_\zeta^2 = \frac{\hat{\sigma}_x^2 \sigma_\zeta^2}{\hat{\sigma}_x^2 + \sigma_\zeta^2}$, where $\gamma \equiv \frac{\sigma_\zeta^2}{\hat{\sigma}_x^2 + \sigma_\zeta^2}$.*

The proof is almost identical to the one of Lemma 1 and is omitted. Note that in the case where $\hat{\xi}_1 = \hat{\mu}_1^i$, we have $E\hat{\zeta}_1^i = 0$. However, since ζ_1 is normally distributed, $\hat{\xi}_1 = \hat{\mu}_1^i$ occurs with probability zero. As a result, we generically have $E\hat{\zeta}_1^i \neq 0$.

We finally determine a voter's posterior opinion about the probability that the incumbent politician is competent. Conditional on the incumbent being competent, voter i treats $\hat{\xi}_1$ as

[11] In fact, it will turn out that G is a linear function.

[12] Note that it would not be rational to use the prior belief about x_1^* since this would mean neglecting useful information.

drawn from the distribution $\hat{x}_1^i$. Similarly, conditional on the incumbent being incompetent, the voter treats $\hat{\xi}_1$ as drawn from $\hat{x}_1^i + \hat{\zeta}_1$. Denote by f_c^i the density function associated with $\hat{x}_1^i$, and by f_{ic}^i the density function associated with $\hat{x}_1^i + \hat{\zeta}_1^i$. The subscripts c and ic stand for competent and incompetent, respectively. Using this notation, the posterior probability that the incumbent politician is competent is determined as follows.

Lemma 3 (Posterior $\boldsymbol{\alpha}$) *(i): Voter i's posterior opinion about the probability that the incumbent is competent is given by*

$$\hat{\alpha}^i = Pr[\text{competent} \left| \hat{\xi}_1 \right.] = \frac{\alpha f_c^i \left(\hat{\xi}_1 \right)}{\alpha f_c^i \left(\hat{\xi}_1 \right) + (1-\alpha) f_{ic}^i \left(\hat{\xi}_1 \right)}. \tag{3}$$

(ii): $\hat{\alpha}^i$ is a strictly decreasing function of $\left| \hat{\xi}_1 - \hat{\mu}_1^i \right|$.

To understand the logic of (3), suppose that $\hat{\xi}_1$ would be a discrete random variable. Then, $f_c^i \left(\hat{\xi}_1 \right)$ and $f_{ic}^i \left(\hat{\xi}_1 \right)$ would denote the probabilities that ξ_1 takes on its inferred value in the case of the competent and the incompetent politician, respectively. Thus, (3) would reflect a standard updating formula. Lemma 3 shows that the same logic applies if f_c and f_{ic} refer to continuous random variables, provided they are well-behaved as it is true for normal random variables.

Part (ii) of Lemma 3 will be used for deriving a politician's best-response choice of g_1 below. It shows that the larger the distance between a politician's signal $\hat{\xi}_1$, as inferred by the voter, and a voter's posterior opinion $\hat{\mu}_1^i$, the lower the probability that the voter assigns to the event that the incumbent politician is competent. This manifests how a voter judges the competence of an incumbent politician through the lens of his (posterior) opinion.

4.3 Voters' Reelection Decision

The politician in office in the second period is either the incumbent from the first period or a newly elected challenger. In either case, he sets $g_2 = \xi_2$ at stage *P2* of the game. This follows from our assumption that, conditional on being reelected, a politician's objective is to maximize welfare[13] and from Assumption 3.

We consider now a voter i's reelection decision at stage $V1$ of the game. A voter i considers his expected utility in case of reelection of the incumbent and compares it to the expected utility obtained in case of the election of a challenger. He casts his vote for the incumbent if and only if expected utility his higher under the incumbent than under the challenger. There is no strategic voting since there is a continuous population of voters.

Expected utility in the case of reelection of the incumbent is determined as follows. From the perspective of voter i, the incumbent is competent with probability $\hat{\alpha}^i$ (see Lemma 3). In the case of the competent incumbent, $g_2 = x_2^*$ since $\xi_2 = x_2^*$. Voter i does not observe x_2^* but substitutes his belief x_2^i. Using (2), $EV_2^i = -E\left(x_2^i - x_2^i - \varepsilon_2\right)^2 = -\sigma_\varepsilon^2$. In case that the incumbent is incompetent, $g_2 = \xi_2 = x_2^i + \hat{\zeta}_1^i$, according to Assumption 2 and Lemma 2. Using (2) again, we have $EV_2^i = -E\left(x_2^i + \hat{\zeta}_1^i - x_2^i - \varepsilon_2\right)^2 = -\left[\left(E\hat{\zeta}_1^i\right)^2 + \hat{\sigma}_\zeta^2 + \sigma_\varepsilon^2\right]$.[14] Overall, expected utility from reelecting the incumbent is given by

$$EV_2^i = -\sigma_\varepsilon^2 - \left(1 - \hat{\alpha}^i\right)\left[\left(E\hat{\zeta}_1^i\right)^2 + \hat{\sigma}_\zeta^2\right]. \tag{4}$$

The logic behind (4) is that the utility loss due to the variance of ε_1 is realized for both politician types. In contrast, the loss due to the fact that an incompetent politician's signal is noisy arises only with probability $1 - \hat{\alpha}^i$ from the perspective of voter i.

Expected utility from a challenger is determined very similarly. There $\hat{\alpha}^i$ has to be replaced by α and $\hat{\zeta}_1^i$ by ζ_2. Since $E\zeta_2 = 0$, we obtain, in analogy to (4),

$$EV_2^i = -\sigma_\varepsilon^2 - (1 - \alpha)\,\sigma_\zeta^2. \tag{5}$$

A voter reelects the incumbent if and only if expected utility as given by (4) exceeds expected utility as given by (5). Rearranging directly leads to the condition stated in the below lemma.

[13]The aim of this assumption is to simplify the analysis. We could obtain almost identical results if we were allowing for rent extraction in a way similar to Persson and Tabellini (2000, Ch. 4.5). To see this, suppose that there is an upper bound on the amount of rents that a politician can extract. Suppose further that a competent politician makes better use of the remaining government budget by better promoting welfare due to superior information (see Rogoff, 1990). In such a model, rent seeking would not affect a politician's choice of g_t.

[14]Here we have used the fact that, for any random variable z, $E\left(z^2\right) = (Ez)^2 + Var\,(z)$.

Lemma 4 (Reelection Decision) *Voter i reelects the incumbent if and only if*

$$\frac{1-\hat{\alpha}^i}{1-\alpha}\ \frac{\left(E\hat{\zeta}_1^i\right)^2+\hat{\sigma}_\zeta^2}{\sigma_\zeta^2}\le 1. \tag{6}$$

To understand condition (6), consider first the limit case in which a voter would not learn anything about the incumbent's bias ζ_1 from observing the incumbent's action at stage P1. This would be the case if the function G were constant for all levels of the signal ξ_1. In this case, $E\hat{\zeta}_1^i = E\zeta_1 = 0$ and $\hat{\sigma}_\zeta^2 = \sigma_\zeta^2$. Thus, condition (6) simplifies to $\hat{\alpha}^i \geq \alpha$. This means that a voter casts his ballot for the incumbent if and only if it is more likely that the incumbent is competent than that a challenger is competent.

If voters update their beliefs about an incumbent's bias ζ_1 by observing g_1 at stage P1, reelection of the incumbent is compatible with $\hat{\alpha}^i < \alpha$. Thus, it is possible that a voter prefers to reelect an incumbent politician even if he believes that the probability that the incumbent is competent is lower than the probability that a challenger would be competent. Lemma 4 shows that this requires that $\left(E\hat{\zeta}_1^i\right)^2 + \hat{\sigma}_\zeta^2$ is sufficiently smaller than σ_ζ^2. This is the case if $\left|\hat{\xi}_1 - \hat{\mu}_1^i\right|$ is small and g_1 provides a relatively sharp signal about the incumbent's bias ζ_1, such that $\hat{\sigma}_\zeta^2$ is small (see Lemma 2). In this case, reelecting the incumbent, a voter expects a relatively small variance associated with g_2 relative to the variance associated with g_2 when being set by a challenger. This comes at a benefit since voters are risk averse over g_2.

4.4 Politician Behavior

A politician's behavior at stage *P2* has already been discussed at the beginning of the last subsection. At stage *P1*, a politician chooses g_1 such that he maximizes the probability of getting reelected. This entails maximizing the probability that (6) holds for the median voter, that is the voter associated with the median of μ_1^i, denoted by μ_1^m. (See Appendix B for a discussion of sufficient conditions for the voter associated with μ_1^m being pivotal.)

From Lemma 3(ii), $\hat{\alpha}^m$ strictly decreases in $\left|\hat{\xi}_1 - \hat{\mu}_1^m\right|$. Furthermore, from Lemma 2, $\left(E\hat{\zeta}_1^m\right)^2$ strictly increases in $\left|\hat{\xi}_1 - \hat{\mu}_1^m\right|$. It follows that a politician maximizes the probability that (6) holds for the median voter by setting g_1 such that $\hat{\xi}_1 = E\left[\hat{\mu}_1^m \,|\xi_1\right]$ because this maximizes $\hat{\alpha}^m$ and minimizes $E\left(\hat{\zeta}_1^m\right)^2$. A politician determines $\hat{\xi}_1$ via his choice of g_1 as

$\hat{\xi}_1 = G^{-1}(g_1)$. Using this, it follows that the condition $\hat{\xi}_1 = E[\hat{\mu}_1^m | \xi_1]$ is equivalent to $g_1 = G(E[\hat{\mu}_1^m | \xi_1])$. Substituting for $E[\hat{\mu}_1^m | \xi_1]$ from Lemma 1 and using Assumption 3, we state this as follows.

Lemma 5 (Politician's Behavior) *The incumbent politician chooses* $g_1 = G((1-\beta)\mu_1^m + \beta\xi_1)$.

This lemma establishes that, given voters' belief G about a politician's behavioral strategy, setting $g_1 = G((1-\beta)\mu_1^m + \beta\xi_1)$ is a best response to this belief and to voters' reelection decision as characterized by (6).

4.5 Strategic Beliefs of Sophistication of Degree k

Lemma 4 shows voters' reelection strategy given a belief G about how a politician reacts to his signal ξ_1. Lemma 5 characterizes a politician's *actual* best use of the information revealed by ξ_1. This should be understood as a best response to voters' belief G. Under level-k beliefs, G relates to a politician's actual behavior. In particular, as will become clear below, voters' beliefs and a politician's actual behavior are mutually consistent up to one order. Thus, there is a fixed-point argument involved here that we are now in a position to address.

Level-k beliefs are defined recursively. We start with a *baseline belief* about how a politician sets g_1 as a function of ξ_1. This baseline belief corresponds, by definition, to $k = 0$. This baseline belief entails voters believing that a politician maximizes expected welfare, given his signal ξ_1. Denoting level-k beliefs by G_k, we have $G_0(\xi_1) = \xi_1$. This follows from (2) and Assumption 3.

Consider now a politician's best response to the belief G_0. Using Lemma 5, it follows that a politician maximizes the probability of getting reelected by choosing $g_1 = G_0[E[\hat{\mu}_1^m | \xi_1]] = (1-\beta)\mu_1^m + \beta\xi_1$. It follows that a politician's best response deviates from voters' belief.

The intuition for this fact is as follows. Suppose that the median voter believes that a politician chooses $g_1 = \xi_1$. This implies that, at stage *V1*, the median voter judges an incumbent competent if g_1 comes close to $\hat{\mu}_1^m$. In other words, the median voter judges competence through the lens of his *posterior* opinion about x_1^*, that is $\hat{\mu}_1^m$. A politician anticipating this has an incentive to set g_1 equal to his expectation of $\hat{\mu}_1^m$, that is $E[\hat{\mu}_1^m | \xi_1]$, rather than equal to ξ_1. Using Lemma 1 and Assumption 3, if follows that $E[\hat{\mu}_1^m | \xi_1] = (1-\beta)\mu_1^m + \beta\xi_1$.

By definition, a level-1 belief G_1 entails that voters anticipate a politician's incentive to deviate from G_0. Specifically, $G_1(\xi_1) := G_0((1-\beta)\mu_1^m + \beta\xi_1) = (1-\beta)\mu_1^m + \beta\xi_1$ (because $G_0(\xi_1) = \xi_1$). A politician's best response to this belief is again determined by Lemma 5 and we obtain $g_1 = G_1[E[\hat{\mu}_1^m|\xi_1]] = (1-\beta^2)\mu_1^m + \beta^2\xi_1 \neq G_1(\xi_1)$. Thus, a politician also has an incentive to deviate from level-1 beliefs. Proceeding with this recursion, we define sophistication-k beliefs as follows.

Definition 1 *Beliefs of strategic sophistication of level k are defined by the recursion*

$$G_k(\xi_1) = G_{k-1}((1-\beta)\mu_1^m + \beta\xi_1), \tag{7}$$

where $G_0(\xi_1) = \xi_1$.

Level-k beliefs, defined in this way, provide a special case of the cognitive hierarchy model of Camerer et al. (2004). Our main simplification is the assumption that all voters share the same level of k. In the original version of the model, several levels of k coexist and are distributed according to a Poisson distribution. Second, we assume that the strategic sophistication of politicians is higher than those of voters.[15]

The following Lemma provides a direct analytical expression for G_k and states a politician's best response to this belief. The proof is straightforward and is omitted.

Lemma 6 *Under beliefs of sophistication of degree k, voter i's belief about a politician's behavioral strategy is given by $G_k(\xi_1) = \left(1-\beta^k\right)\mu_1^m + \beta^k\xi_1$. A politician's best response to this belief is given by $g_1 = \left(1-\beta^{k+1}\right)\mu_1^m + \beta^{k+1}\xi_1$.*

Confirming our earlier claim, G_k is indeed continuous and strictly increasing, in fact linear, in ξ_1 for finite k. Hence, $\hat{\xi}_1 = G^{-1}(g_1)$ is well-defined for finite k. The solution of the political game for infinite k is obtained as a limit case.

Under level-k beliefs, voters' beliefs about a politician's behavioral strategy and a politician's actual behavioral strategy are mutually consistent up to one order. Both converge when k approaches infinity. As mentioned before, the evidence discussed in Camerer et al. (2004) suggests that k takes on a value of one or two, in practice.

[15]This assumption strikes us as natural. Technically, it does not matter whether we assume that politicians have a degree of strategic sophistication of $k+1$ (when voters' is k) or any number higher than this.

4.6 The Political Equilibrium

We start the discussion of the equilibrium with a definition of populism.

Definition 2 (Populism) *A politician's choice is populist if it does not only depend on his signal ξ_t but also on the prior belief of the median voter μ_t^m.*

Our main positive result, which characterizes the outcomes in an indirect democracy in the first office period, is the following.

Proposition 1 (Equilibrium First Period) *Suppose voters hold beliefs of degree of sophistication k. (i): If k is finite, then there exists a unique equilibrium in which*

$$g_1 = \left(1 - \beta^{k+1}\right) \mu_1^m + \beta^{k+1} x_1^* \tag{8}$$

in case of the competent politician and

$$g_1 = \left(1 - \beta^{k+1}\right) \mu_1^m + \beta^{k+1} \left(x_1^* + \zeta_1\right) \tag{9}$$

in case of the incompetent politician. (ii): If k is infinite, there exists a unique equilibrium that is obtained as a limit case for $k \to \infty$. This equilibrium is perfectly populist and both politician types set $g_1 = \mu_1^m$.

Proposition 1 shows that, for finite k, the prevailing equilibrium is *separating*. This means that: (i) a politician's choice depends on his signal ξ_1 and different values of the signal lead to different policy choices; (ii) both politician types choose different levels of g_1 with probability one. The difference in policy choices shrinks with a higher level of k. For the limit case of full rationality, that is an infinite k, the prevailing equilibrium is a *pooling* equilibrium. In particular, both politician types choose an identical policy action that does not depend on the signal ξ_1.

For finite k, g_1 is equal to a weighted average of the politician's signal about x_1^* and the median voter's prior belief about x_1^*. Remember that the signal of the competent politician is equal to x_1^* while the signal of the incompetent politician is equal to $x_1^* + \zeta_1$. Any equilibrium involves pandering to the median voter's belief and thus a populist policy choice.

The degree to which policy making is populist is the higher, the lower β^k. In the limit case where $k = \infty$, (part (ii) of Proposition 1), policy making is perfectly populist and neither politician type makes use of his signal. Thus, β^k can be understood as indicating the susceptibility to populism.

As $0 < \beta < 1$, the susceptibility to populism increases with k. Conversely, limited strategic sophistication of voters' beliefs prevents that the policy choice is perfectly populist. This result may seem rather surprising at first. It shows that the effect of the level of strategic sophistication of voter's beliefs on the equilibrium outcome is far from trivial. To understand the intuition of this result, remember that for $k = 0$, voters $expect$ the politician to behave in a non-strategic way and to maximize voters' utility. Thus, voters do not expect the politician to pander to the median voter. While a politician does have an incentive to pander to the opinion of the median voter, this incentive is limited by the fact that he is not expected to do so.

If k is larger than zero, voters $expect$ a competent incumbent to pander to the median voter (see Section 4.5). The effect of this is to strengthen the politician's incentive to pander to the median voter. The higher k, the higher the number of orders by which the politician's incentive to pander to the median voter's opinion is strengthened. In the limit case of an infinite k, the median voter's opinion is the only determinant of the politician's policy choice.

More formally, it can be shown that it is a logical impossibility that g_1 depends on ξ_1 if k is infinite. To see this, suppose that g_1 would indeed depend on ξ_1. Then voters would be aware of this in equilibrium. They also understand that $\xi_1 = x_1^*$ in the case of the competent politician, but they do not observe x_1^* and, thus, voter i substitutes $\hat{\mu}_1^i$ for x_1^*. It follows that a politician who wants to appear competent to the median voter will not actually want to let his policy depend on ξ_1 but rather on $E\left[\hat{\mu}_1^m \left|\xi_1\right.\right] = (1-\beta)\,\mu_1^m + \beta\xi_1$ (see Lemma 1 and Assumption 3). Here ξ_1 enters only with a weight β, which lies between zero and one.

For an infinite k, voters are aware of this incentive not to let g_1 be a function of ξ_1 but only of $(1-\beta)\,\mu_1^m + \beta\xi_1$. This expression, however, still depends on ξ_1. Thus, the same argument as above can be repeated and we find that g_1 can in fact only depend on $(1-\beta^2)\,\mu_1^m + \beta^2\xi_1$ etc. This argument can be iterated an infinite number of times. Because $0 < \beta < 1$, β^k decreases in k. Hence, ξ_1 must necessarily vanish and g_1 cannot depend on ξ_1. If k is finite, in contrast, this argument can be repeated only a finite number of times, which increases in k. With each iteration, g_1 depends less on ξ_1 and more on μ_1^m.

Apart from k, a crucial determinant of the weight β of a politician's signal for his policy choice is σ_ε^2 (see Lemma 1). If σ_ε^2 is low, β is close to one and populism vanishes. To understand this, recall that voters receive the signal $x_1^* + \varepsilon_1$ before making their election decision. They use this signal to judge an incumbent politician's competence. If σ_ε^2 is very low, voters observe the ex ante optimal policy level, that is x_1^*, almost perfectly, and they know it. As a result of getting a very precise signal, the median voter's prior μ_1^m has very little influence on his posterior belief about x_1^*. A politician's aim is to be judged competent through the lens of the median voter's $posterior$ belief. If this posterior belief depends only very little on the prior μ_1^m, a politician's incentive to pander to the median voter's prior belief is comparatively low. As a result, the policy is to a larger degree determined by the politician's signal, which he uses to predict the median voter's posterior belief. In the opposite case, where σ_ε^2 is large, β is comparatively low. Thus, voters' opinions are highly persistent and μ_1^m has a high weight in influencing policy. Overall, it is important to note, however, that σ_ε^2 (and hence β) affect the equilibrium policy choice only if k is finite, that is under limited strategic sophistication of voters' beliefs.

We conclude this section by summarizing the equilibrium outcome in the second period. The result follows directly from the discussion at the beginning of Section 4.3.

Proposition 2 (Equilibrium Second Period) *In the second office period, we have $g_2 = x_2^*$ in case of a competent politician. Furthermore, $g_2 = x_2^* + \zeta_1$ in case of a re-elected incompetent politician and $g_2 = x_2^* + \zeta_2$ in case of an incompetent challenger.*

4.7 Bounded Rationality and Welfare

Using Proposition 1 and 2, it is straightforward to characterize welfare in case of our baseline model that corresponds to an indirect democracy. We do so by using the concept of a loss function defined as $L_t = EV_t^{FB} - EV_t^{EQ}$ for period t. L_t is defined as the difference between expected utility as achieved when g_t is set to its ex ante welfare-maximizing level x_t^* and expected utility as achieved in the equilibrium of the political game. The first-best utility value ex ante results if $g_t = x_t^*$ (see (2)), which leads to $EV_t^{FB} = -\sigma_\varepsilon^2$. We obtain:

Proposition 3 (Welfare Indirect Democracy) *Under indirect democracy, welfare is* *characterized by*

$$L_1^{ID} = \left(1-\beta^{k+1}\right)^2 (x_1^* - \mu_1^m)^2 + \beta^{2(k+1)} (1-\alpha) \sigma_\zeta^2. \quad (10)$$

$$L_2^{ID} = [1-\alpha-\Delta_\alpha] \sigma_\zeta^2, \quad (11)$$

where $\Delta_\alpha \geq 0$.

Consider the first period. The welfare loss from indirect democracy is equal to a weighted average of the distortion $x_1^* - \mu_1$ associated with the median voter's belief and the variance of the incompetent politician's bias. The first term arises from pandering. The second term arises from the fact that no equilibrium entails full pandering for finite k. In this case, politicians will always partially base their policy choice upon their signal ξ_1. Since the signal of the incompetent politician is noisy, the fact that g_1 depends on this signal increases the variance of g_1. This comes at a cost to risk averse voters. It is interesting to note that the weights $\left(1-\beta^{k+1}\right)^2$ and $\beta^{2(k+1)}$ do not add to one if and only if k is finite. We will come back to this in the next section (see Proposition 6).

In the second period, pandering does not arise since no politician has an incentive to manipulate voters' perception of his competence. As a result, only the noise term σ_ζ^2 contributes to the welfare loss. It can be checked that $\Delta_\alpha \geq 0$ follows from the fact that the probability that a competent politician holds office in the second period exceeds α (see the proof of Proposition 3). Thus, the probability that a competent politician holds office in the second period is higher than that a competent politician holds office in the first period. This confirms a well-known result that elections help to mitigate an adverse selection problem.

An important insight from Proposition 3 is that welfare may be higher under limited strategic rationality than under perfect rationality. For an infinite k, no politician would make his policy choice dependent on his signal, as we have shown in Proposition 1. In the case of the competent politician, and also in the case of the incompetent politician if σ_ζ^2 is sufficiently small, it is desirable that he puts some weight on his signal.[16] Since this only happens for finite k, welfare is always higher for some finite k then for an infinite k. This can be seen by minimizing (10)

[16] The welfare maximizing weight trades off the gain in information through the signal against the welfare loss that arises from the fact that the variance of g_1 increases because the signal of the incompetent politician is noisy.

over β^{k+1}. The minimizing value is given by

$$\beta^{k+1} = \frac{(x_1^* - \mu_1^m)^2}{(x_1^* - \mu_1^m)^2 + (1-\alpha)\,\sigma_\zeta^2} \tag{12}$$

The expression on the right-hand side may take on values between zero and one. It follows from (12) that full rationality, that is an infinite k, does never minimize (10), except for the uninteresting case where σ_ζ^2 is infinite. Overall, our finding is that the level of k crucially affects the trade-off associated with the costs and benefits of populism.

5 Comparing Constitutional Regimes

We turn to a comparison of welfare under our baseline case of indirect democracy to the case of direct democracy and to non-accountable agents. These three institutions can be ordered in terms of the degree to which decision making is delegated from voters to their agents. Decision making can either be delegated to completely *independent agents* such as experts; it can be delegated to politicians who want to get reelected and, thus, are only partially independent (*indirect democracy*); or it may not be delegated at all (*direct democracy*).[17]

We first consider direct democracy. We follow Maskin and Tirole (2004) by modeling direct democracy as a political institution where $g_t = \mu_t^m$, that is it is the median voter who directly chooses g_t. The idea is that in a direct democracy voters have the right to ask for referenda and that this would lead to a strong link between policy making and the opinion of the median voter.[18] In this simple benchmark model of direct democracy the median voter is the only relevant actor and there are no strategic elements involved. The following proposition follows directly from inserting g_t into (2) and taking expectations.

Proposition 4 (Welfare under Direct Democracy) *Under direct democracy,* $L_t^{DD} = (x_t^* - \mu_t^m)^2$.

The loss function is again defined as the deviation of expected utility from its first-best level. Comparing welfare under direct democracy to the case of indirect democracy in Proposition

[17] Relating to Maskin and Tirole (2004), the higher the degree of delegation of decision making the lower the *accountability* of decision makers in a constitutional regime.

[18] In New Zealand, Switzerland, and some U.S. states, a referendum can be initiated by voters by means of a citizen petition.

3, we see that both are equivalent for the case of full rationality, that is an infinite k. In both cases there is no regard for a politician's signal in equilibrium. Thus, in both cases there is no effective role for politicians and their advisor teams.

Now we are turning to delegation of policy making to independent agents such as judges, bureaucrats, or experts. In order to facilitate the comparison to indirect democracy, our subsequent assumptions parallel those about politicians. Henceforth, we will refer to independent agents as $agents$, for brevity. Exactly as in the case of politicians, we assume that the agents receive a signal, denoted by ξ_t^a, about x_t^*. For the competent agent we have $\xi_t^a = x_t^*$ whereas, for the incompetent agent, we have $\xi_t^a = x_t^* + \nu_t$.[19] The random variable ν_t reflects a noise term with an expected value of zero and a variance σ_ν^2. The probability that an agent is competent is π. By definition, agents are non-accountable to voters, that is they determine policies in both periods and cannot be ousted after the first period.

In one important aspect, our assumptions about agents deviate from the assumptions made about politicians. We assume that independent agents are fully benevolent. Thus, they set $g_t = \xi_t^a$. We make this assumption as the case of benevolent agents could easily be considered as an ideal, if unfeasible, benchmark for government. Here we are interested in the question under which conditions this ideal benchmark would actually be desirable in a world characterized by imperfect knowledge and limited strategic sophistication of beliefs.

The welfare loss under the independent agent regime is given in the following proposition. The proof is very similar to the one of Proposition 3 and is omitted.

Proposition 5 (Welfare Independent Agents) *In the case of independent agents,* $L_t^{IND} = (1 - \pi)\,\sigma_\nu^2$.

Comparing the outcomes for the three constitutional regimes in the first period, indirect democracy can be understood as a mix of direct democracy and governance by independent agents. To see this, denote ξ_1^p the signal of a politician and assume that $\xi_1^a = \xi_1^p \equiv \xi$. Furthermore, assume that the likelihood that an expert or a politician is incompetent is equal, that is $\alpha = \pi$. Then $g_1^{IND} = \xi$, where IND stands for independent agents. Furthermore,

[19] In the interpretation of independent agents as "experts", the notion "incompetent" may sound rather odd at first. What we have in mind is that if experts disagree, at most one expert opinion can be right. Thus, experts may be wrong even if they are highly trained. Combating crime provides one salient example where experts disagree substantially (see Levitt, 1998, and Buscaglia, 2008), climate change provides another one (see McKibbin and Wilcoxen, 2002; Weitzman, 2007; Stern 2008).

$g_1^{DD} = \mu_1^m$, where DD refers to direct democracy. From Proposition 1 it follows that $g_1^{ID} = \left(1 - \beta^{k+1}\right) g_1^{DD} + \beta^{k+1} g_1^{IND}$, where ID refers to indirect democracy.

This weighted-average nature of indirect democracy makes it attractive to risk averse voters in the sense that $L_1\left(g_1^{ID}\right) < \left(1 - \beta^{k+1}\right) L_1\left(g_1^{DD}\right) + \beta^{k+1} L_1\left(g_1^{IND}\right)$, for finite k. This follows from the fact that L_1 is strictly convex. The fact that the loss associated with g_1^{ID} is lower than a weighted average of the losses associated with either g_1^{DD} or g_1^{IND} is also the reason why the weights associated with the two terms in L_1^{ID} in Proposition 3, namely $\left(1 - \beta^{k+1}\right)^2$ and $\beta^{2(k+1)}$, add to less than one for finite k. We summarize this finding as follows.

Proposition 6 (Comparative Advantage of Indirect Democracy) *Suppose that* $\alpha = \pi$ *and* $\xi_1^a = \xi_1^p$. *If* k *is finite, then* $g_1^{ID} = \left(1 - \beta^{k+1}\right) g_1^{DD} + \beta^{k+1} g_1^{IND}$ *and* $L_1\left(g_1^{ID}\right) < \left(1 - \beta^{k+1}\right) L_1\left(g_1^{DD}\right) + \beta^{k+1} L_1\left(g_1^{IND}\right)$.

The comparative advantage of indirect democracy identified here crucially relies on bounded rationality, that is a finite k. For perfect strategic rationality, where k is infinite, the value of the loss function is identical in the case of direct and indirect democracy. Hence, the comparative advantage disappears. In this sense, we find that bounded rationality makes indirect democracy relatively more attractive compared to the case of full strategic rationality.

Looking beyond level-k beliefs, a comparison of the loss functions in Propositions 3, 4, and 5 shows that the elements that crucially affect which constitutional regime is optimal are: the distortion associated with the median voter's belief, $|\mu_t^m - x_t^*|$; the variance of the incompetent politician's bias, σ_ζ^2; and the corresponding variance of the incompetent agent's bias, σ_ν^2. In the following corollary we point out the comparative statics.

Corollary 1 (Constitutional Comparison) *(i): Independent agents are optimal if* σ_ν^2 *is small relative to* $|\mu_t^m - x_t^*|$ *and* σ_ζ^2. *(ii): Direct democracy is optimal if* $|\mu_t^m - x_t^*|$ *is small relative to* σ_ζ^2 *and* σ_ν^2. *(iii): In case that neither of these conditions applies, the weighted-average nature of indirect democracy may make it optimal.*

6 Conclusion

In this paper we have analyzed politicians' behavior when both voters and politicians are imperfectly informed about how a policy affects voters' welfare. A key novel aspect of our analysis is taking into account voters' limited strategic sophistication when forming beliefs about politicians' behavior. The sophistication of these beliefs is implied by the order k of what we have dubbed level-k beliefs.

Taking voters' strategic sophistication into account leads to a number of interesting and novel insights. In particular, limited strategic sophistication of voters $weakens$ politicians' incentive to pander to public opinion because politicians know that voters $expect$ them to pander less than fully. Thus, pandering is limited even if politicians are exclusively office-motivated and do not care about voters' welfare per se. We have shown that this feature increases the attractiveness of indirect democracy relative to direct democracy or delegation of policy making to non-accountable agents and indirect democracy may often be preferable to the latter two.

In future research, it would be interesting to investigate how the optimal term length of politicians depends on the speed with which voters obtain feedback about the ex-post optimal policy choice. Second, it would be important to explore which forms of indirect democracy may be most desirable, either presidential or parliamentary, majoritarian or proportional. Third, it may be of interest to consider the roles of the media and of education policy in shaping voters' opinions and beliefs. Specifically, introducing heterogeneous levels of sophistication and political awareness across voters and modeling the first as a function of education policy and the second as a function of actions of the media could lead to interesting insights. Finally, the setup of this paper may also apply to decision making in corporations and to principal-agent relationships more generally.

Appendices

A Proofs

Proof of Lemma 1

Voters observe $x_1^* + \varepsilon_1$. From the point of view of voter i, $x_1^* + \varepsilon_1$ is a realization of the random variable $x_1^i + \varepsilon_1$. The voter aims to update his belief about x_1^i. The random variables x_1^i and $x_1^i + \varepsilon_1$ are jointly normally distributed with $E\left[x_1^i\right] = \mu_1^i$, $Var\left[x_1^i\right] = \sigma_x^2$, $E\left[x_1^i + \varepsilon_1\right] = \mu_1^i$, $Var\left[x_1^i + \varepsilon_1\right] = \sigma_x^2 + \sigma_\varepsilon^2$. Furthermore, $Cov\left[x_1^i, x_1^i + \varepsilon_1\right] = \sigma_x^2$. Inserting this in the formulas for conditional expectations and variances for jointly-normal random variables (see e.g. Hogg and Craig, 1995, p. 148) yields the result.

Proof of Lemma 3

Proof of (i). We omit the time subscript as well as the superscript i when there is no danger of confusion. Let, for clarity, ξ denote the random variable whose realization is $\hat{\xi}$. The idea of the proof is to derive the posterior probability $\hat{\alpha}$ for the case that the random variable ξ falls into the (small) interval $I_\delta\left(\hat{\xi}\right) := \left[\hat{\xi} - \delta,\ \hat{\xi} + \delta\right]$ and to consider the limit $\delta \to 0$.

Denote by C the event that a politician is competent and by IC the complementary event. Using the definition of conditional probabilities, it follows that

$$P\left(C\left|\xi \in I_\delta\right.\right) = \frac{\alpha P\left(\xi \in I_\delta\left|C\right.\right)}{\alpha P\left(\xi \in I_\delta\left|C\right.\right) + (1-\alpha)\, P\left(\xi \in I_\delta\left|IC\right.\right)}. \tag{A.1}$$

Note that the denominator is equal to $P\left(\xi \in I_\delta\right)$. In order to consider the limit of (A.1) for the case where $\delta \to 0$, it is useful to rewrite it as

$$P\left(C\left|\xi \in I_\delta\right.\right) = \left[1 + \frac{1-\alpha}{\alpha}\frac{\int_{\hat{\xi}-\delta}^{\hat{\xi}+\delta} f_{ic}\left(\xi\right) d\xi}{\int_{\hat{\xi}-\delta}^{\hat{\xi}+\delta} f_{c}\left(\xi\right) d\xi}\right]^{-1}, \tag{A.2}$$

where the two normal densities f_c and f_{ic} are defined as in the main text preceding Lemma 3. Since normal densities are well-behaved, it follows from standard arguments using the definition of the Riemann integral that $\lim_{\delta\to 0} \frac{\int_{\hat{\xi}-\delta}^{\hat{\xi}+\delta} f_{ic}(\xi) d\xi}{\int_{\hat{\xi}-\delta}^{\hat{\xi}+\delta} f_{c}(\xi) d\xi} = \frac{f_{ic}\left(\hat{\xi}\right)}{f_c\left(\hat{\xi}\right)}$. Substituting this into (A.2) for $\delta \to 0$

and rearranging yields (3).

Proof of (ii). We consider the case that $\hat{\xi} - \hat{\mu}_1 \geq 0$. (Similar arguments apply to the symmetric case $\hat{\xi} - \hat{\mu}_1 < 0$.) Write $\hat{\alpha} = \left[1 + \frac{1-\alpha}{\alpha} \frac{f_{ic}(\hat{\xi})}{f_c(\hat{\xi})}\right]^{-1}$. The posterior $\hat{\mu}_1$ is taken as given here. Therefore, it is sufficient to show that $\frac{f_{ic}(\hat{\xi})}{f_c(\hat{\xi})}$ increases with $\hat{\xi}$. f_c is the normal density describing the distribution of $\xi_c \equiv \hat{x}_1$, while f_{ic} is the normal density associated with $\xi_{ic} \equiv \hat{x}_1 + \hat{\zeta}_1$. Using the formula for the normal density, we have $f_c(\xi) = \frac{1}{\sqrt{2\pi Var(\xi_c)}} \exp\left[-\frac{(\xi - E\xi_c)^2}{2Var(\xi_c)}\right]$ and $f_{ic}(\xi) = \frac{1}{\sqrt{2\pi Var(\xi_{ic})}} \exp\left[-\frac{(\xi - E\xi_{ic})^2}{2Var(\xi_{ic})}\right]$. It then follows that

$$d\left[\frac{f_{ic}(\xi)}{f_c(\xi)}\right]/d\xi = \frac{\sqrt{Var(\xi_c)}}{\sqrt{Var(\xi_{ic})}} \left[\frac{\xi - E\xi_c}{Var(\xi_c)} - \frac{\xi - E\xi_{ic}}{Var(\xi_{ic})}\right] \exp\left[\frac{(\xi - E\xi_c)^2}{2Var(\xi_c)} - \frac{(\xi - E\xi_{ic})^2}{2Var(\xi_{ic})}\right].$$

We are interested in the case that $\xi = \hat{\xi} \geq \hat{\mu}_1$. Lemma 2 implies then that $\hat{\xi} \geq E\xi_{ic} \geq E\xi_c = \hat{\mu}_1$ (note that $E\xi_{ic} = (1-\gamma)\hat{\mu}_1 + \gamma\hat{\xi}_1$). Hence, $\hat{\xi} - E\xi_c \geq \hat{\xi} - E\xi_{ic}$. Furthermore, $Var(\xi_{ic}) > Var(\xi_c)$. Thus $\left[\frac{\xi - E\xi_c}{Var(\xi_c)} - \frac{\xi - E\xi_{ic}}{Var(\xi_{ic})}\right] > 0$ and hence $d\left[\frac{f_{ic}(\hat{\xi})}{f_c(\hat{\xi})}\right]/d\hat{\xi} > 0$.

The case where $\hat{\xi} < \hat{\mu}_1$ is symmetric and analyzed by following the same steps.

Proof of Proposition 1

Proof of part (i). This follows directly from Lemma 5 and 6.

Proof of part (ii). Clearly, part (i) implies that there is a unique limit for $k \to \infty$. This limit is indeed an equilibrium for appropriate off-equilibrium beliefs about a politician's type if he deviates from setting $g_1 = \mu_1^m$. For instance, consider the belief, identical across voters, that a politician setting $g_1 \neq \mu_1^m$ is incompetent with probability one and that $E\hat{\zeta}_1^i$ is sufficiently high. Given this off-equilibrium belief, setting $g_1 = \mu_1^m$ maximizes a politician's probability of getting reelected from Lemma 4. Thus, a politician does indeed not want to deviate from $g_1 = \mu_1^m$. Furthermore, voters do not have any incentive deviate from (6).

Proof of Proposition 3

Denote by $g_{t,c}$ the level of g set in period t by the competent politician and let $g_{t,ic}$ refer to the incompetent politician. Let λ_t denote the probability that a politician is competent in period t.

Then

$$EV_t = -\left[\lambda_t E\left[(g_{t,c} - x_t^* - \varepsilon_t)^2\right] + (1-\lambda_t)\, E\left[(g_{t,ic} - x_t^* - \varepsilon_t)^2\right]\right]. \quad \text{(A.3)}$$

Consider the first period. Clearly, $\lambda_1 = \alpha$. Using this an inserting for $g_{1,c}$, $g_{1,ic}$ from Proposition 1 into (A.3), we obtain

$$EV_1^{EQ} = -\left(1-\beta^{k+1}\right)^2 (\mu_1^m - x_1^*)^2 - (1-\alpha)\,\beta^{2(k+1)}\sigma_\zeta^2 - \sigma_\varepsilon^2.$$

EV_1 is maximized for $g_1 = x_1^*$ which yields $EV_1^{FB} = -\sigma_\varepsilon^2$. Inserting this and the above expression into the definition of L_t yields (10).

We turn next to the second period. We show first that it is more likely that a competent politician gets reelected than that an incompetent politician gets reelected. To establish this, we show that the probability that (6) holds for the median voter is lower for an incompetent incumbent than for a competent incumbent. We first prove this for finite k. By Lemma 3, $\hat{\alpha}^m$ is a strictly decreasing function of $\left|\hat{\xi}_1 - \hat{\mu}_1^m\right|$. By Lemma 6, the median voter's belief is that $g_1 = G_k(\xi_1) = \left(1-\beta^k\right)\mu_1^m + \beta^k \xi_1$. Hence, $\hat{\xi}_1 = \frac{g_1}{\beta^k} - \frac{1-\beta^k}{\beta^k}\mu_1^m$. By Proposition 1, $g_1 = \left(1-\beta^{k+1}\right)\mu_1^m + \beta^{k+1}\xi_1$. In the case of the competent politician, $\xi_1 = x_1^*$. Inserting this into the expression for g_1, then inserting g_1 into the expression for $\hat{\xi}_1$ and using Lemma 1 yields that $\hat{\xi}_1 - \hat{\mu}_1^m = -\beta\varepsilon_1 \equiv \phi_c$. Similarly, it follows for the incompetent politician that $\hat{\xi}_1 - \hat{\mu}_1^m = \beta(\zeta_1 - \varepsilon_1) \equiv \phi_{ic}$. The variance of ϕ_{ic} is equal to $\beta^2\left(\sigma_\zeta^2 + \sigma_\varepsilon^2\right)$, whereas the variance of ϕ_c is equal to $\beta^2\sigma_\zeta^2$ and thus strictly smaller than the variance of ϕ_{ic}. It follows that the probability that $\left|\hat{\xi}_1 - \hat{\mu}_1^m\right| \geq A$, for any $A \in (0,\infty)$, is strictly greater for the incompetent than for the competent politician. Hence, the probability that $\hat{\alpha}^m \leq B$, for any $B \in (0,1)$, is strictly greater for the incompetent than for the competent politician.

From Lemma 2, $E\hat{\zeta}_1^m = \gamma\left(\hat{\xi}_1 - \hat{\mu}_1^m\right)$. Thus, the above arguments also imply that the probability that $\left(E\hat{\zeta}_1^m\right)^2 \geq C$, for any $C \in (0,\infty)$, is strictly greater for the incompetent incumbent than for the competent incumbent. Since σ_ζ^2 does not differ across types, this establishes that the probability that (6) holds for the median voter is strictly smaller in case of an incompetent incumbent than in case of a competent incumbent for finite k.

If k is infinite, $\hat{\xi}_1$ cannot be inferred and no information about the incumbent's type is observed since we have a pooling equilibrium. Hence, the probability of getting reelected must be equal for both types. Overall, we have shown that it is more likely that a competent politician gets reelected than that an incompetent incumbent gets reelected.

There are three events in which the politician in the second period is competent: (1) A competent incumbent gets reelected; (2) a competent incumbent gets ousted and replaced by a competent politician; (3) an incompetent incumbent gets ousted and replaced by a competent politician. Denote the probability that a competent politician gets reelected by ρ_c and the probability that an incompetent politician gets reelected as ρ_{ic}. Denote the event that the second period politician is competent by C_2. We have then

$$\begin{aligned} Pr\left[C_2\right] &= \alpha\rho_c + \alpha^2\left(1-\rho_c\right) + \alpha\left(1-\alpha\right)\left(1-\rho_{ic}\right) \\ &= \alpha\left[1+\left(1-\alpha\right)\left(\rho_c-\rho_{ic}\right)\right] \geq \alpha. \end{aligned} \tag{A.4}$$

The last inequality follows from the fact that it is more likely that a competent politician gets reelected than that an incompetent politician gets reelected, as shown above. Overall, we have now established that $\lambda_2 = \alpha + \Delta_\alpha$ for some $\Delta_\alpha \geq 0$ (see (A.3)).

Inserting the expressions for $g_{2,c}$, $g_{2,ic}$ given in Proposition 2 into (A.3) yields

$$EV_2 = -\left(1-\alpha-\Delta_\alpha\right)\sigma_\zeta^2 - \sigma_\varepsilon^2.$$

Again, $EV_2^{FB} = -\sigma_\varepsilon^2$. Inserting this and the above expression into the definition of L_t yields (11).

B Sufficient Conditions for the Median Voter Theorem to Apply

A voter i is pivotal in our model if and only if the following two conditions hold: (i) If the pivotal voter casts his ballot for the incumbent then at least half of voters prefer the incumbent; (ii) if the the pivotal voter casts his ballot for the challenger, then at least 50 percent of voters prefer the challenger. If the two conditions hold for the i associated with the median of μ_1^i (that is μ_1^m), we say that the median voter theorem applies.

We first show that a voter i reelects the incumbent if and only if $\hat{\xi}_1$ falls within a finite interval. The reelection decision is determined by (6). The posterior probability $\hat{\alpha}^i$ is a strictly decreasing function of $\left|\hat{\xi}_1 - \hat{\mu}_1^i\right|$as stated in Lemma 3($ii$). Furthermore, $\left(E\hat{\zeta}_1^i\right)^2$ is a strictly increasing function of $\left|\hat{\xi}_1 - \hat{\mu}_1^i\right|$ from Lemma 2. Since $\hat{\alpha}^i$ enters negatively on the left-hand side of (6) while $\left(E\hat{\zeta}_1^i\right)^2$ enters positively, it follows that (6) holds if and only if $\left|\hat{\xi}_1 - \hat{\mu}_1^i\right|$ is sufficiently small. Using Lemma 1, we can state that voter i reelects the incumbent if and only if

$$\hat{\xi}_1 \in \left[(1-\beta)\,\mu_1^i + \beta\,(x_1^* + \varepsilon_1) \pm \delta^{crit}\right] \equiv I_{re}^i, \tag{B.1}$$

where δ^{crit} is a strictly positive real number that is common across voters and is a function of σ_x^2, σ_ε^2, σ_ζ^2. We dub I_{re}^i voter i's reelection interval.

The median voter theorem applies if (i) $\hat{\xi}_1 \in I_{re}^m$ implies that $\hat{\xi}_1 \in I_{re}^i$ for at least half of voters and, conversely, (ii) $\hat{\xi}_1 \notin I_{re}^m$ implies that $\hat{\xi}_1 \notin I_{re}^i$ for at least half of voters. The following lemma states two sufficient conditions for this to hold.

Lemma 7 (Median Voter) *The median voter is pivotal if either of the following conditions are fulfilled:*

(i) $\mu_1^i = \mu_1^m$ *for at least half of voters;*

(ii) $\mu_1^{\max} - \mu_1^{\min} \leq 2\delta^{crit}/\,(1-\beta)$

where $\mu_1^{\min} \equiv \min\{\mu_1^i\}$ *and* $\mu_1^{\max} \equiv \max\{\mu_1^i\}$.

Proof. *Proof of (i).* If $\mu_1^i = \mu_1^m$ for at least half of voters then it follows immediately that $\hat{\xi}_1 \in I_{re}^m$ implies that $\hat{\xi}_1 \in I_{re}^i$ for at least half of voters. Conversely, $\hat{\xi}_1 \notin I_{re}^m$ implies that $\hat{\xi}_1 \notin I_{re}^i$ for at least half of voters.

Proof of (ii). Using Lemma 1 and (B.1), it can be checked that the condition $\mu_1^{\max} - \mu_1^{\min} \leq 2\delta^{crit}/\,(1-\beta)$ is equivalent to $\max I_{re}^{\min} \geq \min I_{re}^{\max}$, where the superscripts min and max refer to the i with the lowest and highest μ_1^i, respectively. If this holds then $\hat{\xi}_1 \in I_{re}^m$ implies that either $\hat{\xi}_1 \in I_{re}^{\max}$ or $\hat{\xi}_1 \in I_{re}^{\min}$. The reason is that every point in I_{re}^m is contained in $I_{re}^{\min}$ or $I_{re}^{\max}$ since I_{re}^m is situated between the latter two und the latter two overlap in at least one point. Furthermore, whenever $g \in I_{re}^m$ and $g \in I_{re}^{\min}$, then also $g \in I_{re}^i$ for all i with $\mu_1^{min} \leq \mu_1^i \leq \mu_1^m$. Similarly, whenever $g \in I_{re}^m$ and $g \in I_{re}^{\max}$, then also $g \in I_{re}^i$ for all i with $\mu_1^m \leq \mu_1^i \leq \mu_1^{\max}$.

Hence, in either case, if $\hat{\xi}_1 \in I_{re}^m$ then $\hat{\xi}_1$ belongs to the reelection interval of at least half of voters.

Consider now the case that $\hat{\xi}_1 \notin I_{re}^m$. It follows that either $\hat{\xi}_1 \notin I_{re}^i$ for all i with $\mu_1^{min} \leq \mu_1^i \leq \mu_1^m$ or for all i with $\mu_1^m \leq \mu_1^i \leq \mu_1^{\max}$. In either case, $\hat{\xi}_1 \notin I_{re}^i$ holds for at least 50 percent of voters. Overall, this establishes that m is pivotal. ■

Obviously, according to part (i), the voter associated with μ_1^m is pivotal if a majority of voters share his belief. The second condition in Lemma 7 limits the range of μ_1^i. In particular, it requires, that the upper limit of the reelection interval for the voter with the lowest μ_1^i is at least as great as the lower limit of the reelection interval for the voter with the highest μ_1^i. It is possible to derive further sufficient conditions for the median voter theorem to apply.

C Indirect Democracy for Biased ζ_1

As mentioned in the main text, politicians may be understood as a representative sample of the general population if the political selection process is not biased in favor of the elite or any other particular group. If politicians are representative for the general population, then we would expect that the incompetent politician's signal is related to the distribution of voters' beliefs. (For the competent politician, this does not apply, since he observes the truth.) A simple way of capturing this is assuming that $E\zeta_t$ is related to $\mu_t^m - x_t^*$.

In the following, we consider the limit case where, from an objective point of view, $E\zeta_t = \mu_t^m - x_t^*$. By objective we mean from the point of view of the economic theorist analyzing the problem. In contrast, we need to assume that a politician believes that $E\zeta_t = 0$ for himself. Otherwise, he could make use of the information about $E\zeta_t$ to unbias his belief about x_t^*. Second, we also assume that voters believe that $E\zeta_t = 0$. More precisely, we assume here that a majority of voters hold beliefs that are identical to the beliefs of the median voter. In this case, it is indeed appropriate to assume that a majority of voters believe that $E\zeta_t = 0$. Otherwise, their beliefs about $E\zeta_t$ would be inconsistent with their own beliefs about x_t^*.

In this case, all positive results in Section 4 continue to hold. However, the welfare expressions in Proposition 3 are modified. In particular, we obtain

$$L_1 = \left[\alpha\left(1-\beta^{k+1}\right)^2 + 1 - \alpha\right](x_1^* - \mu_1^m)^2 + \beta^{2(k+1)}(1-\alpha)\sigma_\zeta^2. \tag{C.1}$$

$$L_2 = [1-\alpha(1+\Delta_\alpha)]\left[(x_2^* - \mu_2^m)^2 + \sigma_\zeta^2\right], \tag{C.2}$$

As to be expected, the terms $(x_t^* - \mu_t^m)^2$ have a stronger influence on the welfare loss than in the baseline case. In particular, in the case of Proposition 3, the coefficient for $(x_1^* - \mu_1^m)^2$ is smaller than in the case of (C.1). The expression $(x_2^* - \mu_2^m)^2$ does not appear at all in Proposition 3. The overall conclusion is that if the incompetent politician's signal is biased towards the beliefs that are prevalent among voters, indirect democracy becomes more similar to direct democracy as defined in Section 5.

D Indirect Democracy in the Case of Non-Degenerate Beliefs of Politicians

Our analysis is based on the assumption that politicians receive a signal ξ_t which they use as a "point estimate" for x_t^* in the sense of classical statistics. This introduces an asymmetry between voters and politicians since the former are Bayesian and hold non-degenerate prior beliefs x_t^i.

It is straightforward to turn a politician into a Bayesian in our framework by assuming that his prior belief about x_t^* is given by $x_t^p \sim N(\xi_t, \sigma_x^2)$, where the superscript p indexes a politician. This would affect the analysis insofar as an incumbent politician would be able to update the probability that he is competent. More important, an incumbent would partially learn about his bias ζ_1 (under the hypothesis that he is incompetent). An incumbent politician's updating would parallel Lemma 2 and 3.

If an incumbent politician learns about his bias ζ_1, Assumption 2 implies that he can partially "unbias" his signal ξ_2 in period 2. This leads to a further incumbent advantage from the perspective of voters since this reduces the expected level of bias in the second period arising from an incompetent politician if this politician is a reelected incumbent. Since voters would take this into account, the reelection condition (6) would get somewhat more complicated. Loosely speaking, the fact that an incumbent can partially unbias his signal reduces the second term on the left hand side of (6). In spite of this modification, the logic of our main results remains entirely valid.

References

Alesina, Alberto, and Tabellini, Guido (2007), "Bureaucrats or Politicians? Part I: A Single Policy Task," *American Economic Review*, 97(1), 169-179.

Alesina, Alberto, and Tabellini, Guido (2008), "Bureaucrats or Politicians? Part II: Multiple Policy Tasks," *Journal of Public Economics*, 92, 426-447.

Atkinson, Anthony B., and Stiglitz, Joseph E. (1980), *Lectures on Public Economics,* McGraw-Hill, Maidenhead, England.

Besley, Timothy (2006), *Principled Agents? The Political Economy of Good Government*, Oxford University Press, Oxford.

Besley, Timothy, and Case, Anne (1995), "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits," *Quarterly Journal of Economics,* 110(3), 769-798.

Blanes i Vidal, Jordi, and Leaver, Clare (2008), "Pandering Judges," Working Paper.

Buscaglia, Edgardo (2008), "The Paradox of Expected Punishment: Legal and Economic Factors Determining Success and Failure in the Fight against Organized Crime," *Review of Law & Economics*, 4(1), 290-317.

Canes-Wrone, Brandice, Herron, Michael C., and Shotts, Kenneth W. (2001), "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science,* 45(3), 532-550.

Camerer, Colin F., Ho, Teck-Hua, and Chong, Juin-Kuan (2004), "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119(3), 861-898.

Caplan, Bryan (2001), "What Makes People Think like Economists? Evidence on Economic Cognition from the 'Survey of Americans and Economists on the Economy'," *Journal of Law and Economics*, 44, 395-426.

Dixit, Avinash K., and Weibull, Jörgen W. (2006), "Political Polarization," Working Paper.

Hogg, Robert V., and Craig, Allen T. (1995), *Introduction to Mathematical Statistics,* 5th edition, Prentice Hall: New Jersey.

Holmström, Bengt (1999), "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169-182.

Levitt, Steven D. (1998), " Juvenile Crime and Punishment," *Journal of Political Economy*, 106(6), 1156-1185.

Maskin, Eric, and Tirole, Jean (2004), "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034-1054.

McKibbin, Warwick J., and Peter J. Wilcoxen (2002), "The Role of Economics in Climate Change Policy," *Journal of Economic Perspectives,* 16(2), 107-129.

Persson, Torsten, and Guido Tabellini (2000), *Political Economics: Explaining Economic Policy,* MIT Press, Cambridge, Mass.

Prat, Andrea (2005), "The Wrong Kind of Transparency," *American Economic Review,* 95(3), 862-877.

Rogoff, Kenneth (1990), "Equilibrium Political Budget Cycles," *American Economic Review,* 80(1), 21-36.

Schultz, Christian (2008), "Information, Polarization and Term Length in Democracy," *Journal of Public Economics,* 92, 1078-1091.

Stern, Nicholas (2008), "The Economics of Climate Change," *American Economic Review: Papers & Proceedings*, 98(2), 1-37.

Weitzman, Martin L. (2007), "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature,* 45(3), 703-724.