Learning, public good provision, and the information trap*

Aleksander Berentsen[†], Esther Bruegger[‡] and Simon Loertscher[§]

October 15, 2007

Abstract

We consider an economy with uncertainty about the true production function for a public good. By using Bayes rule the economy can learn from experience. We show that it may learn the truth, but that it may also converge to an inefficient policy where no further inference is possible so that the economy is stuck in an information trap. We also show that our results are robust with respect to small experimentation.

Keywords: Public economics, learning, size of government.

JEL-Classification: D72, H10, D83.

*The paper has benefitted from very insightful comments of two anonymous referees and the editor. We also want to thank Joshua Gans, Roland Hodler, Andy McLennan, Torsten Persson, Fabienne Peter, Ronny Razin, Kurt Schmidheiny, Christian Schultz, Martin Wagner, Volker Wieland and seminar participants at the University of Bern, SSES 2004 in Basel, PCS 2005 in New Orleans, EPCS 2005 in Durham and AEA in Boston 2006 for most valuable discussions and comments. Very special thanks go to Roberto Schonmann. Financial support by the WWZ-Forum Basel is gratefully acknowledged. A previous version of the paper circulated under the title "Learning, voting and the information trap". Any remaining errors are ours.

†Economics Department, University of Basel, Petersgraben 51, CH-4003 Basel. eMail: aleksander.berentsen@unibas.ch

‡NERA Economic Consulting, 1166 Avenue of the Americas, 34th Floor, New York, NY 10036. eMail: estherbruegger@gmx.net

§Economics Department, University of Melbourne, Economics & Commerce Building, Victoria 3010, Australia. eMail: simonl@unimelb.edu.au

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

Immanuel Kant (1784)

1 Introduction

Broadly speaking, there are two theories regarding the effects of government activity on the economy.[1] Some economists emphasize the crucial role of government in securing property rights, enforcing contracts, providing national security and, perhaps, guaranteeing a moderate minimum income for every one. These proponents do not deny that some government activity is better than none and they would probably argue that at a small scale, public production exhibits very large marginal productivity. These marginal products, however, then decline quickly and eventually become negative. Other economists believe that government is most productive if it operates on a large scale because of increasing returns. According to this view, operating on a small scale, the marginal product of government activity is moderate as it merely serves to appease the poor, yet fails to exhaust their full economic potential.[2]

The question which of these two theories is right cannot be answered by a priori arguments.[3] However, it is crucial to know whether in the long-run efficient policies are

[1]See, e.g., Hayek (1944), Hazlitt (1946) and Friedman (1962, 1997) or Rosenstein-Rodan (1943), Myrdal (1975) and Sachs (2005).

[2]This hypothesis is consistent with Acemoglu and Robinson (2000).

[3]Blendon, Benson, Brodie, Morin, Altman, Gitterman, Brossard, and James (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, Blinder, and Poterba (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. Bartels (1996) notes that the "[political ignorance of the American voter is one of the best-

chosen, i.e., whether experience will eventually lead the economy to learn the truth. This is the question we address in this paper. For this purpose, we construct a model with uncertainty about how the economy functions. The decision maker does not know which of two possible production functions for a public good is the true one. In any period, for any given belief, the decision maker chooses a tax rate to maximize his short-run expected utility and, after observing the realized level of production of the public good, he updates his beliefs using Bayes' rule. We show that in the long-run the true production function may be learned, but the economy may also converge to an inefficient policy where no further inference is possible so that the economy is stuck in an information trap. We also show that this result is robust with respect to experimentation.

The paper relates to two strands of literature. Our main result is related to the well-known fact that impatient Bayesian learners can optimally fail to learn the true parameter values (see, e.g., Easley and Kiefer, 1988).[4] In this strand of literature, the most important predecessor to our paper is McLennan (1984) who studies learning by a monopolistic seller who faces two linear demand functions intersecting at some price and who is uncertain about which of the two is true.[5] The paper is also related to Laslier, Trannoy, and Van Der Straeten (2003) who study voting over unemployment

documented features of contemporary politics, but the political significance of this political ignorance is far from clear."

[4]Insofar as incomplete learning is concerned, a very similar phenomenon obtains in models of herding such as Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). However, since there is no repeated decision taking by the same agent(s) in these herding models, there is also no scope for learning over time, and eventually experimentation.

[5]In every period, the seller in McLennan's model observes whether there is a sale or not and updates his beliefs accordingly. Among other things, he shows that with positive probability the seller ends up charging the price where the two demand functions intersect, at which point no further learning is possible. A more detailed discussion of the relationship between our model and McLennan's is deferred to the end of Section 3.

benefits when households do not know the (a fortiori unobserved) distribution of skills of the unemployed. They uncover a possibility of inefficiency that is quite similar to our finding. An important contrast is that in our model the dynamics are not monotone.[6]

The paper also relates to the political economy literature on heterogenous social beliefs that are consistent with either multiple equilibria or long-run divergence in beliefs, such as Piketty (1995), Spector (2000) and Alesina and Angeletos (2005). In contrast to Piketty and Spector, in our model all households share the same information and beliefs, but as in their model they are eventually hindered from learning the truth. An important difference between our model and the one of Alesina and Angeletos is that their equilibria can be ranked unambiguously only from the point of view of the median household.[7]

The remainder of the paper is structured as follows. Section 2 introduces the model. Section 3 analyzes the dynamic learning process. Section 4 extends the model by introducing experimentation by the policy maker. Section 5 concludes. All the proofs are in the Appendix.

[6]The dynamics of our model are more similar to those in Baron (1996), who analyzes voting over public goods programs by a legislature when there is uncertainty about which legislators can make proposals in future periods. As in our model, the economy "hops" towards its absorbing state, which in his model is given by the complete information bliss point of the median voter. In contrast to Baron's model, we have two absorbing states, one of which can be Pareto inefficient.

[7]Moreover, the sources of multiplicity are quite different. In their model, multiplicity stems from differences in social beliefs about which fraction of income is fair or merited, whereas in ours it arises from incomplete information and incomplete learning.

2 The model

There is a continuum of individuals whose total mass is normalized to one. Individual income y_i is distributed according to the density function $f(y_i)$. The mean income is one. The support of the distribution is $[y^{inf}, y^{sup}]$ with $0 < y^{inf} < y^{sup} < \infty$. 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. Individual i's utility is $u_i = c_i + H(g)$. Note that households are identical with respect to their valuation of the public good. Since mean income is one, the government's budget constraint is $g = \tau$, where $0 \leq \tau \leq 1$ is a flat tax rate. Accordingly, individual i's consumption is $c_i = (1 - \tau)y_i$.[8]

We assume that $H(g)$ is twice differentiable, strictly concave and increasing in g for g close to zero. This assumption makes sure that for every household there is a unique bliss point tax rate. Using the budget restrictions $g = \tau$ and $c_i = (1 - \tau)y_i$, we can replace c_i and g and write i's utility as a function of the tax rate only, $u_i(\tau) = (1 - \tau)y_i + H(\tau)$. Note that $H(\tau)$ is concave in τ. We denote by τ^i individual i's optimal tax rate, which is implicitly defined by $\frac{\partial H}{\partial \tau}(\tau^i) = y_i$. Since $H(\tau)$ is concave, τ^i is decreasing in y_i. Thus, the single crossing property is satisfied (see Gans and Smart, 1996; Persson and Tabellini, 2000, ch. 2, condition 2.4). Note that our households have different preferences regarding the role of the government in the economy along the lines of Alesina and Angeletos (2005). The differences arise because the flat tax affects household consumption differently due to the different income levels.

[8]Alternatively, one could assume that individual i's utility is of the form $u_i = c_i + v(h)$, where $v(.)$ is a function of the public good produced and satisfies Inada conditions. The detailed arguments why such a model is completely equivalent to the one we analyze are available upon request.

We further assume that H and the income distribution satisfy $y^{sup} < \frac{\partial H}{\partial g}(0)$ and $\frac{\partial H}{\partial g}(y) < y^{inf}$. This amounts to assuming that the public good is not a perfect substitute for private consumption. Therefore, even the poorest person wants at least a bit private consumption (i.e. a tax rate smaller than one), and even the richest individual wants some public production. Put differently, only the tax rates $\tau \in P \equiv [\tau^I, \tau^{II}]$ with $\tau^I \equiv H'^{-1}(y^{sup}) > 0$ and $\tau^I \equiv H'^{-1}(y^{inf}) < 1$ will be Pareto efficient.

2.1 The decision maker's tax policy

Denote by $\tau^m \in P$ the tax rate chosen by the decision maker, where τ^m maximizes the period expected utility of household m with income y^m. Throughout the paper we abstract from modeling the political process that yields τ^m. The decision maker, therefore, can be a benevolent or malevolent dictator or a democratically elected president. He might care only for the richest (poorest) household by choosing $\tau^m \in \tau^{II}$ ($\tau^m \in \tau^I$) or for any other household in between. The model is quite general in that it encompasses all political processes that yield a tax policy with a tax rate $\tau^m \in P$. As a special case, it includes a political process that yields a tax rate that maximizes the utility of the median income household.[9]

2.2 Two production functions

We now introduce two production functions $H_A(\tau)$ and $H_B(\tau)$ satisfying the assumptions made above. These are supposed to represent the two distinct, commonly held views on

[9] In an earlier version of the paper (Berentsen, Bruegger, and Loertscher, 2005), we constructed a model with electoral competition along the lines of Persson and Tabellini (2000, ch. 3). In the median voter equilibrium, both parties choose the tax rate that maximizes the median income household's expected utility.

the effect of government activity on the economy described in the Introduction. Without loss of generality, we assume that $H_A(\tau)$ is the true production function.

Let $P_A \equiv [\tau_A^I, \tau_A^{II}]$ and $P_B \equiv [\tau_B^I, \tau_B^{II}]$ 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 household m under H_A and H_B, i.e., $\frac{\partial H_A}{\partial \tau}(\tau_A^m) = y^m \quad \text{and} \quad \frac{\partial H_B}{\partial \tau}(\tau_B^m) = y^m$. Note that $\tau_A^m \in P_A$ and $\tau_B^m \in P_B$ and observe that τ_A^m can be called the Kantian policy since it is the policy that would be chosen if the true production function were known to the decision maker. Without loss of generality we assume that $\tau_A^m < \tau_B^m$.

We assume also that the two functions cross at most once for $\tau > 0$. If they do not cross, or cross only at some $\tau \notin [\tau_A^m, \tau_B^m]$, the updating problem is fairly simple, as we shall see in Corollary 1 below. Henceforth, with the exception of this corollary, we focus on the case where the two production functions cross, as depicted in the bottom line of Figure 1.

The production function representing the view that the optimal size of government is small has a shape like H_A in the bottom panels in Figure 1. It is very steep when τ is close to zero, but then flattens quickly and eventually decreases in τ. The production function reflecting the view that government is most efficient if large has a shape similar to H_B in the bottom panels in Figure 1, which is not very steep at the origin but flattens much slower than H_A. If our sketch of these two opposing views is correct, then the two functions H_A and H_B will have to intersect at some point, which we denote by $\tilde{\tau}$.[10]

The production of the public good is exposed to uncertainty. If τ_t is the tax rate in

[10]Conceptionally, $\tilde{\tau}$ corresponds to the price in McLennan (1984) where the demand functions intersect.

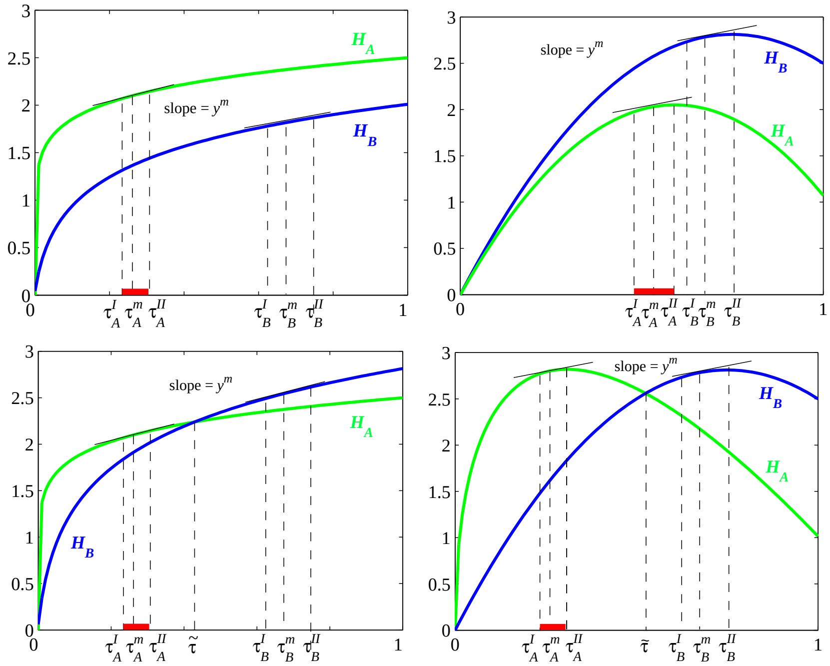

Figure 1: Four variants.

period t, then the decision maker observes the outcome

$$h_t(\tau_t, \varepsilon_t) = H_A(\tau_t) + \varepsilon_t \,, \tag{1}$$

where ε_t is an error term drawn randomly in every period. This error term ε_t captures factors influencing the policy outcome other than the policy itself. The error terms are normally and independently distributed with mean 0 and variance σ^2; we denote its probability density function by $\phi(\varepsilon_t)$.[11] 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 decision maker has the initial belief α_1 that the production function H_A is the true one, with $0 < \alpha_1 < 1$. In every period t, he uses the observed outcomes h_t to

[11] The normality assumption is only sufficient. As becomes clear from the proof of Proposition 2, all our results will hold for any distribution $f(\varepsilon_t)$ that has full support and whose likelihood ratio $l(\varepsilon_t) \equiv \frac{f(H_B(\tau_t)+\varepsilon_t)}{f(\varepsilon_t)}$ is monotone in ε_t and takes on values from zero to infinity.

update his beliefs α_{t+1}. 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 known to the decision maker. Since his belief in period t that H_A is the true is α_t, the expected level of the public good in period t for tax rate τ_t is

$$H_t(\tau_t) \equiv \alpha_t H_A(\tau_t) + (1 - \alpha_t) H_B(\tau_t) \,. \tag{2}$$

3 Dynamics and long-run equilibria

We now derive the long-run equilibrium in our model. We assume that in every period t the decision maker maximizes myopically the expected utility of household m, $u_m(\tau_t) = (1-\tau_t)y^m + H_t(\tau_t)$. This assumption is a good approximation if periods are long compared to the patience of households. The learning problem we explore captures the problem of a decision maker who faces uncertainty about which of two models of reality is the correct one and whose actions affect both his current period payoff and his future beliefs. In Section 4 we endow the decision maker with some forward-looking ability. We do this by allowing him to do experiments.

Proposition 1 characterizes the decision maker's optimal policy.

Proposition 1 *In every period t, the decision maker chooses $\tau_t^m \in [\tau_A^m, \tau_B^m]$, where τ_t^m is implicitly defined by $H_t'(\tau_t^m) \equiv y^m$.*

Figure 2 depicts the decision maker's choice. His initial belief α_1 is such that the expected production function in period 1 is H_1. The policy implemented in period 1 is τ_1^m.

Figure 3 illustrates the impact of the error term on the decision maker's belief and on the equilibrium tax rate in the next period. After implementing τ_1^m, the shock ε_1

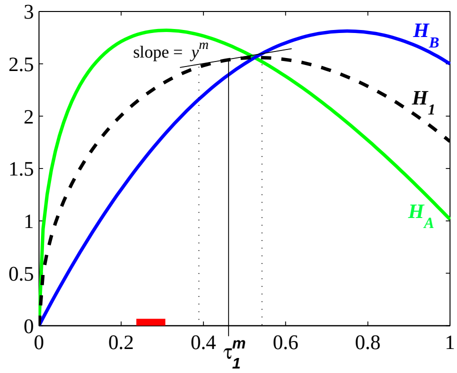

Figure 2: Equilibrium outcome in period 1.

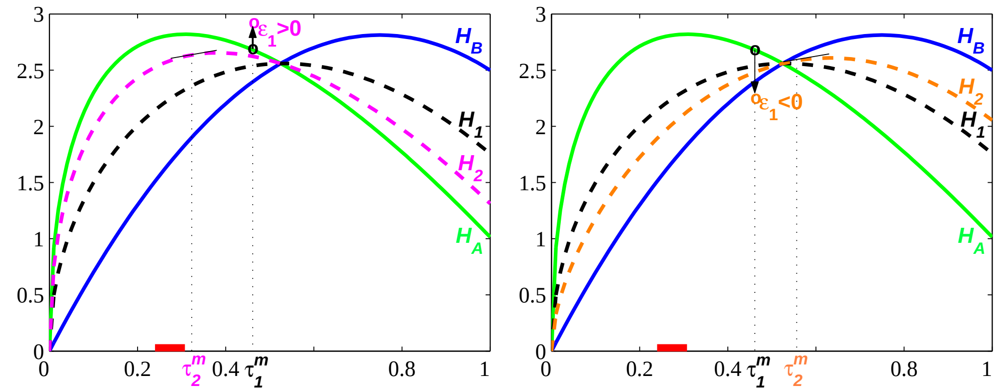

Figure 3: Inferences and outcome in period 2, as a function of ε_1.

materializes. If $\varepsilon_1 > 0$, the outcome is better than expected under H_1, and therefore, the updated belief is $\alpha_2 > \alpha_1$ 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 $\alpha_2 < \alpha_1$ yielding H_2 as shown in the right-hand panel. In both cases, the expected production function H_2 is the basis for equilibrium in period 2.

Note that only a strict subset of the feasible tax rates are implemented in equilibrium, i.e., $\tau_t^m \in [\tau_A^m, \tau_B^m] \subset [0, 1]$. This property is illustrated in Figure 4.

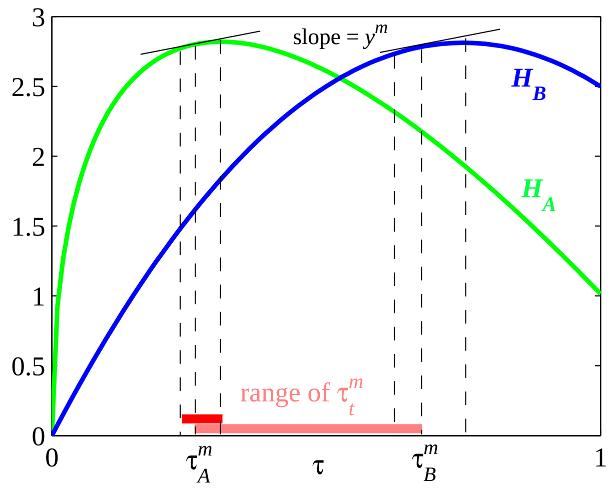

Figure 4: Range of equilibrium tax rates.

3.1 An informal discussion

The decision maker's problem is a problem of inference. Recall that $\mathcal{H}_t \equiv \{(h_i, \tau_i)\}_{i=1}^{t}$ is the history up to date t. Accordingly, let $\Pr(H_A \big| \mathcal{H}_t)$ denote the conditional probability that H_A is true given history $\mathcal{H}_t$. Denote by $\Pr(h_t \big| H_A, \tau_t)$ the probability of observing h_t given that H_A is true and given that policy τ_t is implemented. Then, by Bayes rule

$$\Pr(H_A|\mathcal{H}_t) \;=\; \frac{\Pr(H_A|\mathcal{H}_{t-1})\Pr(h_t|H_A,\tau_t)}{\Pr(H_A|\mathcal{H}_{t-1})\Pr(h_t|H_A,\tau_t)+(1-\Pr(H_A|\mathcal{H}_{t-1}))\Pr(h_t|H_B,\tau_t)}. \tag{3}$$

Since households are rational, they use Bayes rules (3) to update their beliefs; i.e., $\alpha_{t+1} = \Pr(H_A|\mathcal{H}_t)$. Since the probability of observing h_t is higher under the true production function H_A than under 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 (3) 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 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 and is unique.

3.2 The information trap

We now state our main result:

Proposition 2 *Let the two production functions cross at some $\tilde{\tau} \in (\tau_A^m, \tau_B^m)$. Then, the only policies that can be implemented in a long-run equilibrium are $\tilde{\tau}$ and τ_A^m. Formally, a random variable $\tau_\infty \in [0, 1]$ exists such that (i) τ_t^m converges to τ_∞ almost surely as t becomes arbitrarily large, and (ii) the support of τ_∞ is contained in $\{\tilde{\tau}, \tau_A^m\}$.*

The content of Proposition 2 is that the policy converges to a random variable τ_∞ whose distribution has positive probability on, at most, the two real numbers $\tilde{\tau}$ and τ_A^m.[12] Observe that the proposition does not establish that convergence to the information trap $\tilde{\tau}$ occurs with positive probability. For reasons outlined below, we cannot prove such a result analytically. Instead, we use numerical methods to approximate the probabilities that the economy converges to $\tilde{\tau}$ or τ_A^m. Our simulations, reported in Section 3.3, strongly suggest that convergence to $\tilde{\tau}$ occurs with positive probability for a wide range of initial conditions.

The reason why there may be a range 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 is very difficult to learn anything. Hence, the economy becomes stuck with its current beliefs so that the policy will not change anymore. Consequently, one can speak of an information trap around $\tilde{\tau}$.

[12] This is equivalent to saying that the process of beliefs converges to a random variable whose support is contained in $\{\tilde{\alpha}, 1\}$.

If $\tau_A^{II} < \tilde{\tau} < \tau_B^I$ the Pareto sets of H_A and H_B are disjoint and so $\tilde{\tau}$ is Pareto inefficient. In this case, the support of τ_∞ may contain a Pareto inefficient policy. The conditions for this to happen is that H_A and H_B are sufficiently different. In this case, learning the truth is particularly relevant as failure to do so implies that the policy implemented in the long-run equilibrium can be Pareto inefficient.

Proposition 2 has a corollary that follows almost immediately.

Corollary 1 *If the two production functions do not cross on $[\tau_A^m, \tau_B^m]$, then τ_t^m converges almost surely to τ_A^m as t goes to infinity.*

Corollary 1 establishes that if the two models never make the same predictions over the relevant interval $[\tau_A^m, \tau_B^m]$, the economy will ultimately learn the truth. Note also that, from the proof of Proposition 2, it is clear that all our results go through if H_B is convex for τ close to zero as long as H_B is concave for all $\tau \geq \tau_A^m$.[13]

We now derive a lower bound for the probability that the economy learns the true production function. Denote by p the prior probability that the long-run policy converges to $\tilde{\tau}$. Accordingly, $1 - p$ is the probability that in the long-run the policy τ_A^m is implemented. Though we have not been able to derive the probability p analytically, we are able to derive a lower bound on $1 - p$, which we denote by $\gamma(\tilde{\alpha}, \alpha_1)$.

Proposition 3 *For $\alpha_1 > \tilde{\alpha}$, $\gamma(\tilde{\alpha}, \alpha_1) = \frac{\alpha_1 - \tilde{\alpha}}{\alpha_1(1-\tilde{\alpha})}$, which satisfies $\frac{\partial \gamma}{\partial \alpha_1} > 0$ and $\frac{\partial \gamma}{\partial \tilde{\alpha}} < 0$. Otherwise, $\gamma(\tilde{\alpha}, \alpha_1) = 0$.*

Proposition 3 provides an analytical lower bound for the probability that the long-run equilibrium policy is efficient. A priori there is no reason to assume that this lower bound will be tight. However, the numerical simulations reported in the next subsection hit

[13]This guarantees, in particular, that the function $s \equiv H_A - H_B$ satisfies $s'(\tau) < 0$ for all $\tau \in [\tau_A^m, \tau_B^m]$.

this lower bound almost exactly for certain parameter constellations and never fall short of it. We think that these are good reasons to believe that this lower bound is, to some extent, tight. Consequently, it makes sense to analyze its determinants in some more detail.

The fact that $\frac{\partial\gamma}{\partial\alpha_1} > 0$ is very intuitive since one naturally expects a decision maker who is initially better informed to be more likely to adopt the correct belief in the long-run. The sign of the derivative $\frac{\partial\gamma}{\partial\tilde{\alpha}} < 0$ 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.

We now explain why we cannot establish analytically that the economy converges to the information trap with positive probability. In order to do so, we relate our model to McLennan's (1983) model. The two models are very similar in that both assume that there is a policy ($\tilde{\tau}$ in our model, a in McLennan's; for simplicity, we discuss both models using our notation) that, once taken, will inhibit any further inference. The main difference between the two resides in the nature of the random variable, which is binary (sale, no sale) in McLennan's and continuous in our model. The simpler structure allows McLennan to derive the result that under certain restrictions the seller's belief α never jumps over $\tilde{\alpha}$. That is, if he starts with $\alpha_1 < \tilde{\alpha}$, his long-run belief α_∞ will be either 0 or $\tilde{\alpha}$ and if he starts with $\alpha_1 > \tilde{\alpha}$ it will be either $\tilde{\alpha}$ or 1. No such result obtains in our model because for any given $\alpha_t \in (0, \tilde{\alpha})$ there is always a positive probability that a shock occurs such that $\alpha_{t+1} > \tilde{\alpha}$.[14] A more detailed discussion is deferred to Appendix B.

3.3 Numerical results

The simulation results are collected in the two tables below for two different constellations of production functions.[15] For all simulations in the paper, we assume that the decision maker implements the optimal tax rate of the median income household. Figure 5 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

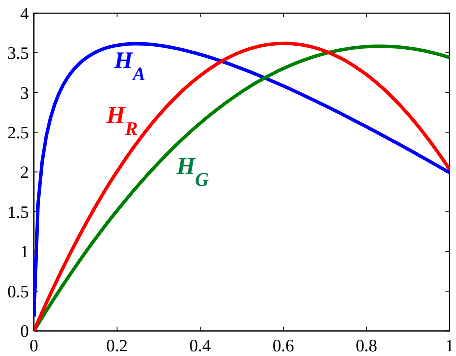

Figure 5: The functions used for the simulations reported in Tables 1 and 2.

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.[16] For

---

[14]To see this, observe that $\alpha_{t+1}(\varepsilon_t) = \frac{\alpha_t}{\alpha_t + (1-\alpha_t)l(\varepsilon_t)}$, where $l(\varepsilon_t)$ is the likelihood ratio that as a function of the shock ε_t can take any value between zero and infinity. Consequently, α_{t+1} is a random variable with support $(0, 1)$.

[15]All matlab files are available on request.

[16]Note that, as claimed in Proposition 2, all draws either converge to τ_A^m or to $\tilde{\tau}$.

H_A and H_G $\tilde{\alpha} = 0.47$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$	γ
$\alpha_1 = 0.1$	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
$\alpha_1 = 0.9$	1	1	0.91	0.92	0.90

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

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 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.

Three further remarks are in order. First, the efficiency potential γ has some bite indeed. For $\sigma = 2$, γ is quite close to the numerical results both in Table 1 and 2. Thus, γ is not merely a theoretical lower bound. Second, the difference between the numerical results and the efficiency potential for $\sigma = 2$ and $\alpha_1 = 0.5$ in Table 1 and for $\alpha_1 = 0.8$ in Table 2 is not statistically significant.[17] Third, consider the columns for $\sigma = 1$ in Table 1 and $\sigma = 0.5$ in Table 2 to see that the probability of convergence to the good policy does not increase monotonically in the initial belief α_1.[18] The intuition for this behavior

[17]The standard errors are 0.03 and 0.05. The complete simulation data and the tables augmented with standard errors are available on request.

[18]In Table 1, the probability decreases from 0.26 (0.04) to 0.21 (0.04), while in Table 2 it decreases from 0.36 (0.05) to 0.21 (0.04), where standard errors are in parentheses. Thus, the difference between 0.36 and 0.21 in Table 2 is statistically significant whereas the difference between 0.26 and 0.21 in Table

H_A and H_R $\tilde{\alpha} = 0.52$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$	γ
$\alpha_1 = 0.1$	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.

seems to be that starting from very bad initial beliefs, i.e. α_1 close to zero, increases in α_1 may well increase the likelihood of adopting a good policy in the long-run. However, as α_1 becomes larger, it gets closer to $\tilde{\alpha}$ and thereby increases the probability of adopting a bad policy in the long-run. Witness, in particular, that the minima in these columns are reached for the α_1 closest to, and to the left of, $\tilde{\alpha}$, which are, respectively, $\alpha_1 = 0.4$ and $\alpha_1 = 0.5$.

4 Experimentation

So far, we have assumed that the decision maker behaves myopically by maximizing the period expected utility of household m. This is in particular questionable if the economy is stuck in the information trap since the decision maker knows that reaching $\tilde{\tau}$ is *bad*: it prevents learning with probability one. It is thus of particular relevance to check whether the result that the economy may end up in an information trap breaks down when we

1 is not.

endow the decision maker with some forward-looking ability. We do this by allowing him to do experiments.

A plausible, and feasible, way of modelling foresighted and experimenting behavior is the following. In any period t, let the decision maker choose $\tau_t^m - \Delta$ and $\tau_t^m + \Delta$, where $\Delta > 0$ measures the degree of foresightedness and τ_t^m is still the myopic bliss point tax rate of household m. The implemented policy is determined by flipping a fair coin.[19] The larger Δ, the greater the degree of foresightedness and/or the larger the willingness to experiment.

The question that interests us is whether in the long-run beliefs converge towards $\alpha_\infty = \tilde{\alpha}$ with positive probability. To answer this question, we again use simulations. The results are collected in Table 3. The true production function H_A and the alternative H_G are the same as in Table 1 above. We set $\sigma = 1$ and let α_1 increase from 0.1 to 0.9, while Δ increases from 0.00001 to 0.1. (The first column with $\Delta = 0$ is a reprint from Table 1.) For each pair (α_1, Δ), we ran 100 draws. Table entries give the number of draws for which the process ended with beliefs $\alpha_\infty = 1$. In this case, the good policy τ_A^m is implemented in the long-run.[20]

We observe the following. First, if Δ is sufficiently large, in particular larger than 0.1, then the truth is always learned in the long-run. Note that $\Delta = 0.1$ implies a difference of 10 percentage points. Even with $\Delta = 0.01$, convergence to the Pareto

[19] Alternatively, one could have a percentage experimentation, according to which positions would be $(1 - \Delta)\tau_t^m$ and $(1 + \Delta)\tau_t^m$. The assumption that randomization is fifty-fifty is made for convenience. We expect the results to be robust to other distributions as long as these are not too skewed towards the true production function.

[20] Obviously, this is true only in an approximate sense because given our modelling assumptions experimentation continues even with $\alpha_\infty = 1$.

H_A and H_G $\tilde{\alpha} = 0.47$ $\sigma = 1$	$\Delta =$ 0	$\Delta =$ 0.00001	$\Delta =$ 0.0001	$\Delta =$ 0.001	$\Delta =$ 0.01	$\Delta =$ 0.1
$\alpha_1 = 0.1$	0.21	0.43	0.68	0.95	1	1
$\alpha_1 = 0.2$	0.26	0.43	0.66	0.95	0.99	1
$\alpha_1 = 0.3$	0.24	0.43	0.68	0.94	0.99	1
$\alpha_1 = 0.4$	0.21	0.36	0.67	0.94	0.99	1
$\alpha_1 = 0.5$	0.25	0.48	0.74	0.94	0.99	1
$\alpha_1 = 0.6$	0.57	0.69	0.85	0.97	1	1
$\alpha_1 = 0.7$	0.76	0.79	0.90	0.99	1	1
$\alpha_1 = 0.8$	0.94	0.93	0.97	1	1	1
$\alpha_1 = 0.9$	0.91	0.97	0.99	1	1	1

Table 3: Simulation results for the model with experimentation.

efficient policy is still almost universal. However, as Δ becomes smaller, the probability of convergence to $\alpha_\infty = 1$ decreases, too. For example, for $\Delta = 0.0001$ and $\alpha_1 \leq 0.4$, less than 70 out of the 100 draws converged to $\alpha_\infty = 1$. $\Delta = 0.0001$ is admittedly a small policy difference. Nonetheless, it represents a positive amount of experimentation and reflects at least some degree of forward-looking behavior. Thus, our information trap result does not break down if we allow for a small degree of non-myopic behavior or experimentation. Second, the non-monotonicity in α_1 observed above carries over to the model with experimentation as can be seen from the column with $\Delta = 0.00001$.

5 Conclusions

We consider an economy where there is uncertainty about how the economy functions. In every period, the decision maker implements a tax policy. Observations of policies and economic outcomes are used to update the decision maker's beliefs, which then serve as the basis for decision making in the following period. We show numerically that the economy can end up in an information trap where no further learning is possible. This result is robust with respect to the introduction of small experimentation.

Putnam (1993) has raised the question why some 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 initially identical societies may differ in the long-run and more specifically, why some countries may adopt Pareto inferior policies even in the long-run. Our explanation is that decision makers face uncertainty and that uncertainty can only be unravelled by experience. Initially identical countries may end up with different outcomes because in combination with bad luck the equilibrium may impede further inferences, so that the uncertainty is never abolished.[21] 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.

Appendix

A Proofs

Proof of Proposition 1 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

[21] Among other things, we have shown that initial beliefs may be crucial for the long-run political outcome. This may help to 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).

standard techniques for the transformation of random variables.[22] 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 $\frac{\partial H(\tau^i)}{\partial \tau} = y_i$ 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, then the density $\omega(\tau_t^i)$ of $\Omega(\tau_t^i)$ is given by $\omega(\tau_t^i) = f(\kappa(\tau_t^i)) \mid \frac{dy_i}{d\tau_t^i} \mid$, where $\mid \frac{dy_i}{d\tau^i} \mid$ 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.

In any period t the household m's optimal tax rate under the expected production function $H_t(\tau_t)$ defined in (2) 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 $\tau_t^m > \tau_B^m$. ■

Proof of Proposition 2 We prove Proposition 2 by showing that the decision maker's belief α_t converges 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 $s(\tau) \equiv H_A(\tau) - H_B(\tau)$ for $\tau \in [\tau_A^m, \tau_B^m]$. 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 $H_A'(\tau_A^m) = H_B'(\tau_B^m)$. Note that for $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$, $s(\tilde{\tau}) = 0$. Therefore, $s(\tau_A^m) > 0$ and $s(\tau_B^m) < 0$.

Let us also define the function $\tau^m(\alpha_t)$, which is the tax rate solving the equation in Proposition 1 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

[22] See, e.g., Hogg and Craig (1995).

theorem, we have $\frac{\partial \tau_t^m}{\partial \alpha_t} = \frac{-s'(\tau_t^m)}{\alpha_t H_A''(\tau_t^m)+(1-\alpha_t)H_B''(\tau_t^m)} < 0$, 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 $w(\alpha_t) \equiv s(\tau^m(\alpha_t))$, 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 $\frac{\partial w}{\partial \alpha_t} = s'\tau^{m\prime} > 0$ 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 $-\infty < w(0) < 0 < w(1) < \infty$.

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, and let

$$\Pr(h_t|H_A) = \phi(h_t - H_A(\tau_t)) = \phi(\varepsilon_t) \quad \text{and} \quad \Pr(h_t|H_B) = \phi(h_t - H_B(\tau_t)) = \phi(w(\alpha_t)+\varepsilon_t),$$

be the respective probabilities of observing h_t when H_A and when H_B is true, where $\phi(.)$ is the density of the normal with mean zero and variance σ^2.[23] After history $\mathcal{H}_t = \{(h_i,\tau_i)\}_{i=1}^t$, the period $t+1$ belief can be written as

$$\alpha_{t+1} = \frac{1}{1+\frac{\Pr(H_B)\Pr(h_1|H_B)\Pr(h_2|H_B)...\Pr(h_t|H_B)}{\Pr(H_A)\Pr(h_1|H_A)\Pr(h_2|H_A)...\Pr(h_t|H_A)}} = \frac{1}{1+\frac{(1-\alpha_1)\phi(w(\alpha_1)+\varepsilon_1)\phi(w(\alpha_2)+\varepsilon_2)...\phi(w(\alpha_t)+\varepsilon_t)}{\alpha_1\phi(\varepsilon_1)\phi(\varepsilon_2)...\phi(\varepsilon_t)}}. \tag{4}$$

Define

$$N_{t+1} \equiv \frac{(1-\alpha_1)\phi(w(\alpha_1)+\varepsilon_1)\phi(w(\alpha_2)+\varepsilon_2)\cdot\ldots\cdot\phi(w(\alpha_t)+\varepsilon_t)}{\alpha_1\phi(\varepsilon_1)\phi(\varepsilon_2)\cdot\ldots\cdot\phi(\varepsilon_t)}, \tag{5}$$

such that (4) becomes $\alpha_{t+1} = \frac{1}{1+N_{t+1}}$. This defines the function $\alpha_t = \alpha(N_t)$ 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

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

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 $\lim_{N_t \to \infty} r(N_t) = w(0) < 0$. Thus, for $\tilde{\tau} \in [\tau_A^m, \tau_B^m]$, there is a unique $\tilde{N}$ such that $r(\tilde{N}) = 0$. In light of these new definitions,

$$N_{t+1} = N_1 \cdot \frac{\phi(r(N_1)+\varepsilon_1)}{\phi(\varepsilon_1)} \cdot \frac{\phi(r(N_2)+\varepsilon_2)}{\phi(\varepsilon_2)} \cdot ... \cdot \frac{\phi(r(N_t)+\varepsilon_t)}{\phi(\varepsilon_t)} = N_t \cdot \frac{\phi(r(N_t)+\varepsilon_t)}{\phi(\varepsilon_t)}. \quad (6)$$

Notice that (6) is a non-linear stochastic first-order difference equation.

Observe first that if the sequence takes either the value 0 or the value $\tilde{N}$, it will take this value forever. This becomes immediate for $N_t = 0$ by inserting $N_t = 0$ into (6). For $N_t = \tilde{N}$, $r(\tilde{N}) = 0$ implies that $\frac{\phi(r(N_t)+\varepsilon_t)}{\phi(\varepsilon_t)} = \frac{\phi(\varepsilon_t)}{\phi(\varepsilon_t)} = 1$, implying in turn $N_{t+1} = \tilde{N}$. If N_t is infinity, N_{t+1} will be too, since $\lim_{N_t \to \infty} r(N_t)$ is a finite negative number.

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

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

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

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

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 $i = 0, 1, \ldots$. Note that δ can be chosen arbitrarily small. Now define the variable $\overline{\varepsilon}_{t_\delta+i}$ by

$$\overline{\varepsilon}_{t_\delta+i} \equiv \frac{\sigma^2}{r(N_{t_\delta+i})} \cdot \ln \frac{N_{t_\delta+i}}{C + \delta} - \frac{1}{2} r(N_{t_\delta+i}) \,. \tag{7}$$

Note that $\overline{\varepsilon}_{t_\delta+i}$ is a shock such that $N_{t_\delta+i+1} = C + \delta$. Assume that $C < \tilde{N}$. Then, the variable $\overline{\varepsilon}_{t_\delta+i}$ is negative and finite for all $N_{t_\delta+i} \in [C - \delta, C + \delta]$, because all terms in (7) are finite. Therefore, for every $N_{t_\delta+i} \in [C - \delta, C + \delta]$, $\Pr(\varepsilon_{t_\delta+i} < \overline{\varepsilon}_{t_\delta+i}) = \Phi(\overline{\varepsilon}_{t_\delta+i}) > 0$, which means that the probability to draw an $\varepsilon_{t_\delta+i} < \overline{\varepsilon}_{t_\delta+i}$ 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. □

From Slutzki'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 2. ■

Proof of Corollary 1 It is clear that if $H_A(\tau) \neq H_B(\tau)$ for all $\tau \in [\tau_A^m, \tau_B^m]$, then the function $s(\tau) \equiv H_A(\tau) - H_B(\tau)$ is never equal to zero in the relevant interval. Consequently, the functions $\omega(\alpha_t)$ and $r(N_t)$ defined in the proof of Proposition 2 will also be non-zero in the relevant range. Therefore, equation (6) has a unique fixed point, which is $N_{t+1} = N_t = 0$, corresponding to $\alpha_{t+1} = \alpha_t = 1$. ■

Proof of Proposition 3 From Proposition 2 we know that α_t either converges to 1 or to $\tilde{\alpha}$. What we need to characterize in order to prove Proposition 3 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 p be the probability of convergence towards $\tilde{N}$. Then $E[N_\infty] = (1-p) \cdot 0 + p \cdot \tilde{N} = p \cdot \tilde{N}$, which implies $p \leq \frac{N_1}{\tilde{N}}$ and hence $(1-p) \geq 1 - \frac{N_1}{\tilde{N}}$, where it will be recalled that $N_1 = \frac{1-\alpha_1}{\alpha_1}$ and hence $E[N_1] = N_1$. Now

$$1 - \frac{N_1}{\tilde{N}} = 1 - \frac{\frac{1-\alpha_1}{\alpha_1}}{\frac{1-\tilde{\alpha}}{\tilde{\alpha}}} = \frac{\alpha_1 - \tilde{\alpha}}{\alpha_1(1-\tilde{\alpha})} \equiv \gamma(\tilde{\alpha}, \alpha_1).$$

Thus, $1 - p \geq \gamma(\tilde{\alpha}, \alpha_1)$ holds. Since $1 - p \geq 0$, zero is the lower bound when $\alpha_1 < \tilde{\alpha}$. ■

B Relation McLennan's (1984) model

One important consequence of the arguments outlined in Subsection 3.2 in the text is that we cannot use McLennan's (1984, p. 343-4) arguments to establish analytically that $\tilde{\alpha}$ is reached with positive probability. To see this, observe first that our Proposition 3 is actually a statement conditional on H_A being true. Without this condition, it would read: The support of τ_∞ is $\{\tau_B^m, \tilde{\tau}, \tau_A^m\}$, or in terms of beliefs, the support of α_∞ is $\{0, \tilde{\alpha}, 1\}$. Second, denote by $p_0(\alpha)$, $\tilde{p}(\alpha)$ and $p_1(\alpha)$ the unconditional probability of converging to the absorbing state $0, \tilde{\alpha}$ and 1, respectively. Since $p_0(\alpha)$, $\tilde{p}(\alpha)$ and $p_1(\alpha)$ are probabilities and because all paths converge,

$$p_0(\alpha) = 1 - \tilde{p}(\alpha) - p_1(\alpha). \tag{8}$$

Then because of the elementary property of Bayesian updating that the expected posterior is equal to the prior,

$$p_0(\alpha)0 + \tilde{p}(\alpha)\tilde{\alpha} + p_1(\alpha) = \alpha. \tag{9}$$

In contrast to McLennan, who has the additional restriction that for, say, $\alpha < \tilde{\alpha}$ the only absorbing states are $\{0, \tilde{\alpha}\}$, the system of the two equations (8) and (9) with three unknowns is indeterminate. Without additional restrictions, neither $\tilde{p} = 0$ nor $\tilde{p} > 0$ can be ruled out. Therefore, it is not possible to prove that $\tilde{p}(\alpha) > 0$ along the lines in McLennan (p. 343-4), as suggested by one careful reader.

References

Acemoglu, D., and J. A. Robinson (2000): "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective," *Quarterly Jour-*

nal of Economics, 115, 1167–1199.

Alesina, A., and G.-M. Angeletos (2005): "Fairness and Redistribution," *American Economic Review*, 95(4), 960–80.

Banerjee, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797–817.

Baron, D. P. (1996): "A Dynamic Theory of Collective Goods Programs," *The American Political Science Review*, 90(2), 316–330.

Bartels, L. M. (1996): "Uninformed Votes: Information Effects in Presidential Eleections," *American Journal of Political Science*, 40(1), 194–230.

Bauer, P. T. (1981): *Equality, the Third World, and Economic Delusion.* Harvard University Press.

Berentsen, A., E. Bruegger, and S. Loertscher (2005): "Learning, voting and the information trap," Discussion Paper 05.16, University of Bern.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

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.

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.

Friedman, M. (1962): *Capitalism and Freedom.* University of Chicago Press.

——— (1997): "If Only the U.S. Were as Free as Hong Kong," *Wall Street Journal*, July 8, A14.

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.

Gans, J. S., and M. Smart (1996): "Majority Voting with Single-Crossing Preferences," *Journal of Public Economics*, 59(2), 219–237.

Hayek, F. A. (1944): *The Road To Serfdom.* Routledge.

Hazlitt, H. (1946): *Economics in One Lesson.* Harper & Brothers.

Hogg, R., and A. T. Craig (1995): *Introduction to Mathematical Statistics.* Englewood Cliffs, N.J. : Prentice Hall.

Kant, I. (1784): "An Answer to the Question: What is Enlightenment?," *http://www.english.upenn.edu/ mgamer/Etexts/kant.html.*

Laslier, J.-F., A. Trannoy, and K. Van Der Straeten (2003): "Voting under ignorance of job skills of unemployed: the overtaxation bias," *Journal of Public Economics*, 87(3-4), 595–626.

McLennan, A. (1984): "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7(3), 331–47.

Myrdal, G. (1975): "The Equality Issue in World Development: Nobel Memorial Lecture," *Scnadinavian Journal of Economics*, pp. 413–432.

Persson, T., and G. Tabellini (2000): *Political Economics*. MIT Press, Masschusetts.

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.

Rosenstein-Rodan, P. (1943): "Problems of Industrialization of Eastern and South-Eastern Europe," *Economic Journal*, 53, 202–211.

Sachs, J. (2005): *The End of Poverty: Economic Possibilities for Our Time*. The Penguin Press.

Spector, D. (2000): "Rational Debate and One-Dimensional Conflict," *Quarterly Journal of Economics*, 115(1), 181–200.