

POLICY GAMBLES

Sumon Majumdar and Sharun W. Mukand¹

¹Majumdar: Dept. of Economics, Queens University; Kingston, Ontario. email: sumon@qed.econ.queensu.ca. Mukand: Dept. of Economics, Tufts University. email: sharun.mukand@tufts.edu. This paper has benefitted from comments by Jon Eaton, Ami Glazer, David Laibson, Gib Metcalf, Dilip Mookherjee, Stephen Morris and Dani Rodrik and seminar participants at Boston University, the University of Chicago, Queen's University, Penn State and the 2001 Canadian Economics Association meeting in Montreal.

Abstract

This paper develops a theory of policy making, that examines the incentives for experimentation with new policies and the scrapping of adopted policies. We demonstrate that a government which cares about its reputation out of electoral concerns, takes socially inefficient policy gambles that may result in two kinds of inefficiencies — first, a government may inefficiently experiment by undertaking a new policy initiative that it (and the voter) knows is unlikely to succeed, and second, the government may prefer to *not* learn from experience and instead persist with an adopted policy despite publicly observable evidence of its failure. Furthermore, these inefficiencies are systematically related to the electoral cycle. Early on in its term a government is likely to enact policies that are either too conservative *or* too radical, while later on in its term the government is likely to show inefficient policy persistence.

Keywords: Learning, Policy Persistence, Policy Experimentation, Leadership, Reputation

JEL Classification: D72, O20, P16.

1 Introduction

Policy making is an uncertain process, with policymakers often lacking a clear blueprint on the appropriate choice to be made. This uncertainty is rife in all arenas of policymaking – be it the impact of tax cuts, the extent of privatization, the timing of deregulation or even the most effective way to deal with an external threat. Given the endemic nature of this uncertainty, simple policy experimentation can generate useful information about the “appropriateness” of a particular course of action. Indeed many successful policy innovations started out as experiments – be it the deregulation of the airline industry in the U.S., temporary capital controls in Chile, “welfare reform” under President Clinton or privatization in Thatcherite Britain. This suggests that an issue of central importance in the political economy of policy making is the following: what factors influence a government’s incentive to engage in policy experimentation *and* learn from the information so generated? In addressing the issue this paper takes a first step towards developing a framework to analyze a government’s incentives for learning and policy experimentation.

Even a cursory examination of the experience with policy experimentation suggests inefficiencies – with governments being inefficiently conservative on some occasions as well as inefficiently reckless on others. For instance, merely learning about the (in)appropriateness of a particular policy is not useful, if on observing failure, the policymaker does not adapt and *change* course. Nevertheless, a striking aspect of the history of policymaking is the apparent unwillingness of leaders to learn from previous experiments. Political leaders are typically reluctant to change course mid-way, even if the policy is publicly perceived to be failing.¹ In the 1980s, Presidents Sarney of Brazil and Alan Garcia of Peru persisted with “heterodox” reform packages long after it was obvious to most observers that it was a failed experiment (Dornbusch and Edwards, 1990). Similarly, governments in transition economies such as Russia and Ukraine persisted with a drastic form of “big-bang” privatization, despite awareness that a mid-course correction was needed

¹In her sweeping survey of decision making by leaders through history, *The March of Folly*, Tuchman (1984) puts it most pithily, “Persistence of error is the problem.....to recognize error, to cut losses, to alter course, is the most repugnant option in government. For a chief of state admitting error is almost out of the question.”

(Aslund, 2002; Freedland, 2000). More recently, Domingo Cavallo, the architect of Argentina's experiment with a currency board, undertook desperate measures to persist with it in 2001, despite it being evident that retaining the currency board was likely to engender a crisis – as it eventually did.² Thus, the puzzle that comes up is why on observing failure, do politicians not learn from experience, but rather feel compelled to continue with a policy that no one is optimistic about. In many of these situations, it seems as if the only remaining special interest lobbying in favor of the policy is the policy maker himself.

This reluctance to learn from the results of an experiment makes it even more puzzling to understand why often the very same governments inefficiently gambled with the experiment in the first place. Such inefficiency in policy experimentation is suggested in the initiation of not only the (failed) “heterodox” experiments mentioned above, but also the experience of several of the experiments with large-scale economic reform in Eastern Europe and Africa. Indeed, the recklessness of such experimentation is documented by Weyland (2002, p111) in a number of case studies. For instance, he describes how newly elected President Perez of Venezuela disregarded information supplied by his economic advisers when he embarked on a recklessly bold experiment with neo-liberal reform. Likewise, on assuming office President Collor of Brazil defied his economic advisers and faced down public opposition when he initiated drastic privatization in the nineties. Weyland further documents similar instances of gambles in the initiation of neo-liberal policy packages in Argentina and Peru since “rather than covering their bases, the initiators of neo-liberal plans put all eggs in one basket by linking their fate to the uncertain outcomes of drastic reform”. Similarly, a newly elected President Clinton took a gamble in introducing legislation that involved a radical overhaul of health care, aware that ensuring its passage was difficult and that a more incremental approach had a better chance of success (Blendon et. al., 1995). Likewise, within days of assuming office Gustavo Noboa of Ecuador disregarded the advise of members of his economic team and initiated in the teeth of public opposition, a drastic experiment with dollarization (Cohen, 2000). These examples suggest that an additional puzzle is why governments

²Indeed Cavallo was willing to undertake several drastic measures such as manipulating tariffs, taxing financial transactions and giving export subsidies to prevent the visible collapse of the currency board (Wijnholds, 2003).

may ignore information and instead choose to experiment with an unpopular policy. Thus it seems that at times some governments have a proclivity towards inefficient experimentation, while on other occasions they appear to be inefficiently conservative. In order to see why both types of inefficiencies may occur, we focus on a simple mechanism – a political leader’s electoral concerns.

In this paper we develop a framework that helps examine the impact of electoral pressures on a government’s incentives to engage in policy experimentation as well as learn from it.³ The government faces a choice between maintaining the safe, status quo policy or experimenting with a new, untried policy that may generate higher, though uncertain returns. If the experiment is perceived to be unsuccessful, the government always has the option of costlessly reverting back to the tried and tested status quo policy. The key aspect to observe is that policy experimentation results in learning not just by the political leader, but also by the citizen-voter. Through a policy’s success or failure, the public learns not only about the appropriateness of the policy itself, but also about the incumbent’s competence at identifying appropriate policies in the first place.

This results in inefficiencies of two kinds. While a policy experiment’s poor performance generates valuable information for the policymaker about its (in)appropriateness and the need for its scrapping, he may fear that doing so will be interpreted as a sign that the government was not sure of its choice of policy in the first place. Thus, although the leader learns, he is afraid to publicly use this information in effecting a change in policy. The adverse reputational impact of a policy reversal gives the incumbent an incentive to ignore useful information produced by experimentation and inefficiently persist with its initial policy choice.

However, this raises the puzzling question of policy adoption: why would a leader who is so concerned with re-election as to inefficiently persist with a failed policy, be interested in inefficiently experimenting with a new policy of uncertain merit? Once again, we argue that reputational concerns are crucial and may give an incentive to either gamble recklessly and experiment with a

³Political theorists of the state such as Hecló (1974), Skocpol (1985) and Hall (1993) emphasize the “autonomy” of learning and policy experimentation from political pressures. In contrast, Besley (2001), Tomassi (2002) and Mukand and Rodrik (2003) suggest that political factors may be an important influence on experimentation and learning.

new policy or alternatively, show an inefficient degree of conservatism in maintaining the status quo. Strikingly, our theory implies that a politician is likely to be inefficiently conservative in experimenting precisely when the costs of such experimentation are low or the payoffs are high. In such cases the reputational gains from successful experimentation are low, and this makes the incumbent more hesitant in launching an experiment. Our analysis suggests that both these inefficiencies in policy choice can be clearly related to the electoral cycle. While later on in their tenure, governments become too conservative and inefficiently persist with policies, earlier on in their tenure, governments may be either too conservative or too reckless in engaging in policy experimentation.

Our benchmark model suggests that politicians are typically not rewarded for changing policies too often, because it signals incompetence. However, there are surely occasions where a politician who shows “flexibility” is electorally rewarded. In an extension, we show that for such flexibility to be rewarded, ideological considerations are likely to be an important part of the answer.

Given the preoccupation of politicians with their public reputation, it is hardly surprising that reputational models have been influential in the political agency literature (see Rogoff, 1990). More in the spirit of the present paper is the literature which captures the effect of uncertainty on the mapping from policies to outcomes in electoral models, as in Harrington (1993) and Canes-Wrone, Herron and Shotts (2001). In an attempt to signal ability, governments may enact policies that ‘pander’ to voter beliefs rather than their welfare. Similarly, Maskin and Tirole (2001) analyze constitutional design issues to show that signaling preferences may result in governmental pandering. Our framework instead focuses on a government’s incentives to engage in policy experimentation and change course in response to dynamic learning by both itself and the electorate.

Our paper is clearly also related to the small but influential literature which addresses the puzzle of inefficient policy persistence. For instance, Alesina and Drazen (1991) show how a ‘war of attrition’ between different groups can endogenously result in a costly delay in policy enactment. A second mechanism, due to Fernandez and Rodrik (1991), emphasizes the role of individual specific uncertainty in preventing the adoption of economic reforms. Finally, Coate and Morris (1999) argue that policies persist since, once implemented, a policy increases effectiveness

of the lobbying efforts of its beneficiaries. Our framework is also closely related to the literature on reputational decision making of managers, as in Scharfstein and Stein (1990), Ottaviani and Sorensen (2000) and especially Prendergast and Stole (1995). We elaborate on this literature in greater detail in Section 4A.

The rest of the paper is organized as follows. Section 2 gives a basic description of the model. The equilibrium of the model with electoral concerns is analyzed in section 3. Section 4 further discusses aspects of our model and related literature. Section 5 concludes.

2 Description of the Model

We begin by giving a brief outline of the model. Consider a government which has just been elected into office. It's choice of policy is governed by national welfare considerations as well as its own future re-election concerns. The incumbent assesses the situation and then decides whether to continue with the safe, status quo policy or gamble and experiment with an uncertain course of action. If enacted, the government and the public learn about the impact of the new policy on the economy. In light of this observed success or failure, the government has the option of persisting with its new policy initiative or not. The electorate then votes on whether or not to re-elect the government. We now elaborate on this structure.

Policies: For simplicity, we will restrict government policy choices to two alternatives: either to stick with the policy already in place, which we call the status quo policy a_S , or to take a new policy initiative a_N . Both policies affect a publicly observable outcome, say, national income. An important feature of the technology of policy making is that different policies are appropriate for different situations and economic environments. For example, reforms which are successful in one country may not be suitable for another. Accordingly, we assume that the success of a policy is contingent on the underlying structure of the economy, which maybe one of two types, S and N , with a_S and a_N being the 'appropriate' policies for the two environments respectively. If the underlying state is S , then enacting the new policy initiative a_N causes a net loss in welfare. However if the state is N , then adoption of the new policy is successful with probability p and

increases total national welfare. Thus, a_N is a new policy whose appropriateness for the particular economy is ex-ante unclear. Governments that are more able, are assumed to be better equipped at recognizing the suitability of policies for their respective economic environments.

We normalize the gain (over and above what can already be achieved) to continuing with the status-quo policy in either state as 0. Thus the status-quo policy is assumed to be one whose efficacy for the economy is already well-understood. There is a cost c to enacting the new policy initiative, while the potential gain from it in national income is denoted by Δ . This gain is however stochastic: it occurs only in state N , and even then with probability p . For economies of type S , the probability of success from policy a_N is zero. Thus, given the assumed technology, a success with the policy a_N yields an immediate output gain of Δ , as well as valuable information that it is appropriate for the economy i.e. that the underlying state is N , and therefore the policy should be continued with in the second period too. We make the following assumption to ensure that the expected net payoff from enacting a_N in state N is positive.

Assumption 1: $p\Delta - c > 0$

If however the government adopts policy a_N and the state is S , then there is no gain to national income, and there is a net loss in welfare since a cost c has been incurred.⁴ Alternatively, c could also be interpreted as the output from the policy a_S ; enacting the new policy in place of the status quo then has an opportunity cost of c , and is incurred regardless of the success or failure of a_N .

Although simple, the above payoff structure is applicable in a wide variety of contexts and has the added advantage of simplifying the analysis. In the context of many economic policy initiatives, success is often crucially dependent on external and internal market conditions; correct understanding of the conditions is important in making policy decisions. Similarly, the cost c can be any and every cost that governments incur when they move away from the status quo. These can include the cost of making compensatory transfers to interest groups that lose out, or something as simple as the cost of training the bureaucracy to effectively administer the new policy. Similarly,

⁴One could equivalently assume that policy a_N yields a gain G with probability p in state N ; in all other cases, it causes a loss in income of L . In our framework, this would mean $c = L$, and $\Delta = G - L$.

if the issue being studied is conflict, then the gain Δ from taking the new policy may well be the ‘peace dividend’. In this case, the state of the world could be the morale or preparedness of the enemy which would determine the success or failure of, for instance, an aggressive policy position.

Politicians and Voters: The government is assumed to be run by an elected politician. Prospective politicians differ in ability and can be one of two types: either high ability H , or of low ability L . Politicians differ in their capability to acquire or process information about the appropriateness of policies for the economy. In order to keep the analysis simple, we assume that a high ability incumbent receives a perfectly reliable signal about the state of the world i.e. he knows the true state of the world for sure. On the other hand, a low ability politician only receives a signal x from the interval $[\underline{x}, \bar{x}]$.⁵ The probability of receiving a particular signal depends on the state: if the true state is S , then the density function for signal x is given by $\phi_S(x)$, while if the state is N , then the density is $\phi_N(x)$. We make the following assumptions on these densities.

Assumption 2: $g(x) = \frac{\phi_N(x)}{\phi_S(x)}$ is strictly increasing in x , with $g(x) \rightarrow 0$ as $x \rightarrow \underline{x}$, and $g(x) \rightarrow \infty$ as $x \rightarrow \bar{x}$.

This assumption (the monotone likelihood ratio property) implies that a higher value of the signal x corresponds to a greater likelihood that the state is N . If the incumbent’s prior about the state being N is given by π_0 , then on seeing the signal x , a low ability incumbent’s belief that the state of the world is N is given by

$$b(x) = \Pr[N|x] = \pi_0 \phi_N(x) / [\pi_0 \phi_N(x) + (1 - \pi_0) \phi_S(x)] \quad (1)$$

Assumption 2 implies that this posterior belief $b(x)$ is increasing in x . Since this belief will play a central role in the decision-making process of a low ability government, we shall treat it directly as a random variable. Let us denote by $F_i(b)$ the probability that the belief (about the state being N) of a L type is less than or equal to b , when the state is i , where $i \in \{S, N\}$. Thus,

$$F_i(b) = \Pr[b(x) \leq b|i] = \Pr[x \leq g^{-1}\left(\frac{1 - \pi_0}{\pi_0} \frac{b}{1 - b}\right)|i]$$

⁵Although we focus on the model with one type having perfect information about the state of the world, as we show in Appendix II, the model can be extended to a continuum of types each differing in the probability with which they receive a perfect signal. The analysis in this extension is very similar to the model we present here.

Note that since the belief about the state being N is increasing in x , we have $F_N(b) \leq F_S(b)$.

We assume that politicians care both about the welfare of the population, as well as their own future electoral prospects. The latter maybe because being in office gives them some private non appropriable ‘ego rents’. To model this, we assume that the government’s objective function is given by

$$W_P = \gamma(\text{National Welfare}) + (1 - \gamma)(\text{Prob. of Re-election}) \quad (2)$$

where $\gamma \in (0, 1)$, is the relative weight that it puts on national welfare, which is assumed to be the same for both types of politicians.

Let us suppose that at the beginning of its term, public perception about the new government being of high ability is given by λ . At the end of the term, the government faces an election. There is a single representative voter who cares about her welfare (in this case, the net national income over the electoral cycle, after accounting for costs and benefits from policy) and thus would like to elect the most able government i.e. one which is more likely to identify ‘appropriate’ policies for the economy. This focus on a representative voter is deliberate, since we wish to explicitly minimize inefficiency in government decision-making arising due to voter heterogeneity and ideology. To this end, we also assume that all voters share the same prior π_0 on the effectiveness of the policy a_N for the economy i.e. all believe the state to be N with probability π_0 .

At the end of the term, a challenger is randomly drawn from the population and public perception about his ability is given by ω , where ω is distributed over $[0, 1]$ according to a distribution with the cdf given by $G(\omega)$. This perception may be formed on the basis of the challenger’s performance in other arenas or through an unmodeled ‘charisma factor’. Thus, if the voter’s end of term assessment about the incumbent’s ability is R , then the incumbent’s ex-ante probability of being re-elected is $G(R)$.⁶

The Timing of Decision Making: At the beginning of its term ($T = 0$), the new government has for a limited time a ‘window of opportunity’ to change the existing policy a_S . It receives

⁶This particular assumption for re-election is not important to the model. As will be clear below, the direction of the results are unchanged so long as we assume that the probability of being re-elected is some increasing function of the voter’s assessment about the government’s ability.




Figure 1: The timing of events

a private signal about the appropriateness of the policies i.e. the state of the world, and then faces the choice of either maintaining the status quo policy a_S or enacting the new policy a_N . If it decides to maintain status quo, then output remains the same. On the other hand, if the new policy is enacted, a cost c is initially incurred and both the government and the electorate learn about its impact on output, which is realized mid-term (at the beginning of $T = 1$). If the new policy turns out to be successful, everyone infers that the state is N , and therefore a_N is the appropriate policy. However, in the case where the gain of Δ is **not** realized, it is not clear whether this failure is due to a random draw or due to the policy being an inappropriate one; the government now faces the important decision of continuing with its policy initiative a_N or reverting back to the initial status quo a_S . If it continues with a_N , it again incurs a cost c , and its impact on output (i.e. Δ or 0) is observed before the election. The decision tree for a government, along with the timing of events is shown in figure 1.

The representative voter makes inferences on the ability of the incumbent based on the sequence of policy choices as well as the realization of their impact on output, and chooses either to re-elect or throw out the government at the end of the term.

2.1 Benchmark case: Socially Efficient Decision Making

The focus of our analysis is to study the impact of electoral imperatives on a government's policy choices. In order to facilitate this examination, we study as a benchmark case, the policy choices of a government which cares *only* about national welfare.

If the incumbent is of high ability, the analysis of its choices is straightforward. Since by assumption a high ability incumbent receives a perfectly informative signal, he will always choose a policy in accordance with his private signal. If he receives a signal that the state is N , he will find it optimal to enact a_N . Furthermore, even in the event that the policy does not achieve success in the short-run i.e. Δ is not realized in the first period, the high ability incumbent will find it optimal to persist with a_N since by assumption 1, $p\Delta - c > 0$. Similarly, if the signal is that the state is S , then the high ability incumbent will choose to maintain the status quo policy a_S .

The low ability incumbent receives a signal x which is only imperfectly correlated with the state of the world. So his efficient policy choices will be a function of the strength of this signal. Recall that under assumption 2, higher values of x make it more likely that the state is N . Hence a low ability incumbent will be willing to experiment with the new policy initiative only if his private signal x is high enough. Furthermore, for extremely high values of x (i.e. $b(x) \approx 1$), it is optimal for the incumbent to persist with a_N even if Δ is not realized in the first period. Thus in analyzing the choices of a low ability incumbent, there will be two cutoff beliefs, b_0 and b_1 , with $b_0 < b_1$. Only if his belief about N is higher than b_0 will he enact policy a_N at $T = 0$; those incumbents with beliefs between b_0 and b_1 will revert back to the status quo policy a_S in the face of failure, while only those with initial beliefs greater than b_1 will persist. These beliefs correspond to two cutoff signals x_0 and x_1 , with $b_0 = b(x_0)$ and $b_1 = b(x_1)$, where $b(\cdot)$ is the posterior belief on receiving signal x , as defined by equation (1). Those with a signal $x > x_0$ enact a_N and only those with signal $x > x_1$ persist even on failure. We now solve for the two cutoff beliefs.

First, consider the event when the realized mid-term output from policy a_N has been 0. The incumbent then updates his initial belief b (about the state being N) to $\frac{(1-p)b}{(1-p)b+1-b}$. He will be indifferent between continuing with a_N and switching back to a_S only if $\frac{(1-p)b}{1-pb}p\Delta - c = 0$. This

gives the efficient cutoff level of belief for $T = 1$ as

$$b_1^{eff} = \frac{c}{p[c + (1 - p)\Delta]}$$

Next, consider an incumbent who is indifferent between enacting a_N at $T = 0$ or sticking to the status quo a_S . He realizes that if the policy a_N does not yield an output gain of Δ mid-term, he will revert back to a_S . On the other hand, if it is a success then it becomes clear that the state is N and he should continue with a_N the next period. Thus if his initial belief is b_0 , the expected output from enacting a_N is given by $b_0p(\Delta + p\Delta - c) - c$. This gives the efficient cutoff level of belief for $T = 0$ as

$$b_0^{eff} = \frac{c}{p(\Delta + p\Delta - c)}$$

Since $p\Delta - c > 0$, we have $b_0^{eff} < b_1^{eff}$. To summarize, if the incumbent's belief about the state being N is greater than b_0^{eff} , he will take the new policy initiative a_N , but only if his belief exceeds b_1^{eff} will he persist with this policy even if it does not find mid-term success.

In this section we analyzed the benchmark case of an incumbent who cares *only* about the welfare of the representative voter. We now relax this extreme assumption of a purely benevolent politician and allow him to be also concerned about his future electoral prospects.

3 Policy Making and the Electoral Imperative

Governments do care about national output and voter welfare. Electoral concerns, however, often weigh heavily on the policy choices they make. Since the representative voter's welfare is a function of the ability of the government in identifying appropriate policies, she will always prefer to re-elect the incumbent if his perceived ability is greater than that of the challenger. Therefore, voters attempt to glean all possible information that they can about the government's ability from its policy choices. An incumbent who chooses policies that boost national income increases his chances of retaining power. In the face of uncertainty, experimentation with a new policy, will result in learning about the appropriateness of alternative courses of action. However, a government that appears to learn from its performance and switches policies will be perceived not to have been




Figure 2: Equilibrium policy decisions by the L type government as a function of his belief b

sure about its choice of policy in the first place, since more able governments do not need to learn as much. Once account is taken of this signaling aspect of policy choices, inefficiencies may arise both in the adoption of new policies as well as in the scrapping of adopted policies.

Let us consider a Bayes-Nash equilibrium of this political game in which the incumbent's objective function is given by (2). It will consist, for each type of incumbent, of a strategy for the initial period i.e. whether or not to enact the new policy a_N . This will be contingent on the private signal that the incumbent receives regarding the state of the world. In the case where it decides to enact a_N , it will also have to decide whether or not to persist with a_N (or to revert back to a_S) after realization of the mid-term output. Voters observe both the sequence of policy choices as well as the outcomes, and based on this information form expectations about the ability of the incumbent. Finally, the strategies of the incumbent and the voters' beliefs have to be consistent with each other in that the voters' beliefs are derived using Bayes' rule and the government's strategy (whenever possible), and the incumbent's strategy is optimal given these beliefs.

We begin by proposing the following equilibrium structure: a high ability government always chooses policies in accordance with its private signal. Recall that it obtains perfect information about the state of the world. So only if its signal indicates that the state is N , will it enact a_N in the first period, but having done so, it will persist with a_N in the second period even in the case of a failure. For a low ability government with a signal x , it will take the policy initiative a_N only if its belief $b(x)$ exceeds a certain cutoff value b_0^* . In the event that a output gain of Δ is realized by the middle of the term, it will continue with a_N in the second period. However, in the case

of failure, it will persist with a_N only if its initial belief $b(x)$ exceeded the cutoff value b_1^* , where $b_1^* > b_0^*$; otherwise, it will revert back to a_S . Decisions for the L type are summarized in figure 2.

In our model, while there are two types of governments, following the realization of beliefs, the ex-ante difference between the types disappears. For example, a high ability government who gets perfect information that the state is N , is identical to a low ability type with the belief $b(x) = 1$, and their decisions will be the same. This feature of the model means that even if there were a continuum of types differing in their probability of getting a perfect signal (as we develop in Appendix II), the equilibrium structure will again involve only two cut-off beliefs b_0^* and b_1^* , with all types with beliefs above b_0^* enacting a_N and only those with beliefs above b_1^* persisting.

We will now examine the conditions under which the above equilibrium arises.

3.1 The Second Period: Policy Persistence in the Face of Failure

We begin with the second period, after payoffs from the first period policy choice have been publicly realized. If the initial choice of policy a_N resulted in an output gain of Δ , everyone infers that the state is N and therefore (under assumption 1), it is appropriate to continue with a_N . On the other hand, if there has been no gain in output, then the incumbent has two choices: either to persist with a_N or to revert back to the original status quo policy a_S . Under the proposed equilibrium, a high ability politician never switches back. Therefore, if a government chooses to abandon the policy a_N and revert back to a_S , then the voter infers that it must be one of low ability. If however, the incumbent chooses to persist with a_N , then the voter is not sure whether it is one of high ability or if it is a low ability type with a belief exceeding b_1^* .

One possibility from persisting with a_N even in the light of mid-term failure is that it may yield an output gain of Δ in the second period. In such an event, the voter would infer that the state is N , and his end of term assessment about the incumbent's ability will be given by $R_S(b_1^*) =$

$$\begin{aligned} & P(H \mid a_N \text{ at } T = 0, \text{ output} = 0 \text{ in 1st pd.}, a_N \text{ at } T = 1, \text{ output} = \Delta \text{ in 2nd pd.}) \\ &= \frac{\lambda}{\lambda + (1 - \lambda)[1 - F_N(b_1^*)]} = \frac{\lambda}{1 - (1 - \lambda)F_N(b_1^*)} \end{aligned} \quad (3)$$

In state N , all high ability governments enact policy a_N and persist with it, but so do low ability

governments with beliefs exceeding b_1^* . As b_1^* rises, the range of beliefs over which a L type will persist, shrinks. Therefore, this reputation $R_S(b_1^*)$ is increasing in b_1^* .

The other possibility from continuing with a_N in the face of mid-term failure is that it may fail in the second period too. Since a_N has probability $1 - p$ of yielding no gain in output even in state N , the voter's personal belief about the state being N after observing two failures of policy a_N is given by $\pi_2 = \frac{(1-p)^2\pi_0}{(1-p)^2\pi_0+(1-\pi_0)}$, and his assessment of the incumbent's ability is $R_F(b_1^*) = P(H | a_N \text{ at } T = 0, \text{ output} = 0 \text{ in pd.1, } a_N \text{ at } T = 1, \text{ output} = 0 \text{ in pd. 2}) =$

$$\frac{\lambda\pi_2}{\lambda\pi_2 + (1-\lambda)[\pi_2(1-F_N(b_1^*)) + (1-\pi_2)(1-F_S(b_1^*))]} \quad (4)$$

Observe that, as expected, this reputation is lower than that under success viz. $R_S(b_1^*)$.

In making its decision on whether or not to persist with a_N in the face of failure, the incumbent considers the consequences of its action both on its electoral prospects as well as on national welfare. If he switches back to a_S , the expected gain in output is 0, and he is immediately identified as being of low ability by the voters (and so his probability of re-election approaches zero). Alternately, for a low ability government with an initial belief b_1^* , persisting with a_N in the face of failure yields an expected gain in output of $\frac{(1-p)b_1^*}{1-pb_1^*}p\Delta - c$, and a gamble over his electoral prospects, with his reputation in the events of success and failure given by (3) and (4). For this incumbent to be indifferent between persisting with a_N or switching back to a_S , we must have:

$$V(b_1^*) = \gamma \left[\frac{(1-p)b_1^*}{1-pb_1^*}p\Delta - c \right] + (1-\gamma) \left\{ \frac{(1-p)b_1^*}{1-pb_1^*}p G(R_S(b_1^*)) + \left(1 - \frac{(1-p)b_1^*}{1-pb_1^*}p\right) G(R_F(b_1^*)) \right\} = 0. \quad (5)$$

Observe that the left-hand side of this equation is increasing in b_1 . Therefore, given suitable conditions on the end points, there exists a unique value of $b_1^* \in (0, 1)$ that satisfies the above equation. This is done in proposition 1 below.

Note that the incumbent's reputation from persisting with a_N (both $R_S(b_1^*)$ and $R_F(b_1^*)$) is always positive. Therefore the first term on the left hand side of the above equation (which is the expected output from persisting with a_N) must be negative i.e. there are beliefs under which *persisting with the policy is inefficient, yet the incumbent finds it optimal to do so*. The intuition for the result is simple. A high ability government has nothing to learn from a policy's

performance, since it already knows the appropriate policy sequence to be adopted. In contrast, consider a government that appears to learn from its performance since it switches policies after a poor realization of national output. Any government that flip flops by changing policies suggests indecision and a lack of confidence in adopting the policy in the first place – thereby revealing itself to be of low ability. However, this perception of low ability may jeopardize the incumbent’s electoral chances. It is this fear of an adverse impact on his re-election prospects that gives the incumbent an incentive to persist with a policy that even he believes is not the optimal one.⁷

This tendency for governments to *not* learn from performance of the policy, but rather persist with it long after it has become publicly discredited, is especially true for those policies whose introduction is closely associated with the incumbent. This is (arguably) true for interventionist heterodox policies observed in Latin America and the cross country policy experience reported by Yatawara (1998).⁸ Perhaps the most compelling instance of such inefficient persistence comes from the Johnson administration’s decision “to commit American ground forces to Vietnam” (Berman, 1988, p475). This was because over time “the war in Southeast Asia had become Lyndon Johnson’s war”. More tellingly, despite awareness that he had gambled and misplayed his hand, Johnson responded to all advice, as well as a steady stream of pessimistic intelligence reports and dismal military results with further escalating the level of military involvement. Indeed it is difficult to avoid the impression that Johnson’s decision to persist was a last ditch gamble to preserve his reputation and save his presidency. As argued by Schandler (1977), “Johnson found it politically

⁷Observe the importance of discontinuous reputational updating in our framework. If the government switches policies, it results in a dramatic drop in the payoff from $R_S(\cdot)$ and $R_F(\cdot)$, all the way to zero. In this model, one type always receives a perfect signal and the other type never does; so the reputation from switching is 0. In a model with a continuum types, each differing in the probability of getting a perfect signal, for (almost) every type there is a positive probability of receiving a signal between b_0^* and b_1^* . Thus in such a model, the mean reputation from switching is not 0, but it is still discretely lower than either of the reputations from persisting i.e. R_S and R_F . So the inefficient persistence result holds in this modified model too.

⁸For instance, see Pastor & Wise (1992) and Stokes (1999) on the heterodox policies pursued by Alan Garcia of Peru. Dornbusch & Edwards (1991), Tommasi & Velasco (1996) and Rodrik (1996) also describe instances of inefficient policy persistence and inefficient adoption. Consistent with our story, Blomberg et al. (2001) also empirically observe that the probability of the fixed exchange rate being maintained increases as an election approaches.

and personally impossible to change publicly the policy he had tenaciously pursued for so long”.

Let us now turn our attention to the incumbent’s first period decision.

3.2 The First Period: Policy Initiatives when No Initiatives are Needed?

We now examine a new government’s first-period incentives to maintain the status quo or experiment with the new policy initiative. To begin with, consider the reputation of an incumbent who maintains the status quo policy a_S . The voters infer that this can occur either because the government is one of high ability and knows for sure that the state is S , or because it is a low ability government whose private belief is below b_0^* . Therefore the voter’s perception of the ability of a government that maintains the status quo is given by

$$R_{\text{status quo}} = P(H | a_S) = \frac{\lambda(1 - \pi_0)}{\lambda(1 - \pi_0) + (1 - \lambda)[\pi_0 F_N(b_0^*) + (1 - \pi_0)F_S(b_0^*)]} \quad (6)$$

Note that this reputation is decreasing in b_0^* . As b_0^* increases, the threshold belief level for a low ability type to enact a_N rises. Thus it becomes more likely that a government which maintains the status quo a_S is of low ability.

On the other side, let us evaluate the impact of experimenting with the policy initiative a_N . Consider first the case when this experiment works and yields an output gain of Δ by the middle of the term. Since success can occur only in state N , the voter learns that the incumbent had initially chosen the ‘appropriate’ policy. Now, this could have been either because the government was a high ability one and had perfectly identified a_N as being appropriate for the economy or because a low ability government had received a signal which resulted in its belief about state N to exceed b_0^* . In this case, the voter’s perception about the incumbent’s ability is given by

$$P(H | a_N \text{ at } T = 0, \text{ output} = \Delta \text{ in 1st pd.}) = \frac{\lambda}{\lambda + (1 - \lambda)(1 - F_N(b_0^*))} = \frac{\lambda}{1 - (1 - \lambda)F_N(b_0^*)} \quad (7)$$

As one would expect, ‘successful’ policy experimentation boosts reputation above λ . However (as $b_0^* < b_1^*$), it is smaller than $R_S(b_1^*)$, the public perception about a government that persists with a_N even in the face of failure and ultimately achieves success. This is related to the dynamic nature

of our game in which the separation of types takes place temporally: as more of the L types drop out over time, anyone who persists with a_N is likely to see his reputation enhanced.

However, the policy a_N may also result in a short-term failure. Recall that only those L types with beliefs greater than b_1^* persist even in the face of failure. Thus, for a low ability incumbent with the marginal belief of $b_0^* < b_1^*$, it is optimal in the second period to revert to the policy a_S in the event of failure. In this case, the government is revealed to be one of low ability and its second period total payoff (both in terms of gain in output as well as re-election prospects) is zero.

Thus for a low ability government with a belief of b_0^* to be indifferent between maintaining the status quo a_S and initiating the new policy initiative a_N in the first period, we need to equate the difference in the expected payoffs from the two options. This gives:

$$\begin{aligned} & \gamma[pb_0^*(\Delta + p\Delta - c) - c] \tag{8} \\ = & (1 - \gamma)\left[G\left(\frac{\lambda(1 - \pi_0)}{\lambda(1 - \pi_0) + (1 - \lambda)[\pi_0 F_N(b_0^*) + (1 - \pi_0)F_S(b_0^*)]}\right) - pb_0^*G\left(\frac{\lambda}{1 - (1 - \lambda)F_N(b_0^*)}\right)\right] \end{aligned}$$

The government cares both about the national income as well as its own prospects for re-election. The left hand side of the above equation is the expected gain in output from enacting policy initiative a_N in the first period, and is increasing in b_0^* . On the right hand side, the first term is the government's probability of remaining in power if it chooses to maintain the status quo a_S ; as noted earlier, this is decreasing in b_0^* . The second term on the right is the expected probability of re-election from enacting a_N in the first period, and is increasing in b_0^* . Thus the difference between these two probabilities is decreasing in b_0^* . Therefore, under suitable conditions on the end points, there exists a unique value of $b_0^* \in (0, 1)$ that satisfies the above equation.

Apart from end point conditions, the other condition that has to be satisfied for the above structure to be an equilibrium is that b_1^* should exceed b_0^* . Note that the two equations (5) and (8) for determining b_1^* and b_0^* are independent of each other. Thus, if the parameters are such that the solution to (5) i.e. b_1^* , is less than the solution to (8) i.e. b_0^* , it means that the payoff to the marginal type (one with belief b_0^*) from persisting with a_N is positive, and therefore *all* incumbents who start with policy a_N will persist with it, even in the face of failure. This positive payoff $V(b_0^*)$ (following failure of a_N) will then have to be incorporated in equation (8) for determining b_0^* . The

equilibrium will now thus involve an out-of-equilibrium action, namely policy switching in the face of failure, and an associated out-of-equilibrium belief. However, since the payoffs are continuous in b , $V(b)$ goes to zero as b approaches the solution to (5) i.e. b_1^* . Therefore, the general direction of comparative static results on the initial period action will remain unchanged in this case; we will of course lose the comparative static results on the second period action.

Assumption 3 below gives a sufficient condition under which b_0^* is less than b_1^* . Note that b_1^{eff} exceeds b_0^{eff} ; thus if the relative weight on output γ is close to unity, then b_0^* will be below b_1^* . Secondly, if the probability of success p is high enough, then a first period failure will cause the updated beliefs on state N to be pessimistic enough so that only for very high initial beliefs b_1^* (and above) that the government will choose to persist with a_N . This will again guarantee that b_1^* exceeds b_0^* . Also, comparing gains from the two periods, success in the first period has the additional informational value of revealing that the state is N and the associated option value $p\Delta - c$ of implementing a_N in the second period too. When this value is high, it will push b_0^* to be low, and lead it to be lower than b_1^* . All of these effects are captured in the assumption below.

Assumption 3: $(2 - p)\frac{\Delta}{p\Delta - c} < \frac{\gamma c}{1 - \gamma}$

The following proposition summarizes our preceding discussion and its proof completes the argument for demonstrating the existence of equilibrium

Proposition 1 *Under assumptions 1 – 3, there exist unique values b_0^* and $b_1^* \in (0, 1)$ with $b_0^* < b_1^* < b_1^{eff}$, such that (i) in the first period, only those low ability governments with beliefs $b(x) \geq b_0^*$, and high ability governments who are informed that the state is N , enact the new policy initiative a_N ; all others maintain a_S . (ii) If an output gain of Δ is not realized in the first period, then in the second period only those low ability governments with beliefs $b(x) \geq b_1^*$, and high ability governments who are informed that the state is N , persist with a_N ; all others switch to a_S .*

Proof. See Appendix I. ■

This proposition is of interest for two reasons. First, it shows that there exists an equilibrium in which governments inefficiently persist with a policy despite public evidence of its likely failure.

The second aspect of interest is that electoral imperatives may also distort the first period policy choice, since there is no reason why b_0^* should equal b_0^{eff} .

An appealing aspect of our analysis is its simplicity. It is therefore germane to emphasize three aspects of our framework that make it easy to demonstrate the existence of equilibrium. First is the fact that we have restricted our analysis to examining the dynamics within a single electoral cycle. By doing so, we have assumed that at the start of the game, the origin of reputation is unmodeled and is identical across high or low ability governments. Second is our assumption that a high ability government is perfectly informed about the underlying state. However, this assumption while important, can be relaxed somewhat (see Appendix II). To see the third distinctive feature of our model consider equation (5), which determines the second period cut-off belief b_1^* ; it is independent of the first-period cut-off belief b_0^* . In fact, this is true even if we extend the model to more periods; the cut-off rule for each period would be represented by an equation analogous to (5). This feature of the model, which is convenient, is due to two reasons. Firstly, for those continuing with a_N , everyone deduces that their belief must be above b_1^* , and therefore their reputation is only a function of b_1^* . Secondly, for those switching back to policy a_S , their reputation is independent of b_0^* . Remarkably, this feature is true even in a more general version of the model where there are a continuum of types (see Appendix II).⁹

3.3 Comparative Statics:

We now examine the government's inefficiency in making its first period policy choice. A government with belief b_0^* , in making its decision on whether or not to experiment with policy a_N weighs a possible gain in output gain against a reputational gamble. This is given by equation (8), the two sides of which are depicted in figure 3. The left-hand side is the potential gain in output, while the right-hand side is the difference in expected reputation from the two courses

⁹In this version of the model that we describe in Appendix II, switching does not cause the reputation to drop to 0, but on switching the electorate deduces that the government's initial signal was not perfect. Since the relative probability of getting an imperfect signal between b_0^* and b_1^* is the same for everybody, b_0^* and b_1^* do not affect the reputation from switching.




Figure 3: Equilibrium decision for a L type in the first period, obtained from equation (8).

of action. Notice that for different values of the parameters, the right hand side of equation (8) may intersect the left hand side either to the left or to the right of b_0^{eff} . Therefore, compared with efficient decision-making, a low ability government may be *more* or *less* likely to experiment with a new policy initiative in the first period. If $b_0^* > b_0^{eff}$, then there are situations when the expected output from implementing a_N is positive, yet the government chooses not to do so (i.e. for $b_0 \in (b_0^{eff}, b_0^*)$). Similarly, when $b_0^* < b_0^{eff}$, the government appears too liberal in the sense that there are situations when the expected output from implementing a_N is negative, yet the government does so (i.e. the interval $b_0 \in (b_0^*, b_0^{eff})$). Thus, depending on the parameters, the government's first period policy choice may be either too conservative (in the sense of choosing to retain the status quo a_S when adopting a new policy would be optimal) or too radical (in the sense of taking a new policy initiative when it is not needed).

The following proposition obtains the condition under which either case occurs, and discusses the effects of some of the parameters on this condition. It also summarizes the comparative static results on the second period decision threshold b_1^* .

Proposition 2 *Assume that the conditions of proposition 1 hold. Then*

[I] $b_0^* \geq b_0^{eff}$ according as

$$G\left(\frac{\lambda(1-\pi_0)}{\lambda(1-\pi_0) + (1-\lambda)\left[\pi_0 F_N\left(\frac{c}{p[\Delta+p\Delta-c]}\right) + (1-\pi_0)F_S\left(\frac{c}{p[\Delta+p\Delta-c]}\right)\right]}\right) \geq \frac{c}{\Delta+p\Delta-c} G\left(\frac{\lambda}{1-(1-\lambda)F_N\left(\frac{c}{p[\Delta+p\Delta-c]}\right)}\right) \quad (9)$$

Thus $b_0^* > b_0^{eff}$ if either

(i) c is sufficiently small, or Δ or p is sufficiently large, or

(ii) $F_N\left(\frac{c}{p[\Delta+p\Delta-c]}\right) < (1-\pi_0)(1-F_S\left(\frac{c}{p[\Delta+p\Delta-c]}\right))$.

[II] The second period cut-off belief for persisting with a_N in the face of failure, b_1^* , is decreasing in γ , λ and Δ , and increasing in c .

Proof. See Appendix I. ■

This proposition demonstrates in a particularly sharp way, the dilemma facing a new government in its policy choice. Some new governments that are uncertain, may wish to experiment with new policy initiatives but do not wish to be found out as ‘experimenting’. Others may believe that the new initiative is unlikely to work, yet may adopt it as a gamble to bolster their reputation and the resulting electoral prospects. Thus, their actions might be either too conservative or too radical, the key being how each is viewed by the voting public.

Under what conditions on the parameters would we expect the incumbent to be inefficiently radical as against inefficiently conservative? First, consider the cost and gains from implementing a_N . As the above proposition shows, when c is small or Δ and p are large, then $b_0^* > b_0^{eff}$. In other words, especially when the costs of adoption are small or potential payoffs large, that it is *less* likely that a randomly drawn government will in fact enact policy a_N . While this might seem somewhat surprising, it accords quite well with our emphasis on the reputational underpinning of inefficient policy gambles. To see this observe that a low value of $\frac{c}{p[\Delta+p\Delta-c]}$ makes it more likely that *all* governments (irrespective of ability or information) have a greater incentive to enact a_N . Thus any electoral gain from successfully implementing the policy initiative is small, and this makes the government more reluctant in adopting it in the first place.

Condition (ii) in the above proposition is one under which the reputation from maintaining the status quo exceeds that from implementing a_N even with success, and therefore a government with re-election concerns will be reluctant to enact the policy initiative a_N . In this case, the probability of a L type government getting a signal above b_0^{eff} i.e. $1 - F_N(\frac{c}{p[\Delta+p\Delta-c]})$, is high in state N , and is low in state S ; therefore, the reputation from achieving success with a_N is not too high, and this reduces the incumbent's desire to gamble by experimenting with the new policy initiative.

In the event of a failure with a_N in the first period, the government faces a choice of whether or not to persist with it in the second period. As the above proposition shows, an increase in Δ or a decrease in the cost c causes b_1^* to fall i.e. the government persists more. Second, observe that as the initial perception about the government's ability λ increases, so does the electoral payoff (both $R_S(\cdot)$ and $R_F(\cdot)$) from continuing; hence, as λ rises, b_1^* decreases. In other words, when λ is high, a L type incumbent will be more reluctant to switch to a_S , since this is more likely to jeopardize his electoral chances. Conversely, observe that when λ is very low (i.e. $\lambda \rightarrow 0$), then b_0^* approaches b_0^{eff} . In particular, if the initial reputation is very low, then the incumbent's policy choice is highly unlikely to affect the chances for re-election. If so, then the government may as well adopt the efficient policy i.e. one that maximizes welfare.

3.3.1 *Opportunity cost of experimentation and the degree of inefficiency*

Taking the new policy initiative entails costs. As mentioned earlier, these costs can be interpreted in various ways – they may be costs of administering a new policy and/or making compensatory transfers to potential ‘losers’, or the output from the status quo policy. This broader connotation of costs is useful in analyzing the welfare impact of changes over time in the opportunity cost of experimenting with the new policy.

In particular, suppose that the first period cost of enacting the new policy is c_0 and the second period cost is c_1 , with c_0 not necessarily being equal to c_1 . The preceding analysis remains unchanged except for incorporating c_1 into equation (5) for determining b_1^* , and incorporating c_0 and $p\Delta - c_1$ into equation (8) for determining b_0^* .

Any decrease in the opportunity cost of first-period experimentation c_0 makes experimentation

more attractive i.e. b_0^* decreases. At the same time, it also lowers the efficient cut-off for experimentation i.e. $b_0^{eff} = \frac{c_0}{p[\Delta+p\Delta-c_1]}$. Thus, it maybe of interest to study the overall effect of a change in the cost of experimentation on the relative degree of inefficiency in policy adoption. To do this, let us rewrite equation (8) as follows: $b_0^* - b_0^{eff} =$

$$\frac{1}{p(\Delta + p\Delta - c_1)} \frac{1 - \gamma}{\gamma} \left[G\left(\frac{\lambda(1 - \pi_0)}{\lambda(1 - \pi_0) + (1 - \lambda)[\pi_0 F_N(b_0^*) + (1 - \pi_0) F_S(b_0^*)]}\right) - p b_0^* G\left(\frac{\lambda}{1 - (1 - \lambda) F_N(b_0^*)}\right) \right]$$

Note that as c_0 becomes smaller, so does b_0^* . The effect of this on the two reputations is that while it raises the reputation from maintaining the status quo policy, $R_{\text{status quo}}$, it lowers that from achieving success with policy a_N , R_s . Thus, from an electoral point of view, it makes maintaining the status quo more attractive and therefore has the effect of increasing $b_0^* - b_0^{eff}$. Thus, if the incumbent was too conservative in enacting policy a_N in the first place, then any lowering of the cost of experimentation exacerbates the degree of inefficiency. Therefore, somewhat surprisingly, a lowering of the cost of experimentation worsens matters, precisely when the Δ is large or p is high (i.e. when $b_0^* > b_0^{eff}$). On the other hand, if the incumbent was too liberal i.e. $b_0^* < b_0^{eff}$, a smaller c_0 leads to reducing the extent of inefficiency, by pushing b_0^* towards b_0^{eff} .

A similar analysis can be done to study the effects of a change in the second period cost c_1 on the degree of inefficient second-period persistence. As c_1 falls, the second period cut-off belief b_1^* falls i.e. more now persist. But this also reduces the following reputations, whether on success or failure. Since inefficient persistence is due to a desire to acquire these reputations, a fall in them means that the effect on output due to a decrease in c_1 is to reduce the degree of inefficiency.

3.3.2 Reputation and Ideology:

In our benchmark model we have suppressed ideological predilections of the government in power. While clearly important, the role of ideological factors in explaining a government's policy decisions should not be exaggerated – many policies engender relatively low dimensional conflict amongst voters (e.g. war, foreign policy, inflation). We now sketch an example to suggest that introduction of ideological considerations into the model may yield interesting insights.

So far, we assumed that the initial priors on the state being N i.e. π_0 , are common between

the incumbent and the electorate. Suppose now we follow Harrington (1993) in assuming that the incumbent’s prior beliefs π^I are different from the citizen-voter’s beliefs π_0 .¹⁰ Furthermore, this difference in priors is assumed to be known to the electorate. For example, a ‘conservative’ government may be commonly identified as being a strong believer in the efficacy of certain types of policies (e.g. the status quo).

The structure of the preceding analysis is unchanged except that $F_i(b) = \Pr[b(x) \leq b|i] = \Pr[x \leq g^{-1}(\frac{1-\pi^I}{\pi^I} \frac{b}{1-b})|i]$ is affected by changes in π^I : as π^I increases, for any received signal x , the government’s belief on the state being N becomes stronger in the sense of first-order stochastic dominance i.e. $F_i(b)$ decreases. Given this, any time a government with a higher π^I maintains the status quo, its reputation is higher; at the same time, its reputation from achieving success with the policy a_N is lower. Together, they imply that a government that is known to be ex-ante more optimistic about the policy a_N (i.e. one with a higher π^I), will in fact be more conservative in adopting it, and the bias in the first period is more likely to be in favor of the status quo. On the other hand as π^I increases, inefficient policy persistence in the second-period is likely to be lessened, since now the reputations from continuing i.e. both $R_S(b_1^*)$ and $R_F(b_1^*)$, are lower.

Politicians’ are not typically rewarded for being ‘flexible’ and our benchmark model per se suggested that this is because it signals incompetence. However, there are occasions where it seems that a politician who shows ‘flexibility’ is electorally rewarded; our above example suggests that ideology maybe part of the answer. Here, capable governments are those that take the correct action in accordance with the appropriate state. Incompetent governments try to imitate them and get their highest electoral reward by appearing to contradict their own prior, or in other words, by showing that given good enough evidence they would be willing to overturn their prior.

Finally, observe that a change in the electorate’s initial perception about the state being N has the opposite effect on the government’s decisions. An increase in π_0 lowers the incumbent’s

¹⁰Piketty (1995) has indicated how such heterogeneity in beliefs can arise out of a learning process. Benabou and Tirole (2002) show that ideology can arise as part of a learning process with time-inconsistent preferences. Members of differing political parties can then have very different beliefs about the probability distribution describing different states of the world, than ordinary citizens.

reputation from maintaining the status quo, while that from success is unchanged (as a success reveals perfectly that the state is N). Thus, a higher π_0 means that the government will experiment more in the first period. In case of a failure in the first period, the government's second period decision weighs off the output losses against the gains to reputation. Again, the reputation from success, $R_S(b_1^*)$, is unaffected by changes in π_0 , but now with a higher π_0 , a second-time failure is looked upon more favorably by the electorate. Thus, as π_0 increases, the reputation from twice failure with a_N , $R_F(b_1^*)$, is higher. This leads to more inefficient persistence as π_0 increases.

4 Discussion

In this section we explore some aspects of the framework that we introduced, including a discussion of the literature on reputational herding, other equilibria and the window of opportunity.

A. Relation to the Literature – Reputational Models: Our framework perhaps shares most with the literature developed in the context of reputational decision making. For instance, Scharfstein and Stein (1989), under the assumption that managers do not know their own type, show that they engage in herd behavior; once a manager takes a particular action, it changes the public prior on the action, and this induces other managers to ignore their private information and follow this particular course of action.¹¹ Observe that this feature is also shared by our framework, where an increase in the electorate's prior π_0 , is more likely to result in a_N being enacted in the first period (see Section 3.3.2). Once agents know their own type, they may choose a policy that contradicts their own public prior in order to boost their reputation – as in Avery and Chevalier (1999) and Prendergast and Stole (1995). Once again observe (in Section 3.3.2) that higher is the incumbent's prior belief π^I , the more likely is he to take a_S and contradict his own (publicly known) prior.

Our framework perhaps shares most with Prendergast and Stole (1995). They demonstrate that a manager's concern for his current reputation can lead to investment decisions that are radical early on and then display increasing conservatism over time. There are important differences however. First, given our preoccupation with governmental decision making, our framework em-

¹¹Their result is generalized by Ottaviani and Sorensen (2000).

phasizes the discrete nature of the trade-off all leaders face: between sticking to the safe, status quo policy and experimenting with a more uncertain course of action. Second, in our analysis, the public observability of the impact of a government’s policies on the economy is given a central role. In the context of an economy this is perhaps appropriate, since not only might the effect of a policy be directly experienced by the voter, but the effects of policy initiatives are also widely reported in the media and other watchdogs and think tanks, which influence public perception. Thus we allow for the public to make an estimate of the ability of the politician as a function not just of policy choices, but also of national output, both of which are observable. Furthermore, the simplicity of our framework, allows us to dispense with myopic decision making on the part of the incumbent. Politicians care not only about their present actions, but also the implications of present policy choices on future perceptions – resulting in the first period policy choice being either too conservative or radical.

B. Other Equilibria: Given that we are considering a dynamic signaling game, there may be the possibility of equilibria other than the one discussed above. However, observe that the above equilibrium involved all possible action-sequences being taken in equilibrium, and therefore it did not involve specifying any “out-of-equilibrium” beliefs. Thus, any other equilibrium would necessarily involve some action-sequence being not undertaken in equilibrium. Consider first, a potential equilibrium where no one undertakes policy a_N in the initial period, with the associated out-of-equilibrium belief the most pessimistic possible i.e. anyone who enacts a_N is considered to be a low type. For a H type government who has received a perfect signal that the state is N , its payoff from enacting a_N is then $2\gamma[p\Delta - c] + (1 - \gamma)G(0)$, while its payoff from maintaining a_S is given by $(1 - \gamma)G(R_{\text{status quo}})$. Now, if $\frac{\gamma}{1-\gamma}$ exceeds $\frac{1}{2[p\Delta - c]}$ (which is implied by Assumption 3), the payoff from enacting a_N exceeds that from maintaining the status-quo. Therefore, under assumption 3, there cannot be an equilibrium where no government enacts a_N . Similarly, under assumption 3, one can rule out equilibria under which a_S is never maintained.

The other possible equilibria that we need to consider involve neither type playing some action in the second period. This is not possible for action a_N as a H type government with a perfect signal that the state is N would prefer (again, under assumption 3) the output gains from a_N viz.

$\gamma[p\Delta - c]$ over the maximum possible loss in reputation, $(1 - \gamma)G(1)$. Suppose on the other hand no one switches in equilibrium; now the payoff from continuing for the marginal type b_0^* would be given by $V(b_0^*)$ from the left-hand side of equation (5), and in this case would be positive. Further, this would have to be incorporated into her payoff (in the case of failure) in the first-period equation i.e. (8). But $V(b_0^*)$ is continuous in b_0^* , and assumption 3 ensures that $b_1^* > b_0^*$ i.e. that the payoff to the marginal type from continuing in spite of failure, can never be positive. Thus, neither type switching under failure cannot be an equilibrium either.

Hence, under assumption 3, the equilibrium discussed in the previous section is the unique equilibrium of the model.

C. Window of opportunity: In our benchmark model, the window of opportunity for enacting the new policy a_N occurs only in the first period. If the incumbent government chooses not to enact a_N at that time, we assumed that the opportunity disappears. Suppose instead the opportunity for enacting a_N remained in both periods so that even if the government decided to maintain the status quo in the first period, it would still have the option to enact the new policy a_N in the second period. Nevertheless, it is reassuring to observe that the equilibrium discussed above is still an equilibrium of this modified game. The reasoning is as above: a H type with a perfect signal that the state is N , will prefer to enact a_N in the initial period itself as the output gain $\gamma[p\Delta - c]$ outweighs the maximum possible loss of reputation $(1 - \gamma)G(1)$. Given this, any government who maintains the status quo in the first period and enacts a_N in the second period, will be identified as an L type. Therefore, anyone who is ever going to enact a_N is better off doing so in the first period itself. Thus, in this modified game, any policy experimentation will happen early on in the electoral cycle.

Suppose instead of at the beginning of the electoral cycle, the window of opportunity facing a government occurred in the middle of the cycle. What impact would this have on the government's incentives to take the new policy? Since this now results in a shorter time horizon within the present electoral cycle, a government does not have to worry about its second period decision of whether or not to persist. In this case, since there is less scope for dynamic sorting, it increases the incentive for an L type to experiment i.e. b_0^* is lower. Thus, conditional on reputation, an

incumbent is likely to experiment more in the middle of the term than if the opportunity occurred at the beginning. Of course, in real life this effect is likely to be compounded by the incumbent's reputation in the middle of the term. If this is high, then it is less likely to take a gamble, while if low, then it may be more inclined to experiment and thereby gamble on its success.

In the above example, the extent of the inefficiency arises from the dynamic nature of the sorting that occurs. If the opportunity to enact a policy occurs relatively late, then inefficient policy initiatives are likely, since low ability governments are less likely to be recognized as such by the voter. Perhaps the sharpest way to see this is to consider a simple variant – where there is *uncertainty* about the gains that accrue in the second period. For instance, suppose c_1 is either 0 or Δ and is not known ex-ante, but becomes clear at the beginning of $T = 1$, *before* the government takes its second-period decision. In this simple case, given the extreme nature of the costs, the second-period decision is simple: if $c_1 = 0$, then all who enacted a_N continue with it, failure or not, while if $c_1 = \Delta$, none continue. In other words, the action taken by the government in the second period reveals no more information about its type. The effect of this for the marginal L type in the first period is to increase his gains from experimenting, since his second-period decision is not informative. This leads to an important insight – anything that results in the optimal follow-up decisions to be uncorrelated with the type of the incumbent, causes the follow-up action to lose its signaling value. As a result while there is less inefficiency in the following periods, it exacerbates the incumbent's desire to experiment in the first period, as it lessens the chances of being caught out in the second period.

5 Conclusion

The appropriateness of many policies for a particular economic or institutional structure is plagued by uncertainty. Given this uncertainty, governmental learning and policy experimentation take on an important role. In this paper we took a first step in analyzing the impact of electoral imperatives on a government's decision to learn through policy experimentation. Once account is taken of a government's electoral concerns, two kinds of inefficiencies can be identified. First,

early on in their tenure, governments have a tendency to gamble by either recklessly experimenting with new policies, or being inefficiently conservative by sticking to the status quo. However, once a government has been in power for some time, it acquires a reputational stake in policies that it previously enacted, and becomes inefficiently conservative.

A number of issues raised in this paper warrant further exploration. Firstly, while our analysis of ideological considerations yielded some interesting insights, it is rather preliminary. We believe that a systematic analysis of the role of voter heterogeneity (in both beliefs and preferences) and ideology in influencing policy experimentation and persistence, would be useful. A second point worth exploring is the impact of electoral imperatives on the government's incentive to experiment with and choose between policies that require a varying length of time to show results (i.e. short versus medium or long term). Of course, tackling such issues systematically would require the development of a more elaborate dynamic structure than we have here. Finally, for some policies, it is possible that the underlying state of the world changes, even if only gradually. If so, the importance of policy experimentation becomes even more acute. Indeed, much can be learned from enriching our framework to analyze policy experimentation with a changing state of the world. We leave this and much else for future work.

References

- Alesina, Alberto and Drazen, Allan. "Why are Stabilizations Delayed?" *American Economic Review*, December 1991, 81(5), 1170-88.
- Aslund, Andres. "Problems with Economic Transformation in Ukraine" in *Transition: The First Decade* edited by M. Blejer and M. Skreb, MIT Press, 2002.
- Benabou, Roland and Jean Tirole. "Belief in a Just World and Redistributive Politics." *mimeo*, Princeton 2002.
- Berman, Larry. "Lyndon B. Johnson: Paths Chosen and Opportunities Lost" in Fred I. Greenstein, ed., *Leadership in the Modern Presidency*. Cambridge, Mass.: Harvard University Press, 1988.

Besley, Timothy. "Political Institutions and Policy Competition." in *The Institutional Foundations of a Market Economy* eds. Kochendorfer-Lucius & Pleskovic, World Bank, 2001.

Blendon, Robert, Mollyann Brodie and John Benson "What Happened to America's Support for the Clinton Health Plan?" *Health Affairs*, 1995, 14(2), 7-23.

Canes-Wrone, B., Michael Herron and K. Shotts. "Leadership and Pandering: A Theory of Executive Policymaking." *American Journal of Political Science* 2001, 45, 532-550

Coate, Stephen and Morris, Stephen. "Policy Persistence." *American Economic Review*, December 1999, 89(5), 1327-36.

Cohen, Benjamin. "Political Dimensions of Dollarization." Federal Reserve Bank of Dallas, 2000.

Cukierman, Alex and Mariano Tommasi. "When Does it Take a Nixon to Go to China?" *American Economic Review*, March 1998, 88(1), 180-97.

Dornbusch, Rudiger and Sebastian Edwards. *The Macroeconomics of Populism in Latin America*. Chicago: University of Chicago Press, 1990.

Fernandez, Raquel and Rodrik, Dani. "Resistance to Reform: Status Quo Bias in the Presence of Individual Specific Uncertainty." *American Economic Review*, 1991.

Freedland, C. *Sale of the Century: Russia's Wild Ride* Random House, 2000

Hall, Peter. "Policy Paradigms, Social Learning and the State: the Case of Economic Policymaking in Britain", *Comparative Politics*, vol 25-3, pp 275-96; 1993

Harrington, Joseph. "Economic Policy, Economic Performance and Elections." *American Economic Review*, March 1993, 83(1), 27-42.

Heclo, Hugh. *Modern Social Politics in Britain and Sweden*. New Haven, Yale, 1974.

Maskin, E. and J. Tirole. "The Politician and the Judge: Accountability in Government", mimeo.

Mukand, Sharun and Dani Rodrik. "In Search of the Holy Grail: Policy Convergence and Economic Performance", mimeo. Harvard, July 2003.

Ottaviani, Marco and Peter Sorensen. "Herd Behavior and Investment: Comment", *American Economic Review*, June 2000, 90(3), 695-704.

Piketty, Thomas. "Social Mobility and Redistributive Politics." *Quarterly Journal of Economics*, August 1995, 110(3), 551-84.

Prendergast, Canice and Stole, Lars. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy*, December 1996, 104(6), 1105-34.

Rodrik, Dani. "Understanding Economic Policy Reform." *Journal of Economic Literature*, March 1996,34(1), pp. 9-41.

Rogoff, K. "Equilibrium Political Budget Cycles", *American Economic Review*, 80:21-36, 1990

Schandler, Herbert Y. *The Unmaking of a President: Lyndon Johnson and Vietnam*. Princeton, N.J.: Princeton University Press, 1977.

Scharfstein, David and Jeremy Stein. "Herd Behavior and Investment", *American Economic Review*, June 1990

Skocpol, Theda. "Bring the State Back In: Structure of Analysis in Current Research" in *Bring the State Back In* edited by Evans, Rueschmeyer and Skocpol, Cambridge, 1985.

Stokes, Susan. "Policy Switches" in *Democracy, Accountability and Representation* ed. Manin, Przeworski & Stokes, Cambridge. 1999

Tommasi, Mariano. "Crisis, Political Institutions and Policy Reform", mimeo. San Andres, 2002.

Tommasi, Mariano and Andres Velasco. "Where Are We in the Political Economy of Reform?", mimeo, *New York University*, 1996.

Tuchman, Barbara W. *The March of Folly: From Troy to Vietnam*. New York: Knopf, 1984.

Weyland, Kurt. *The Politics of Market Reform in Fragile Democracies*. Princeton, 2002.

Wijnholds, J. "The Argentine Drama: A View from the IMF", *mimeo* 2003.

Appendix I: Proofs

Proof of Proposition 1: From our discussion of second-period (first-period respectively) decision making by low ability governments in section 3.1 (section 3.2, resp.), the expected payoff from policy a_N is increasing in b_1 (b_0 resp.) while that from a_S is constant. Therefore, all incumbents with beliefs above b_1^* (b_0^* resp.) persist (choose a_N in the first period, resp.), while those with beliefs less than b_1^* revert back to a_S (choose a_S in the first period, resp.). Note that a H type government who receives perfect information that the state is N (S respectively) has belief $b_0 = 1$ ($b_0 = 0$ resp.) and therefore it is optimal for such a government to enact a_N and persist with it (maintain a_S resp.). Thus, to prove the existence of this equilibrium, we now need to show that there exists solutions to equations (5) and (8), and that $b_0^* < b_1^* < b_1^{eff}$.

To check that solutions to equations (5) and (8) exist, note that the LHS and RHS of both are continuous in b . As $R_S(b)$ and $R_F(b)$ are both increasing in b , and $R_S(b) \geq R_F(b)$, the LHS of (5) is increasing in b . At $b = b_1^{eff}$, $\frac{(1-p)b_1^{eff}}{1-pb_1^{eff}}p\Delta = c$, and so $\text{LHS}(5) > 0$; while at $b = 0$, $\text{LHS}(5) = -\gamma c + (1-\gamma)G(\frac{\lambda\pi_2}{\lambda\pi_2+1-\lambda})$ which is negative by assumption 3. Thus by continuity, a solution to (5) exists, is unique and lies in the interval $[0, b_1^{eff})$.

Similarly, the LHS of (8) is increasing in b , while the RHS is decreasing in b . At $b = 0$, $\text{LHS}(8) = -\gamma c < 1 - \gamma = \text{RHS}(8)$ at $b = 0$; while at $b = 1$, $\text{LHS}(8) = \gamma[p(\Delta + p\Delta - c) - c] > (1-\gamma)[G(\frac{\lambda(1-\pi_0)}{\lambda(1-\pi_0)+1-\lambda}) - 1] = \text{RHS}(8)$. Again by continuity, a solution b_0^* to equation (8) exists and is unique.

Next we need to show that b_0^* is less than b_1^* .

Let us rewrite equations (5) and (8) for determining b_1^* and b_0^* in terms of a general belief b as:

$$\gamma[pb\Delta - c] + pb(1-\gamma)G[\frac{\lambda}{1-(1-\lambda)F_N(b)}] + (\frac{1-pb}{1-p} - pb)G[R_F(b)] - \gamma c \frac{p}{1-p}(1-b) = 0 \quad (5')$$

$$\begin{aligned} &\gamma[pb\Delta - c] + (1-\gamma)pbG[\frac{\lambda}{1-(1-\lambda)F_N(b)}] \\ &+ \gamma pb(p\Delta - c) - (1-\gamma)G[\frac{\lambda(1-\pi_0)}{\lambda(1-\pi_0) + (1-\lambda)[\pi_0 F_N(b) + (1-\pi_0)F_S(b)}] = 0 \quad (8') \end{aligned}$$

The left-hand side of both equations are increasing in b and we know that $b_1^* < b_1^{eff}$. So if we can

show that the LHS of (5') is less than that of (8') for all $b \in [0, b_1^{eff}]$, then any solution to (8') i.e. b_0^* , will be smaller than that for (5') viz. b_1^* .

Thus, we require to show that $(\frac{1-pb}{1-p} - pb)G[R_F(b)] - \gamma c \frac{p}{1-p}(1-b)$ is less than $\gamma pb(p\Delta - c) - (1-\gamma)G[\frac{\lambda(1-\pi_0)}{\lambda(1-\pi_0)+(1-\lambda)[\pi_0 F_N(b)+(1-\pi_0)F_S(b)]}]$ for $b \in [0, b_1^{eff}]$.

$$\begin{aligned} \text{i.e. that } & \left(\frac{1-pb}{1-b} + \frac{p}{1-p} \right) G(R_F(b)) + \frac{1}{1-b} G\left(\frac{\lambda(1-\pi_0)}{\lambda(1-\pi_0) + (1-\lambda)[\pi_0 F_N(b) + (1-\pi_0)F_S(b)]} \right) \\ & < \frac{\gamma}{1-\gamma} \left(\frac{pc}{1-p} + \frac{pb}{1-p}(p\Delta - c) \right) \quad \text{for } b \in [0, b_1^{eff}] \end{aligned}$$

Note that the first part on the left-hand side of the above inequality is increasing in b , while the second part is the product of two terms, one increasing in b and the other decreasing in b . Hence the first part achieves its maximum at $b = b_1^{eff}$, and is bounded above by: $\frac{1-pb_1^{eff}}{1-b_1^{eff}} + \frac{p}{1-p}$, while the second term is bounded above by $\frac{1}{1-b_1^{eff}}$.

Recall that $b_1^{eff} = \frac{c}{p[c+(1-p)\Delta]}$. Thus, the LHS of the above inequality is bounded above by:

$$\frac{1-pb_1^{eff}}{1-b_1^{eff}} + \frac{p}{1-p} + \frac{1}{1-b_1^{eff}} = \frac{p}{1-p}(2-p)\frac{\Delta}{p\Delta - c} \quad (10)$$

On the other hand, the RHS of the inequality is bounded below by $\frac{\gamma}{1-\gamma}\frac{pc}{1-p}$. That the expression (10) is less than $\frac{\gamma}{1-\gamma}\frac{pc}{1-p}$ is stated as Assumption 3. Thus, under assumption 3, for any $b \in (0, b_1^{eff}]$, the LHS of (5') is smaller than that of (8'). Since the LHS of both equations is increasing in b , any solution b_0^* to (8') will be less than that for (5') viz. b_1^* .

Proof of Proposition 2: Observe that at $b_0 = b_0^{eff}$, the LHS of equation (8) is 0. Therefore, whether $b_0^* \geq b_0^{eff}$ depends on whether the RHS of (8) at b_0^{eff} is greater or less than 0. Evaluating the RHS of equation (8) at $b_0^{eff} = \frac{c}{p[\Delta+p\Delta-c]}$ yields condition (9).

(i) The impact of parameters c , Δ and p follows from the fact that the LHS of (9) is decreasing in $\frac{c}{p[\Delta+p\Delta-c]}$ while the RHS is increasing in $\frac{c}{p[\Delta+p\Delta-c]}$.

(ii) $F_N(\frac{c}{p[\Delta+p\Delta-c]}) < (1-\pi_0)(1-F_S(\frac{c}{p[\Delta+p\Delta-c]}))$ implies that $R_{\text{status quo}}$ exceeds R_s , both evaluated at b_0^{eff} . Since $\frac{c}{p[\Delta+p\Delta-c]} < 1$, this implies that the reputation from maintaining the status quo exceeds that from achieving success with a_N . Hence, under this condition, the government is too conservative.

The comparative static results on b_1^* follow from the discussion at the end of section 3.1.

Appendix II: Model with a continuum of types

Consider a version of the model where there are a continuum of possible types of the incumbent. The incumbent now gets either a perfect signal about the state of the world or an imperfect signal $x \in [\underline{x}, \bar{x}]$. Types differ in the probability with which they receive a perfect signal. Denoting this probability by t , we now assume that possible government types t is distributed over the interval $[0, 1]$.

Thus in the state i , $i \in \{S, N\}$, a government of type t may receive a perfect signal about the state (with probability t), or with probability $1 - t$, will receive a signal $x \in [\underline{x}, \bar{x}]$ according to the density function $\phi_i(x)$. Assumption 2 on the density functions $\phi_S(\cdot)$ and $\phi_N(\cdot)$ is still maintained, so that a higher value of the signal x corresponds to a higher belief that the state is N .

Governments wish to maximize a weighted sum of output and expected reputation:

$$\gamma(\text{National welfare}) + (1 - \gamma)(\text{Mean reputation})$$

We will again seek to characterize an equilibrium of the type described earlier i.e. a government that receives perfect signal that the state is S will maintain the status quo policy, while if it receives a perfect signal that the state is N , it will enact policy a_N and continue with it even on failure. On the other hand, if the government receives an imperfect signal x , it will update its belief about the state of the world according to (1); it will experiment with a_N only if its belief exceeds b_0^* . If a_N achieves success, then it will continue with a_N in the second period; however, if there is a failure, it will continue with a_N only if its initial belief exceeded b_1^* .

Assume that the electorate's initial belief about the incumbent's type is characterized by a distribution with density $f(t)$, $t \in [0, 1]$. Denote the mean of this distribution by λ .

REPUTATIONS

Consider the event where the government enacts policy a_N in the first period and it achieves success. In this case, everyone deduces that the state must be N , and the voters' updated distribution about the government's type is given by the density function:

$$f(t | a_N \text{ at } T = 0, \text{ output} = \Delta \text{ in 1st pd.}) = \frac{f(t)[t + (1 - t)(1 - F_N(b_0^*))]}{E[s + (1 - s)(1 - F_N(b_0^*))]}$$

Hence, the expected reputation in such an event is given by:

$$R_S(b_0^*) = ER_{\text{success in 1st period}} = \frac{Et^2 + (Et - Et^2)(1 - F_N(b_0^*))}{Et + (1 - Et)(1 - F_N(b_0^*))} = \frac{\lambda - (\lambda - Et^2)F_N(b_0^*)}{1 - (1 - \lambda)F_N(b_0^*)}$$

Simple differentiation shows that this reputation is increasing in b_0^* .

Similarly, the expected reputation for the incumbent in the event that he sticks with the policy a_N in spite of a midterm failure and it achieves success in period 2 is given by:

$$R_S(b_1^*) = ER_{\text{failure in 1st pd., success in 2nd pd.}} = \frac{\lambda - (\lambda - Et^2)F_N(b_1^*)}{1 - (1 - \lambda)F_N(b_1^*)}$$

In the case that the government sticks with a_N in spite of a midterm failure, and it turns out to be a failure in the second period too, then as before, the electorate updates its belief about the state being N to π_2 . It now believes that either the government did get a perfect signal about the state being N , or if it got an imperfect signal, then this must have been high enough to result in the belief exceeding b_1^* . A similar calculation as above gives the expected reputation in this case:

$$R_F(b_1^*) = ER_{\text{failure in 1st pd., failure in 2nd pd.}} = \frac{Et^2 + (\lambda - Et^2)[\pi_2(1 - F_N(b_1^*)) + (1 - \pi_2)(1 - F_S(b_1^*))]}{\lambda + (1 - \lambda)[\pi_2(1 - F_N(b_1^*)) + (1 - \pi_2)(1 - F_S(b_1^*))]}$$

Again, note that this reputation is increasing in b_1^* .

On the other hand, if a_N turns out to be a failure in the first period, and the government decides to revert back to a_s in period 2, the electorate understands that the government had not received a perfect signal at the beginning. Furthermore, this imperfect signal must have caused the government's belief to lie between b_0^* and b_1^* so that it was optimistic enough to undertake the policy experiment, but not sure enough to continue with it in the face of midterm failure. In this case, the voters' updated distribution about the government's type is given by the density function: $f(t \mid a_N \text{ at } T = 0, \text{ output} = 0 \text{ in 1st pd., } a_S \text{ at } T = 1) =$

$$\frac{f(t)(1 - t)[\pi_2(F_N(b_1^*) - F_N(b_0^*)) + (1 - \pi_2)(F_S(b_1^*) - F_S(b_0^*))]}{E(1 - t)[\pi_2(F_N(b_1^*) - F_N(b_0^*)) + (1 - \pi_2)(F_S(b_1^*) - F_S(b_0^*))]} = \frac{f(t)(1 - t)}{E(1 - t)}$$

and therefore the expected reputation in such an event is given by: $ER_{\text{failure in 1st period, switch}} = \frac{\lambda - Et^2}{1 - \lambda}$. Unlike in the model with one type being perfect, here the reputation from switching does not fall to 0, but importantly, it does not depend on b_0^* or b_1^* and is less than λ . It is even less

than $R_F(b_1^*)$, the reputation from failing twice with policy a_N . This is what leads incumbents to inefficiently persist with a_N in the face of midterm failure.

Finally, we need to calculate the expected reputation in the case that the incumbent decides to maintain the status quo policy in the first period. A similar calculation as above shows that this is given by:

$$ER_{\text{maintain status quo}} = \frac{(1 - \pi)Et^2 + (\lambda - Et^2)[\pi F_N(b_0^*) + (1 - \pi)F_S(b_0^*)]}{(1 - \pi)\lambda + (1 - \lambda)[\pi F_N(b_0^*) + (1 - \pi)F_S(b_0^*)]}$$

Simple differentiation shows that this reputation is decreasing in b_0^* .

FIRST and SECOND PERIOD DECISIONS

Having established the expected reputations from the various possible outcomes, we can use similar equations as before to characterize b_0^* and b_1^* . Note that a government with belief b_1^* is indifferent between persisting with a_N and switching back to a_S , and balances the output loss and reputation gain from the two options: $\gamma \frac{(1-p)b_1^*}{1-pb_1^*} p\Delta - c +$

$$(1 - \gamma) \left\{ \frac{(1-p)b_1^*}{1-pb_1^*} pG(R_S(b_1^*)) + \left(1 - \frac{(1-p)b_1^*}{1-pb_1^*} p\right) G(R_F(b_1^*)) \right\} = (1 - \gamma) G\left(\frac{\lambda - Et^2}{1 - \lambda}\right)$$

The left-hand side is increasing in b_1^* , while the right hand side is independent of it. Therefore, there will exist a unique solution to the above equation. Further, since $R_S(b_1^*)$ and $R_F(b_1^*)$ both exceed the reputation from switching, it implies that $\frac{(1-p)b_1^*}{1-pb_1^*} p\Delta - c < 0$ i.e. some governments persist with a_N even though they believe that it is inefficient to do so.

In the first period, a marginal incumbent with belief b_0^* is indifferent between maintaining the status quo policy a_S , or experimenting with a_N (with the realization that it will revert back to a_S if a_N does not succeed in period 0) :

$$\gamma[pb_0^*(\Delta + p\Delta - c) - c] + (1 - \gamma)[pb_0^*G(R_S(b_0^*)) + (1 - pb_0^*)G\left(\frac{\lambda - Et^2}{1 - \lambda}\right)] = (1 - \gamma)G(ER_{\text{maintain status quo}})$$

While the right-hand side of this equation is increasing in b_0^* , the right hand side is decreasing in b_0^* . Therefore, under suitable conditions on the end points, there will exist a unique solution to the above equation.