

Berentsen, Aleksander; Bruegger, Esther; Loertscher, Simon

Research Report

Learning, voting and the information trap

WWZ Forschungsbericht, No. 03/05

Provided in Cooperation with:

Center of Business and Economics (WWZ), University of Basel

Suggested Citation: Berentsen, Aleksander; Bruegger, Esther; Loertscher, Simon (2005) : Learning, voting and the information trap, WWZ Forschungsbericht, No. 03/05, Universität Basel, Wirtschaftswissenschaftliches Zentrum (WWZ), Basel

This Version is available at:

<https://hdl.handle.net/10419/127500>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



July 2005

Learning, voting and the information trap

WWZ-Research Report 2005/ 03

Aleksander Berentsen, Esther Bruegger,
Simon Loertscher

The Authors:

Aleksander Berentsen, Esther Bruegger and Simon Loertscher are Researchers at
University of Basel
Department of Business and Economics
Petersgraben 51
4003 Basel
Switzerland

A publication of the University of Basel, Department of Business and Economics (WWZ)

Address for information:

WWZ Forum, Petersgraben 51, CH-4003 Basel Fax +41 61 267 33 33

The WWZ Research Papers document the results of research projects at the WWZ.

© WWZ 2005 and the authors, Reproduction for other purposes than the personal use needs the permission of the authors and the WWZ Forum. Please contact the WWZ Forum

Learning, voting and the information trap^{*}

Aleksander Berentsen, Esther Bruegger and Simon Loertscher[†]

April 3, 2005

Abstract

We consider a median voter model with uncertainty about how the economy functions. The distribution of income is exogenously given and the provision of a public good is financed through a proportional tax. Voters and politicians do not know the true production function for the public good, but by using Bayes rule they can learn from experience. We show that the economy may converge to an inefficient policy where no further inference is possible so that the economy is stuck in an information trap.

Keywords: Political economy, Meltzer-Richard model, Bayesian learning.

JEL-Classification: D72, H10, D83.

^{*}This paper has benefitted from very valuable discussions with Roland Hodler, Fabienne Peter, Ronny Razin, Kurt Schmidheiny, Martin Wagner, and, most of all, Roberto Schonmann, and from comments of seminar participants at the University of Bern, at the SSES conference 2004 in Basel, at the PCS meeting 2005 in New Orleans and at the EPCS meeting 2005 in Durham. Financial support by the WWZ-Forum Basel is gratefully acknowledged. Any remaining errors are ours.

[†]Berentsen: Economics Department, University of Basel, Petersgraben 51, CH-4003 Basel. E-mail: aleksander.berentsen@unibas.ch. Bruegger: Economics Department, University of Bern, Gesellschaftsstrasse 49, CH-3012 Bern. E-mail: bruegger@ucla.edu. Loertscher: Economics Department, University of Bern, Vereinsweg 23, CH-3012 Bern. E-mail: sl2331@columbia.edu.

...for after falling a few times they would in the end certainly learn to walk...

Immanuel Kant

1 Introduction

We consider a median voter model with uncertainty about how the economy functions. The distribution of income is exogenously given and the provision of a public good is financed through a proportional tax. Voters and politicians do not know the true production function for the public good, but by using Bayes rule they can learn from experience. We show that the economy may converge to an inefficient policy where no further inference is possible so that the economy is stuck in an information trap.

We introduce uncertainty by making the following two assumptions about the production of the public good. First, there are two production functions, and voters do not know which one is true. Second, the production of the public good is disturbed by exogenous shocks. In each period voters and politicians observe the implemented policy and the associated random output of the public good. Using this information and Bayes rule they update their beliefs about which production function is true. Each voter has a most preferred policy, which depends on these common beliefs and her personal income. In every period, the election outcome and consequently the production of the public good is determined by the median voter.¹

We consider two questions. (1) Does the stochastic process of beliefs and associated policies converge? (2) If so, where do they converge to? We show analytically that the policies converge to a random variable. The support of this random variable includes two policies. Interestingly, one of the policies can be Pareto inefficient. We use numerical methods to approximate the distribution of the random variable. The probability of converging to the inefficient policy increases in the variance of the shocks and in voters' initial beliefs attributed to the wrong production function.

There is a substantial political economy literature that deals with incomplete information. However, most of this literature deals with asymmetric informa-

¹Blendon et al. (1997) conducted an opinion survey showing that there is a substantial gap between economists' and the public's beliefs about how the economy functions. Fuchs et al. (1998) report findings from another survey that there are significant differences even among professional economists about policy questions as well as parameter estimates. This can be regarded as evidence of uncertainty about which is the correct model.

tion in the sense that some types of agents are better informed than others.² Closer related to our work is the seminal paper by Piketty (1995), in which agents have to learn the parameters of the model. In his model, agents have heterogeneous initial beliefs and have access to heterogeneous private information, which is why they end up with heterogeneous beliefs even in the long run. Obviously, at most one of these beliefs can be correct. In contrast, in our model all voters share the same information and beliefs, but are eventually hindered from learning the truth because further inference becomes impossible once they always observe the same outcome (or more precisely, once they always observe outcomes that have the same probability under either hypothesis). Building on Piketty's paper, Spector (2000) considers a cheap talk game in which all agents derive identical utility from a collective decision, but differ with respect to their beliefs. His assumptions are in complete contrast with ours insofar as we assume that all voters have the same information but are affected in different ways from the same policy. Our paper also relates to the literature on Bayesian learning (see, e.g., McLennan, 1984; Easley and Kiefer, 1988), which has established that impatient optimizers may optimally fail to learn the true parameter values. The model of Alesina and Angeletos (2003) is very similar to ours in that different beliefs are consistent with different equilibria, so that different social beliefs and political outcomes are self-reproducing. An important difference is that in their model the equilibria can be ranked unambiguously only from the point of view of the median voter. Moreover, the sources of multiplicity are quite different. It stems from differences in social beliefs about which fraction of income is fair or merited in their model, whereas in ours it arises because of incomplete information and incomplete learning due to bad luck and the median voter theorem.

The remainder of the paper is structured as follows. In section 2, the basic model is outlined. In section 3 we introduce uncertainty and the dynamic learning process of voters, and we show that this process converges. We derive also a lower bound for the probability that in the long run a Pareto efficient policy is adopted. In section 4, we then report simulation results that strongly support the view that the probability of reaching a Pareto inefficient policy is strictly positive for a wide set of initial conditions. Section 5 concludes.

²See, for example, Feddersen and Pesendorfer (1996), Blumkin and Grossmann (2004) or Schultz (2004).

2 The basic model

Our starting point is the model developed by Persson and Tabellini (2000, ch. 3), which builds on the seminal work of Meltzer and Richard (1981).³ We first describe the model without uncertainty.

2.1 The Hotelling-Downs model with a public good

There is a continuum of individuals whose total mass is normalized to one. Individual income y_i is distributed according to the differentiable distribution function $F(y_i)$, where $f(y_i) = F'(y_i)$ denotes the probability density function. The mean income is denoted by y and the median income by y^m . The support of the distribution is $[0, y^{sup}]$ with $y^{sup} < \infty$. Each individual i derives utility from private consumption c_i and from a public good $H(g)$, which is a function of government expenditure g . Therefore, individual i 's utility is

$$u_i = c_i + H(g) \quad (1)$$

Note that individuals differ only with respect to their private consumption, but are identical with respect to their valuation of the public good.

The government's budget constraint is

$$g = \tau y, \quad (2)$$

where $0 \leq \tau \leq 1$ is a flat tax rate. Accordingly, individual i 's consumption is

$$c_i = (1 - \tau)y_i. \quad (3)$$

We make the following assumptions for the production function $H(g)$. Assumption 1: $H(g)$ is twice differentiable and strictly concave. Assumption 2: $\frac{\partial H}{\partial g}(0) > \frac{y^{sup}}{y} > 0$. Assumption 3: $\frac{\partial H}{\partial g}(y) < 0$. These assumptions imply that $H(g)$ has a unique interior maximum in $[0, y]$ which avoids boundary solutions in the voting model we consider below.

Using the budget restrictions (2) and (3), we can rewrite (1) to get i 's utility from policy τ

$$u_i(\tau) = (1 - \tau)y_i + H(\tau). \quad (4)$$

From now on, as in (4), we write $H(\tau)$ instead of $H(\tau y)$ because y is a constant. Note that $H(\tau)$ is also concave in τ . Moreover, because of Assumption 2, we have $\frac{\partial H}{\partial \tau}(0) > y^{sup}$.

³We thank Roland Hodler for suggesting writing down the model in this way.

By τ^i we denote individual i 's optimal tax rate, which is implicitly defined by

$$\frac{\partial H}{\partial \tau}(\tau^i) = y_i. \quad (5)$$

Since $H(\tau)$ is concave, τ^i is decreasing in y_i . Thus, the single crossing property is satisfied (see Persson and Tabellini, 2000, ch. 2, condition 2.4). Denote τ^m the optimal tax rate of the median income voter.

2.2 Pareto efficient and Pareto inefficient policies

The optimal tax rate of the richest individual τ^I is defined by $H'(\tau^I) = y^{sup}$. Assumption 2 implies $\tau^I > 0$; some government activity is better than none for *all* individuals, even for the richest one. The optimal tax rate of the poorest individual τ^{II} is defined by $H'(\tau^{II}) = 0$. Assumption 3 implies $\tau^{II} < 1$, so even the poorest individual will prefer τ^{II} to any higher tax rate. The interval $P \equiv [\tau^I, \tau^{II}] \subset [0, 1]$ contains all Pareto efficient tax rates. Accordingly, the regions $[0, \tau^I)$ and $(\tau^{II}, 1]$ contain Pareto inefficient policies.




Figure 1: Pareto efficient and inefficient policies.

Figure 1 illustrates the set of Pareto efficient tax rates. This is also the interval of *conflictual politics* since voters do not unanimously agree which of these policies are better or worse. This is in contrast to policies $\tau \notin P$, which are considered by all individuals worse than either τ^I or τ^{II} .

2.3 Median Voter Equilibrium

We assume that every individual votes for the policy available which is closest to the policy that maximizes its utility given in (4). We focus on two party competition. Parties (or candidates) are opportunistic and derive utility solely from being in office. Parties simultaneously choose a policy $\tau \in [0, 1]$, which they commit to implement. The parties maximize the number of votes as opposed to maximizing the probability of winning. If both parties get the same number of votes, the winner is drawn by flipping a fair coin. Under the assumptions made, the unique equilibrium of the game is the well-known median voter equilibrium. Both parties choose τ^m as their policy and the winner is chosen randomly.

3 The model with uncertainty

In this section, we introduce uncertainty by making the following two assumptions. First, there are two possible production functions, labelled $H_A(\tau)$ and $H_B(\tau)$, only one of which is true, both satisfying the assumptions of section 2. Second, the production of the public good is disturbed by some factors exogenous to the model. Voters and politicians have some initial beliefs about which production function is the true one. They use the observed outcomes to update their beliefs.

3.1 Uncertainty and its unravelling

Without loss of generality, we assume that $H_A(\tau)$ is the true production function. Let $P_A \equiv [\tau_A^I, \tau_A^H]$ and $P_B \equiv [\tau_B^I, \tau_B^H]$ be the sets of Pareto efficient tax rates associated with the production function H_A and H_B , respectively. Let τ_A^m and τ_B^m be the optimal tax rates for the median voter under H_A and H_B , i.e.,

$$\frac{\partial H_A}{\partial \tau}(\tau_A^m) = y^m \quad \text{and} \quad \frac{\partial H_B}{\partial \tau}(\tau_B^m) = y^m. \quad (6)$$

Note that $\tau_A^m \in P_A$ and $\tau_B^m \in P_B$. Without loss of generality we assume that $\tau_A^m < \tau_B^m$. Furthermore, we assume that the two functions cross exactly once at $\tilde{\tau}$, and that $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$.

The production of the public good is exposed to uncertainty. If τ_t is the tax rate in period t , then voters (and politicians) observe the outcome

$$h_t(\tau_t, \varepsilon_t) = H_A(\tau_t) + \varepsilon_t, \quad (7)$$




Figure 2: Two production functions.

where ε_t is an error term drawn randomly in every period.⁴ It is common knowledge that the error terms are normally and i.i.d. with mean 0 and variance σ^2 ; we denote its probability density function by $\phi(\varepsilon_t)$. Note that without noise, the learning process, described below, would be degenerate since one observation would be sufficient to identify the true production function.

The time line is as shown in Figure 3. In every period t , an election takes

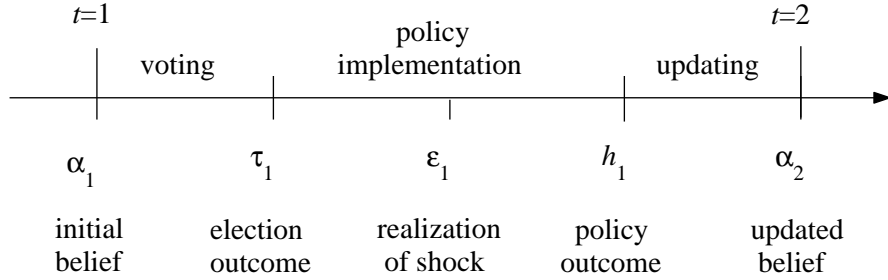


Figure 3: The time line.

place, $t = 1, 2, \dots$. Both implemented policies and the outcomes from these policies are observed ex post. That is, in period $t + 1$, the entire history $\mathcal{H}_t \equiv \{(h_j, \tau_j)\}_{j=1}^t$ of previously implemented tax rates and associated policy outcomes is common knowledge. The beliefs of voters and politicians in period t that H_A is the true production function are denoted by α_t . Then the

⁴The error term ε captures factors influencing the policy outcome except the policy itself.

expected level of the public good in period t for tax rate τ_t is

$$H_t(\tau_t) \equiv \alpha_t H_A(\tau_t) + (1 - \alpha_t) H_B(\tau_t). \quad (8)$$

Proposition 1 *In every period t , both candidates take the position τ_t^m , where τ_t^m is implicitly defined by*

$$H'_t(\tau_t^m) = \alpha_t H'_A(\tau_t^m) + (1 - \alpha_t) H'_B(\tau_t^m) = y^m. \quad (9)$$

Proof. Since H_A and H_B are concave, $H_t(\tau_t)$ is concave. For any concave function and beliefs α_t , the distribution function for τ_t^i can be derived using standard techniques for the transformation of random variables.⁵ Let $\tau_t^i = \kappa(y_i)$ denote the inverse of the function $y_i = H'_t(\tau_t^i)$ derived from the optimality condition (5) of the model without uncertainty. Since $H'_t(\tau_t^i)$ exists, $\frac{dy_i}{d\tau_t^i} = H''_t(\tau_t^i)$. If we denote by $\Omega(\tau_t^i)$ the distribution of τ_t^i , the density $\omega(\tau_t^i)$ of $\Omega(\tau_t^i)$ is given by

$$\omega(\tau_t^i) = f(\kappa(\tau_t^i)) \left| \frac{dy_i}{d\tau_t^i} \right| \quad (10)$$

where $\left| \frac{dy_i}{d\tau_t^i} \right|$ denotes the absolute value of the derivative $\frac{dy_i}{d\tau_t^i} = H''_t(\tau_t^i)$. Consequently, the optimal tax rate of the voter with the median income is the median optimal tax rate, which is denoted by τ_t^m . The median voter theorem applies and the median optimal tax rate will be implemented. \square

Figure 4 depicts the equilibrium outcome, as stated in Proposition 1. Initial




Figure 4: Equilibrium outcome in period 1.

beliefs α_1 are such that the expected production function in period 1 is H_1 , so that the policy implemented in period 1 is τ_1^m .

⁵See e.g. Hogg and Craig (1995).

Figure 5 illustrates the impact of the error term on the beliefs and on the equilibrium tax rate in the next period. After implementing τ_1^m , the shock ε_1 materializes. If $\varepsilon_1 > 0$, the outcome is better than expected under H_1 , and therefore, updated beliefs are $\alpha_2 > \alpha_1$ and the new expected production function H_2 is as shown in the left hand panel. On the other hand, if $\varepsilon_1 < 0$, the outcome is worse than expected under H_1 , and therefore, beliefs are downgraded to $\alpha_2 < \alpha_1$, yielding H_2 as shown in the right hand panel. In both cases, the realization of ε_1 , the expected production function H_2 is the basis for equilibrium in period 2.




Figure 5: Inferences and outcome in period 2, as a function of ε_1 .

Next we show that only a strict subset of the feasible tax rates $\tau \in [0, 1]$ are implemented in equilibrium.

Proposition 2 *Let τ_t^m be the median tax rate in any period t . Then,*

$$\tau_t^m \in [\tau_A^m, \tau_B^m] \quad \forall t.$$

Proof. By Proposition 1, in any period t the median voter's optimal tax rate under the expected production function $H_t(\tau_t)$ defined in (8) is implemented in equilibrium. Since by definition $\frac{\partial H_A}{\partial \tau}(\tau_A^m) = \frac{\partial H_B}{\partial \tau}(\tau_B^m)$ and since $H_A(\tau)$ and $H_B(\tau)$ are both concave, we know that $\frac{\partial H_A}{\partial \tau} > y^m$ and $\frac{\partial H_B}{\partial \tau} > y^m$ for all $\tau < \tau_A^m$. Hence, since $\alpha_t \leq 1$ for all t , $\tau_t^m \geq \tau_A^m$ for all t follows. Symmetric arguments can be applied to rule out $\tau_t^m > \tau_B^m$. \square

Proposition 2 is illustrated in Figure 6.

3.2 An informal discussion of the convergence results

The voters' problem in our model is basically a problem of inference. Recall that $\mathcal{H}_t \equiv \{(h_i, \tau_i)\}_{i=1}^t$ is the publicly observed history up to date t . Accordingly,




Figure 6: Range of equilibrium tax rates.

let $\Pr(H_A|\mathcal{H}_t)$ denote the conditional probability that H_A is true given history \mathcal{H}_t . Denote $\Pr(h_t|H_A, \tau_t)$ the probability of observing h_t given that H_A is true and given that policy τ_t is implemented. Then, by Bayes rule

$$\Pr(H_A|\mathcal{H}_t) = \frac{\Pr(H_A|\mathcal{H}_{t-1}) \Pr(h_t|H_A, \tau_t)}{\Pr(H_A|\mathcal{H}_{t-1}) \Pr(h_t|H_A, \tau_t) + (1 - \Pr(H_A|\mathcal{H}_{t-1})) \Pr(h_t|H_B, \tau_t)}. \quad (11)$$

Since voters are rational, they use Bayes rules (11) to update their beliefs, i.e., $\alpha_{t+1} = \Pr(H_A|\mathcal{H}_t)$, where in the initial period $t = 1$, $0 < \alpha_1 < 1$. Since the probability of observing h_t is higher under the true production function H_A than under the wrong one H_B , α_{t+1} should be expected to converge to 1 as the number of observations gets large. However, recall that the two production functions intersect at $\tilde{\tau}$ which implies that $\Pr(h_t|H_A, \tilde{\tau}) = \Pr(h_t|H_B, \tilde{\tau})$. Inspection of (11) reveals that in this case, $\alpha_{t+1} = \alpha_t$. The observation h_t is equally likely under production function H_A as under H_B . In this case, the learning process comes to a halt. Let $\tilde{\alpha}$ be the belief such that in political equilibrium $\tilde{\tau}$ is implemented. That is, $\tilde{\alpha}$ solves

$$\tilde{\alpha} H'_A(\tilde{\tau}) + (1 - \tilde{\alpha}) H'_B(\tilde{\tau}) = y^m,$$

where $\tilde{\tau}$ is such that $H_A(\tilde{\tau}) = H_B(\tilde{\tau})$. Clearly, $\tilde{\alpha} \in (0, 1)$ exists. Moreover, the fact $\frac{\partial \tilde{\alpha}}{\partial y^m} < 0$ is readily established by noting that for a given belief α , the preferred tax rate of any voter decreases in her income. Therefore, as the median income increases, a higher belief that H_B is true is required for the median voter's preferred tax rate to be $\tilde{\tau}$, and hence, $\tilde{\alpha}$ decreases in y^m .

This raises two important questions: (1) Does the stochastic process of beliefs and policies converge? (2) If so, to what beliefs and policies does it converge? In Section 3.3 we show that the process of beliefs converges to a random

variable whose support consists solely of $\tilde{\alpha}$ and 1 and that this is equivalent to saying that the policy converges to a random variable whose support is $\tilde{\tau}$ and τ_A^m .

Another question is how likely the convergence to the policy $\tilde{\tau}$ is. In Section 4 we use numerical methods to approximate the probability of reaching $\tilde{\alpha}$ and 1 (or equivalently $\tilde{\tau}$ and τ_A^m) as a function of initial conditions such as initial beliefs, the shape of the production functions and the variance of shocks. Our simulations suggest that convergence to $\tilde{\tau}$ occurs for a wide range of initial conditions. This is interesting because $\tilde{\tau}$ can be Pareto inefficient.

3.3 Convergence of the stochastic process

We now state our main results. These are (1) that in the long-run, the equilibrium policy and equilibrium beliefs converge and (2) that they do not necessarily converge to a Pareto efficient policy and the true probability, respectively. We comment on these findings after the proposition and its proof.

Proposition 3 *There exists a random variable $\tau_\infty \in [0, 1]$ such that*

1. $\tau_t^m \rightarrow \tau_\infty$ almost surely as $t \rightarrow \infty$, and
2. the support of τ_∞ is $\{\tilde{\tau}, \tau_A^m\}$.

Proof. We prove Proposition 3 by showing that the voters beliefs α_t converge to a random variable α_∞ almost surely. From Proposition 1 we then get the convergence result for τ_t^m .


We first define the function

$$s(\tau) \equiv H_A(\tau) - H_B(\tau) \quad \text{for } \tau \in [\tau_A^m, \tau_B^m]. \quad (12)$$

The fact that $s'(\tau) < 0$ for $\tau \in [\tau_A^m, \tau_B^m]$ is readily established, using $H'_A(\tau) < H'_B(\tau)$ for $\tau \in [\tau_A^m, \tau_B^m]$, which follows from concavity of both H_A and H_B and the fact that $H'_A(\tau_A^m) = H'_B(\tau_B^m)$, noted in (6). Note that for $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$, $s(\tilde{\tau}) = 0$. Therefore, $s(\tau_A^m) > 0$ and $s(\tau_B^m) < 0$. Figure 7 provides an illustration.

Let us also define the function $\tau^m(\alpha_t)$, which is the tax rate solving equation (9) as a function of the beliefs α_t . So for a given belief α_t we have $\tau_t^m = \tau^m(\alpha_t)$, the unique optimal tax rate of the median voter. Using the implicit function theorem, we have

$$\frac{\partial \tau_t^m}{\partial \alpha_t} = \frac{-s'(\tau_t^m)}{\alpha_t H''_A(\tau_t^m) + (1 - \alpha_t) H''_B(\tau_t^m)} < 0, \quad (13)$$

Figure 7: An illustration of function $s(\tau)$.

since $-s' > 0$ and $\alpha_t H_A'' + (1 - \alpha_t) H_B'' < 0$ by concavity. This is also quite intuitive. As the beliefs that H_A is true increase, the equilibrium tax rate decreases, i.e., is closer to τ_A^m . Finally, let us define

$$w(\alpha_t) \equiv s(\tau^m(\alpha_t)), \quad (14)$$

which gives us the difference between the two production function in equilibrium as a function of the beliefs in period t . The function w is defined on the interval $[0, 1]$. The fact that $\frac{\partial w}{\partial \alpha_t} = s' \tau^{m'} > 0$ follows immediately from the above observations. Moreover, because with $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$, $s(\tilde{\tau}) = 0$, we have $w(\alpha(\tilde{\tau})) = 0$ for a unique $\tilde{\alpha} \in (0, 1)$ and $0 > w(0) > -\infty$ for $0 < w(1) < \infty$.

Let $\alpha_1 = \Pr(H_A)$ and $1 - \alpha_1 = \Pr(H_B)$ be the exogenously given prior beliefs that H_A and H_B are true, respectively. After observing history $\mathcal{H}_1 = (h_1, \tau_1)$, voters apply Bayes rule to get

$$\begin{aligned} \alpha_2 &= \Pr(H_A | \mathcal{H}_1) = \frac{\alpha_1 \Pr(h_1 | H_A)}{\alpha_1 \Pr(h_1 | H_A) + (1 - \alpha_1) \Pr(h_1 | H_B)} \\ &= \frac{\Pr(H_A) \Pr(h_1 | H_A)}{\Pr(H_A) \Pr(h_1 | H_A) + \Pr(H_B) \Pr(h_1 | H_B)}. \end{aligned}$$

After observing history $\mathcal{H}_2 = \{(h_i, \tau_i)\}_{i=1}^2$, they use α_2 and Bayes rule to get

$$\begin{aligned} \alpha_3 &= \Pr(H_A | \mathcal{H}_2) = \frac{\alpha_2 \Pr(h_2 | H_A)}{\Pr(H_A) \Pr(h_2 | H_A) + (1 - \alpha_2) \Pr(h_2 | H_B)} \\ &= \frac{\Pr(H_A) \Pr(h_1 | H_A) \Pr(h_2 | H_A)}{\Pr(H_A) \Pr(h_1 | H_A) \Pr(h_2 | H_A) + \Pr(H_B) \Pr(h_1 | H_B) \Pr(h_2 | H_B)}. \end{aligned}$$

By induction, after observing history $\mathcal{H}_t = \{(h_i, \tau_i)\}_{i=1}^t$, we'll have

$$\alpha_{t+1} = \frac{\Pr(H_A) \Pr(h_1 | H_A) \Pr(h_2 | H_A) \dots \Pr(h_t | H_A)}{\Pr(H_A) \Pr(h_1 | H_A) \Pr(h_2 | H_A) \dots \Pr(h_t | H_A) + \Pr(H_B) \Pr(h_1 | H_B) \Pr(h_2 | H_B) \dots \Pr(h_t | H_B)}$$

or equivalently

$$\alpha_{t+1} = \frac{1}{1 + \frac{\Pr(H_B) \Pr(h_1|H_B) \Pr(h_2|H_B) \dots \Pr(h_t|H_B)}{\Pr(H_A) \Pr(h_1|H_A) \Pr(h_2|H_A) \dots \Pr(h_t|H_A)}}. \quad (15)$$

Since by assumption ε_t is distributed according to the normal with mean zero and variance σ^2 , which we denote as $\phi(\cdot)$, substituting yields⁶

$$\begin{aligned} \Pr(h_t|H_A) &= \phi(h_t - H_A(\tau_t)) = \phi(\varepsilon_t) \quad \text{and} \\ \Pr(h_t|H_B) &= \phi(h_t - H_B(\tau_t)) = \phi(s(\tau_t) + \varepsilon_t). \end{aligned}$$

Thus, using (14) we can write (15) as

$$\alpha_{t+1} = \frac{1}{1 + \frac{(1-\alpha_1)\phi(w(\alpha_1)+\varepsilon_1)\phi(w(\alpha_2)+\varepsilon_2)\dots\phi(w(\alpha_t)+\varepsilon_t)}{\alpha_1\phi(\varepsilon_1)\phi(\varepsilon_2)\dots\phi(\varepsilon_t)}}. \quad (16)$$

Define

$$N_{t+1} \equiv \frac{(1 - \alpha_1)\phi(w(\alpha_1) + \varepsilon_1)\phi(w(\alpha_2) + \varepsilon_2) \cdot \dots \cdot \phi(w(\alpha_t) + \varepsilon_t)}{\alpha_1\phi(\varepsilon_1)\phi(\varepsilon_2) \cdot \dots \cdot \phi(\varepsilon_t)}, \quad (17)$$

such that (16) becomes

$$\alpha_{t+1} = \frac{1}{1 + N_{t+1}}. \quad (18)$$

That is, (18) defines $\alpha_t = \alpha(N_t)$ with $\frac{\partial \alpha(N_t)}{\partial N_t} < 0$. Note also that $\alpha_{t+1} \in (0, 1] \Leftrightarrow N_{t+1} \in [0, \infty)$. Moreover, we can now define a sequence of random variables $\{N_i\}_{i=1}^t$, the initial value of which is exogenously given as $N_1 = \frac{1-\alpha_1}{\alpha_1}$. Finally define $r(N_t) \equiv w(\alpha(N_t))$, where

$$\frac{\partial r}{\partial N_t} = \frac{\partial w}{\partial \alpha_t} \frac{\partial \alpha_t}{\partial N_t} < 0$$

is readily established. It is also easy to see that $r(0) = w(1) > 0$ and that $\lim_{N_t \rightarrow \infty} r(N_t) = w(0) < 0$, both finite. Thus, for $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$, there is a unique \tilde{N} such that

$$r(\tilde{N}) = 0. \quad (19)$$

In light of these new definitions,

$$\begin{aligned} N_{t+1} &= N_1 \cdot \frac{\phi(r(N_1) + \varepsilon_1)}{\phi(\varepsilon_1)} \cdot \frac{\phi(r(N_2) + \varepsilon_2)}{\phi(\varepsilon_2)} \cdot \dots \cdot \frac{\phi(r(N_t) + \varepsilon_t)}{\phi(\varepsilon_t)} \\ &= N_t \cdot \frac{\phi(r(N_t) + \varepsilon_t)}{\phi(\varepsilon_t)} = N_t \cdot e^{-\frac{r(N_t)}{2\sigma^2}(r(N_t)+2\varepsilon_t)}. \end{aligned} \quad (20)$$

⁶Note that for a continuous random variable any single observation has probability zero. Nonetheless, L'Hopital's rule can be used to determine to posterior probability, so that the density rather than the cdf is appropriate.

Notice that (20) is a non-linear stochastic first-order difference equation.

Observe first that if the sequence takes either the value 0, the value \tilde{N} , or is infinity, it will take this value forever. This becomes immediate for $N_t = 0$ by inserting $N_t = 0$ into (20). For $N_t = \tilde{N}$, note that $r(\tilde{N}) = 0$ implies that the exponent in (20) becomes 0 for any ε_t , implying $N_{t+1} = \tilde{N}$. If N_t is infinity, N_{t+1} will be too, since $\lim_{N_t \rightarrow \infty} r(N_t)$ is a finite negative number.

Note also that the sequence $\{N_t\}$ is a martingale. The reason is first that

$$\begin{aligned} E[N_{t+1}] &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} N_{t+1} \cdot \phi(\varepsilon_1, \varepsilon_2, \dots, \varepsilon_t) d\varepsilon_1 d\varepsilon_2 \dots d\varepsilon_t \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} N_1 \cdot \phi(r(N_1) + \varepsilon_1) \cdot \dots \cdot \phi(r(N_t) + \varepsilon_t) d\varepsilon_1 d\varepsilon_2 \dots d\varepsilon_t \\ &= N_1 < \infty, \end{aligned}$$

where the joint normal $\phi(\varepsilon_1, \varepsilon_2, \dots, \varepsilon_t) = \phi(\varepsilon_1) \cdot \phi(\varepsilon_2) \cdot \dots \cdot \phi(\varepsilon_t)$ by independence. Second,

$$\begin{aligned} E[N_{t+1} | \{N_i\}_{i=1}^t] &= N_t \int_{-\infty}^{\infty} \frac{\phi(r(N_t) + \varepsilon_t)}{\phi(\varepsilon_t)} \phi(\varepsilon_t) d\varepsilon_t \\ &= N_t \int_{-\infty}^{\infty} \phi(r(N_t) + \varepsilon_t) d\varepsilon_t = N_t. \end{aligned}$$

The martingale convergence theorem (e.g., Durrett, 2005, p. 233) states that N_t converges almost surely to a limit N_∞ with $E[N_\infty] < \infty$. For the interpretation of our model, it is necessary to evaluate the random variable N_∞ . Lemma 1 states that the martingale either converges towards 0 or towards \tilde{N} .

Lemma 1 *The support of the random variable N_∞ is $\{0, \tilde{N}\}$.*

Proof. From the observation we made above, we know that $\Pr(N_{t+1} = 0 | N_t = 0) = 1$ and $\Pr(N_{t+1} = \tilde{N} | N_t = \tilde{N}) = 1$. We now prove by contradiction that there exists no other value C the martingale N_t can converge to. Note, that the martingale convergence theorem directly states that N_t cannot converge to infinity.

Assume there exists a number $C \in (0, \infty)$ where N_t can converge to. Then, for every $\delta \in \mathbb{R}$ such that $0 \notin [C - \delta, C + \delta]$ and $\tilde{N} \notin [C - \delta, C + \delta]$, there exists a time period t_δ , for which we have $N_{t_\delta+i} \in [C - \delta, C + \delta]$ for $i = 0, 1, \dots$. Note that δ can be chosen arbitrarily small. Now define the variable $\bar{\varepsilon}_{t_\delta+i}$ by

$$\bar{\varepsilon}_{t_\delta+i} \equiv \frac{\sigma^2}{r(N_{t_\delta+i})} \cdot \ln \frac{N_{t_\delta+i}}{C + \delta} - \frac{1}{2} r(N_{t_\delta+i}). \quad (21)$$

Note that $\bar{\varepsilon}_{t_\delta+i}$ is a shock such that $N_{t_\delta+i+1} = C + \delta$. Assume that $C < \tilde{N}$. Then, the variable $\bar{\varepsilon}_{t_\delta+i}$ is negative and finite for all $N_{t_\delta+i} \in [C - \delta, C + \delta]$, because all terms in (21) are finite. Therefore, for every $N_{t_\delta+i} \in [C - \delta, C + \delta]$,

$$\Pr(\varepsilon_{t_\delta+i} < \bar{\varepsilon}_{t_\delta+i}) = \Phi(\bar{\varepsilon}_{t_\delta+i}) > 0, \quad (22)$$

which means that the probability to draw an $\varepsilon_{t_\delta+i} < \bar{\varepsilon}_{t_\delta+i}$ is strictly positive for every $N_{t_\delta+i} \in [C - \delta, C + \delta]$. Thus, with a positive probability we observe an $N_{t_\delta+i+1} > C + \delta$ for every period $t_\delta + i$ because $N_{t_\delta+i+1}$ depends negatively on $\varepsilon_{t_\delta+i}$. This means, that

$$\inf_{N_{t_\delta+i} \in [C-\delta, C+\delta]} \Pr(N_{t_\delta+i+1} \notin [C - \delta, C + \delta]) > 0,$$

which is a contradiction to the assumption of convergence of N_t . Hence, N_t cannot converge to C .

In order to prove non-convergence towards a $C > \tilde{N}$, we define $\underline{\varepsilon}_{t_\delta+i}$ as

$$\underline{\varepsilon}_{t_\delta+i} \equiv \frac{\sigma^2}{r(N_{t_\delta+i})} \cdot \ln \frac{N_{t_\delta+i}}{C - \delta} - \frac{1}{2} r(N_{t_\delta+i})$$

and use the equivalent reasoning as above.

We are now only left to show that the probability of N_t converging to the set union of all C is still 0. By choosing intervals around C with rational endpoints, the probabilities can be summed up for the union set. Since we can choose δ arbitrarily, it is always possible to find an interval with rational endpoints for all C . Therefore, the sum of probabilities over these intervals is 0. This completes the proof of Lemma 1. \square

From Slutski's Theorem we know that if N_t converges to N_∞ with support $\{0, \tilde{N}\}$ almost surely, then α_t converges to α_∞ with support $\{\tilde{\alpha}, 1\}$ almost surely. For the belief $\alpha_t = 1$ the tax rate τ_A^m is implemented, for $\tilde{\alpha}$ it is $\tilde{\tau}$. Therefore, the support of τ_∞ is $\{\tau_A^m, \tilde{\tau}\}$. This completes the proof of Proposition 3. \square

3.4 The efficiency potential

Proposition 3 states that the economy converges to either τ_A^m or $\tilde{\tau}$. If $\tau_A^H < \tilde{\tau} < \tau_B^I$ the Pareto sets of H_A and H_B are disjoint and $\tilde{\tau}$ lies in between them, i.e., is Pareto inefficient. The conditions for this require that H_A and H_B are sufficiently different. From now on we assume that $\tilde{\tau}$ is Pareto inefficient.

Corollary 1 *If $\tau_A^H < \tilde{\tau} < \tau_B^I$, then the economy can converge to a Pareto inefficient policy.*

An interesting question is how likely it is that voters end up with a Pareto efficient policy. However, the distribution of τ_∞ cannot be determined analytically. That means, the probabilities that the political economy converges to $\tilde{\tau}$ and to τ_A^m cannot be derived analytically as a function of initial conditions. In Section 4 we will use numerical simulations to approximate this distribution. Yet we attain an analytical result for the lower bound of the probability that the policy converges to τ_A^m . For that purpose, we define the efficiency potential as this minimal probability, which we denote as ξ . That is,

$$\xi \equiv \inf \Pr \left(\lim_{t \rightarrow \infty} \tau_t \rightarrow \tau_A^m \mid \alpha_1, \tilde{\tau} \right).$$

Proposition 4 $\xi = \max \left\{ 0, \frac{\alpha_1 - \tilde{\alpha}}{\alpha_1(1 - \tilde{\alpha})} \right\}$.

Proof. From Proposition 3 we know that α_t either converges to 1 or to $\tilde{\alpha}$. What we need to characterize in order to prove Proposition 4 is actually the distribution of the random variable N_∞ over $\{0, \tilde{N}\}$, from which we can then deduce the distribution of the random variable α_∞ over $\{1, \tilde{\alpha}\}$.

Corollary 2.11 in Durrett (2005) implies that $E[N_\infty] \leq E[N_1]$. Let μ be the probability of convergence towards \tilde{N} . Then

$$\begin{aligned} E[N_\infty] &= (1 - \mu) \cdot 0 + \mu \cdot \tilde{N} = \mu \cdot \tilde{N} \\ \Rightarrow \quad \mu &\leq \frac{N_1}{\tilde{N}} \quad \Rightarrow \quad (1 - \mu) \geq 1 - \frac{N_1}{\tilde{N}}, \end{aligned}$$

where it will be recalled that $N_1 = \frac{1 - \alpha_1}{\alpha_1}$. As it is a probability, ξ must be nonnegative. It equals the minimum value of $(1 - \mu)$ if $(1 - \mu) > 0$. It follows that

$$\xi = \max \left\{ 0, 1 - \frac{N_1}{\tilde{N}} \right\} = \max \left\{ 0, 1 - \frac{\frac{1 - \alpha_1}{\alpha_1}}{\frac{1 - \tilde{\alpha}}{\tilde{\alpha}}} \right\} = \max \left\{ 0, \frac{\alpha_1 - \tilde{\alpha}}{\alpha_1(1 - \tilde{\alpha})} \right\}.$$

□

Note that unless $\alpha_1 = 1$ (in which case the problem is degenerate), ξ is strictly less than one. Taking first derivatives, we get $\frac{\partial \xi}{\partial \alpha_1} > 0$ and $\frac{\partial \xi}{\partial \tilde{\alpha}} < 0$ for $\xi > 0$. Clearly, these derivatives are only valid for $\alpha_1 > \tilde{\alpha}$. Otherwise, $\xi' = 0$.

The first observation is intuitive, since it is natural to expect voters who are initially better informed to be more likely to converge to the correct belief in the long run. The sign of the derivative $\frac{\partial \xi}{\partial \tilde{\alpha}} < 0$ is also intuitive, but understanding

it requires a moment's reflection. For a given $\alpha_1 > \tilde{\alpha}$, a series of bad shocks is required for the beliefs to be downgraded to $\tilde{\alpha}$. Obviously, as $\tilde{\alpha}$ decreases, a longer series of bad shocks is required for beliefs to be downgraded to $\tilde{\alpha}$. Since a longer series of bad shocks is less likely, the efficiency potential increases as $\tilde{\alpha}$ decreases. As noted above, $\frac{\partial \tilde{\alpha}}{\partial y^m} < 0$. Therefore, the efficiency potential increases in the median income. That is, on average richer countries should be associated with better policies. Note, though, that this prediction of the model hinges on the assumption that H_A is the true production function. Were H_B true, then the efficiency potential would decrease in y^m .

4 Numerical Results

Of course, we are not only interested in determining the efficiency potential, which after all gives us only a minimal probability of reaching the good policy. It is equally interesting to learn something about the probability of implementing the bad policy in the long run. Unfortunately, as explained above, the distribution of α_∞ cannot be calculated explicitly. We therefore have to rely on simulations in order to approximate the probability that beliefs converge to $\alpha_\infty = \tilde{\alpha}$ and $\alpha_\infty = 1$, respectively. This probability is a function of initial beliefs, the noisiness of the production functions, and the production functions, and in particular of the slope $s(\tilde{\tau})$.

4.1 Simulations for different initial beliefs

The simulation results are collected in the two tables below for two different constellations of production functions. Figure 8 shows three functions which are taken as the production function of the public good. For Table 1, we use the blue function (H_A) as the true production function, and the green function (H_G) as the alternative production function. For Table 2, again the blue function (H_A) is the true production function and the red one (H_R) is the alternative. An entry in the table is the share of draws for which the belief converged to 1 for a given combination of initial belief α_1 and noise σ . For every entry we did a hundred draws. One minus the table entry gives the share of draws that converged to the inefficient tax rate.⁷ For example, the 1 in the top left entry of Table 1 means that for $\alpha_1 = 0.1$ and $\sigma = 0.2$ every draw converged to 1, for the blue (true) and green (untrue) production function. Note that the smaller

⁷It is reassuring that all draws either converge to τ_A^m or to $\tilde{\tau}$.

H_A and H_G $\tilde{\alpha} = 0.47$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$	ξ
$\alpha_1 = 0.1$	1	0.99	0.21	0.01	0
$\alpha_1 = 0.2$	1	0.98	0.26	0.02	0
$\alpha_1 = 0.3$	1	0.98	0.24	0	0
$\alpha_1 = 0.4$	1	0.99	0.21	0	0
$\alpha_1 = 0.5$	1	0.97	0.25	0.10	0.12
$\alpha_1 = 0.6$	1	1	0.57	0.48	0.42
$\alpha_1 = 0.7$	1	0.99	0.76	0.67	0.62
$\alpha_1 = 0.8$	1	1	0.94	0.76	0.78

Table 1: Results when H_A is true and H_G is the alternative.

H_A and H_R $\tilde{\alpha} = 0.52$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$	ξ
$\alpha_1 = 0.1$	0.99	0.36	0.01	0	0
$\alpha_1 = 0.2$	1	0.29	0	0	0
$\alpha_1 = 0.3$	1	0.28	0	0	0
$\alpha_1 = 0.4$	1	0.29	0	0	0
$\alpha_1 = 0.5$	1	0.21	0	0	0
$\alpha_1 = 0.6$	1	0.58	0.31	0.29	0.27
$\alpha_1 = 0.7$	1	0.75	0.59	0.58	0.53
$\alpha_1 = 0.8$	1	0.92	0.76	0.71	0.73
$\alpha_1 = 0.9$	1	0.96	0.91	0.90	0.88

Table 2: Results when H_A is true and H_R is the alternative.




Figure 8: The functions used for the simulations reported in Tables 1 and 2.

σ , the higher the probability of reaching τ_A^m . This is intuitive because a smaller variance of the shocks increases the informativeness of the policy outcome.

4.2 Probability of convergence as a function of noise

It is also interesting to see how the long run equilibrium depends on the noise in the production function. For that purpose, we simulated an economy with two given production functions and given initial beliefs, and let only the variance of the error term vary. The results are depicted in Figure 9.

Specifically, we simulated for the functions $H_A(\tau) = \ln(\tau + 0.003) - 4(\tau + 0.003) + 6$ and $H_B = 8.9\tau - 5.5\tau^2$, which gives us a $\tilde{\tau} = 0.543$. We set $\alpha_1 = \frac{1}{2}$ and $y^m = \frac{1}{2}$. Note that the maximal value of H_A is 3.61. Then $\tilde{\alpha} = 0.5744$, and the efficiency potential of this political economy is 0. The results are very




Figure 9: Simulation results when only the variance of the error varies.

intuitive as it is natural to expect that the noisier public production, the harder it is to learn the truth and consequently, the less likely it is to converge to τ_A^m .

As we discuss in the conclusions, this indicates that if it is possible to reduce σ^2 , this would be very effective in increasing the likelihood of convergence to τ_A^m .

4.3 The information trap

The reason why there is an environment around $\tilde{\tau}$ from which the policy can eventually not escape is that the two production functions have very similar values in the neighborhood of $\tilde{\tau}$. The closer one gets to $\tilde{\tau}$, the less distinguishable the true and the false production function become. Once one is close enough to $\tilde{\tau}$, it thus becomes very difficult to learn anything from observation. Hence, the economy becomes stuck with its current beliefs once these are sufficiently close to $\tilde{\alpha}$, as a consequence of which policy will not change anymore. Hence, one can speak of an information trap around $\tilde{\tau}$, because voters cannot gather any new information.

All voters then know and perfectly agree that policy $\tilde{\tau}$ is not Pareto efficient. That is, they are all perfectly aware that their policy is lost somewhere in the middle. So, why do they not just change the policy? As the same policy affects different people in different ways, they do not agree in which direction they should move. Given beliefs $\tilde{\alpha}$, low income voters would prefer tax rates $\tau > \tilde{\tau}$, while rich individuals would prefer smaller tax rates, and the median voter y^m finds $\tilde{\tau}$ optimal. Since once $\tilde{\tau}$ is implemented, it will be implemented forever, there is a kind of prisoner's dilemma flavor associated with this outcome.

The simulation results reported below strongly suggest that the economy can converge to $\tilde{\tau}$ for a wide range of initial conditions. This suggests that even if the median voter (and all his neighbors) could coordinate on some small policy experiments and vote for tax rates slightly higher or lower than $\tilde{\tau}$, the society still faces the problem that it will eventually fall back into the trap. In order to truly escape the trap, some large scale experimentation would be required, like, e.g., implementing $\frac{\tilde{\tau}}{2}$ or $2\tilde{\tau}$, in order to induce beliefs to change substantially enough. So as to make clear how voters could coordinate on such a change, a very different model would have to be developed.

5 Conclusions

Putnam (1993) has raised the question why some democratic governments fail and others succeed. He explains the failure and success of democracies by referring to differences in political institutions and attitudes. We have provided an

alternative explanation why, in general, political outcomes in initially identical societies may differ in the long run and more specifically, why some democracies may adopt Pareto inferior policies even in the long run. Our explanation, which we see largely as complementary to Putnam's, rests on the assumption that voters face uncertainty and that uncertainty can only be unravelled by experience. The basic reason why initially identical countries may end up with different outcomes is that in combination with bad luck the political equilibrium may impede further inferences, so that uncertainty is never abolished.⁸ Since in our model economies may fail to converge to Pareto efficient policies as a consequence of bad shocks, its predictions are consistent with the observations of Easterly (2001), who notes that some countries' meager growth performance may be caused by bad luck.

Are there any lessons to be learnt from our model for policy makers and policy advisors? Though it is arguably hard or even impossible to directly affect people's beliefs, it is not necessarily true that these beliefs cannot be influenced at all. Anything that reduces the variance in the public production function has a positive effect on the probability of converging to a Pareto efficient policy. So, if it is possible to reduce this variance, e.g., by sharpening the predictions of the competing models a society believes in, the long run beliefs of a society, and consequently its policies, may be different. Thus, even in our model societies are not simply doomed to fail.

To be sure, a lot of work remains to be done for the future. The robustness of the model could be checked by allowing for altruistic politicians or abstentionism. It would also be desirable to see how the model behaves if the distribution of income were endogenous. Another extension would be to let voters have different beliefs so as to formalize the idea that "many political conflicts are not only conflicts of interest but also conflicts about how the economy functions", as expressed by Saint-Paul (2000, p. 917).

⁸Among other things, we have shown that initial beliefs may be crucial for the long run political outcome. This may help better understand the economic and political difficulties former colonies face who may have been endowed with bad initial beliefs at the time of independence, as emphasized, e.g., by Bauer (1981).

References

- Alesina, A. and G.-M. Angeletos (2003). Fairness and redistribution. UCLA Department of Economics, Levine's Bibliography.
- Bauer, P. T. (1981). *Equality, the Third World, and Economic Delusion*. Harvard University Press.
- Blendon, R. J., J. M. Benson, M. Brodie, R. Morin, D. E. Altman, D. Gitterman, M. Brossard, and M. James (1997). Bridging the gap between the public's and economists views of the economy. *Journal of Economic Perspectives* 11(3), 105–118.
- Blumkin, T. and V. Grossmann (2004). Ideological polarization, sticky information, and policy reforms. Working Paper.
- Durrett, R. (2005). *Probability: Theory and Examples*. Curt Hinrichs.
- Easley, D. and N. M. Kiefer (1988). Controlling a stochastic process with unknown parameters. *Econometrica* 56(5), 1045–64.
- Easterly, William, S. (2001). *The Elusive Quest for Growth*. MIT Press.
- Feddersen, T. J. and W. Pesendorfer (1996). The swing voter's curse. *American Economic Review* 86(3), 408–424.
- Fuchs, V. R., A. S. Blinder, and J. M. Poterba (1998). Economists' views about parameters, values and policies: Survey results in labor and public economics. *Journal of Economic Literature* 36(3), 1387–1425.
- Hogg, R. and A. T. Craig (1995). *Introduction into mathematical statistics*. Englewood Cliffs, N.J. : Prentice Hall.
- McLennan, A. (1984). Price dispersion and incomplete learning in the long run. *Journal of Economic Dynamics and Control* 7(3), 331–47.
- Meltzer, A. H. and S. F. Richard (1981). A rational theory of the size of government. *Journal of Political Economy* 89(5), 914–27.
- Persson, T. and G. Tabellini (2000). *Political Economics*. MIT Press, Massachusetts.
- Piketty, T. (1995). Social mobility and redistributive politics. *Quarterly Journal of Economics* 110(3), 551–84.

- Putnam, R. D. (1993). *Making Democracy Work*. Princeton University Press.
- Saint-Paul, G. (2000). The "New Political Economy": Recent Books by Allen Drazen and by Torsten Persson and Guido Tabellini. *Journal of Economic Literature* 38(4), 915–25.
- Schultz, C. (2004). Information, polarization, and accountability in democracy. Working Paper.
- Spector, D. (2000). Rational debate and one-dimensional conflict. *The Quarterly Journal of Economics* 115(1), 181–200.