



Democracy, populism, and (un)bounded rationality

Johannes Binswanger ^{a,*}, Jens Prüfer ^b

^a Department of Economics, CentER, Netspar, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands

^b Department of Economics, CentER, Tilec, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands

ARTICLE INFO

Article history:

Received 21 December 2010

Received in revised form 7 December 2011

Accepted 16 February 2012

Available online 24 February 2012

JEL classification:

D72

D78

D83

Keywords:

Strategic sophistication

Democracy

Populism

Sophistication-*k* beliefs

Accountability

Experts

ABSTRACT

In this paper we aim to understand how bounded rationality affects performance of democratic institutions. We consider policy choice in a representative democracy when voters do not fully anticipate a politician's strategic behavior to manipulate his reelection chances. We find that this limited strategic sophistication affects policy choice in a fundamental way. Under perfect sophistication, a politician does not make any use of his private information but completely panders to voters' opinions. In contrast, under limited sophistication, a politician makes some use of private information and panders only partially. Limited sophistication crucially determines how welfare under representative democracy compares to welfare under alternative political institutions such as direct democracy or governance by experts. We find that, under limited strategic sophistication, representative democracy is preferable to the other institutions from an ex ante perspective.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

How does bounded rationality affect the performance of democratic institutions? We analyze policy choice in a representative democracy when voters have limited cognitive abilities to anticipate a politician's strategic behavior. In particular, voters may not fully acknowledge a politician's incentive to bias a policy choice towards their opinions, that is to engage in populism. A politician's incentive to bias a policy choice in this way originates from the fact that it may make voters believe that he is competent and thus improve his chances to get reelected. We compare political outcomes and the resulting level of welfare under representative democracy with direct democracy and a form of government where decision making is delegated to experts. We show that the welfare-ranking of the three institutions is crucially affected by bounded rationality.

The existing political economy literature that aims to understand the working of democratic institutions is almost exclusively based on the assumption of perfect rationality (Besley, 2006). As put by Besley (2006, p. 172) "going forward it would be interesting to understand better what the differences are between behavioral models of politics and the postulates of strict rationality." In this paper we study how democratic institutions work under bounded rationality. We thus contribute to a small but growing literature on how deviations from perfect rationality shape policies.

We focus on bounded rationality in the form of limited strategic sophistication. In particular, our analysis is based on a model of limited strategic thinking, dubbed *k*-thinking. This model was introduced by Stahl and Wilson (1994, 1995) and Nagel (1995). Since then, a sizeable literature has developed that explores *k*-thinking theoretically and empirically, including Ho et al. (1998),

* Corresponding author. Tel.: +31 13 466 2920; fax: +31 13 466 3042.

E-mail address: j.binswanger@uvt.nl (J. Binswanger).

Costa-Gomes et al. (2001), Crawford (2003), Camerer et al. (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberry (2007a), and Goldfarb and Yang (2009), among others. The literature has found strong experimental support for k -thinking.

As a result of their empirical success, models of k -thinking have been used to study the performance and design of institutions when agents are boundedly rational rather than infinitely sophisticated, as assumed in standard analysis. For instance, Camerer et al. (2004) discuss speculation and price setting from the perspective of k -thinking. This is important for the analysis of financial market regulation and central bank policy. Crawford and Iriberry (2007b) and Crawford et al. (2009) study the performance of auctions, an important institution for an efficient allocation of resources. By focusing on the political domain, our paper contributes to studying the performance of social institutions under level- k thinking.

Apart from bounded rationality, the framework we consider is a fairly standard one. Our framework is reminiscent, for instance, of Maskin and Tirole (2004). In particular, a representative voter is endowed with an opinion about which policy maximizes his expected utility. We will refer to the representative voter simply as “the voter”. Importantly, the voter's opinion may be wrong. There is a politician who observes the voter's opinion. Moreover, the politician receives a signal indicating which policy maximizes the voter's utility from an ex-ante point of view. 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. As in career concern models (for instance, Holmström, 1999; Prat, 2005), 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. The voter's aim is to (re)elect the politician whom he believes to be most competent.

In order to get reelected, a politician has an incentive to pander to the voter's opinion. This incentive to pander may potentially be mitigated by the voter receiving a (noisy) signal about which policy has been optimal before the election takes place. The voter uses this signal to judge the incumbent politician's competence. Crucially, however, the voter's judgment is also (rationally) influenced by his prior opinion about which policy is optimal. This induces an incentive for a politician to pander to the voter's opinion because it will make him look more competent in the voter's eyes and increase the chance of getting reelected.

Beliefs of limited strategic sophistication refer to the voter's limited ability to anticipate the politician's strategic incentives to pander to the voter's opinion. Since this inability may concern different orders of strategic behavior, these beliefs will be defined in a recursive way. We refer to them as sophistication- k beliefs or, simply, k -beliefs.

The recursive definition of k -beliefs works as follows. If a voter is fully naive, he has a 0-belief. This means that he believes that a politician chooses a policy in a way to maximize the voter's welfare. This belief reflects the case where the voter takes the constitutional role of politicians in a democracy, as managers of the state on behalf of the people, literally. As we will show, under this naive belief, a politician faces an incentive not to fully maximize the voter's welfare but to partially distort the policy towards the voter's opinion. If the voter has a 1-belief, he anticipates the politician's incentive to deviate from the voter's 0-belief. This, however, induces a second-order incentive for the politician to deviate from the voter's 1-beliefs. If the voter also anticipates the politician's second-order incentive to deviate, he is endowed with a 2-belief etc. Iterating forward in this way, we finally end up with perfect rationality as the limit case of an ∞ -belief.

We take the voter's level of k as given. The empirical evidence suggests values for k of one or two (Camerer et al., 2004; Crawford and Iriberry, 2007b). 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 behavioral literature, or rather infinitely high turns out to crucially matter for our results. First, we find that the higher the voter's degree of sophistication, the *stronger* the incentive of a politician to pander. The intuition for this result is that a more sophisticated voter *expects* the politician to pander more. As a result, the politician will indeed pander more. Under perfect sophistication (or rationality), the voter expects the politician to fully pander and the politician does so, in turn. This entails that the politician will not make use of any private information. Overall, the politician's incentive to pander is crucially affected by limited strategic sophistication. Thus, a representative democracy yields different policy results under bounded rationality, compared to the case of perfect rationality.

Under a broad range of circumstances, the fact that the politician will not make use of any private information under perfect rationality is deteriorating for welfare. Under these circumstances, welfare is higher under limited strategic sophistication than under perfect rationality. This challenges the common wisdom that deviations from full rationality are mostly detrimental for welfare.

Our analysis proceeds with comparing the desirability of representative democracy in comparison to other political institutions in the case of bounded rationality. In particular, we compare representative democracy to direct democracy and to the case where policy making is fully delegated to independent experts. We show that the welfare-ranking of the three institutions is crucially affected by beliefs of limited strategic sophistication. We also show that, from the ex ante perspective of a constitutional designer, beliefs of limited strategic sophistication give representative democracy an edge over direct democracy and over delegation to experts.

The remainder of this paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces a model of representative democracy. In Section 4, we formally introduce beliefs of limited strategic sophistication and solve the model. In Section 5, we compare representative democracy to the case of direct democracy and to the case of delegation of policy making to independent experts. We discuss supporting evidence as well as the robustness of our findings in Section 6 and conclude in Section 7. Proofs are contained in the Appendix A if not stated otherwise.

An earlier version of this paper (Binswanger and Prüfer, 2009) presents results for the model under more general assumptions. It also contains several extensions, showing that the results from the baseline setup presented here are fairly robust.

2. Related theoretical literature

In this section, we discuss the related theoretical literature, deferring the discussion of empirical issues until Section 6. 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'. The crucial difference between our analysis and Maskin and Tirole's work is that our analysis is based on k -beliefs. Canes-Wrone et al. (2001) consider a setup where both voters and politicians are imperfectly informed about which of two policy options is optimal. Again, our analysis differs because of our focus on bounded rationality. Furthermore, Canes-Wrone et al. do not compare the welfare properties of different political institutions. Dixit and Weibull (2007) study the possibility of polarization in voters' beliefs in a setup that is reminiscent of ours, but assuming that voters are perfectly rational. Alesina and Tabellini (2007) analyze the optimal allocation of responsibilities between politicians and bureaucrats under perfect rationality.

Several existing papers explore deviations from standard preferences or beliefs on political outcomes. Hillman (2010) analyzes the role of "expressive behavior." This means behavior that is adopted in order to create a positive self-image or to please others, in a way that runs against one's own material self-interest. People may engage in this type of behavior since it spends "expressive utility". Hillman shows how expressive behavior may lead to disadvantageous policies that no one would choose if responsible for the outcome. Bénabou and Tirole (2006) explore how people may be motivated to distort their beliefs about the importance of hard work for an advantageous economic career. These beliefs, in turn, determine the size of the welfare state and may lead to multiple equilibria. Sieg (2001) explores whether patterns of political business cycles that arise under rational expectations may also emerge if agents are boundedly rational adaptive learners. He finds that rationality is not a necessary prerequisite for political business cycles.

3. A model of representative democracy

3.1. The political game: voters, politicians, and timing

3.1.1. Voters

We consider an economy populated by a representative voter, whom we simply refer to as "the voter".¹ There are two periods (referring to two office periods of politicians). In each period, the voter's utility is determined by

$$V = -(g - x^* - \varepsilon)^2. \quad (1)$$

The variable $g \in \mathbf{R}$ denotes a policy action that is set by the office-holding politician. Neglecting ε , the utility maximizing level of g is given by $x^* \in \mathbf{R}$. A crucial assumption in our framework is that x^* is unobserved. For technical reasons, we assume that x^* is drawn at the beginning of each period by nature from a normal distribution with mean Ex^* and variance σ_x^2 . The mean may vary across periods and is unknown. The variable ε is a normally distributed random variable with an expected value of zero and a variance of σ_ε^2 . We assume that ε is identically and independently distributed over time and independent of all other random variables in the model. As is the case for x^* , ε is also unobserved. The distribution of ε is common knowledge.

As we will discuss in more detail below, nature first draws x^* , before ε is realized. The policy action g is to be set *after* x^* has been determined but *before* ε is realized. Thus, x^* determines the ex ante optimal policy. It specifies how, from an ex ante point of view, a choice of g translates into voters' utility. In contrast, ε represents a short-term shock to x^* and determines the ex post optimal level of g . Whereas x^* and ε are not observed in isolation, the voter observes the sum $x^* + \varepsilon$ after g has been set. This allows him to learn, although imperfectly, about x^* .

From an ex-ante perspective, the voter's utility in a period is given by the expected value of V , that is by

$$EV = -E[(g - x^* - \varepsilon)^2]. \quad (2)$$

There are two essential features of Eq. (1) or (2). First, x^* determines a unique interior optimum for g from an ex ante point of view. Second, there is risk aversion with respect to g .

As already stated, x^* is not observed and Ex^* is unknown. However, the voter has a *prior belief* about x^* . Specifically, his prior belief about x^* is given by x , which is a normally distributed random variable with mean μ and variance σ_x^2 . To clearly distinguish x and μ from k -beliefs introduced below, we shall speak of the former as the voter's *opinion* about x^* .

3.1.2. Politicians

The policy action g is chosen and implemented by an incumbent politician. A politician observes the voter's prior opinion μ (and σ_x^2) about x^* . The politician does not observe x^* directly. Rather, he receives a signal ξ that is informative about x^* . For

¹ See Binswanger and Prüfer (2009) for the case of voters with heterogeneous beliefs.

simplicity, we assume that the politician does not have any prior about x^* before receiving the signal ξ . Thus, the signal ξ acts as the politician's prior.

There are two politician types that we dub competent and incompetent, respectively. The ex ante probability that a politician is competent is denoted by α and is common knowledge. The competent politician obtains a prior $\xi = x^*$ that corresponds to the truth. An incompetent politician obtains a prior $\xi = x^* + \zeta$, where ζ is a normally distributed random variable with mean zero and variance σ_ζ^2 . We assume that ζ is independent of all other random variables in the model and that it is independently drawn in periods one and two.² Furthermore, the distribution of ζ is common knowledge. As it is common in the literature on career concerns (see Holmström, 1999, or Prat, 2005), we assume that a politician does not observe his type. Since ξ acts as a politician's prior, we have $E[x^*|\xi] = \xi$ for both politician types.

We adopt an assumption of Alesina and Tabellini (2007) and Glaeser et al. (2005), that the politician has lexicographic preferences. In particular, a politician cares first about reelection. Conditional on being reelected, he maximizes welfare of the voter. The second objective reflects that the politician's private economic preferences are the same as the voter's as in a "citizen-candidate" model (Osborne and Slivinski, 1996; Besley and Coate, 1997). However, the politician's primary goal is to stay in office.

3.1.3. Timing of the political game

For clarity, we introduce a time index $t = 1, 2$ that refers to first and second office periods, respectively. The timing of the game is as follows.

- Office period one:
 - Stage 1: Nature draws x_1^* and determines the type of the incumbent politician and his prior ξ_1 .
 - Stage 2: The incumbent politician chooses the policy g_1 , which becomes public knowledge.
 - Stage 3: Nature draws ε_1 and sends the signal $x_1^* + \varepsilon_1$ to the voter.
 - Stage 4: The voter decides whether to reelect or to oust the incumbent politician.
- Office period two:
 - Stage 1: Nature draws x_2^* and determines the prior ξ_2 of the politician in office.
 - Stage 2: The politician chooses g_2 .
 - Stage 3: Nature draws ε_2 .

Note that the voter has only one move in the entire game: he decides whether to cast his vote for the incumbent politician or for a challenger. We solve the game by backward induction for a *sophistication- k equilibrium*, a solution concept that we define below and that boils down to a special case of a perfect Bayesian equilibrium for $k = \infty$.³

3.2. Discussion of the model

In this subsection, we comment on key assumptions made above. Elections and representative democracy serve various purposes. First, elections are a mechanism of preference aggregation (Downs, 1957; Riker, 1982). Second, elections help to discipline and monitor politicians and thus alleviate problems of moral hazard (Barro, 1973; Ferejohn, 1986). Third, elections allow for selecting political leaders according to their ability to choose appropriate policies (Rogoff, 1990; Canes-Wrone et al., 2001). In our analysis, we focus on the latter aspect in an environment of limited strategic sophistication. In the interest of transparency and to keep the analysis tractable, we exclude the first two aspects by assuming that voters are homogeneous and that there is no scope for moral hazard.

Concerning a politician's information, we have assumed that the signal ξ directly acts as the politician's *prior* belief. Assuming, instead, that the politician would have a prior before obtaining ξ would simply complicate the analysis. In that case, the politician's best information about x^* when choosing a policy would be a weighted average of his prior and the signal ξ . This would mean carrying around more notation without leading to additional interesting insights. Thus, we directly take ξ as the politician's prior and refer to it as such. An immediate consequence of this is that, upon learning ξ , a rational politician will not apply Bayes' rule but base his actions directly on ξ .

In our model, a politician gets a signal about the ex ante optimal policy while the voter does not. In this sense, a politician is better informed than the voter (although the signal is noisy in the case of the incompetent politician). That politicians get information that voters do not get reflects that politicians have access to various bodies of advisers. Moreover, politicians have a much stronger incentive to be well informed about policies than voters since a single voter has only a negligible influence on policy choice (Maskin and Tirole, 2004), whereas a politician's career depends crucially on his information. Finally, voters may also behave "expressively" (Hillman, 2010), which has an effect similar to being prone to distorted or missing information (although "expressive" behavior may not be seen as irrational).

² In Binswanger and Prüfer (2009), we consider the case that an incumbent politician who stays in office in the second period keeps the draw of ζ from the first period for the second one.

³ As is usual in dynamic games of incomplete information, multiple perfect Bayesian equilibria exist in this game, i.e. multiple equilibria exist in the case of *unlimited* strategic sophistication. Since, in this paper, our focus is on the role of limited strategic sophistication, we do not provide a fully-fledged analysis of the complete set of perfect Bayesian equilibria that arise under unlimited sophistication. Interestingly, though, the sophistication- k equilibrium provides a natural equilibrium selection criterion among standard perfect Bayesian equilibria under unlimited sophistication: there exists a perfect Bayesian equilibrium under unlimited strategic sophistication that provides the limit case where boundedly rational agents would converge to if going through an infinite k -thinking/learning process. See also footnote 11.

We view the assumption that a politician does not observe his own *type* as natural since it is generally very difficult to provide objective evidence that a certain political platform is “wrong” (i.e. far off the welfare maximizing one). Thus, there is very limited information a politician may rely on to determine his *type*. Alternatively, one may take this assumption as reflecting that politicians are overconfident. Thus, every politician may view himself as competent and believe that he got the perfectly revealing signal himself (while other politicians may get noisy ones).⁴

As an example of a situation the model may apply to, consider the allocation of funds for combating crime. Specifically, suppose that there is a given budget to be spent for combating crime. 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 (for instance, prison infrastructures). In this example, x^* 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 corresponds to a short-term shock to the “threat of crime” and may originate from a sudden rise in youth unemployment, a sudden increase in immigration or the like.⁵

4. Analysis of representative democracy

4.1. Preliminary considerations

We start the analysis by establishing two preliminary results, namely about the voter's reelection decision and about the voter's posterior opinion about x_1^* after having obtained the signal $x_1^* + \varepsilon_1$. Equipped with these two preliminary results, we will then turn to the decisive element of our analysis: bounded rationality in the form of beliefs of limited strategic sophistication.

To study the reelection decision, we first need to consider a politician's choice in the second office period. A politician's primary goal is to maximize the chance of staying in office. However, in the second period, there is no further office period ahead. Therefore, the politician's objective is to maximize voters' expected welfare (Eq.(2)).⁶ As a result, the politician sets $g_2 = \xi_2$. Hence, in case of a competent politician, $g_2 = x_2^*$ and, using Eq. (2), $EV_2 = -\sigma_\varepsilon^2$. In case of an incompetent politician, $g_2 = x_2^* + \xi_2$ and $EV_2 = -\sigma_\varepsilon^2 - \sigma_\xi^2$. Thus, expected utility is higher in the case of the competent politician.

This observation directly determines the voter's reelection decision. The voter wants to maximize the probability that the politician in the second period is competent. Thus, a voter reelects an incumbent if and only if the posterior probability that the incumbent is competent, *after having observed* g_1 and $x_1^* + \varepsilon_1$, is higher than the prior probability that an unknown challenger is competent, given the voter's belief about the strategy of a competent politician. Denote the voter's posterior probability that an incumbent is competent by $\hat{\alpha}$. Then the voter reelects the incumbent if and only if $\hat{\alpha} > \alpha$.⁷

We now turn to the second preliminary piece of analysis, namely the voter's posterior opinion about x_1^* after having obtained the signal $x_1^* + \varepsilon_1$. It is characterized by the following lemma.

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

It is noteworthy that we have assumed that voters are perfectly rational when it comes to Bayesian updating. While it may be interesting to explore deviations from Bayesian updating, we abstain from this by focusing on the implications of limited strategic sophistication in an otherwise unperturbed environment. We do so for the sake of transparency.

4.2. Beliefs of limited strategic sophistication

Now we turn to the first office period. Because almost all variables that appear below refer to the first period – and because the results referring to the first period are the ones of main interest – we find it convenient to simplify notation and drop the time index when there is no danger of confusion. Thus, if not explicitly stated otherwise, μ refers to μ_1 etc.

For the reelection decision, the voter infers the posterior likelihood $\hat{\alpha}$ that the incumbent politician is competent (see previous subsection). To do so, he essentially has two pieces of information: first, he observes the politician's choice of g ; second, obtaining nature's signal about x^* , he forms a posterior opinion $\hat{\mu}$ (as determined by Lemma 1).

The voter uses g and $\hat{\mu}$ to infer the politician's prior ξ . For this, the voter must realize that g is a *function* of ξ . In other words the voter understands that the politician uses ξ when choosing g . Since the politician's prior *as inferred by the voter* need not a priori be identical with the prior that the politician actually has obtained, we denote the *inferred* prior by $\hat{\xi}$.

We assume that the voter believes that his posterior opinion $\hat{\mu}$ is the best available or trustworthy information about x^* , i.e. he assumes that g does not provide any information about x^* beyond $\hat{\mu}$. This assumption reflects that the voter is overconfident about

⁴ For evidence on overconfidence among executives, see Malmendier and Tate (2005).

⁵ We assume that x^* 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.

⁶ As mentioned before, this is best understood as the politician promoting his own private economic well-being as in a “citizen-candidate” model.

⁷ For an explicit expression of $\hat{\alpha}$ see Binswanger and Prüfer (2009) where we analyze a more involved version of the model presented here. Explicitly determining $\hat{\alpha}$ is not relevant for the following arguments.

his own opinion.⁸ Overconfidence has been found to be a very common property of human judgment (Alpert and Raiffa, 1982; Biais et al., 2005; Malmendier and Tate, 2005). In their survey of behavioral finance, De Bondt and Thaler (1995) call overconfidence “perhaps the most robust finding in the psychology of judgment.”

Due to this assumption, the voter will judge the politician's competence as follows. Having inferred $\hat{\xi}$, the voter will conclude that it is likely that the incumbent is competent and reelect him if $\hat{\xi}$ comes close to $\hat{\mu}$. If the distance between $\hat{\mu}$ and the inferred $\hat{\xi}$ is large, the voter will conclude that it is unlikely that the incumbent is competent and hence oust him. In other words, the voter rewards conformity of politicians, in the sense that he rewards $\hat{\xi}$ being close to $\hat{\mu}$.

It is crucial to realize that, at this point, it is an open issue how the voter infers $\hat{\xi}$ from g . It is a fundamental observation that the voter can only infer $\hat{\xi}$ if he has a belief about how the politician uses his prior ξ to choose g . Formally, we may state that the voter has a belief that the politician chooses g according to a function G , such that $g = G(\xi)$. Observing g , he thus infers $\hat{\xi}$ as the inverse of the G function (we neglect existence of the inverse of G for the moment). The k -beliefs of limited strategic sophistication represent a specification of the belief function G .

4.2.1. Strategic belief of sophistication of level zero

In models of sophistication- k thinking (or k -thinking, for short), the corresponding k -beliefs are defined in a recursive way. By construction, the recursion starts with an initial belief that is fully non-strategic, i.e. naive (Camerer et al., 2004; Crawford and Iriberry, 2007b). The only belief of the voter that qualifies as fully non-strategic and naive is the belief that the politician maximizes the voter's welfare. Every other belief contains some strategic element and can thus not qualify as an initial belief.

The naive $k = 0$ -belief takes the constitutional role of politicians in a democracy, as managers of the state on behalf of the people, literally. It is important to notice, though, that initializing k -beliefs as non-strategic/naive does not mean that any individual has these naive beliefs. Rather the iteration of k -beliefs reflects some learning process that starts with a naive belief and develops from there (Crawford and Iriberry, 2007b). Evidence from experiments points to the fact that most individuals iterate by one or two levels (Camerer et al., 2004; Crawford and Iriberry, 2007b).

According to 0-beliefs, the politician sets $g = \xi$. This maximizes Eq. (2) given that the politician's best information about x^* is ξ . Making use of the G function introduced above, we write 0-beliefs as $G_0(\xi) \equiv \xi$ (that is G_0 is the identity function).

A strategic politician plays best response to the voter's belief. The politician anticipates the logic of the voter's reelection decision. In particular, the closer the $\hat{\xi}$ that the voter infers from g comes to the voter's posterior opinion about x^* , $\hat{\mu}$, the higher the voter's estimated likelihood that the incumbent politician is competent. In other words, the lower the distance

$$|\hat{\xi} - \hat{\mu}|,$$

the higher $\hat{\alpha}$.

However, the politician does not know $\hat{\mu}$ when choosing g because $\hat{\mu}$ depends on $x^* + \varepsilon$ (see Lemma 1). The politician's only relevant piece of information is his prior ξ . Thus, the politician maximizes the expected chance of getting reelected by minimizing

$$|g - E[\hat{\mu}|\xi]|.$$

Knowing that $\hat{\mu} = (1 - \beta)\mu + \beta(x^* + \varepsilon)$ from Lemma 1 (and given that $E[x^*|\xi] = \xi$ and $E[\varepsilon|\xi] = 0$), the politician sets

$$g = E[\hat{\mu}|\xi] = (1 - \beta)\mu + \beta\xi. \tag{3}$$

Thus, if the voter holds the 0-belief $g = G_0(\xi) = \xi$, the politician deviates from this belief. He finds it in his interest to distort the policy choice towards the voter's prior opinion, μ , about x^* . The “weight” of this distortion is given by $1 - \beta$ under 0-beliefs.

4.2.2. Strategic belief of sophistication of level k

A voter's 1-belief is defined as a best response to the best response of a politician to a voter's 0-belief. Thus,

$$G_1(\xi) \equiv (1 - \beta)\mu + \beta\xi. \tag{4}$$

We consider again a strategic politician's best response to this belief. Since the voter believes that $g = G_1(\xi)$ and since the voter infers $\hat{\xi}$ as the inverse of G_1 , we have, using Eq. (4),

$$\hat{\xi} = \frac{1}{\beta}g - \frac{1 - \beta}{\beta}\mu. \tag{5}$$

⁸ Technically speaking, we assume that the posterior $\hat{\mu}$ is only a function of μ and $x^* + \varepsilon$ (see Lemma 1). Without overconfidence, the voter would also embrace that g potentially reveals some additional information about x^* . As a result, the posterior $\hat{\mu}$ would be a function of μ , $x^* + \varepsilon$ and g . Since g depends on the politician's best estimate of $\hat{\mu}$, this introduces an intricate fixed point argument that would complicate the analysis significantly. However, it would still lead to the same qualitative results as in the current model. Specifically, there would still be pandering because $\hat{\mu}$ would still depend on μ , if only to a lesser extent. Apart from this technical consideration, we perceive the assumption of overconfident voters, who do not acknowledge that the policy implemented by a politician may reveal information about the true state of the world, as very realistic.

Exactly in the same way as under 0-beliefs, the incumbent wants to minimize the distance $|\hat{\xi} - \hat{\mu}|$ in expectations. He does so by setting g such that $\hat{\xi} = E[\hat{\mu}|\xi]$. Using Eq. (5), Lemma 1, and the fact that $E[x^* + \varepsilon|\xi] = \xi$, we obtain that

$$\frac{1}{\beta}g - \frac{1-\beta}{\beta}\mu = (1-\beta)\mu + \beta\xi.$$

Rearranging, we obtain

$$g = (1-\beta^2)\mu + \beta^2\xi. \quad (6)$$

We infer that, under 1-beliefs, the weight of the politician's pandering to the voter's opinion is $1-\beta^2$. Under 0-beliefs it is $1-\beta$ (see Eq. (3)). Since β lies between zero and one (see Lemma 1), the politician's distortion of the policy choice towards the opinion of the voter is *greater* under the 1-belief than under the 0-belief. The intuition for this is the following. Under the 1-belief, a voter already *expects* the politician to pander to his belief by a weight of $1-\beta$ (see Eq. (4)). The politician takes this into account and this drives him to pander according to the voter's expectations and even somewhat more (a pandering effect of second order), to maximize the chances of getting reelected.

It is now straightforward to continue the recursion that leads to a general belief of strategic sophistication of degree k . Suppose that the voter holds a belief of sophistication of level k that has already been defined (for instance, for $k=0$ or $k=1$). Formally, this means that the voter believes $g = G_k(\xi)$. A strategic politician knows that the voter infers $\hat{\xi} = G_k^{-1}(g)$.⁹ From arguments made above, the politician's best response is to set g such that $\hat{\xi} = E[\hat{\mu}|\xi]$. This means that $G_k^{-1}(g) = (1-\beta)\mu + \beta\xi$. Starting with the initial belief $G_0(\xi) = \xi$ it is straightforward to determine the expressions for G_1 (as we have done above), G_2 , G_3 etc. The result is given in the following lemma (the calculations are straightforward and the proof is omitted).

Lemma 2. *The voter's k -belief about the politician's strategy is given by $G_k(\xi) = (1-\beta^k)\mu + \beta^k\xi$.*

We summarize a politician's best response to the voter's k -belief in the following lemma.

Lemma 3. *A politician's best response to the voter's k -belief $g = G_k(\xi) = (1-\beta^k)\mu + \beta^k\xi$ is given by $g = (1-\beta^{k+1})\mu + \beta^{k+1}\xi$.*

In Section 6, we show that it is actually not crucial for our results whether the voter's k -type is lower or higher than the politician's. However, as we discuss there, we perceive the assumption made above, that politicians can outguess voters and not vice versa, as a natural benchmark. The reason is that politicians are likely to be better trained in strategic thinking than voters because this is crucial for their career. For the remainder of this section and in Section 5, we will continue assuming that the exogenously given levels of sophistication are k for the voter and $k+1$ for the politician.

4.2.3. Strategic belief of infinite sophistication

If we continue iterating level- k beliefs to the limit of infinity, we end up with the benchmark of perfect strategic sophistication, that is perfect rationality. This belief would then support a perfect Bayesian equilibrium of the political game (see below). The intriguing finding is that under perfect rationality, the voter *expects* the politician to perfectly pander to his prior opinion (see Lemma 2). As a result, the politician will do so. Thus, in the perfect rationality case, the politician will not make any use of private information when choosing a policy but will only adhere to the voter's opinion.

4.3. The equilibrium of the political game

Before stating the equilibrium outcome, we wish to define our solution concept of the political game, which we dub sophistication- k equilibrium. We provide a definition that is fairly informal and targeted to the particular game that we consider here.

Definition 1. (Sophistication- k equilibrium) A sophistication- k equilibrium is a strategy combination and a set of beliefs about the state of nature and about the behavior of the other player, such that at each node of the game between a level- k player (the voter) and a level- $k+1$ player (the politician):

1. The strategies for the remainder of the game are Nash given the beliefs and strategies of the other player;
2. The level- k player holds a k -belief about the behavior of the other player;
3. The $k+1$ player anticipates the belief of the level- k player;
4. The beliefs about the state of nature are rational and determined by Bayes' law.

As we discuss in Section 6, the politician's level of sophistication could be higher than $k+1$ without changing the results quantitatively. We also show that for other finite realizations of the voter's and the politician's k -types (which can be identical or different), our results hold qualitatively.

Before we state our main positive result, we also wish to define what we mean by populism.

⁹ It can be checked, using Lemma 2, that the existence of the inverse function G_k^{-1} is actually assured as long as k is finite.

Definition 2. (Populism) A politician's choice, g , is populist if it does not only depend on his own information, ξ , but also on the prior opinion of the voter, μ . The extent of populism depends on the weight that a politician puts on μ .

Proposition 1. (Sophistication- k equilibrium) (i): Under limited strategic sophistication, that is if k is finite, there exists a unique sophistication- k equilibrium. In particular, the policy outcome for the first period is

$$g = (1 - \beta^{k+1})\mu + \beta^{k+1}x^* \quad (7)$$

in case of the competent politician and

$$g = (1 - \beta^{k+1})\mu + \beta^{k+1}(x^* + \zeta) \quad (8)$$

in case of the incompetent politician. (ii): Under perfect rationality, that is if k is infinite, there exists a unique equilibrium that is obtained as a limit case for $k \rightarrow \infty$. This equilibrium is perfectly populist and both politician types set $g = \mu$.

Proposition 1 shows that, under limited strategic sophistication, the policy outcome is a weighted average of the politician's own information and the voter's opinion. Thus, while the policy choice is always populist, it is only partially so since the politician does make use of his information.¹⁰ The higher the degree of sophistication k , the higher the weight a politician puts on the voter's opinion as opposed to his own information. Because of this, the difference between the policy choice of the two politician types shrinks with a higher k . When k reaches infinity, the policy choice is perfectly populist. The outcome is the same as if the voter had chosen the policy himself. This corresponds to the institution of direct democracy (see the following section).

It is noteworthy that the perfect Bayesian equilibrium that is obtained in the limit case of an infinite k , requires an appropriate off-equilibrium belief that supports this equilibrium. Suppose the voter believes that if the politician chooses $g \neq \mu$, the probability that the incumbent is competent is lower than α . In this case, the voter would oust the politician. Hence, the politician has no incentive to deviate from the equilibrium stated in Proposition 1 (ii).

Since dynamic games of incomplete information often have a large number of perfect Bayesian equilibria, equilibrium selection criteria are an important issue. The fact that the perfect Bayesian equilibrium in Part (ii) of Proposition 1 represents the limit case of a sophistication- k equilibrium provides a natural equilibrium selection criterion. It provides the limit that boundedly rational agents would converge to if going through a k -thinking/learning process.¹¹

Overall, comparing the two cases in Proposition 1, we observe that the institution of representative democracy works fundamentally different under limited strategic sophistication compared to under perfect strategic rationality. The result that the politician's incentive to pander *increases* with strategic sophistication appears quite surprising at first, but it becomes intuitive once the logic of sophistication- k belief is understood.

In the following corollary, we point out the effects of two further important parameters of the model: σ_ε^2 and α .

Corollary 1. (Comparative statics) (i): If the intermediate shock's variance, σ_ε^2 , grows, the extent of populism increases. (ii): Both the frequency and the extent of populism are independent of α , the share of competent politicians in the economy.

With respect to part (i), recall that $\beta = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2}$ (see Lemma 1). Thus, if σ_ε^2 is low, β is close to one and populism vanishes. The intuition is that the voter receives the signal $x^* + \varepsilon$ before making his election decision. He uses 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^* , almost perfectly. As a result of getting a very precise signal, the voter's prior μ has very little influence on his posterior opinion about x^* . A politician's aim is to be judged competent through the lens of the voter's *posterior* opinion. If this posterior opinion depends only very little on the prior μ , a politician's incentive to pander to the voter's prior opinion is comparatively low. As a result, the policy is to a larger degree determined by the politician's own information, which he uses to predict the voter's posterior opinion. In the opposite case, where σ_ε^2 is large, β is comparatively low. Thus, the voter's opinion is highly persistent and μ has a high weight in influencing policy. 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 the voter.

With respect to part (ii) of Corollary 1, note that, according to Proposition 1, any policy of any politician is populist, independent of the politician's type. When comparing Eqs. (7) to (8), it is evident that an incompetent politician's policy g is more noisy than a competent politician's. However, according to our definition, the extent of populism depends on the weight that a politician puts on the voter's prior opinion, μ , and this weight, $1 - \beta^{k+1}$, is equal across politician types.

4.4. Bounded rationality and welfare

Using Proposition 1, it is straightforward to characterize welfare. We do so by using the concept of a loss function L . L is defined as the difference between the voter's expected utility as achieved when g is set to its first-best level – that is the ex ante welfare-

¹⁰ For empirical evidence on this, see Section 6.

¹¹ It is straightforward to show that, under perfect rationality, any policy choice $g = \mu + \Delta$ (where Δ is an arbitrary real number) can be supported as an equilibrium by appropriate off-equilibrium beliefs. In all of these equilibria, the policy action does not depend on the politician's information. As mentioned in footnote 2, we do not provide a fully-fledged analysis of the complete set of perfect Bayesian equilibria that arise under unlimited strategic sophistication. While this would be very interesting, it is beyond the scope of this paper and would require an analysis in its own right. In this paper, we focus on the role of limited strategic sophistication and compare the outcome to the limit case $k \rightarrow \infty$ that would arise if agents went through an infinite number of k -thinking/learning rounds.

maximizing level x^* – and expected utility as achieved in the equilibrium of the political game. Formally, $L = EV^{FB} - EV^{EQ}$ in each period. Our welfare measure is thus entirely based on the voter's utility and the politician's utility does not appear. This is an assumption commonly made in the literature (see, e.g., Maskin and Tirole, 2004). The idea is that the weight of a politician is very small compared to the whole electorate. Technically, this is approximated by the electorate (the “voter”) having unit mass, whereas the politician has measure zero.

The first-best utility value ex ante amounts to $EV^{FB} = -\sigma_\xi^2$ (see Eq. (2)). Using this, we obtain:

Proposition 2. (Welfare representative democracy) Under representative democracy, welfare is characterized by

$$L^{RD} = (1 - \beta^{k+1})^2 (x^* - \mu)^2 + \beta^{2(k+1)} (1 - \alpha) \sigma_\xi^2 \quad \text{in period 1,} \quad (9)$$

$$L^{RD} = [1 - \alpha - \Delta_\alpha] \sigma_\xi^2 \quad \text{in period 2,} \quad (10)$$

where $\Delta_\alpha \geq 0$.

Consider the first period. The welfare loss from representative democracy is equal to a weighted average of the distortion associated with the voter's opinion, $|x^* - \mu|$, and the variance of the incompetent politician's prior, σ_ξ^2 . The first term arises from pandering. The second term arises from the fact that no equilibrium entails full pandering for finite k but the politician will always partially base his policy choice upon his private information. Since the prior of the incompetent politician is noisy, the fact that g depends on this prior increases the variance of g . This comes at a cost to risk averse voters. It is interesting to note that the weights, $(1 - \beta^{k+1})^2$ 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 in the context of Proposition 5.

In the second period, pandering does not arise since no politician has an incentive to manipulate the voter's perception of his competence. As a result, only the noise term σ_ξ^2 contributes to the welfare loss. It can be checked that $\Delta_\alpha \geq 0$. This follows from the fact that the probability that a competent politician holds office in the second period exceeds α (see the Proof of Proposition 2). 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 (Besley, 2006).

An interesting and important insight from Proposition 2 is that welfare is higher under limited strategic sophistication than under perfect rationality provided that the politician's private information is sufficiently valuable on average (that is $(1 - \alpha)\sigma_\xi^2$ is not too large). Thus, under a broad range of circumstances, bounded rationality leads to an increase in welfare compared to perfect rationality. Within the behavioral economics literature, this is a rather unusual finding. After all, behavioral distortions such as loss aversion, hyperbolic discounting etc. are usually detrimental for welfare, not welfare-enhancing (see Della Vigna and Malmendier, 2004; Gabaix and Laibson, 2006; Köszegi and Rabin, 2006; Laibson, 1997).

That limited strategic sophistication can be beneficial for welfare originates from the fact that smart players have too big an incentive to engage in detrimental “outsmarting games.” A smart voter expects a politician to pander. As a result, the politician will pander, but the politician will outsmart the voter by pandering even more than the voter expects. Now, if the voter becomes even smarter, he will expect even more pandering, resulting the politician pandering even more etc. In the limit of full rationality, both players have exhausted every opportunity to outsmart each other. This is only compatible with a situation of perfect populism. As long as the politician's private information is sufficiently valuable, on average, this is a bad outcome since it means wasting valuable information.

The fact that the lack of perfect strategic sophistication is beneficial for welfare is reminiscent of the economic benefits of trust (Fehr, 2009). Intuitively, mutual trust also reduces the propensity to engage in “outsmarting games” and, thus, has an effect similar to limited strategic sophistication. However, as emphasized by Fehr (2009), trust is not primarily a phenomenon of cognitive limitation, but of social preferences and betrayal aversion. Thus, while trust also limits the propensity for strategic behavior, this mainly originates from a different source than limited sophistication.

More formally, welfare under limited strategic sophistication can be compared to welfare under perfect rationality as follows. Welfare is always higher for some finite k than for an infinite k . This can be seen by minimizing Eq. (9) with respect to β^{k+1} . Assuming an interior solution – and thus neglecting that $\beta^{k+1} < \beta$ since $k \geq 0$ – the minimizing value is given by

$$\beta_{min}^{k+1} = \frac{(x^* - \mu)^2}{(x^* - \mu)^2 + (1 - \alpha)\sigma_\xi^2}. \quad (11)$$

The expression on the right-hand side is independent of k and may take on values between zero and one. It follows from Eq. (11) that perfect rationality, that is an infinite k , does never minimize Eq. (9), except for the uninteresting case where σ_ξ^2 is infinite, that is when the incompetent politicians' information would be infinitely noisy, on average.

When is highest welfare achieved under $k=0$? If the welfare-maximizing k is 0, then the nonnegativity constraint on k for minimizing Eq. (9) is binding. In particular, $k=0$ leads to the highest welfare if

$$\beta \leq \frac{(x^* - \mu)^2}{(x^* - \mu)^2 + (1 - \alpha)\sigma_\xi^2}. \quad (12)$$

This is the case if the right-hand side of Eq. (12) is relatively large, i.e. comes close to one (recall that $\beta < 1$). In the limit case where there are no incompetent politicians, i.e. if $\alpha = 1$ and/or $\sigma_\xi^2 = 0$, the right-hand side of Eq. (12) is, in fact, equal to 1. By

continuity, $k = 0$ also minimizes Eq. (9) if α is close to one or σ_ξ^2 is small. However, $k = 0$ does not lead to the highest value of welfare if α is sufficiently smaller than one or σ_ξ^2 is relatively large.

Intuitively, a competent politician's signal reveals the welfare-maximizing level of g . Thus, any other determinant of g – such as μ – that prevents g from exclusively depending on this truthful signal lowers welfare. In contrast, if there are incompetent politicians and their signals are sufficiently noisy, then dependence of g on these signals increases its variance and makes policy choice more erratic. Since the voter is risk averse, this comes at a cost. In this case, some degree of populism, i.e. higher k , is beneficial as it lowers the variance of g and this may increase welfare.

5. Comparing constitutional regimes

We now turn to the comparative institutional analysis. Our aim is to show that limited strategic rationality crucially affects the welfare ranking of different political institutions. In particular, we compare welfare under our baseline case of representative democracy to two other important benchmark cases: direct democracy, and governance by independent experts. The three institutions can be ordered in terms of the degree to which decision making is delegated from voters to their agents, and thus the degree to which agents are accountable to voters. Decision making can either be delegated to completely *independent experts*; it can be delegated to politicians who want to get reelected and, thus, are only partially independent (*representative democracy*); or it may not be delegated at all (*direct democracy*).¹²

5.1. Direct democracy

We first consider direct democracy. We follow Maskin and Tirole (2004) by modeling direct democracy as a political institution where

$$g = \mu, \quad (13)$$

that is the voter directly chooses g himself. The idea is that in a direct democracy voters have the right to ask for referenda and that this leads to a strong link between policy making and the opinion of voters (Gerber, 1996; Besley and Coate, 2008).¹³

In reality μ is likely to depend on political institutions. In a direct democracy, the political discourse is likely to make voters better informed than in representative regimes. Furthermore, in regimes with direct-democratic elements, there is more scope for experts to influence voters' opinions. For simplicity, and in the interest of transparency, our formal analysis does not take this into account. However, it is important to keep this omission in mind when interpreting our results.

There are no strategic elements involved in decision making. The following proposition follows directly from inserting g into Eq. (2) and taking expectations.

Proposition 3. (*Welfare under direct democracy*) Under direct democracy, the welfare loss in each period amounts to $L^{DD} = (x^* - \mu)^2$.

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 representative democracy in Eq. (9), we see that both are equivalent for the case of perfect rationality, that is an infinite k .

5.2. Independent experts

Now we turn to delegation of policy making to independent experts. For consistency, we assume that experts have the same lexicographic preferences as politicians. However, because experts are non-accountable to voters—that is they determine policies in both periods and cannot be ousted after the first period—the first goal of politicians, to get reelected, is not relevant for experts. Hence, we assume that experts solely maximize welfare, which makes them similar to benevolent dictators.

The regime of independent agents is often viewed as an ideal form of governance. What we have in mind is, for instance, that policy choices are mainly determined by academic experts or experts with a previous career in academia. Specifically, one could imagine that those agents who are responsible for implementing policy – politicians or bureaucrats – are bound by the constitution to directly follow the advice of academic experts. In practice, such an institution does not exist in pure form, although top-level central bankers or supreme court judges may come close.

Importantly, however, complaints of academic economists that politicians are not listening to their advice are commonplace. Implicit in this complaint is the judgment that governance by experts would be desirable. Furthermore, it is often implicitly assumed that the policy goal of academic experts is to maximize welfare, broadly defined as a weighted average of the utility of different groups of citizens. In our setting, this would be equal to maximizing Eq. (2). We therefore believe that the view that experts are benevolent is, if not fully true, an important benchmark case and thus of interest for a comparative analysis of different governance institutions.

To facilitate the comparison to representative democracy, we assume that experts obtain a prior about x^* , denoted by ξ^{EXP} , similar to the case of politicians. There is a competent expert type for which we have $\xi^{EXP} = x^*$. For the incompetent expert type, we have $\xi^{EXP} = x^* + \nu$. The random variable ν reflects a noise term with an expected value of zero and a variance σ_ν^2 . The probability

¹² 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.

¹³ In New Zealand, Switzerland, and some U.S. states, a referendum can be initiated by voters by means of a citizen petition.

that an expert is competent is π . As mentioned, we assume that experts are fully benevolent. In other words, their objective is to maximize Eq. (2).¹⁴

Consequently, an expert sets

$$g = \xi^{EXP}. \quad (14)$$

The welfare loss under the expert regime is given in the following proposition. The proof is very similar to the one of Proposition 2 and is omitted.

Proposition 4. (Welfare experts) *In the case of experts, the per period loss amounts to $L^{EXP} = (1 - \pi)\sigma_v^2$.*

5.3. Comparison of institutions (ex ante and ex post)

Comparing the outcomes for the three constitutional regimes in the first period, representative democracy can be understood as a mix of direct democracy and governance by experts.

For a comparison of institutions, it is instructive to take on a neutral ex ante perspective, according to which the quality of information (that is the prior about x^*) is the same for politicians under representative democracy and for experts. Denote ξ^P the prior of a politician and assume that $\xi^{EXP} = \xi^P \equiv \xi$. Furthermore, assume that the likelihood that an expert or a politician is incompetent is equal, that is $\alpha = \pi$. Then $g^{EXP} = \xi$, where EXP stands for experts. Furthermore, $g^{DD} = \mu$, where DD refers to direct democracy. It follows from Proposition 1 that $g^{RD} = (1 - \beta^{k+1})g^{DD} + \beta^{k+1}g^{EXP}$ in the first period, where RD refers to representative democracy.

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

Proposition 5. (Ex ante superiority of representative democracy) *Consider an ex ante perspective where $\alpha = \pi$ and $\xi^{EXP} = \xi^P$. If k is finite, then $g^{RD} = (1 - \beta^{k+1})g^{DD} + \beta^{k+1}g^{EXP}$ and $L(g^{RD}) < (1 - \beta^{k+1})L(g^{DD}) + \beta^{k+1}L(g^{EXP})$.*

Proposition 5 implies that – in the context of our specific model – representative democracy yields greater expected welfare than the convex combination of direct democracy and delegation of decision making to independent experts.¹⁵ At the heart of this ex ante superiority lies the fact that, under limited strategic sophistication, policy is determined as a weighted average of the voter's opinion and the politician's information. This makes the policy outcome very balanced. To understand this, note that, on the one hand, if policy making is determined strictly by voters themselves – as in the case of direct democracy or also under representative democracy in the case of perfect strategic rationality – there is the risk that voters err to a large extent. On the other hand, if policy making is determined by non-accountable experts, there is the risk that they are wrong. The institution of representative democracy builds, in fact, a kind of portfolio of those two institutions, from an ex ante perspective. This happens precisely because the politicians pander to voters, but only partially so. However, this only happens if the policy choice is indeed a weighted average of voters' opinions and politicians' information, that is only if there is limited strategic sophistication. For an infinite k , the ex ante superiority disappears. In that case, welfare under representative and under direct democracy is identical in our model. Furthermore, for an infinite k , there is no natural way to rank the latter two institutions with respect to independent agents from an ex ante perspective.

Moving from the neutral ex ante to the ex post perspective, where parameter values are already realized, we compare the loss functions in Propositions 2, 3, and 4. We find that the elements that crucially affect which constitutional regime is optimal are: the distortion associated with the median voter's belief, $|\mu - x^*|$; the variance of the incompetent politician's bias, σ_v^2 ; and the corresponding variance of the incompetent expert's bias, σ_e^2 . We summarize the effects of these three elements on the relative optimality of constitutional regimes in the following Corollary.

Corollary 2. (Constitutional comparison) (i): *Experts are optimal if σ_e^2 is small relative to $|\mu - x^*|$ and σ_v^2 .* (ii): *Direct democracy is optimal if $|\mu - x^*|$ is small relative to σ_e^2 and σ_v^2 .*

6. Discussion

6.1. Empirical evidence

The main positive prediction of our k -belief model of representative democracy is that, when a politician can get reelected, he partially panders to the voter's opinion. If no reelection is possible (i.e. in the second period in our model), the politician does not

¹⁴ The term “incompetent expert” is rooted in the fact that experts often disagree. If so, at most one expert opinion can be right. In other words, experts, even if benevolent and highly trained, may be wrong. Combating crime provides one example where experts disagree substantially (see Levitt, 1998, and Buscaglia, 2008), climate change provides another one (see McKibbin and Wilcoxon, 2002; Weitzman, 2007; Stern, 2008), to mention just a few domains.

¹⁵ It is important to recall that, in reality, we probably have $\alpha \neq \pi$ and $\xi^{EXP} \neq \xi^P$. It is an empirical question, however, in which direction these values differ.

pander. Evidence for this pattern is provided by Heith (1998) and Heith (2004). There, the author reports results from her studies of archive documents that shed light on how U.S. presidents make use of *public opinion research*. As put by Heith (1998, p. 165), “few things are more important to the modern White House than public opinion.” Results of public opinion research are found to be crucial for agenda setting and have a strong influence on policy making. On the other hand, “the permanent campaign does not dominate decision making” (Heith, 2004, p.134). In sum, this means that public opinion is an important ingredient, but not an exclusive determinant of presidential decision making. This is fully consistent with the prediction of our model that, before re-election, policies are determined as a weighted average of the voter's opinion and the politician's private information.

A very interesting result reported in Heith (1998) is that public opinion research is far more important in a president's first term than in the second. In fact, they are most important in the first and second year of the first term, when agendas are set. This is consistent with the prediction of our model that, in the second period, there is no pandering.

List and Sturm (2006) provide evidence that politicians pander when facing reelection by showing that, in their last term, U.S. governors tend to be less populist with regard to environmental policy. McArthur and Marks (1988) find that representatives in the U.S. Congress change their voting behavior after not getting reelected, i.e. when being “lame duck.” Going beyond elected officials, Blanes i Vidal and Leaver (2011) provide evidence that even appointed judges pander to voters' opinions when they face a threat of losing decision making power.

Turning to the case of direct democracy, our analysis predicts that the policy outcome in a direct democracy directly reflects the voter's opinion, while the outcome in a representative democracy does partially so, provided limited strategic sophistication. In line with this pattern, Gerber (1999) finds that in U.S. states with voter initiatives, measures of public opinion and policy outcomes are more closely correlated than in states without voter initiatives.

Besley and Coate (2003) provide evidence on the comparison of elected vs. appointed officials by analyzing the behavior of electricity regulators. They find that elected officials lead to more pro-consumer regulations than appointed officials. This is consistent with the idea that elected officials pander to their electorate while appointed officials are less prone to do so. In a similar spirit, Lim (2011) provides evidence that elected judges pander to their electorate and that their behavior differs from appointed judges.

6.2. Different k -types for voters and politicians

Here, we briefly discuss how our results change if the constellation of the voter's and politician's levels of k is different from the baseline case in Section 4. For the analysis of more general constellations of k -beliefs, it is useful to imagine that each level of k corresponds to a (hypothetical) belief type (Crawford and Iriberri, 2007b). A (hypothetical) 0-type voter has a naive belief, namely that the politician maximizes the voter's utility. A 1-type politician plays best response to the belief of the 0-type voter. A 1-type voter has a belief that is a best response to the 1-type politician. A 2-type politician plays best response to the beliefs of a 1-type voter. A 2-type voter has a belief that is a best response to the strategy of the 2-type politician etc.

The fact that the iteration starts with a 0-type, does not mean that this type actually exists or is empirically plausible. Rather, the iteration reflects a *learning* process that starts at the 0-type belief and stops at a specific finite level of k because of bounded rationality. In experimental analysis, when allowing for heterogeneous types, most subjects are found to be either 1- or 2-types (Camerer et al., 2004; Crawford and Iriberri, 2007b).

In our derivation of k -beliefs in Section 4, we have implicitly assumed that the politician's level of sophistication is higher than the voter's by one degree. For instance, if the voter is of 1-type, then the politician is of 2-type. We view this as a plausible assumption since a politician has a high incentive to get trained in strategic behavior. In contrast, a single voter has a much lower incentive to learn about a politician's strategic behavior since a voter has a negligible effect on the outcome of an election.¹⁶ In this context, it is noteworthy that a politician of type $k + x$, $x > 1$ would choose the same actions as a politician with type $k + 1$. The reason is that both types play best response to a voter of type k , and this is all that matters.

Although we view the baseline assumption that the politician's degree of sophistication is higher than the voter's as plausible, it is not crucial for our results. Suppose, instead, that the politician and the voter share the same k . Then, based on the reasoning outlined above, the politician plays best response to a (hypothetical) voter's “ $k-1$ ”-type, and the voter's belief is a best response to the politician's (actual) k -type. Since the politician's choice of g is a best response to a voter's “ $k-1$ ”-belief, we can use Lemma 3 to infer that $g = (1 - \beta^k)\mu + \beta^k\xi$, rather than $g = (1 - \beta^{k+1})\mu + \beta^{k+1}\xi$. Comparing the two expressions, it becomes immediately clear that any comparative static result with respect to k that we derived assuming that the belief type of the politician is $k + 1$, will qualitatively be the same when the belief type of the politician is k .

Based on the same reasoning, it is also straightforward to infer that if the belief type of the voter is k and the belief type of the politician is $k-1$, then the politician's action is $g = (1 - \beta^{k-1})\mu + \beta^{k-1}\xi$. This is simply a response to the belief of a (hypothetical) “ $k-2$ ”-type voter. Results for other possible k -type constellations can be inferred along the same lines. In sum, we conclude that the assumption that the politician's level of sophistication is higher than the voter's is not crucial for our results.

7. Conclusion

The debate about the best form of government goes back to ancient times. Despite the long history in search of optimal political institutions, the issue is still not settled, yet. As *The Economist* wrote recently, when commenting about current governance

¹⁶ Compare this to the “Behavioral IO” literature where it is common to assume that consumers suffer from a behavioral bias while firms or their managers are fully rational; see Della Vigna and Malmendier (2004) and Gabaix and Laibson (2006).

and other political problems in California: “the main culprit has been direct democracy [...] This citizen legislature has caused chaos.”¹⁷ Moreover, that newspaper identified a late trend in the entire Western world towards more direct democracy: “With technology making it ever easier to hold referendums and Western voters ever more angry with their politicians, direct democracy could be on the march.”

In this paper, we have addressed the topical issue of the optimal political governance regime using novel insights from behavioral economics. In particular, we have performed a comparative analysis of political institutions under bounded rationality and limited strategic sophistication.

In our model of representative democracy, we have shown that limited strategic sophistication crucially affects politicians' incentives to engage in populism. Under limited sophistication, a policy is determined as a weighted average between politicians' private information and voters' opinions. In contrast, under unlimited sophistication, the policy outcome is perfectly populist and no politician has an incentive to make use of any private information, however valuable this information may be. A surprising finding is that the higher the level of strategic sophistication, the more populist is the policy choice of politicians. The reason is that a more sophisticated voter *expects* politicians to pander more. As a result, politicians will indeed pander more.

In our model, limited strategic sophistication crucially affects the desirability of various political institutions: by comparing representative democracy to direct democracy and to delegation of policy making to independent experts, we have argued that, from an ante perspective, representative democracy is preferable to the other two institutions if there is limited strategic rationality. In a direct democracy, where policy making is determined by voters themselves, there is the risk that voters err to a large extent. On the other hand, there is a risk that independent experts may be wrong as well. In our setup, the institution of representative democracy builds, in fact, a *portfolio* of the two institutions direct democracy and independent experts and thus diversifies the mentioned risks. This happens precisely because politicians pander to voters under limited strategic sophistication, but the pandering is only partial.

Although the model presented in this paper is stylized and rests on specific assumptions, it contributes to the ongoing discussion about the optimal form of democratic institutions: it sheds light on the transmission channels through which limited strategic sophistication may impact policy choice and makes clear predictions about the degree of populism depending on voters' sophistication.

In future research, it would be interesting to consider the role of the media, taking into account their potential to shape voters' beliefs. Furthermore, one may analyze the role of education policy. In particular, one may consider endogenizing the level of strategic sophistication and take into account heterogeneous levels of sophistication. Furthermore, the setup of this paper may also apply to decision making in corporations and to principal-agent relationships more generally.

Acknowledgments

We would like to thank Jordi Blanes i Vidal, 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, EEA 2009, and ISNIE 2009 annual meetings for very helpful comments. We are also very grateful to two anonymous referees and the editor, Heinrich Ursprung, for valuable suggestions on how to improve the paper. All errors are our own.

Appendix A

Proof of Lemma 1. At stage 3 of period one, the voter observes $x^* + \varepsilon$ as the realization of the random variable $x + \varepsilon$. The voter aims to update his belief about x . The random variables x and $x + \varepsilon$ are jointly normally distributed with $E[x] = \mu$, $Var[x] = \sigma_x^2$, $E[x + \varepsilon] = \mu$, $Var[x + \varepsilon] = \sigma_x^2 + \sigma_\varepsilon^2$. Furthermore, $Cov[x, x + \varepsilon] = \sigma_x^2$. Substituting this into the formulas for conditional expectations and variances for jointly-normal random variables (see, for instance, [Hogg and Craig, 1995](#), p. 148) yields the result. *Q.E.D.*

Proof of Proposition 1. Part (i): This follows directly from [Lemma 2](#) and [Lemma 3](#).

Part (ii): Clearly, part (i) implies that there is a unique limit for $k \rightarrow \infty$. This limit is indeed an equilibrium for appropriate off-equilibrium beliefs about a politician's type if he deviates from setting $g = \mu$. For instance, consider the belief that a politician setting $g \neq \mu$ is incompetent with probability one. Under this belief, the voter would never reelect a politician setting $g \neq \mu$ (see [Section 4.1](#)). Given such an off-equilibrium belief, a politician setting $g = \mu$ maximizes his probability of getting reelected. Hence, he does not want to deviate from $g = \mu$. *Q.E.D.*

Proof of Proposition 2. Denote by g_c the level of g set by the competent politician and let g_{iC} refer to the incompetent politician. Let λ denote the probability that a politician is competent in a given period. Then, using Eq. (2),

$$EV = - \left[\lambda E \left[(g_c - x^* - \varepsilon)^2 \right] + (1 - \lambda) E \left[(g_{iC} - x^* - \varepsilon)^2 \right] \right]. \quad (15)$$

Consider the first period. Clearly, $\lambda = \alpha$. Using this, and inserting for g_c , g_{iC} from [Proposition 1](#) into Eq. (15), we obtain

$$EV^{EQ} = - \left(1 - \beta^{k+1} \right)^2 (\mu - x^*)^2 - (1 - \alpha) \beta^{2(k+1)} \sigma_\zeta^2 - \sigma_\varepsilon^2.$$

¹⁷ April 23rd, 2011, p.12

EV is maximized for $g = x^*$, which yields $EV^{FB} = -\sigma_e^2$. Inserting this and the above expression into the definition of L yields Eq. (9).

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 if k is finite. When making his reelection decision, the voter considers how the inferred prior of the politician, $\hat{\xi}$, compares to the voter's best available information about x^* , that is $\hat{\mu}$. In particular, the smaller $|\hat{\xi} - \hat{\mu}|$, the higher the posterior probability $\hat{\alpha}$ that the voter assigns to the event that the incumbent is competent. Since a politician wants to maximize $\hat{\alpha}$ (see Section 4.1), the politician chooses g such that $\hat{\xi} = E[\hat{\mu}|\xi]$ (see Section 4.2). Since the prior ξ is more informative in case of the competent politician than in case of the incompetent politician, the likelihood that the former achieves a small value of $|\hat{\xi} - \hat{\mu}|$ is higher. Hence, the likelihood that the voter infers a high value of $\hat{\alpha}$ is higher in case of a competent incumbent than in case of an incompetent incumbent. In turn, the likelihood that $\hat{\alpha} \geq \alpha$ is higher in case of the competent politician. Hence, a competent incumbent will be reelected more often than an incompetent politician (see Section 4.1 for the fact that an incumbent gets reelected if and only if $\hat{\alpha} \geq \alpha$).¹⁸

If k is infinite, both politician types choose $g = \mu$ (see Proposition 1). Thus, there is no way to make any updating about the likelihood that the incumbent is competent. Hence, the reelection chances are the same for both politician types. Overall, the chance that the competent type gets reelected is (at least weakly) higher than the chance that the incompetent type gets reelected for any value of k , including infinity.

We now derive Eq. (10). 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 challenger; (3) an incompetent incumbent gets ousted and replaced by a competent challenger. 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 . Then we have

$$\begin{aligned} Pr[C_2] &= \alpha\rho_C + \alpha^2(1-\rho_C) + \alpha(1-\alpha)(1-\rho_{IC}) \\ &= \alpha[1 + (1-\alpha)(\rho_C - \rho_{IC})] \geq \alpha. \end{aligned} \quad (16)$$

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 derived above. Overall, we have now established that, in period two, $\lambda = \alpha + \Delta_\alpha$ for some $\Delta_\alpha \geq 0$ (see Eq. (16)).

Using that $g_C = x^*$, $g_{IC} = x^* + \zeta$ in the second period (see Section 4.1) and inserting this into Eq. (15) yields a second period expected utility of

$$EV = -(1-\alpha-\Delta_\alpha)\sigma_\zeta^2 - \sigma_e^2.$$

Again, $EV^{FB} = -\sigma_e^2$. Inserting this and the above expression into the definition of L yields Eq. (10).

References

- Alesina, A., Tabellini, G., 2007. Bureaucrats or politicians? Part I: a single policy task. *The American Economic Review* 97, 169–179.
- Alpert, M., Raiffa, H., 1982. A progress report on the training of probability assessors. In: Kahneman, D., Slovic, P., Tversky, A. (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, U.K., and New York, NY, pp. 294–305.
- Barro, R., 1973. The control of politicians: an economic model. *Public Choice* 14, 19–42.
- Bénabou, R., Tirole, J., 2006. Belief in a just world and redistributive politics. *Quarterly Journal of Economics* 121, 699–746.
- Besley, T., 2006. Principled agents? *The Political Economy of Good Government*. Oxford University Press, Oxford.
- Besley, T., Coate, S., 1997. An economic model of representative democracy. *Quarterly Journal of Economics* 112, 85–114.
- Besley, T., Coate, S., 2003. Elected versus appointed regulators. *Journal of the European Economic Association* 1, 1176–1206.
- Besley, T., Coate, S., 2008. Issue unbundling via citizens' initiatives. *Quarterly Journal of Political Science* 3, 379–397.
- Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. *Review of Economic Studies* 72, 287–312.
- Binswanger, J., Prüfer, J., 2009. Democracy, populism, and (un)bounded rationality. CentER Discussion Paper No. 2009-10. Tilburg University, Tilburg.
- Blanes i Vidal, J., Leaver, C., 2011. Are tenured judges insulated from political pressure? *Journal of Public Economics* 95, 570–586.
- Buscaglia, E., 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, 290–317.
- Camerer, C.F., Ho, T., Chong, J., 2004. A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119, 861–898.
- Canes-Wrone, B., Herron, M.C., Shotts, K.W., 2001. Leadership and pandering: a theory of executive policymaking. *American Journal of Political Science* 45, 532–550.
- Costa-Gomes, M.A., Crawford, V.P., 2006. Cognition and behavior in two-person guessing games: an experimental study. *The American Economic Review* 96, 1737–1768.
- Costa-Gomes, M.A., Crawford, V.P., Broseta, B., 2001. Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69, 1193–1235.
- Crawford, V.P., 2003. Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. *The American Economic Review* 93, 133–149.
- Crawford, V.P., Iriberrí, N., 2007a. Fatal attraction: salience, naivete, and sophistication in experimental 'hide-and-seek' games. *The American Economic Review* 97, 1731–1750.
- Crawford, V.P., Iriberrí, N., 2007b. Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75, 1721–1770.
- Crawford, V.P., Kugler, T., Neeman, Z., Pauzner, A., 2009. Behaviorally optimal auction design: examples and observations. *Journal of the European Economic Association* 7, 377–387.

¹⁸ See Binswanger and Prüfer (2009) for a more formal version of the arguments in this paragraph.

- De Bondt, W.F., Thaler, R.H., 1995. Financial decision making in markets and firms: a behavioral perspective. In: Jarrow, R.A., Maksimovic, V., Ziemba, W.T. (Eds.), *Handbooks in Operations Research and Management Science*, Vol. 9. Elsevier, Amsterdam, pp. 385–410.
- Della Vigna, S., Malmendier, U., 2004. Contract design and self-control: theory and evidence. *Quarterly Journal of Economics* 119, 353–402.
- Dixit, A.K., Weibull, J.W., 2007. Political polarization. *Proceedings of the National Academy of Science of the United States of America* 104, 7351–7356.
- Downs, A., 1957. *An Economic Theory of Democracy*. Harper Collins, New York.
- Fehr, E., 2009. On the economics and biology of trust. *Journal of the European Economic Association* 7, 235–266.
- Ferejohn, J., 1986. Incumbent performance and electoral control. *Public Choice* 50, 5–25.
- Gabaix, X., Laibson, D., 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121, 505–540.
- Gerber, E.R., 1996. Legislative response to the treat of popular initiatives. *American Journal of Political Science* 40, 99–128.
- Gerber, E.R., 1999. *The Populist Paradox: Interest Group Influence and the Promise of Direct Legislation*. Princeton University Press, Princeton.
- Glaeser, E.L., Ponsetto, G.A.M., Shapiro, J.M., 2005. Strategic extremism: why republicans and democrats divide on religious values. *Quarterly Journal of Economics* 120, 1283–1330.
- Goldfarb, A., Yang, B., 2009. Are all managers created equal? *Journal of Marketing Research* 46, 612–622.
- Heith, D.J., 1998. Stuffing the white house public opinion apparatus. *Public Opinion Quarterly* 62, 165–189.
- Heith, D.J., 2004. *Polling to Govern: Public Opinion and Presidential Leadership*. Stanford University Press, Stanford.
- Hillman, A.L., 2010. Expressive behavior in economics and politics. *European Journal of Political Economy* 26, 403–418.
- Ho, T.H., Camerer, C., Weigelt, K., 1998. Iterated dominance and iterated best response in experimental *p*-beauty contests. *The American Economic Review* 88, 947–969.
- Hogg, R.V., Craig, A.T., 1995. *Introduction to Mathematical Statistics*, fifth ed. Prentice Hall, New Jersey.
- Holmström, B., 1999. Managerial incentive problems: a dynamic perspective. *Review of Economic Studies* 66, 169–182.
- Kőszegi, B., Rabin, M., 2006. A model of reference-dependent preferences. *Quarterly Journal of Economics* 121, 1133–1165.
- Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112, 443–477.
- Levitt, S.D., 1998. Juvenile crime and punishment. *Journal of Political Economy* 106, 1156–1185.
- Lim, C.S.H., 2011. Turnover and Accountability of Appointed and Elected Judges. Working Paper. Stanford University.
- List, J.A., Sturm, D.M., 2006. How elections matter: theory and evidence from environmental policy. *Quarterly Journal of Economics* 121, 1249–1281.
- Malmendier, U., Tate, G., 2005. CEO overconfidence and corporate investment. *Journal of Finance* 60, 2661–2700.
- Maskin, E., Tirole, J., 2004. The politician and the judge: accountability in government. *The American Economic Review* 94, 1034–1054.
- McArthur, J., Marks, S.V., 1988. Constituent interest vs. legislator ideology: the role of political opportunity cost. *Economic Inquiry* 26, 461–470.
- McKibbin, W.J., Wilcoxon, P.J., 2002. The role of economics in climate change policy. *Journal of Economic Perspectives* 16, 107–129.
- Nagel, R., 1995. Unraveling in guessing games: an experimental study. *The American Economic Review* 85, 1313–1326.
- Osborne, M.J., Slivinski, A., 1996. A model of political competition with citizen candidates. *Quarterly Journal of Economics* 111, 65–96.
- Prat, A., 2005. The wrong kind of transparency. *The American Economic Review* 95, 862–877.
- Riker, W.H., 1982. Liberalism against Populism: A Confrontation between the Theory of Democracy and the Theory of Social Choice. Freeman, San Francisco.
- Rogoff, K., 1990. Equilibrium political budget cycles. *The American Economic Review* 80, 21–36.
- Sieg, G., 2001. A political business cycle with boundedly rational agents. *European Journal of Political Economy* 17, 39–52.
- Stahl, D., Wilson, P., 1994. Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization* 25, 309–327.
- Stahl, D., Wilson, P., 1995. On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* 10, 218–254.
- Stern, N., 2008. The economics of climate change. *The American Economic Review* 98 Papers & Proceedings, 1–37.
- Weitzman, M.L., 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45, 703–724.