

*ISSN 1471-0498*



**DEPARTMENT OF ECONOMICS  
DISCUSSION PAPER SERIES**

**POLITICAL ACCOUNTABILITY AND POLICY  
EXPERIMENTATION: WHY TO ELECT LEFT-HANDED  
POLITICIANS?**

**Tim Willems**

Number 647  
March 2013

Manor Road Building, Manor Road, Oxford OX1 3UQ

# Political Accountability and Policy Experimentation: Why to Elect Left-Handed Politicians\*

Tim Willems<sup>†</sup>

June 12, 2013

## Abstract

This paper analyzes how political accountability affects the incentives of policy makers to learn through experimentation with new policies. It shows that when voters face an inference problem on the competence level of policy makers, reelection concerns reduce experimentation incentives to the benefit of the status quo. Whether experimentation in representative democracies is *suboptimally* low, depends on society's degree of risk aversion relative to that of the decision maker. If the level of experimentation is suboptimal, taking decisions by direct democracy, or electing risk-loving politicians, could improve welfare. Interestingly, risk-lovers also seem to be overrepresented among Presidents of various countries.

*JEL-classification:* D72, D83

*Key words:* policy experimentation, learning, political economy, reform, status quo bias, career concerns

---

\*I thank Tim Besley, John Feddersen, Bård Harstad, James Hollway, Ian Jewitt, Ruixue Jia, Sharun Mukand, Rick van der Ploeg, Germán Vera Concha, Sweder van Wijnbergen, Sam Wills, and participants of the 2013 RES Conference for useful comments and discussions.

<sup>†</sup>Nuffield College and Department of Economics, University of Oxford, Manor Road Building, Manor Road, Oxford OX1 3UQ, UK. E-mail: tim.willems@economics.ox.ac.uk.

# 1 Introduction

*The country needs and, unless I mistake its temper, the country demands bold, persistent experimentation. It is common sense to take a method and try it: if it fails, admit it frankly and try another. But above all, try something.*

*Franklin D. Roosevelt, 1932*<sup>1</sup>

As pointed out by Franklin D. Roosevelt in the above quotation, the experimental approach to learning can be a powerful way to improve upon existing policies. But under what conditions will the amount of policy experimentation be socially optimal? And how does a country's institutional structure affect its policy makers' incentives to experiment with new, untried (but potentially beneficial) policies? In an attempt to answer these questions, this paper builds a political economy model in which there is scope for "active" learning on what policy is appropriate.

While doing so, this paper models another relevant learning process, namely the inference problem that citizens face on the competence level of the official in charge. In such an environment of two-sided learning, this paper analyzes what happens to experimentation incentives when public officials are being made accountable (*i.e.* when they are subjected to reelection, as in representative democracies). It points out that, in the presence of an incumbency advantage, a desire for reelection makes public officials less willing to experiment with new policies. The reason is that policy makers will then fear failure of an experiment, as that will result in electoral defeat. Here, incompetent policy makers inflict a negative externality upon competent ones (even if they are not in office) - decreasing their incentives to experiment with new policies already under risk neutrality. This could explain why some well-known policy experiments were implemented by non-accountable regimes, whose decision makers did not have to worry about reelection (*cf.* the experiments with capitalism in China, Augusto Pinochet's neoliberal reforms in Chile, as well as the market experiments implemented by Suharto in Indonesia (but do see footnote 2)). It might

---

<sup>1</sup>As quoted in Callander (2011, p. 643).

simultaneously explain why economies in need for reform often *appoint* technocratic government officials (for example of the kind we have recently seen in Greece and Italy), who are not always interested in reelection *ex ante* (although they do sometimes change their mind *ex post* - see the case of Mario Monti in Italy).

Whether the amount of policy experimentation in representative democracies is *suboptimally* low, depends on the ranking of the degree of risk aversion of the decision maker relative to that of society. If the decision maker is less risk averse than society, he would experiment too much in the absence of accountability and subjecting him to reelection is actually welfare improving. Some might say that Augusto Pinochet was less risk averse than Chilean society, which resulted in excessive policy experimentation under his regime, while the same might have applied to China (where the Tiananmen Square protests of 1989 showed that Deng Xiaoping may have overshoot the optimal level of experimentation at the time).<sup>2</sup>

However, if the decision maker and society share the same degree of risk aversion, representative democracies will experiment too little - and even more so if society is less risk averse than the decision maker. Since a group is generally believed to be less risk averse than its individual members (*cf.* Samuelson (1964), Arrow and Lind (1970) and Chambers and Echenique (2010)), many would see the latter case as most relevant. In those circumstances, welfare can be improved by taking decisions via direct democracy (such that career-concerned politicians no longer determine whether to experiment or not), or by replacing the incumbent with an atypical policy maker who is *less* risk averse than society.

In representative democracies, society can thus be made better off by appointing decision makers who are "non-representative" in a sense that they do not share the social objective function, but have a greater appetite for risk instead. The reason is that their relative love for risk counteracts the inclination of accountable politicians

---

<sup>2</sup>Needless to say, there are many additional problems with non-accountable regimes in practice. Consequently, they typically produce inferior economic outcomes and often fail to implement welfare-enhancing reforms, for example due to noncongruent interests of policy makers (see Rodrik (2000, Section III) for a discussion, Giuliano, Mishra, and Spilimbergo (forthcoming) for empirical evidence, and Maskin and Tirole (2004) for a model). In this paper, I only consider the utopian welfare-maximizing social planner and use that perfect benchmark to compare the efficiency of representative democracies (which are likely to have their own imperfections) against.

to experiment too little due to reelection concerns. Hereby, the message of the present paper relates to that of Rogoff (1985), who showed that it can be optimal to appoint a non-representative "conservative" central banker who places too large a weight on inflation-rate stabilization (relative to employment stabilization) from a social welfare point of view.

Finally, this paper provides an explanation for why there seems to exist a bias against reforms to the benefit of the status quo: by maintaining the status quo, accountable politicians are able to increase the expected duration of their time in office.<sup>3</sup> Different from Fernandez and Rodrik (1991, who obtain a status quo bias in a median voter framework), this paper's bias results from the principal-agent problems underlying representative democracies.

On a more general level, this paper intends to contribute to the literature on policy experimentation. As for example noted by Besley (2000, p. 5) and Callander (2011, p. 643, 658), the large interest in experimentation by policy makers has not yet been matched by a comparable academic interest. Consequently, models of policy experimentation are rare. The present paper hopes to contribute in this dimension by developing a simple and tractable model of policy experimentation that could also be used to analyze other questions than the one being considered here.

## 2 Related literature

This paper relates to several existing literatures. Firstly, it is related to earlier studies on policy experimentation. An important question in this area has always been how countries can learn from each other's experiences. Experimentation then entails a positive externality: countries that did not incur the costs of experimentation, can learn from the experiment as well (assuming that the outcome is publicly observable). As Bolton and Harris (1999) show, this produces free-riding behavior - leading to a level of experimentation that is suboptimally low. On the other hand, Bolton and

---

<sup>3</sup>One of my colleagues who has been both an MP as well as a minister in The Netherlands (Rick van der Ploeg) always claims that "unfortunately, the only way to survive in politics is by changing as little as possible". This paper's model suggests that there indeed lies a truth in this notion.

Harris (1999) also identify an encouragement effect: if an individual country can bring forward the time at which the information generated by the experimentation of other countries becomes available, experimentation may be excessive.

In a related contribution, Callander and Harstad (2012) show that policy makers can get an incentive to mitigate free-riding behavior by selecting experiments that are less informative to other states, even if these experiments are less attractive to themselves. They then analyze how centralization (for example via the installation of a federal system) can correct this inefficiency.

Mukand and Rodrik (2005) add another element: they point out that reforms that succeed in one country, may fail completely in another. Mukand and Rodrik show that this setup can imply that countries in the neighborhood (by the relevant metric) of flourishing leaders do very well (since they can successfully imitate policies that have proved to work), distant countries will experiment themselves (which leads to moderate outcomes with high variance), while countries in the "near-periphery" bear the cost and are worst off: they are located too far away from the leader to be able to imitate their policy successfully, but they are close enough so that an honest policy maker will still choose to mimic the leader in fear of being held as corrupt if he does not do so (in the Mukand-Rodrik model, imitating an existing, successful policy is a way to signal that the policy is not corrupt). This implies that countries in the "near-periphery" are disciplined inefficiently, as a result of which they will end up with policies that are inappropriate for their local conditions.

Majumdar and Mukand (2004) consider a setup in which voters face an inference problem on the ability of policy makers to identify which state the world is in. They show that this can explain why policy makers sometimes persist in certain policies, even when outcomes suggest that they are not working. In their model, changing direction signals a lack of knowledge on what policy to implement, as a result of which policy makers may refrain from doing so. The present paper on the other hand takes a setup in which politicians differ in their ability to *implement* reforms, which has different implications as we will see below. However, as I will argue in footnote 11, intuition suggests that addition of the Majumdar-Mukand channel would only strengthen the findings that I will arrive at.

Callander (2011) in turn analyzes whether "trial and error"-learning actually generates optimal outcomes and investigates how the learning process is affected by uncertainty about voter preferences. Where the addition of noise typically obstructs learning, Callander finds that it facilitates learning in the policy context. The reason is that this uncertainty leads to a greater variety of policies being tried - thereby generating more information, which improves the quality of future policy making.

Finally, the present paper also links to the theoretical literature analyzing the impact of institutions on government performance, such as Maskin and Tirole (2004). They build a simple model with which they capture the virtues and drawbacks of political accountability. They focus upon a setup in which the public official may have different preferences than the electorate and analyze whether accountability can discipline these "noncongruent" policy makers. Maskin and Tirole (2004) however do not touch upon the relationship between a country's institutional structure and policy experimentation incentives, which is the focus of the present paper.

### **3 Model**

This section describes the model setup. In the model, learning is two-sided: voters want to learn whether officials are competent in implementing policy changes, while officials want to learn what policy is appropriate for their country. By choosing their policy, officials also determine how much information is revealed about their competence level. The resulting interaction between the election process and policy outcomes, lies at the heart of this paper. Throughout the paper, I define "experimentation" as the act of taking an action that entails static losses, but could generate dynamic gains via the information it produces.

A key challenge in models of optimal experimentation is always to maintain analytical tractability, as this is easily lost when there is scope for active learning. In that case, one has to resort to numerical methods in order to get answers (see Beck and Wieland (2002) for a general example and Willems (2012) for a model of price setting). To avoid this, the present paper models the learning problem in a very simple way - thereby maintaining tractability (despite the complicating presence of

two-sided learning), while trying to minimize loss of generality.

Each period  $t$ , the model is characterized by an unobserved state of the world,  $\Omega_t$ , which essentially governs "how the world works". This state may vary over time and is either given by  $A$  or  $B$  (so  $\Omega_t \in \{A, B\}$ ). The model consists of two periods (1 and 2) and at the beginning of period 1, decision makers and the electorate share the prior belief that  $\mathcal{P}^- = \Pr[\Omega_1 = A] = p$ . Throughout I will assume that  $p > \frac{1}{2}$ , so  $A$  is the most likely state for period 1.

Decision making officials ( $\mathcal{O}$ ) can be of two (time-invariant) types when it comes to implementing reforms: competent ( $\mathcal{C}$ ) or incompetent ( $\mathcal{I}$ ). Each period, officials have the option to implement one of three possible policies: first of all, they can choose to maintain the status quo ( $s$ ). I assume that continuation of the latter is relatively easy and does not require any particular skills from the policy maker (as Samuel Johnson once said, "To do nothing is in every man's power"). Consequently, this policy action pays off irrespective of the official's competence level. Moreover, and as in Majumdar and Mukand (2004), I assume that this is a safe policy option whose efficacy for the economy is well understood, working equally well in both states of the world. It yields a certain social pay-off (indicated by  $W$ ) of  $\sigma \in (0, 1)$ , so for period  $t$ :

$$W(s|\Omega_t = A, \mathcal{O} = \mathcal{C}) = W(s|\Omega_t = A, \mathcal{O} = \mathcal{I}) = W(s|\Omega_t = B, \mathcal{O} = \mathcal{C}) = W(s|\Omega_t = B, \mathcal{O} = \mathcal{I}) = \sigma$$

Alternatively, the country can experiment by implementing a new, untried policy - a more challenging option that does require competence from the incumbent official in order to be able to pay off. Suppose that there are two reform options,  $a$  and  $b$ . These represent the risky actions and pay off 1 only if they are implemented in the correct state of the world, by a competent policy maker. If they are implemented in the wrong state, or if they are implemented by an official who is incompetent in administering the reform process, their pay-off is zero:

$$\begin{aligned} W(a|\Omega_t = A, \mathcal{O} = \mathcal{C}) &= W(b|\Omega_t = B, \mathcal{O} = \mathcal{C}) &= 1 \\ W(a|\Omega_t = B) &= W(a|\mathcal{O} = \mathcal{I}) = W(b|\Omega_t = A) = W(b|\mathcal{O} = \mathcal{I}) &= 0 \end{aligned}$$

Given the assumption that  $p = \Pr[\Omega_1 = A] > \frac{1}{2}$ , policy makers who choose to experiment will always do so by implementing action  $a$  in period 1. Action  $b$  can however be chosen in period 2 if the information generated by the first-period experiment (more on which below) suggests that  $\Omega_2$  will be equal to  $B$ .

The fact that the pay-off structure is binary ("Arrow-Debreu security-like") is not essential: this just gives analytical convenience by equalizing pay-offs with probabilities (see Maskin and Tirole (2004)) and is only meant to capture the idea that the "right" risky option improves upon the status quo ( $W(j|\Omega_t = J, \mathcal{O} = \mathcal{C}) > \sigma$ ), while the status quo dominates the "wrong" risky option ( $W(j|\Omega_t \neq J) < \sigma$ ). This implies that experimentation is riskier than maintaining the status quo.

The decision maker is office-loving, for example due to the perks and prestige associated with his position: he derives ego rents  $R$  from each period he is in office (but the special case of  $R = 0$  yields the same main results so its presence is not essential). In addition, a decision maker obtains personal utility from his country doing well - for example because this makes him more popular in the streets. The period utility function of a decision maker in office is therefore given by:

$$u_t = (1 - \theta) R + \theta W_t, \tag{1}$$

where  $\theta \in (0, 1)$  is the relative weight that the decision maker puts on social welfare generated in period  $t$ ,  $W_t$ .

Once the first-period social pay-off  $W_1$  has materialized, it becomes publicly observable to all agents in the model. Because the pay-off of the experimental action varies with the state  $\Omega$ , the outcome of a first-period experiment produces information on "how the world works". The decision maker can exploit this information when making his period 2 decision. More specifically, I assume that first-period experimentation generates information on the second-period state in the form of a signal  $\omega \in \{A, B\}$  that brings a degree of confidence  $q$  (implying that this information is expected to be correct with probability  $q$ , so  $\mathbb{E}\{\Pr[\omega = J|\Omega_2 = J]\} = q$ ). One can think of this signal as combining information from Bayesian inference on the outcome of the first-period experiment, with the possibility that the state changes

over time (*i.e.* that  $\Omega_2 \neq \Omega_1$ ). Due to the reduced form nature of  $q$ , one does not have to specify the probability of this switch occurring explicitly. Without loss of generality, I assume that the reliability of the signal does not depend upon the outcome on the first-period experiment, such that  $q_{W_1=1} = q_{W_1=0} = q$  - which eases the analysis by making it less taxonomical.<sup>4</sup> Maintaining the safe status quo  $s$  produces no information on the state (as its pay-off does not vary with  $\Omega$ ), so  $q$  also captures the expected "informational return" to experimentation.

All agents in the model share the prior belief  $\mathcal{X}^-$  that the incumbent is incompetent when it comes to implementing reforms. This prior belief is described by  $\mathcal{X}^- = \Pr[\mathcal{O} = \mathcal{I}] = x$ . The type of the official is unknown to both the electorate as well as to the official himself (the same symmetrical structure as in Holmström (1999)): it can only be inferred from the outcome of his actions. The assumption that the policy maker does not know his own type eliminates any signaling behavior on his behalf, which would obscure the analysis by adding an extra strategic dimension (although this could be an interesting extension). In addition, it seems realistic to assume that a freshly appointed policy maker is initially uncertain on his own level of competence in implementing a new policy, especially if his success also depends upon his relationship with new, unknown staff members assisting him (*cf.* Persson and Tabellini (2000) and Besley (2006), who proceed along the same lines).

Note that the binary pay-off structure implies that the expected first-period pay-off from following the risky strategy equals  $(1 - \mathcal{X}^-)p$ . After all, the first-period action  $a$  will only pay off 1 if the policy maker is competent (to which agents attach prior probability  $(1 - \mathcal{X}^-)$ ) and if  $\Omega = A$  (probability  $p$ ). I make the following assumption on the relationship between  $(1 - \mathcal{X}^-)p$  and  $\sigma$ :

**Assumption 1:**  $(1 - \mathcal{X}^-)p < \sigma$ .

This restricts the analysis to the interesting case in which a decision maker would never experiment in a static setup. It captures the defining notion that experimen-

---

<sup>4</sup>Note that observing  $W_1 = 1$  after implementing action  $a$  in period 1 perfectly reveals that  $\Omega_1 = A$ . But since there is a probability that the aggregate state switches and  $\Omega_2 \neq \Omega_1$ , this information is still noisy with respect to  $\Omega_2$  - thereby justifying a  $q_{W_1} < 1$ .

tation creates static losses, but (as we will see below) there could be dynamic gains.

The prior belief that the official is incompetent (given by  $\mathcal{X}^- = \Pr[\mathcal{O} = \mathcal{I}] = x$ ) is updated after the period 1 pay-off has realized. I assume that the updating rule is such that the belief that the official is incompetent goes up after a first-period policy failure, which is intuitive (and consistent with Bayes' rule, which is the rule that I will employ for concreteness). If the official chooses to maintain the status quo  $s$ , no information on his type is produced (as the status quo pay-off does not vary with  $\mathcal{O}$ ) and the posterior just equals the prior, hence:  $\mathcal{X}_{W_1=\sigma}^+ = \Pr[\mathcal{O} = \mathcal{I}|W_1 = \sigma] = \mathcal{X}^- = x$ , where the subscript " $W_1 = \sigma$ " indicates that this is the posterior after having maintained the status quo in the first period.<sup>5</sup>

If the decision maker implements a risky policy, beliefs on his type can be updated. In particular, if it turns out that the first-period policy has paid off (*i.e.*  $W_1 = 1$ ), Bayesian updating implies that  $\mathcal{X}_{W_1=1}^+ = \Pr[\mathcal{O} = \mathcal{I}|W_1 = 1] = 0$  (since a policy would never pay off if the official were incompetent). If the first-period policy has failed to pay-off on the other hand, this can either be due to incompetence of the official, or due to the fact that a competent official unluckily implemented the wrong policy (*e.g.*  $a$  while  $\Omega = B$ ).<sup>6</sup> Applying Bayes' rule then yields:

$$\mathcal{X}_{W_1=0}^+ = \Pr[\mathcal{O} = \mathcal{I}|W_1 = 0] = \frac{x}{px + (1-p)} (> \mathcal{X}^- = x), \quad (2)$$

Note that  $\mathcal{X}_{W_1=0}^+ > \mathcal{X}_{W_1=\sigma}^+ > \mathcal{X}_{W_1=1}^+$ , so beliefs on the official's level of incompetence are indeed decreasing in the success of the first-period policy.

Conditional upon experimentation in the first period, the model structure implies

---

<sup>5</sup>The assumption that maintaining the status quo produces absolutely no information on the official's type might be too extreme. However, one can easily verify that this paper's results generalize to the more realistic setup in which the status quo option produces *less* information on the official's type than the experimental action, which seems plausible. Consider for example a setting in which the status quo pays off with probability  $z_S$  if maintained by an incompetent official, while the experimental action implemented by an incompetent policy maker pays off with probability  $z_E$ , where  $z_S > z_E$  (capturing the notion that maintaining the status quo requires less competence than implementing a new policy). The limiting case that I consider in the paper (where  $z_S = 1$  and  $z_E = 0$ ) just delivers analytical convenience without loss of generality.

<sup>6</sup>So "competence" does *not* refer to the official's ability to infer which state of the world we are in (see Majumdar and Mukand (2004) for a model where it does). In the present paper, "competence" rather refers to the policy maker's ability to administer the reform process.

that the expected social pay-off from implementing one of the risky options in period 2 equals  $(1 - \mathcal{X}_i^+) q$ , with  $i \in \{W_1 = 0, W_1 = 1\}$ . I will assume that:

**Assumption 2:**  $(1 - \mathcal{X}_i^+) q > \sigma$  for all  $i$ .

This assumption rules out the uninteresting scenario in which the risky options are always dominated in expectation by the safe policy  $s$ . Together, Assumptions 1 and 2 capture the notion that experimentation entails static losses, but that there could be dynamic gains.

If an election takes place at the beginning of period 2, the incumbent is challenged by another candidate who is believed to be incompetent with probability  $x$  (the parameter  $x$  thus also represents the inverse of the perceived quality of the pool from which a challenger can be drawn). If the incumbent is voted out after having experimented in period 1, the generated signal on  $\Omega_2$  ( $\omega$ ) is publicly observable such that his successor knows the correct state of nature with probability  $q$  as well.<sup>7</sup> As a result, a decision maker who is subject to reelection will always be voted out at the beginning of period 2 if his first-period policy failed. After all, both the incumbent and the challenger then have the same information on  $\Omega_2$ , but the electorate's posterior belief that the incumbent is incompetent equals  $\mathcal{X}_{W_1=0}^+$ , which is greater than  $x$  (the probability with which the challenger is believed to be incompetent).

The fact that a politician loses office *with certainty* after a policy failure is extreme. This is however just meant to capture the more general idea that a policy maker's *probability* of losing the next election goes up after a failure, which seems realistic. The fact that he will be voted out with probability 1 only delivers analytical convenience without loss of generality.

Finally, I will first follow the literature in assuming that:

**Assumption 3:** Whenever candidates are perceived as being equally competent by voters, the incumbent is reelected.

---

<sup>7</sup>I thus assume that the incumbent is not able to take the signal  $\omega$  home with him after being voted out. This seems realistic as such information is often present with civil servants who stay in office, laid down in documents that are handed over to the next government, etc. etc. I will however return to this issue in Section 6.

As said, this tie-breaking rule is relatively standard in the literature and is a way to capture the incumbency advantage<sup>8</sup> (but it will be relaxed in Section 6.1 below). The existence of the incumbency advantage is well-documented in reality (*cf.* Gelman and King (1990)). Explaining its origin is not the goal of the present paper, and the results that are to follow are general in a sense that they do not depend upon its exact nature. To prevent possible loss of generality, I therefore leave its roots unspecified - especially since there is no consensus on this yet. It could be motivated by the existence of government turnover costs (*e.g.* due to the destruction of the incumbent's experience), a preference of voters for a known candidate over an unknown one, or because the incumbent is preferred along some other dimension (which is why he ended up in power in the first place; see footnote 10 for an elaboration on this point and see Ashworth and Bueno de Mesquita (2008) for a model with such a selection effect).

Summarizing, this paper consists of the following sequence of events:

1. Nature determines states  $\{\Omega_1, \Omega_2\}$  as well as whether the policy maker is competent or not.
2. The policy maker decides which first-period policy to implement.
3. First period pay-off materializes; the prior belief on the policy maker's level of competence  $\mathcal{X}^-$  is updated to  $\mathcal{X}^+$ ; if the policy maker chose to experiment in period 1, a public signal  $\omega$  on  $\Omega_2$  is released.
4. Election: the incumbent is reelected or replaced.
5. The policy maker decides which second-period policy to implement.
6. Second period pay-off materializes.

---

<sup>8</sup>See *e.g.* Jehiel (1999, p. 659), Prato and Strulovici (2011, p. 12), and Samuelson and Zeckhauser (1988, p. 9). Meyer (1991) has shown that such a form of favoring the former winner can also be the optimal strategy.

## 4 Decision maker's behavior

At this stage, we are able to analyze what the introduction of political accountability does to policy experimentation incentives. To show that this paper's argument does not rely upon risk aversion, I start by assuming that both society as well as the decision making official are risk neutral (I will analyze the impact of risk aversion in Section 5).

### 4.1 Efficient benchmark: non-accountable official

First consider the case of a decision making official who is not accountable ("NA"). He cares about social welfare as in equation (1). The lack of accountability implies that the decision maker does not have to worry about reelection as he will remain in office for the full two periods - no matter what. His lifetime utility  $U$  is then given by (I abstract from discounting for notational simplicity):

$$U^{NA} = 2(1 - \theta)R + \theta [W_1(\alpha_1|\Omega_2, \mathcal{O}) + W_2(\alpha_1, \alpha_2|\Omega_2, \mathcal{O})],$$

where  $W_1(\alpha_1)$  expresses that first-period social welfare is a function of the first-period action  $\alpha_1$ , while  $W_2(\alpha_1, \alpha_2)$  captures that the second period's level of social welfare depends upon both period's actions (due to the fact that the first period action  $\alpha_1$  affects the availability of information in period 2). The official's maximization problem thus reads:

$$\max_{\alpha_1, \alpha_2 \in \{a, b, s\}} 2(1 - \theta)R + \theta \mathbb{E} \{W_1(\alpha_1|\Omega_1, \mathcal{O}) + W_2(\alpha_1, \alpha_2|\Omega_2, \mathcal{O})\} \quad (3)$$

$\theta > 0$  implies that any policy that solves the official's problem (3), also maximizes expected social welfare. As a result, we can use this outcome as an efficient benchmark against which we can compare the efficiency of other institutional setups.

To determine whether or not he is better off by first-period experimentation, the official has to evaluate two possible strategies from his prior perspective:

E. Experiment by choosing a risky action in period 1. Note that this will always

be action  $a$ , since this action dominates  $b$  for period 1 given that  $p > \frac{1}{2}$ . I will refer to this option as " $E$ " (for "experimental").

S. Maintain the status quo in period 1. This option will be referred to as " $S$ " (for "safe" or "status quo").

Since experimentation entails static losses (Assumption 1), a decision maker will never implement one of the experimental actions in period 2 if he maintained the status quo in period 1 (because the latter strategy does not generate any information on the state). Consequently, the official's lifetime utility from maintaining the status quo in period 1 equals:

$$U_S^{NA} = 2(1 - \theta)R + \theta 2\sigma \quad (4)$$

For the experimental option  $E$  on the other hand, the official's expected lifetime utility reads:

$$\mathbb{E}\{U_E^{NA}\} = 2(1 - \theta)R + \theta(1 - x)(p + q), \quad (5)$$

where the expectations operator  $\mathbb{E}\{\cdot\}$  only incorporates information that is available in the prior stage (which is the relevant stage here, as the decision maker is considering whether or not he should experiment in period 1).

The first term of equation (5) captures the certain ego rents that the official enjoys from being in power. On top of these, the experimental strategy can pay off if the official is competent (to which he attaches prior probability  $(1 - x)$ ), while  $\theta$  captures the degree to which he derives personal utility from social welfare. " $p$ " represents the expected social pay-off from following the risky strategy in period 1, while the official also takes into account that first-period experimentation produces information on the second-period state  $\Omega_2$ . He expects this information to be correct with probability  $q$ , in which case he can secure a certain social pay-off of 1 in the second period. With probability  $(1 - q)$  he expects the information to be wrong, in which case the pay-off is zero.

Comparing (4) and (5) shows that a non-accountable official will experiment in period 1 if and only if:

$$\sigma \leq \frac{(1 - x)(p + q)}{2} \equiv \bar{\rho}^{NA} \quad (6)$$

Equation (6) has a nice intuitive interpretation: it shows that experimentation by taking action  $a$  in the first period is more attractive when we are relatively certain that  $a$  is the correct action (high  $p$ , which implies low uncertainty on the optimal policy  $p(1-p)$ ), when action  $a$  is highly informative on the state of the world (high  $q$ ), when the experimental strategy is easy to implement (*i.e.* it is unlikely to fail due to incompetence; low  $x$ ), or when the status quo performs relatively poorly (low  $\sigma$ ). The latter effect suggests that policy experiments should be more common in countries where the status quo generates poor outcomes (conditional on decision makers believing that there exists a policy that would generate better outcomes). The reason is that a low  $\sigma$  reduces the opportunity cost of experimentation. This seems to be in accordance with reality: China-type Special Economic Zones, Chile-style macro-experiments, as well as randomized controlled trials all seem more widespread in less developed countries than they are in developed ones (while reforms in more developed countries tend to be implemented when the status quo is yielding poor results (*cf.* the reforms in Britain and New Zealand in the 1980s)).

## 4.2 Representative democracy

In a representative democracy ("RD"), the decision maker (now referred to as "politician") does have to care about reelection as he can be voted out at the beginning of the second period. At that point, he loses his direct pleasure from being in office ( $R$ ), as well as any pay-offs related to his country doing well (as he is no longer responsible and therefore no longer admired for his country's performance). This leaves him with his outside option, the value of which I set to zero.

His lifetime utility is then given by:

$$U^{RD} = (1 + \mathbf{1}_{\text{Re}})(1 - \theta)R + \theta [W_1(\alpha_1|\Omega_1, \mathcal{O}) + W_2(\alpha_1, \alpha_2|\Omega_2, \mathcal{O}) \cdot \mathbf{1}_{\text{Re}}]$$

Here,  $\mathbf{1}_{\text{Re}}$  is an indicator function that equals 1 if the politician is reelected at the beginning of period 2, and 0 otherwise. All other terms are defined as before.

The politician's problem also consists of choosing whether to go for the experimental policy  $E$ , or to stick to the safe status quo  $S$ . When he sticks to the status

quo, he is reelected: after having implemented  $s$  in the first period, he, nor any other politician, can do better in expected value terms than again implementing  $s$  in period 2 (as no information on the second-period state  $\Omega_2$  is produced). Since maintenance of the status quo also does not produce any information on the official's type, this makes the electorate formally indifferent between all candidates. But because the incumbent has an electoral advantage over his challengers, he is reelected (whatever the exact reason for this incumbency advantage may be - recall the discussion following Assumption 3 and see Section 4.3 below for an interpretation in terms of a model of "agenda setting").<sup>9</sup> His lifetime utility from following option  $S$  therefore equals:

$$U_S^{RD} = 2(1 - \theta)R + \theta 2\sigma \quad (7)$$

The expected lifetime utility from option  $E$  now however reads:

$$\mathbb{E}\{U_E^{RD}\} = [1 + (1 - x)p](1 - \theta)R + \theta(1 - x)(p + pq) \quad (8)$$

Under representative democracy, an experimenting politician has a probability  $(1 - x)p$  of surviving into the second period. After all, the first-period experiment will only pay-off if the state is correct (which has probability  $p$ ), while the official also needs to be competent (probability  $(1 - x)$ ). If the first period policy fails, the incumbent is voted out (recall the discussion in Section 3). Conditional upon survival the politician receives  $R$  in the second period, while  $(1 - x)pq$  is the expected probability that a competent official is reelected and receives correct information on  $\Omega_2$  after which he can implement the right period 2 action (securing a unit pay-off to society). As before,  $\theta$  captures the degree to which the politician enjoys personal utility from his country doing well.

A comparison of equations (7) and (8) then shows that an accountable politician will only decide to experiment in period 1 when:

$$\sigma \leq \frac{(1 - x)(p + pq) - [1 - (1 - x)p] \frac{1 - \theta}{\theta} R}{2} \equiv \bar{\rho}^{RD} \quad (9)$$

---

<sup>9</sup>I will moreover return to this issue in Section 6.1, where Assumption 3 is relaxed.

The interpretation of equation (9) is very much like that of (6), but now the official is not certain whether he will survive into the second period if he chooses to experiment, while the direct utility flow from being in office ( $R$ ) also shows up. This utility flow is being put at risk if the politician experiments in period 1, so the larger the value he attaches to being in office relative to the value he attaches to social welfare (as measured by  $(1 - \theta) / \theta$ ), the less willing he becomes to experiment.

### 4.3 Implications of accountability

Now that we have established conditions under which accountable and non-accountable officials are willing to experiment, we can compare the two policy rules and see how accountability affects experimentation incentives.

Because  $0 < p, x < 1$ , and  $R \geq 0$ , the RHS of inequality (9) is always smaller than that of (6) - even for  $R = 0$  and/or  $\theta = 1$ . That is:  $\bar{\rho}^{RD} < \bar{\rho}^{NA}$ , with  $\bar{\rho}^{NA}$  being the cut-off value for  $\sigma$  below which a non-accountable official would start to experiment; see equation (6). This shows that an accountable politician is less eager to experiment with new policies than a decision maker who is not subject to reelection (*i.e.* there is a "reelection bias"). For an accountable politician, the status quo has to be really bad (low  $\sigma$ ) before he becomes willing to experiment, or experimentation has to be really attractive (high  $p$  and  $q$ ; low  $x$ ).

Hence, for those parameter values such that  $\sigma \in (\bar{\rho}^{RD}, \bar{\rho}^{NA}]$ , an official who faces no risk of losing office would experiment in the first period, while an accountable politician would not - despite the fact that experimentation is socially optimal in these circumstances (this follows from the fact that the non-accountable official's policy maximizes social welfare; recall the discussion following (3)).

What happens is that the mere existence of incompetent policy makers, combined with the inference problem voters face, "crowds out" desirable experimentation, even if competent policy makers hold office: under accountability the policy maker is afraid that his first-period experimental action turns out to be wrong, after which voters will regard him as "incompetent" (even if this is not true) and vote him out. Instead, maintaining the status quo suffices to get reelected - for example because the

incumbent is preferred along some other dimension, which is why he found himself in power in the first place.<sup>10</sup> If the incumbency advantage is indeed driven by such a selection effect, one can think of this paper's model as being a model of agenda setting: as long as an incumbent expects to have an advantage over his challengers in some dimension  $X$ , he will attempt to free-ride on this advantage by trying to keep all other dimensions (such as "competence") out of the spotlights. This can be achieved by maintaining the status quo, as that action does not produce any new information along the competence dimension (also see Zwiebel (1995), who relates a similar idea in a corporate setup to the (now outdated) saying that "no one ever got fired for buying IBM" - the erstwhile status quo action when it came to buying new computer equipment). Related channels operate when the incumbency advantage is caused by government turnover costs, more media attention for the incumbent during the election campaign, or by a preference of voters for a known candidate over an unknown one.

Because the risky action decreases expected tenure, career-concerned politicians fail to internalize the full positive informational spillover that results from experimentation, as it may accrue to a different policy maker. Note that this result does not rely upon the existence of risk aversion, as it already arises in a setup where both

---

<sup>10</sup>Consider the simple setup in which voters determine who to elect at the beginning of period 1 based upon two dimensions. The first dimension captures voter beliefs on competence, while the second dimension represents ideas on, say, immigration policy (but this could also be a candidate's characteristic, such as "likeability" or "oratory skills"). Voters can choose from  $N \geq 2$  candidates, but because none of these initially has experience in office, they are all perceived as being identical along dimension 1 - even more so because their performance will also depend upon their relationship with new, unknown staff members. Consequently, the second dimension becomes pivotal and the candidate whose immigration policy plans are most appealing to voters (or: the candidate who is most likeable/the best orator), ends up in office. This incumbent subsequently realizes that if he manages to avoid making any mistakes in the first dimension (such that beliefs on his competence level do not fall, which can be achieved by maintaining the status quo), the second dimension will again be pivotal in the period 2 election. This then implies that the incumbent has a good chance of being reelected, as voters selected (so apparently liked) him along this dimension in the first period. More generally, the exact nature of the incumbency advantage does not matter for this paper's findings: a more exogenous incumbency advantage (such as government turnover costs or more media attention for the incumbent during the election campaign), or assuming that voters prefer a known candidate over an unknown one, produces the exact same results.

decision maker and society are risk neutral.<sup>11</sup>

Hereby, this paper's model can explain why there seems to exist a bias towards maintaining the status quo at the expense of implementing reforms. This "status quo bias" has been studied extensively since the seminal contribution of Fernandez and Rodrik (1991). In a median voter framework, they point out that when there is individual-specific uncertainty (individual voters not knowing whether they will win or lose as a result of the reform), expected pay-off maximization can imply that the median voter blocks an efficiency-enhancing reform *ex ante*, even though he would welcome it *ex post* (already under risk neutrality). In practical democracies however, most policies are decided upon via representation, which is not always well-approximated by a median voter analysis (Cukierman and Spiegel, 2003). By modeling the principal-agent relationship underlying representative democracies, this paper has shown that a status quo bias can still arise in such a setup: when reforms entail a certain amount of policy experimentation (which does not seem particularly unrealistic, especially if one considers reforms that have not been implemented before in the respective country), accountable politicians who have a desire for reelection will choose to maintain the status quo too often from a social welfare point of view (already in the absence of individual-specific uncertainty, which drives the Fernandez-Rodrik mechanism). As in Fernandez and Rodrik (1991), this argument does not rely upon risk aversion, irrationality, or hysteresis due to sunk costs. Different from that paper however, this paper's status quo bias results from the principal-agent problems underlying representative democracies.

---

<sup>11</sup>Intuition moreover suggests that the reelection bias would strengthen if voters would not allow politicians to learn from past experiences. This for example happens in Majumdar and Mukand (2004) where a policy change makes politicians look bad (as it may suggest that they do not know what they are doing). This makes the experimental action even less attractive, as voters effectively do not allow policy makers to use the information generated by experimentation in that setup. This would only add to the reelection bias resulting from the present paper's model.

## 5 The impact of risk aversion

So far the analysis has assumed that both society as well as decision makers are risk neutral, to illustrate that this paper's channel does not rely upon risk aversion ("RA"). In practice, both society ("Soc") and decision making officials ("DM") are probably risk averse - individual decision makers even more so than society (*cf.* Samuelson (1964), Arrow and Lind (1970) and Chambers and Echenique (2010)). This section therefore analyzes the impact of risk aversion on the incentives to experiment. To do so, I first consider the case where  $RA_{DM} > RA_{Soc}$ , after which I turn to that of  $RA_{DM} < RA_{Soc}$ .

**Case I:**  $RA_{DM} > RA_{Soc}$  First suppose that the decision maker is more risk averse than society. In particular, let us assume that the decision maker is risk averse, while society continues to be risk neutral (but conclusions are robust to the more general case in which the decision maker is *more* risk averse than society - the latter not necessarily being risk neutral). The period utility function of the decision maker in office now reads:

$$u_t = \Psi((1 - \theta)R + \theta W_t),$$

where  $\Psi(\cdot)$  represents a concave function.

Then, by taking similar steps as before, the non-accountable official's expected lifetime utility from implementing the risky option can be shown to equal:

$$\mathbb{E}\{U_E^{NA}\} = (1 - x)(p + q)\Psi((1 - \theta)R + \theta) + [2 - (1 - x)(p + q)]\Psi((1 - \theta)R) \quad (10)$$

Since  $\Psi(\cdot)$  is concave, Jensen's inequality implies that risk aversion reduces the decision maker's expected utility from the risky experimental strategy (RHS of (10)  $<$  RHS of (5)), which makes the official less willing to experiment with new, untried policies. Note that this channel is active already *without* accountability (equation (10) refers to a non-accountable official), so it works independently of the "desire for reelection channel" focused at in Section 4.

Since society is still assumed to be risk neutral, it has the same appetite for experimentation as before. However, risk averse decision makers are less willing to experiment, so in this case experimentation would already be suboptimally low in the absence of accountability. Making the decision maker accountable would only enhance this bias by the mechanism set out in Section 4.3.

**Case II:**  $RA_{DM} < RA_{Soc}$  Performing a similar analysis on a decision maker who is less risk averse than society, shows that there will be too much experimentation without accountability. In those circumstances, introducing political accountability is actually welfare improving, as the bias against experimentation that comes with it counters the inclination of the decision maker to experiment too much.

For representative democracies, this channel has the interesting implication that voters would be better off if they were governed by a policy maker whose degree of risk aversion is lower than that of society. After all, this character trait counteracts the reelection bias as a result of which such a politician is more likely to implement socially desirable policy experiments. (Of course provided that he is not that risk loving that he would overshoot the optimal level of experimentation.)

This result relates to the seminal work by Rogoff (1985) on why it can be optimal to appoint a central banker who does not share the social objective function, but places more weight on inflation stabilization than society does. In a political context, this paper points out that it can be optimal for representative democracies to elect decision makers who do not share the social objective function either, but have non-representative risk attitudes instead.

What can we say about the relative degree of risk aversion of policy makers in practice? Do we actually observe such a bias? Unfortunately, information on risk attitudes of elected policy makers is hard to obtain, but to generate some tentative evidence one could look at an interesting proxy: hand preference. The reason is that the psychology literature has established that left-handers tend to be more eager to engage in experimentation and risky activities than right-handers (see *e.g.* Halpern and Coren (1990) and McNamara *et al.* (1994), who also report that left-handers are more likely to experiment with drugs).

Evidence from this proxy indeed suggests that relative risk lovers (*i.e.*: left-handers) tend to do well in politics: despite the fact that left-handers only form about 10 percent of the population, 5 out of the last 7 U.S. presidents were left-handed by birth (Jimmy Carter and George W. Bush are the two exceptions; Ronald Reagan was born left-handed but was taught to write with his right hand). In fact, we can go back as far as 1929<sup>12</sup> and even though this is likely to produce a biased picture (since left-handedness used to be seen as a disability that had to be suppressed; *cf.* Reagan's experience), still 7 out of the last 14 Presidents wrote with their left hand. More comprehensive evidence - for example from other countries and from other correlates of risk aversion - is needed to draw definite conclusions, but it is a striking observation that may require an explanation: in both cases, the p-values of these statistics equal 0.00018.

With respect to other correlates of risk attitudes, psychological studies such as Lilienfeld *et al.* (2012) have established that many U.S. Presidents had a psychopathic lack of fear, which is associated with a greater willingness to take risks as well. Simultaneously, Presidents with this "disorder" initiated more new projects and are viewed as the most successful ones *ex post*, so voters indeed seem to appreciate such a character trait - exactly as hypothesized by Franklin D. Roosevelt in the introductory quote for this paper. Hübler (2012) furthermore reports a negative relationship between height and risk aversion and, interestingly, U.S. Presidents have also been taller than average (see Persico, Postlewaite and Silverman (2004); this height-difference already existed well before the introduction of modern media, so it seems to be more than a simple "appearance effect"). Finally, across countries, left-handers are overrepresented among the current 20 Latin American Presidents as well: although I haven't been able to classify all of them, there seem to be at least 5 left-handers (p-value  $\approx 0.043$ ).<sup>13</sup> By combining the information from the U.S. and Latin America, we end up with (at least) 12 left-handers out of 34 Presidents, which already has a p-value of 0.000066 - similar to that of flipping a supposedly fair coin

---

<sup>12</sup>See <http://www.handresearch.com/news/barack-obama-left-handed-us-president.htm>.

<sup>13</sup>See the pictures of them signing a document at <http://greatlefties.blogspot.co.uk/2012/03/latin-american-leaders.html>. Unfortunately, I have not been able to find the corresponding information for European leaders.

and obtaining 22 "Heads" in 25 trials.

This paper's model suggests a possible (but by no means "the one and only") explanation: being left-handed (as well as being tall and having a psychopathic lack of fear) correlates with being less risk averse and more eager to experiment, which is a character trait that voters appreciate in policy makers. The electorate could infer such risk attitudes from earlier political careers: both Clinton and Obama for example showed as Governor respectively Senator that they were not afraid to change the status quo (Obama also played this card via his 2008 Presidential campaign slogan: "*Change We Can Believe In*", emphasis mine).<sup>14</sup> To the extent that voters find these observable pieces of information credible,<sup>15</sup> they may signal to them that those types are also more likely to undertake welfare-improving policy experiments when elected President.

If, on the other hand, there is some other channel responsible for the apparent overrepresentation of risk-loving policy makers, then the present paper suggests that this channel entails a positive externality, as it counters the reelection bias.

Since voters are not perfectly able to infer risk characteristics of candidates from their behavior, one may also want to rely upon self-selection. This could be achieved by increasing the risk in the politician's compensation scheme. As shown by Bandiera *et al.* (2011), firms (who rely upon a combination of inference by a hiring committee and offering a particular payment structure) are able to match risk-seeking CEOs with firms that benefit most from such a characteristic, which suggests that countries might be able to do so as well.

---

<sup>14</sup>Arguably, both Clinton and Obama also lived rather risky personal lives (*cf.* the former's sexual escapades, while Obama experimented with both marijuana and cocaine; at one instance he even wanted to use heroin, but ended up not doing so because he did not like the dealer who was trying to sell it to him). As for example argued in Ludwig (2002), there is a tight link between the personal lives of politicians and policy outcomes.

<sup>15</sup>After all, voters may recognize that especially something like a campaign slogan can just be cheap talk, particularly given the fact that there is a time-consistency problem present here: once elected, this paper's reelection bias may make the fresh President want to stick to the status quo too often, free-riding on his incumbency advantage - despite his earlier promises (some have suggested to me that this is somewhat in line with what Obama did during his first term). In this respect, it might be interesting to analyze the possibility of a separating equilibrium, in which risk averse candidates are not able to mimic the behavior of risk loving ones. I leave this for future work.

## 6 What if there is too little experimentation?

While the reelection bias is countered by appointing policy makers who are less risk averse than society, there is no guarantee that this eliminates the entire bias. Moreover, when a representative democracy finds itself with a decision maker who happens to be (weakly) more risk averse than society, the amount of policy experimentation will be suboptimally low. This raises the question what else can be done to correct this bias (apart from abolishing accountability, which entails costs along other dimensions; recall footnote 2).

The problem essentially stems from a discrepancy between the objective function of the politician and that of society. Most notably, the former is distorted by a non-zero probability of being voted out after first-period experimentation (while maintaining the status quo suffices to get reelected by the incumbency advantage). Since this makes the objective of the decision maker differ from that of society, it brings the inefficiencies. Realigning objectives would therefore solve the problem.

However, affecting the probability of reelection after first-period experimentation, is practically difficult. One interesting way in which this could be achieved, is by letting the incumbent "privatize" (hide) the information on the second-period state  $\Omega_2$  generated by the first-period experiment. Then, after having experimented in period 1, the incumbent has an informational advantage over any challengers, which increases his reelection probability after a first-period policy failure. It is however not obvious whether governments are able to hide information for successors in this way (recall footnote 7), while this may also bring costs along other dimensions as it distorts the democratic process.

The remainder of this section describes two alternative ways via which the inefficiencies could be removed.

### 6.1 Tie-breaking rule

As discussed under Assumption 3 in Section 3, I have so far assumed that the incumbent is able to exploit a slight electoral advantage over his challengers - for example due to a selection effect as in Ashworth and Bueno de Mesquita (2008), as a result of

government turnover costs, or because of the mechanism underlying Meyer (1991). As a result, the incumbent is reelected when voters are indifferent between him and a challenger on the competence dimension, and maintaining the status quo brings him another term in office.

Although quite standard in the literature (see *e.g.* Jehiel (1999)), this tie-breaking rule is not necessarily the optimal one. In fact, if the electorate is able to commit to vote according to a tie-breaking rule that is somewhat tougher on policy makers who chose to maintain the status quo, it can remove the bias against experimentation.

To see this, suppose that the electorate commits to a rule that reelects the incumbent with probability  $\phi$  when he is seen as equally competent as his challengers. Notwithstanding any superiorities that the incumbent may have along other dimensions than competence, and notwithstanding any of the other aforementioned reasons for why reelecting the incumbent in case of indifference might be optimal, I will impose that  $\phi$  is strictly less than 1. By following similar steps as before, it follows that an incumbent will then choose to experiment when:

$$\sigma \leq \frac{(1-x)(p+pq) - [\phi - (1-x)p] \frac{1-\theta}{\theta} R}{1+\phi} \equiv \bar{\rho}_\phi^{RD}$$

The question subsequently becomes: can we set  $\phi$  such that  $\bar{\rho}_\phi^{RD} = \bar{\rho}^{NA}$ ? (With  $\bar{\rho}^{NA}$  still characterizing the efficient cut-off for experimentation, defined in (6).) After some algebra, it is possible to show that the answer is affirmative, namely by setting:

$$\phi = \frac{\left(\frac{1}{2} + \frac{1-\theta}{\theta} R\right) p + \left(p - \frac{1}{2}\right) q}{\frac{p+q}{2} + \frac{1-\theta}{(1-x)\theta} R}$$

One can get more intuition for this expression by analyzing its comparative statics. Doing so, tells us that:

$$(i) \frac{\partial \phi}{\partial q} < 0, (ii) \frac{\partial \phi}{\partial p} > 0, (iii) \frac{\partial \phi}{\partial x} \leq 0, (iv) \frac{\partial \phi}{\partial R} < 0, (v) \frac{\partial \phi}{\partial \theta} \geq 0$$

The reason for the sign of (i) is that when  $q$  increases, the signal generated by first-period experimentation becomes more informative on the optimal period 2 action.

This makes experimentation more attractive, but more so for a non-accountable official who lives in the perfect benchmark world. The reason is that an accountable politician runs the risk of no longer being in charge when the opportunity to exploit the signal arises. Consequently, accountable politicians who maintain the status quo have to be "penalized" with a lower reelection probability in case of indifference  $\phi$ , to induce them to experiment up to the efficient level.

An increase in  $p$  reduces the probability of first-period failure, thereby making both accountable and non-accountable officials equally more willing to experiment. Consequently, there is less need to stimulate experimentation for accountable officials via a lower reelection probability for the incumbent after having maintained the status quo, and  $\partial\phi/\partial p > 0$ .

The reason that  $\partial\phi/\partial x \leq 0$  in (iii), is because an increase in  $x$  makes failure of the experimental action due to lack of competence more likely. This makes all policy makers more reluctant to experiment, but for a good reason - namely the fact that it is less likely to improve social welfare relative to the status quo. But when  $R > 0$ , it also makes a policy maker more reluctant to try the experimental action for a bad reason - namely a personal fear of subsequent electoral defeat, after which he can no longer enjoy his private ego rents  $R$ . Consequently,  $\partial\phi/\partial x = 0$  when  $R = 0$  but  $\partial\phi/\partial x < 0$  when  $R > 0$ . In the latter case, the lower reelection probability in case of indifference again induces the policy maker to experiment up to the efficient level.

The intuition for the signs of (iv) and (v) is very similar: when  $R$  increases or when  $\theta$  decreases, the policy maker cares more about holding on to power as opposed to generating social welfare, as a result of which voters should be tougher on him when he fails to experiment, as that may be driven by a personal fear of losing his job. This can be achieved by lowering  $\phi$ .

One has to keep in mind however, that a tie-breaking rule that penalizes officials who maintain the status quo may also bring costs - for example when the selection effect present in Ashworth and Bueno de Mesquita (2008) is important, or when government turnover costs are large. Analyzing this trade-off in a model that contains both channels, might therefore be an interesting avenue for future research.

## 6.2 Direct democracy

Finally, an alternative way to realign objectives would be to make decisions by direct democracy, via initiatives and referenda. Then "society" is the decision maker, while accountable politicians are only involved with the implementation of society's choices. In such a setup the socially optimal level of experimentation can be achieved, as the decision to experiment or not is then no longer in the hands of officials who are afraid to lose their positions if the experiment fails (or have a higher degree of risk aversion than society). It is however unlikely that this entails a free lunch: in case of large uncertainty on the distributional impact of the reform this could lead to a status quo bias in the traditional sense of Fernandez and Rodrik (1991), while the quality of decisions taken via direct democracy may deteriorate if the electorate is poorly informed about the optimal action.

The trade-off between having decisions taken by society via direct democracy ("DD") or by accountable politicians (who may be better informed, but are also afraid to lose their job if the experiment fails) can easily be analyzed within this paper's framework. Taking the risk neutral setup for clarity, let us suppose that politicians know the correct value of  $p$  by which Nature determines the initial state of the world  $\Omega_1$ . "Society" on the other hand is more uncertain on the state (and thereby on the optimal action), expressed by their belief that  $\Pr[\Omega_1 = A] = \hat{p} < p$  (but still  $\hat{p} > \frac{1}{2}$ , such that the variance  $\hat{p}(1 - \hat{p}) > p(1 - p)$ ). Simultaneously assume that society is also less able to learn from first-period experiences, captured by the fact that society only learns the correct state at the beginning of the second period with probability  $\hat{q} < q$ . The condition under which society decides to experiment is then given by:

$$\sigma \leq \frac{(1-x)(\hat{p} + \hat{q})}{2} \equiv \bar{\rho}^{DD} \quad (11)$$

Note that  $\bar{\rho}^{DD} < \bar{\rho}^{NA}$  (the optimal cut-off for experimentation, given by (6)) so society's inferior knowledge on the optimal action also forms a drag on experimentation. Whether this drag is better or worse than the drag caused by the reelection motive of an accountable politician, can be analyzed by comparing  $\bar{\rho}^{RD}$  (the cut-off

for experimentation in representative democracies; given by (9)) to  $\bar{\rho}^{DD}$ . In particular,  $\bar{\rho}^{DD} > \bar{\rho}^{RD}$  (which implies that direct democracy dominates representative democracy) when:

$$\left[ \frac{1 - (1 - x)p}{1 - x} \right] \left( \frac{1 - \theta}{\theta} \right) R > (p - \hat{p}) + (pq - \hat{q}) \quad (12)$$

The LHS of (12) relates to what a politician puts at risk by experimenting. It shows that the drag caused by the reelection motive increases when:

- i) the experimental action has a rather low probability of success (due to large uncertainty on the state (low  $p$ ) and/or a high probability of failure due to incompetence (high  $x$ )), or when
- ii) decision makers derive a lot of utility from being in office (high  $(1 - \theta)R$ ) relative to the utility they derive from social welfare (low  $\theta$ ).

These factors make an accountable politician less willing to experiment, while society's desire for experimentation is not affected by them.

The RHS of (12) on the other hand captures society's disadvantages. When politicians have better information than society on the initial optimal action ( $p \gg \hat{p}$ ), or when they are more able to learn from past experiences ( $pq \gg \hat{q}$ ),<sup>16</sup> representative democracy is likely to dominate direct democracy.

So if the informational advantage of politicians over society is sufficiently small, or if politicians place too much weight on staying in office relative to increasing social welfare (high  $(1 - \theta)R/\theta$ ), taking decisions by direct democracy would bring the amount of policy experimentation closer to its optimal level. Since  $(1 - \theta)R/\theta$  tends to be high in countries with bad institutions, this paper's analysis suggests that such countries may benefit from taking decisions by direct democracy.

---

<sup>16</sup>Note that this condition corrects for the fact that accountable politicians are generally less willing to experiment than society as they take into account that they will only survive into the second period if the first-period action does not fail. This is captured by the  $p$ -term in the condition.

## 7 Conclusion

*Growth is, of course, a form of change, and growth is impossible when change is not permitted. And successful change requires a large measure of freedom to experiment. A grant of that freedom costs a society's rulers their feeling of control (...) The great majority of societies, past and present, have not allowed it. Nor have they escaped from poverty.*

*Rosenberg and Birdzell (1986, p. 34)*

As expressed in the above quotation, the experimental way of learning is not only a powerful approach, but also a risky route for policy makers. This paper has formalized this risk and has analyzed what it does to experimentation incentives.

In particular, this paper has pointed out that, in an environment where voters face an inference problem on the ability of public officials (competent or incompetent?), a desire for reelection introduces a status quo bias at the expense of policy experimentation. In contrast to Fernandez and Rodrik (1991), who obtain a status quo bias in a median voter framework, this paper's bias results from the principal-agent problems underlying representative democracies.

Whether the amount of policy experimentation in representative democracies is suboptimally low, depends on the ranking of the degree of risk aversion of the decision maker relative to that of society. When society is less risk averse than policy makers, the reelection bias against experimentation is enhanced - leading to a level of experimentation that certainly is too low from a social welfare point of view. In the reverse case where policy makers are less risk averse than society, a non-accountable official will experiment too much. In those circumstances, social welfare would be improved by subjecting the policy maker to reelection.

For representative democracies this simultaneously implies that voters would like to appoint "non-representative" policy makers who are believed to be less risk averse than society, as they are more likely to implement socially desirable policy experiments because their relative love for risk counteracts the reelection bias. There is some tentative evidence that relative risk-lovers are indeed overrepresented among

policy makers. Although we should welcome this observation in the light of the present paper's model, it is of course by no means clear that this is due to demand factors: it could also be the case that risk-lovers are more likely to start a political career, or that there is yet another channel in play. If it turns out that the apparent overrepresentation of risk-loving politicians is driven by some other factor, then this paper suggests that the other factor entails a positive externality, as it counters the reelection bias. Hopefully future research can shed more light on this issue.

When the level of policy experimentation is suboptimally low, taking decisions by direct democracy can increase welfare when politicians place too much weight on staying in office relative to increasing social welfare (which is likely to be the case in countries with bad institutions), provided that society has sufficient information on the optimal action. The reason is that, under direct democracy, the decision to experiment or not is no longer in the hands of career-concerned officials who are afraid to lose their positions if the experiment fails (the root of the problem).

This paper also leaves several issues for future work. It might for example be interesting to turn this paper's two-period framework into a multi-period one. Such a model could then be used to analyze the impact of term limits on experimentation incentives.<sup>17</sup> While term limits remove the reelection motive (which could eliminate the reelection bias), they may also shorten the politician's time horizon, as a result of which it is not clear a priori what term limits would do to experimentation incentives (given that the benefits of experimentation lie in the future).

In addition, it would also be interesting to test this model's prediction that (holding all else constant) the introduction of political accountability reduces policy experimentation incentives. This is however difficult due to at least two factors. Firstly, it is empirically hard to determine whether a particular reform should be classified as containing an experimental component or not (and if so: to what extent?). Secondly, it was already noted in footnote 2 that the lack of political accountability is often accompanied by other problems, such as noncongruent interests of policy makers. This makes it hard to identify the causal effect of political accountability. Exploiting

---

<sup>17</sup>Essentially, one can interpret this paper's setup as one in which a term limit of two periods is in place.

within-country variation as in Besley and Case (1995), so as to control for the general institutional setup, might be a way to get around this issue.

Finally, although this paper is written in the political sphere, the exact same mechanisms may apply to a corporate setup: there, career-concerned managers may refrain from efficiency-enhancing policy changes as they fear that an implementation failure will hurt their career path (*cf.* Jehiel (1999)). Examining the consequences of such a bias for the broader economy could form another avenue for future exploration.

## 8 References

Arrow, Kenneth J. and Robert C. Lind (1970), "Uncertainty and the Evaluation of Public Investment Decisions", *American Economic Review*, 60 (3), pp. 364-378.

Ashworth, Scott and Ethan Bueno de Mesquita (2008), "Electoral Selection, Strategic Challenger Entry, and the Incumbency Advantage", *Journal of Politics*, 70 (4), pp. 1006-1025.

Bandiera, Oriana, Andrea Prat, Luigi Guiso, and Raffaella Sadun (2011), "Matching Firms, Managers and Incentives", NBER Working Paper No. 16691.

Beck, Günter W. and Volker Wieland (2002), "Learning and Control in a Changing Economic Environment", *Journal of Economic Dynamics and Control*, 26 (9-10), pp. 1359-1377.

Besley, Timothy (2000), "Political Institutions and Policy Competition", mimeo, London School of Economics.

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

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

Bolton, Patrick and Christopher Harris (1999), "Strategic Experimentation", *Econometrica*, 67 (2), pp. 349-374.

Callander, Steven (2011), "Searching for Good Policies", *American Political Science Review*, 105 (4), pp. 643-662.

Callander, Steven and Bård Harstad (2012), "Experimentation in Federal Systems", mimeo, Northwestern University.

Chambers, Christopher P. and Federico Echenique (2010), "When Does Aggregation Reduce Risk Aversion?", mimeo, California Institute of Technology.

Cukierman, Alex and Yossi Spiegel (2003), "When Is the Median Voter Paradigm a Reasonable Guide for Policy Choices in a Representative Democracy?", *Economics & Politics*, 15 (3), pp. 247-284.

Fernandez, Raquel and Dani Rodrik (1991), "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty", *American Economic Review*, 81 (5), pp. 1146-1155.

Gelman, Andrew and Gary King (1990), "Estimating Incumbency Advantage without Bias", *American Journal of Political Science*, 34 (4), pp. 1142-1164.

Giuliano, Paola, Prachi Mishra, and Antonio Spilimbergo (forthcoming), "Democracy and Reforms: Evidence from a New Dataset", *American Economic Journal: Macroeconomics*.

Halpern, Diane F. and Stanley Coren (1990), "Laterality and Longevity: Is Left-Handedness Associated With a Younger Age at Death? In: Stanley Coren (Ed.), *Left-handedness: Behavioral Implications and Anomalies*, Oxford: North-Holland.

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

Hübler, Olaf (2012), "Are Tall People Less Risk Averse than Others?", IZA Discussion Paper No. 6441.

Jehiel, Philippe (1999), "Information Aggregation and Communication in Organizations", *Management Science*, 45 (5), pp. 659-669.

Lilienfeld, Scott O. *et al.* (2012), "Fearless Dominance and the U.S. Presidency: Implications of Psychopathic Personality Traits for Successful and Unsuccessful Political Leadership", *Journal of Personality and Social Psychology*, 103 (3), pp. 489-505.

Ludwig, Arnold M. (2002), *King of the Mountain: The Nature of Political Leadership*, Kentucky: The University Press of Kentucky.

Majumdar, Sumon and Sharun W. Mukand (2004), "Policy Gambles", *American Economic Review*, 94 (4), pp. 1207-1222.

- Maskin, Eric and Jean Tirole (2004), "The Politician and the Judge: Accountability in Government", *American Economic Review*, 94 (4), pp. 1034-1054.
- McNamara, Patrick, Dana Blum, Karen O'Quin and Steven Schachter (1994), "Markers of Cerebral Lateralization and Alcoholism", *Perceptual and Motor Skills*, 79 (3), pp. 1435-1440.
- Meyer, Margaret A. (1991), "Learning from Coarse Information: Biased Contests and Career Profiles", *Review of Economic Studies*, 58 (1), pp. 15-41.
- Mukand, Sharun W. and Dani Rodrik (2005), "In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance", *American Economic Review*, 95 (1), pp. 374-383.
- Persico, Nicola, Andrew Postlewaite and Dan Silverman (2004), "The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height", *Journal of Political Economy*, 112 (5), pp. 1019-1053.
- Persson, Torsten and Guido Tabellini (2000), *Political Economics: Explaining Economic Policy*, Cambridge: MIT Press.
- Prato, Carlo and Bruno Strulovici (2011), "Direct Democracy, Political Delegation, and Responsibility Substitution", mimeo, Northwestern University.
- Rodrik, Dani (2000), "Institutions for High-Quality Growth: What They Are and How to Acquire Them", *Studies in Comparative International Development*, 35 (3).
- Rogoff, Kenneth (1985), "The Optimal Degree of Commitment to an Intermediate Monetary Target", *Quarterly Journal of Economics*, 100 (4), pp. 1169-1189.
- Rosenberg, Nathan and L.E. Birdzell (1986), *How the West Grew Rich: The Economic Transformation of the Industrial World*, London: I.B. Tauris & Co. Ltd.
- Samuelson, Paul A. (1964), "Discussion", *American Economic Review*, 59 (1), pp. 88-96.
- Samuelson, William and Richard Zeckhauser (1988), "Status Quo Bias in Decision Making", *Journal of Risk and Uncertainty*, 1 (1), pp. 7-59.
- Willems, Tim (2012), "Actively Learning by Pricing: A Model of an Experimenting Seller", mimeo, University of Oxford.
- Zwiebel, Jeffrey (1995), "Corporate Conservatism and Relative Compensation", *Journal of Political Economy*, 103 (1), pp. 1-25.