Choosing Leaders: Learning from Past Decisions in a Changing Environment*

Matias Iaryczower and Andrea Mattozzi†

November 23, 2011

Abstract

We study the joint problem of a politician, who sets policy in a changing environment, and a voter, who is learning about the ability of the politician to decide whether to reelect him or not. If the voter can observe policy decisions but not the optimal policies, policy choices are uninformative about ability. If the voter can observe optimal policies with noise and enough decisions by the politician, then (i) the voter appoints good types and fires bad types with probability close to one, and (ii) the executive chooses policy as if he had no career concerns. In this case consistent policy records are indicative of ability when past information depreciates fast.

*We thank Juliana Bambaci, Steve Coate, and seminar participants at Cornell, Maryland, Bologna, Roma, the MPSA & APSA annual meetings, the Priorat Workshop in Theoretical Political Science, and the 11^{th} Conference of the Society for the Advancement of Economic Theory for useful comments to previous version of this paper.

†Matias Iaryczower is Assistant Professor of Politics at Princeton, Princeton, NJ 08544. Andrea Mattozzi is Associate Professor of Economics at MOVE and Universitat Autònoma de Barcelona, Barcelona, Spain 08193. Emails: miaryc@princeton.edu, andrea.mattozzi@movebarcelona.eu.

Constant development is the law of life, and a man who always tries to maintain his dogmas in order to appear consistent drives himself into a false position. - Mohandas K. Gandhi.

1 Introduction

Presidents, Governors, company CEOs and central bank chairmen are appointed to make decisions on behalf of others. While their tenets, principles and ideologies are typically well understood *before* they get to office, their ability as decision-makers is often unknown. The decisions they make while in office then stick with them as a label, defining voters' and shareholders' perceptions about their ability, and influencing their reelection chances and career prospects.

While executives typically start their tenure in office with a clear idea of the initial policies they want to implement, the bulk of their job is to revise their policy strategies in response to changes in the environment. Do we maintain a restrained fiscal approach, or is it time to actively pump the economy? Do we maintain a large military presence in a region of conflict, or scale back the operations and resort to diplomacy?

At the core of these problems is that the optimal policy in any given period evolves over time, but it typically not independent of the optimal policy in previous periods. In this context past information will generally be useful to evaluate current policy alternatives, but might also mislead the policy-maker when change is required. As a result of these two countervailing forces, the information accumulated in past periods can either completely swamp current information – making executives unresponsive to new events – or depreciate with time and end up being too uninformative to influence current decisions.

The problem is made all that much harder when policy-makers have career concerns, as for example in the case of a President who wants to be reelected, or a Governor that is considering running for President after his term in office. In this case, the actions taken in office can affect the career prospects of the decision maker, and the decision maker has to consider how his actions are going to be interpreted by the principal who can decide on his political future.

In this paper we tackle the joint problem of a career minded decision-maker, who is learning about the changing environment in order to choose policy, and a principal, who is learning about the ability of the decision maker in order to decide whether to keep/promote him or dismiss him altogether. Can voters make convincing inferences about the ability of the President based on his decisions on when to be active or passive in military interventions abroad, or on when to stimulate the economy or be fiscally prudent? Can senators evaluate the ability of the Federal Reserve chairman based on his decisions on when to pursue an expansionary monetary policy and when to be passive, or of judges based on their conviction rates?

To capture the tradeoff between experience and adaptability to a changing environment, we consider the following model. A policy-maker (the politician) makes T up or down decisions in an office A. The optimal t-period policy is unobservable, and evolves according to a Markov process. Before making each policy decision d_t, the politician observes an imperfect private signal about the optimal t-period policy, with better politicians having a more precise signal. A principal (the voter) observes the politician's decisions in office A, and possibly also imperfect information about the optimal policy in each period, and then decides whether to appoint the politician for a second office B, or instead appoint an untested politician. The politician cares about making "good" policy decisions in office A (according to his own perceptions of the different types of errors), but also about getting appointed for office B (has career/electoral concerns).

In this environment, all past information is useful for current decisions, but the politician's knowledge does not increase with experience in office as it would if the optimal policy was given, and time invariant. The transition probabilities of the optimal policy process can be asymmetric, so that one alternative is highly persistent, but the other is short lived. Thus for example, we can capture a situation in which a passive monetary policy is optimal most of the time, but short spells of expansionary monetary policy are optimal some of the time. We can also capture of course a situation in which long spells of expansionary and passive policy are optimal.

We consider two alternative informational environments for the voter. In the first case, the voter can observe the choices of the politician, but cannot observe the process of optimal policies. For example, the voter might be able to observe the unemployment rate, but in the absence of a deeper understanding of the economy, be ignorant about whether increasing or decreasing the deficit would reduce or increase unemployment further. We say that this voter is *uninformed*. We show that when the voter is uninformed, and the politician puts enough weight on career concerns, in equilibrium policy decisions are independent of – and uninformative about – the ability of the politician. As a result, elections become ineffective instruments to select good politicians to office. Furthermore, posturing completely dominates policy-making, as both high and low ability politicians choose the same given sequence of policies, disregarding all private information about the optimal policy process.

While the uninformed-voter setting is a reasonable approximation in various problems, in many applications it is natural to assume that the voter can at least imperfectly observe the optimal policy in each period. This could be due for example to the effect of the media. To capture this, we endow the voter with an imperfect signal about the realization of the optimal policy in each period. We say that in this case there is partial *transparency* about the optimal policy, and parametrize the extent of transparency in the policy-making environment by the precision of the voter's signal.

In contrast to the case in which the voter is uninformed, we show that for an arbitrarily small level of transparency, if the voter can observe a sufficiently large number of policy decisions, (i) the voter appoints high ability politicians and dismisses low ability politicians with probability close to one, and (ii) the politician sets policy as if he had no career concerns (implying efficiency when he has the same preferences as the voter). Thus, while a very noisy observation of the optimal policy process would not be enough for the voter to choose the right action in the absence of delegation, this limited information is enough to sort out the agency relation and possibly lead to large gains in efficiency. This suggests that in this setting, even an unsophisticated media can greatly improve the quality of political choices taken in office, and may help voters to select and promote only high ability politicians.[1]

When the voter is uninformed, policy decisions are independent of the ability of the politician. With transparency, instead, if the voter can observe a large enough number of policy choices, the politician will choose policy ignoring career concerns. As a result, politicians of different ability levels generate different policy outcomes. A natural question in this setting is then whether observing frequent policy reversals tells us something about the ability of the decision-maker. Does flip-flopping on policy choices reflect poor decision-making, or consistent track records virtue? We show that under some conditions, consistent records are indicative of ability. This happens for example when past information depreciates fast enough. In this case, politicians with consistent records would get reappointed (or reelected) and politicians who flip-flop thrown out of office. In general, however, the extent of policy reversal is not monotonically tied to ability.

The rest of the paper is organized as follows. We review the related literature in Section 2, and present the model in Section 3. We present the results in Section 4. We begin in section 4.1, by characterizing the optimal strategy of a politician with no career or electoral concerns. We then consider career concerns with an uninformed voter (Section 4.2) and with transparency (Section 4.3). We conclude in Section 5. All proofs are in the Appendix.

2 Related Literature

Impetuous Youngsters and Jaded Old Timers. Prendergast and Stole (1996) (PS) consider the problem of a politician (a manager) who chooses investments on a project in each of T periods. The project has unknown profitability, which is unchanging in time, but the manager receives a private signal of its return in each period, with more talented managers receive more precise signals. The manager cares about current profits and the current market's perception of ability. The main

[1]A similar logic holds for an almost incompetent bureaucracy or independent monitoring agency in other principal-agent relationships (ministers/President, investors/fund managers, CEOs/shareholders.

result of the paper is that managers will initially overreact to new information but eventually become unwilling to respond to new information suggesting that their previous behavior was wrong. The model has obvious similarities with our setting, but also important differences. The first and most fundamental difference is that the optimal t-period optimal investment policy is unchanging in time (this is akin to perfect persistence in our setting). In this setting, politician gradually learns the truth as he accumulates signals drawn from the same true state.[2] In our setting, instead, the optimal policy evolves stochastically in time, so past information depreciates. A second important difference is that in PS the manager cares (myopically) about his reputation at the end of the current period, while in our setting politician cares about how the entire decision record affects the voter's perception of his ability. Third, while in PS the reward to reputation is continuous, in our model this is bunched in classes (politician is either appointed for office B or not). The distinction is similar to what Alesina and Tabellini (2007) call bureaucrats and politicians (here PS consider the bureaucrats case, while we consider politicians).

Cheap Talk with Career Concerns. These models build on the seminal contribution of Crawford and Sobel (1982). In the basic setting, there are two players, an expert (the sender) and a decision-maker (the receiver). The optimal policy for both the expert and the decision-maker (typically not the same) are a function of an unobservable state of the world, about which the expert is privately informed. After observing a signal about the state, the expert sends a report to the decision-maker, who then chooses policy. Differently than in CS, the expert's ability (the precision of her private information) is private information. Moreover, the expert has career

[2] Because of this, the change in beliefs is larger at the beginning, when the manager is more uncertain about profitability, and larger the more talented the manager is. After many periods, the manager is more certain about the state, and therefore the next signals typically confirm what he knows already. Thus beliefs and optimal actions (if in a continuum, as here) should change little, at least if the manager is talented. As a result, to be perceived as talented, the manager wants to change a lot early (impetuous youngster) and do not change much later on (jaded old timers). Because the manager cares about the profits, under some conditions there exists a separating equilibrium where the distortion accentuates the impetuous youngsters and jaded old timers type of behavior.

concerns, so that her payoff depends on the receiver's perception of her ability.[3] The overwhelming message is that the expert will generally not report truthfully. This is true with a generic continuous information structure (Ottaviani and Sorensen (2006a), Ottaviani and Sorensen (2006b)), and in a repeated version of the cheap talk game with reputation (Sobel (1985), Benabou and Laroque (1992), Morris (2001)). The closest paper here is Li (2007). In this model, an agent delivers an initial report and a final report about the state of the world based on two private signals of increasing quality, after which the true state becomes publicly observable.[4] The agent is one of two types: smart or mediocre (this is private information). Agent's types differ in the precision of their signal and possibly also in the slope of signal quality improvement. The agent is paid more the more able he is perceived to be. The main result of the paper is that inconsistent reports signal high ability in equilibrium when a smart agents signals improve faster than those of a mediocre one.[5]

Informational Cascades. In our model, past information can overwhelm current information, which is then not used for current decisions. This is akin to what happens in models of informational cascades (Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992)). In the basic setting, a group of individuals decide in sequence without knowing a payoff relevant state of the world. Individuals condition their decisions on their own private information and the history of predecessors' decisions. Each individual's decision only affects her own payoffs. In equilibrium, the informa-

[3]Note that differently than in our setting, here the informed party cannot choose policy, but only send a message to the decision-maker.

[4]The model has two stages, and the same game is repeated in the second stage. However the agents wage does not depend on his second-stage performance and thus has no further reputational concerns at the end of his career. Since there is no conflict of interest between the principal and the agent, there is an equilibrium in the second stage in which the agent reports truthfully.

[5]The characterization of equilibria depend on the initial precision of the good agent p_0. When p_0 is low, in general there is a full revelation equilibrium. If the mediocre agent improves faster than the smart agent, there is a cutoff x_L such that if $p_0 \geq x_L$ there exists a pooling equilibrium, and moreover in any equilibrium the second report is uninformative. When instead the smart agent improves faster than the mediocre agent there exists a cut off value x_H such that if $p_0 \geq x_H$ then there is a unique partial revelation equilibrium in which the smart agent reports truthfully and the mediocre agent sends an untruthful consistent report with positive probability. It is in this case where inconsistent reports signal high ability in equilibrium.

tion in the publicly observed action history can overwhelm one individual's private information, causing all private information from this point on to be ignored. The new actions have no informational content and subsequent individuals face the same problem, with everyone rationally herding on the same action. Moscarini, Ottaviani, and Smith (1998) consider a similar model, with the difference that the state of the world changes stochastically over time, exactly as in our model. They show that because of the depreciation of past information, only temporary informational cascades can arise in this setting. Moreover, when the persistence of the state process is sufficiently low (past information depreciates fast) no cascade ever arises.

Selection Through Elections. Our paper also relates to papers studying how well elections serve as a selection device to sort out politicians with desirable characteristics.[6] Banks and Sundaram (1998) study the optimal retention rule for voters in a model that incorporates both moral hazard and adverse selection. Canes-Wrone, Herron, and Shotts (2001) consider a model in which elected officials have the same preferences as the electorate, and the incumbent attempts to signal talent (e.g., more precise information). With some probability, the voter can observe the state, and thus evaluate the accurateness of the decision. Otherwise, she is uninformed. They conclude that elected officials will pander (choose the popular, ex ante preferred action) only under some limited conditions. Canes-Wrone and Shotts (2007), however, show that elected officials will be more inclined to pander when there is uncertainty regarding their congruence with the electorate. Pandering also comes out in Maskin and Tirole (2004), who compares "judges" and politicians. Maskin and Tirole assume that the official values office per se, and also has a legacy motivation. When the office-holding motive is strong, politicians want to pander. The distortions are larger when the public is poorly informed about what the optimal action is, and when feedback about the quality of the decision is limited (and in this case non-elected offi-

[6] A different although related literature studies elections as disciplining device. Here the focus is not on adverse selection but on how elections can either induce effort by the elected politicians (moral hazard) or induce them to choose the policies preferred by the voter (see for example Barro (1973), Ferejohn (1986), Banks and Sundaram (1998), and Alesina and Tabellini (2007)).

cials are preferred). In our setting this conclusion holds in the extreme case in which the voter is completely uninformed, but breaks down with even limited transparency about optimal policies if the evaluation horizon is long enough.

Transparency. Prat (2005) introduces a distinction between transparency about consequences and transparency about actions. He shows that while transparency on consequences is beneficial, transparency on action can have detrimental effects. In his model, an agent selects a policy that is payoff relevant for a principal. The principal's optimal policy depends on the realization of a binary state, which is unobservable for the principal, but about which the agent is privately but imperfectly informed. The precision of the agent's signal depends on his type, which can be either good or bad. The payoff of the agent is the principal's belief that the agent is good. The principal is better off if the agent's action matches the state, and the larger is her belief that the agent is good (a reduced form of a two-period career concerns model). In this setting, if the principal can observe the agent's action, the agent has an incentive to disregard useful information and act according to how an able agent is expected to act a priori. Fox and Van-Weelden (2011) show that under some conditions (when the prior on the state of the world is sufficiently strong, and the costs are sufficiently asymmetric) the principal is made better off observing the action, but worse off observing the consequences of that action.[7]

Other Papers. The focus and type of results in this paper are very different from that of the "testing experts" literature (see for example Olszewski and Sandroni (2008, 2009)). Most importantly, in these papers the principal is uninformed about the entire data generating process, and tries to device a test that would be able to reject a false theory proposed strategically by an uninformed agent. In our case only the type and history of signals of politician is not known to the voter. Our paper is also very different from others in which the principal can compare the behavior of multiple experts (see for example Meyer (1991) and Bernhardt (1995) among many others).

[7]See also Levy (2005) for the analysis of transparency in committees.

3 The Model

There are two agents, the politician, and the voter, and two stages. In the first stage, the politician makes a sequence of T policy decisions in office A, $d_t \in \{0,1\}$, $t = 1,\ldots,T$. We refer to a complete sequence of policy decisions as a policy, and write this as $h_T(d)$, where for any x and $t = 1, 2, \ldots$, $h_t(x) \equiv (x_1, x_2, \ldots, x_{t-1})$. In the second stage, the voter observes the policy $h_T(d)$ and possibly also some additional information Z, and decides whether to appoint the politician for an office B or to instead elect an untested politician.

In each period t, there is an optimal t-period policy $\omega_t \in \{0,1\}$, which we refer to as a *passive* or *active* policy. While the realization of ω_t is unobservable, it is common knowledge that $\Pr(\omega_1 = 1) = p_1 \geq 1/2$ and that $\Pr(\omega_{t+1} = j|\omega_t = j) = \gamma_j \geq 1/2$. Given $b \in (0,1)$, the politician suffers a loss of b if he chooses an active policy when a passive policy is optimal ($d_t = 1$ when $\omega_t = 0$), suffers a loss of $(1-b)$ if he chooses a passive policy when an active policy is optimal ($d_t = 0$ when $\omega_t = 1$), and has a payoff of zero if $d_t = \omega_t = 0$ or $d_t = \omega_t = 1$. Thus given information $\mathcal{I}$, a politician with no career concerns (whose sole motivation is given by policy considerations) chooses $d_t = 1$ if and only if $\Pr(\omega_t = 1|\mathcal{I}) \geq b$.

While the politician cannot observe the optimal t-period policy, we assume that she can observe a signal s_t, such that $\Pr(s_t = j|\omega_t = j) = \theta$, $j \in \{0,1\}$. The parameter θ represents the *ability* of the politician, and is assumed to be known by the politician but not by the voter. There are two types of politician: experts ($\theta = e$) and amateurs ($\theta = a$), where $1/2 < a < e < 1$. The probability that the politician is an amateur is $\pi \in (0,1)$.

The politician cares about both the consequences of his decisions in office A (with weight $1-\delta$) and about the expected payoff of being appointed to office B by the voter (this has a weight δ). The payoff of a politician of type θ at time t with a history of decisions $h_t(d)$ and signals $h_t(s)$ when choosing a continuation $c_t(d) \equiv (d_t, \ldots, d_T)$ is

$$V_t(c_t(d); h_t(d), h_{t+1}(s), \theta) = \delta T y(c_t(d); h_t(d)) + (1-\delta)\sum_{m=t}^{T} E\left[u(d_m, \omega_m)|h_{t+1}(s), \theta\right]$$

where $\delta \in [0,1]$ is the career or reelection concern, and $y = 1$ if politician is appointed to office B and $y = 0$ otherwise.[8] The voter has a payoff of one if she appoints an expert politician and a payoff of zero if she appoints an amateur politician.

Let $H_t(d)$ and $H_t(s)$ denote the set of possible decision and signal histories at date $t = 1, \ldots, T+1$, and let $\Theta \equiv \{a, e\}$ be the set of types. For any $t = 1, \ldots, T$, $\theta \in \Theta$, $h_t(d) \in H_t(d)$ and $h_{t+1}(s) \in H_{t+1}(s)$, let $\sigma_{\theta,t}(h_t(d), h_{t+1}(s)) = \Pr(d_t = 1|\theta, h_t(d), h_t(s), s_t)$, and let $\sigma_\theta = \sigma_{\theta,1} \times \ldots, \times \sigma_{\theta,T}$. We denote the strategy of the agent as the mapping $\sigma \equiv (\sigma_a, \sigma_e)$. We refer to σ_θ as type θ's choice strategy (the restriction of the politician's strategy σ to its type θ).

Conditional on information $\mathcal{I}$ at time T, each type θ's choice strategy σ_θ induces a distribution over policies $h_T(d)$, call this $f(\cdot|\sigma_\theta; \mathcal{I})$, given by:

$$f(h_T(d)|\sigma_\theta; \mathcal{I}) = \prod_{t=1}^{T} \left[\sum_{\omega_t} \Pr(\omega_t|\omega_{t-1}, \mathcal{I}) \sum_{s_t} \Pr(s_t|\omega_t; \theta) \sum_{d_t} \Pr(d_t|h_t(d), h_t(s), s_t; \sigma_\theta) \right]$$

where $\Pr(d_t = 1|h_t(d), h_t(s), s_t; \sigma_\theta)$ is given by $\sigma_{\theta,t}(h_t(d), h_{t+1}(s))$, and $\omega_0 \equiv \emptyset$ for convenience.

In the paper we consider two scenarios regarding the information $\mathcal{I}_T$ that the voter has at time T. In the first case, $\mathcal{I}_T = \emptyset$, the voter has no information beyond the decision record itself. We say that in this case the voter is *uninformed*. In the second case, the voter receives a sequence of signals z_t, $t = 1, \ldots, T$, where $\Pr(z_t = j|\omega_t = j) = q > 1/2$, for $j \in \{0, 1\}$. We say that in this case the voter is *q-informed* or that the decision making environment is transparent (or q-transparent).

A pure strategy of the voter is a rule $y(\cdot)$ mapping the set of possible decision records $h_T(d)$ and (if the voter is q-informed) signals $h_T(z)$ to a decision of whether to appoint the politician ($y = 1$) or not ($y = 0$). Thus $y(h_T(d), h_T(z)) = 1$ indicates that the voter appoints the politician at office B after decision record $h_T(d)$ and signals $h_T(z)$. A mixed strategy $\varphi(h_T(d), h_T(z))$ is a probability of appointing the politician conditional on $(h_T(d), h_T(z))$.

[8] Note that we have multiplied the career concerns term by T. This is equivalent to assuming that the politician cares about the proportion of correct decisions, and avoids changing the weight of career concerns because of the number of decisions taken in office A.

Given the politician's strategy σ, and conditional on $\mathcal{I}_T$, after observing the policy $h_T(d)$ the voter has a belief over types $\beta^v_\sigma(h_T(d);\mathcal{I}_T) \equiv \Pr(\theta = a|h_T(d);\sigma,\mathcal{I}_T)$, and appoints the politician if and only if

$$\beta^v_\sigma(h_T(d);\mathcal{I}_T) \leq \pi \Leftrightarrow \frac{f(h_T(d)|\sigma_e;\mathcal{I}_T)}{f(h_T(d)|\sigma_a;\mathcal{I}_T)} \geq 1.$$

Thus in equilibrium $\varphi(h_T(d), h_T(z)) > 0$ only if $\beta^v_\sigma(h_T(d);\mathcal{I}_T) \leq \pi$, and similarly $\varphi(h_T(d), h_T(z)) < 1$ only if $\beta^v_\sigma(h_T(d);\mathcal{I}_T) \geq \pi$.

4 Results

4.1 Policy Choices with No Career Concerns

We begin by characterizing the optimal strategy of a politician with no career or electoral concerns. In this setting, the politician's optimal policy strategy is independent of the voter's own beliefs and strategy. As a result, each decision is independent of all previous and future decisions (the decision problems are related, through information, but the decisions themselves are not). The politician therefore chooses an active policy in period t if and only if the probability that this policy is optimal is above the threshold b, that is

$$\Pr(\omega_t = 1|h_{t+1}(s);\theta) = \sum_{\omega_{t-1}} \Pr(\omega_t = 1|\omega_{t-1}, s_t;\theta)\Pr(\omega_{t-1}|h_t(s);\theta) > b$$

Note that the t-period posterior belief is a function of the entire history of signals, including s_t and the t-period history $h_t(s)$. However, the history $h_t(s)$ enters only through its effect on $\Pr(\omega_{t-1}|h_t(s);\theta)$. Hence, for a type θ-politician, the t-period *prior* belief $p^\theta_t(h_t(s)) \equiv \Pr(\omega_t = 1|h_t(s);\theta)$ is a sufficient statistic for $h_t(s)$. We then let $\beta_\theta(s_t, p^\theta_t(h_t(s)) \equiv \Pr(\omega_t = 1|h_t(s), s_t;\theta)$ denote the politician's t-period posterior beliefs as a function of the current period signal s_t and the prior p^θ_t. With this notation, a type θ politician chooses $d_t = 1$ in period t if and only if

$$\beta_\theta(s_t, p^\theta_t) = \frac{\Pr(s_t|\omega_t = 1,\theta)p^\theta_t}{\Pr(s_t|\omega_t = 1,\theta)p^\theta_t + \Pr(s_t|\omega_t = 0,\theta)(1-p^\theta_t)} > b \tag{1}$$

The fundamental tradeoff between experience and adaptability follows from the fact that the information in the t-period prior can overwhelm t-period information. This happens if either $\beta_\theta(1, p_t^\theta) \leq b$ or $\beta_\theta(0, p_t) \geq b$, or equivalently, if

$$p_t^\theta \leq \frac{(1-\theta)b}{(1-\theta)b + \theta(1-b)} \equiv \underline{p}(\theta, b) \quad \text{or} \quad p_t^\theta \geq \frac{\theta b}{\theta b + (1-\theta)(1-b)} \equiv \overline{p}(\theta, b)$$

Therefore, with no career concerns, the politician's dominant strategy $\hat{\sigma} = (\hat{\sigma}_a, \hat{\sigma}_e)$ boils down to choosing an active policy independently of t-period information whenever the t-period prior is sufficiently favorable for the active policy, i.e., $p_t^\theta \geq \overline{p}(\theta, b)$, to be passive independently of the t-period information whenever his t-period prior is sufficiently unfavorable to taking action, i.e., $p_t^\theta \leq \overline{p}(\theta, b)$, and to follow the t-period information whenever p_t^θ is in the interval $P(\theta, b) \equiv (\underline{p}(\theta, b), \overline{p}(\theta, b))$:

$$\hat{\sigma}_{\theta,t}(h_t(d), h_t(s), s_t) = \begin{cases} 1 & \text{if } p_t^\theta \geq \overline{p}(\theta, b) \text{ or } p_t^\theta \in P(\theta, b) \text{ and } s_t = 1 \\ 0 & \text{if } p_t^\theta \leq \underline{p}(\theta, b) \text{ or } p_t^\theta \in P(\theta, b) \text{ and } s_t = 0. \end{cases} \tag{2}$$

Two simple facts will be useful in our analysis. First, note that while in general the politician's t-period decision does not match his t-period signal, when he makes policy decisions according to $\hat{\sigma}$, then $d_t = s_t$ whenever the politician switches from a passive to an active policy or vice-versa.

Remark 1 *Suppose the politician sets policy according to the no-career-concern strategy $\hat{\sigma}$. Then $d_t \neq d_{t-1}$ implies $d_t = s_t$.*

On the flip side, if the politician is sufficiently good relative to the persistence of the state, then "old" information depreciates too fast, and becomes essentially useless for (t-period) decision-making.

Remark 2 *There exists $\hat{\theta}(\gamma_0, \gamma_1) \in (0, 1)$ such that $p_t^\theta \in P(\theta, b)$ for all $t = 1, \ldots, T$ whenever $\theta > \hat{\theta}(\gamma_0, \gamma_1)$. Furthermore, $\hat{\theta}(\gamma_0, \gamma_1)$ is increasing in γ_0 and γ_1.*

This result is relevant because, as we show in Section 4.3, better politicians have more volatile priors, but $P(\theta, b)$ also increases with the type of the politician (as

$1-\underline{p}(\theta,b)$ and $\overline{p}(\theta,b)$ are increasing in θ). The lemma shows that the interval $P(\theta,b)$ increases faster with θ than the support of the politician's t-period prior, and as a result, high ability politicians always have beliefs contained within $P(\theta,b)$. Thus if the politician is sufficiently good relative to the persistence of the optimal policy process, all of his decisions are based only on current period information. Moreover, the speed at which the support of p_t^θ grows with θ relative to $P(\theta,b)$ is decreasing in the memory of the optimal policy. As a result, for relatively low ability levels, the support of p_t^θ is contained in $P(\theta,b)$ if past information depreciates fast, but not if past information depreciates at a slower rate. The threshold ability for which the prior is always in $P(\theta,b)$ is therefore increasing in γ_0 and γ_1.

A valid concern regarding the characterization in (2) as a positive statement is that if decision-makers care about their reputation, posturing can taint the informative content of their policy decisions. In the next sections we reconsider the problem of a decision maker with career concerns under two alternative assumptions about the transparency of the decision-making environment. First, we focus on the case in which the voter is completely uninformed about the optimal policy process. With this, we aim to capture a situation in which the voter is unable to disentangle the effect of the politician's actions from other factors affecting observable outcomes. For example, the voter might be able to observe the unemployment rate, but in the absence of a deeper understanding of the economy, be ignorant about whether increasing or decreasing the deficit would reduce or increase unemployment further. We then move on to the case in which the the environment is at least somewhat transparent, in that the voter observes the optimal policy with noise.

4.2 Career Concerns with an Uninformed Voter

We begin with the case in which the voter is uninformed. We show that when career concerns are sufficiently important, in equilibrium the politician's decisions are completely uninformative about his ability to process information. Moreover, posturing completely dominates policy-making, in the sense that both types of politician choose a single sequence of policies ignoring all of their private information about the optimal policy process. We establish the result in two steps. First, we show that if the weight that the politicians put on political prospects is sufficiently large relative to policy considerations, their decisions in office are uninformative about their ability.

Lemma 1 *Suppose the Principal is uninformed. There exists a* $\underline{\delta} \in (0, 1)$ *such that if* $\delta > \underline{\delta}$*, then in any equilibrium the politician's decisions are uninformative; i.e.,* $\beta^P_\sigma(h_T(d); \emptyset) \equiv \Pr(\theta = a|h_T(d); \sigma, \emptyset) = \pi$ *for all policies* $h_T(d)$ *that have positive probability in equilibrium.*

When the voter is uninformed about the optimal policy in each period, low ability politicians can mimic the statistical properties of the experts' behavior, with the voter being unable to call their bluff. Since in equilibrium the voter cannot commit to appointment decisions that are not sequentially rational, her retention decisions can only be contingent on the observed policy choices. But then if observed policy choices were to be informative, there would be some policies after which the voter would appoint the politician, and others for which she would not appoint the politician. If career or reelection concerns are sufficiently important, the payoff of getting hired dominates the additional flexibility of being free of posturing.[9]

In principle, Lemma 1 allows equilibrium strategies in which the expert politician conditions policy choices on the information he observes, $h_t(s)$, and the amateur politician merely "mimics" the statistical properties of the expert's strategy. However we show in Lemma 2 that this cannot be the case: if policy decisions are uninformative in equilibrium, politicians must *pool* on a policy that is unresponsive to the politician's information. It follows that the inability of the voter to determine which policy is optimal induces even the best politicians to disregard their information completely.

[9]Key to this result is that the voter only decides whether to reelect/appoint the politician to office or not, and cannot commit to punishing the politician for having chosen certain policies come election time. We believe that this coarseness of the set of instruments at the disposition of the voter is indeed one of the main characteristics of a political setting. In other environments, however, it might be natural to assume that the principal has a wider array of instruments at her disposal. We return to this point in our concluding remarks.

Lemma 2 *Suppose the principal is uninformed. If policy choices are uninformative in equilibrium, the politician chooses a policy $h_T(d)$ that is not contingent on his information $\{h_t(s)\}_t$ or ability type: for all t and all $h_t(s)$, $\sigma_j((h_t(s), s_t = 1), h_t(d)) = \sigma_j((h_t(s), s_t = 0), h_t(d)) = \sigma(h_t(d))$ for $j = e, a$.*

The proof builds on the observation that if the expert politician's policies are responsive to his private information, then at some point t mimicking by the amateur politician must entail mixing over $d_t = 1$ and $d_t = 0$. If we assume that the voter appoints the politician whenever she is indifferent, then his payoffs are determined solely by policy concerns, and will therefore prefer to deviate and play his no career concern action, as dictated by $\hat{\sigma}_a(h_{t+1}(s))$. This still leaves open the possibility that there might be a non-pooling equilibrium with mimicking if the voter uses different appointment probabilities after different policies of the politician. But because the politician's payoffs are a function of the (private) belief with which he evaluates the t^{th} decision, there is no way to induce him to mix in the first place.

Lemmas 1 and 2 directly imply Theorem 1.

Theorem 1 *Suppose the voter is uninformed. Then there exists $\underline{\delta} \in (0, 1/2)$ such that if $\delta > \underline{\delta}$, in equilibrium both the expert and amateur politician implement a policy $h_T(\tilde{d})$ that is independent of the politician's private information.*

Note then that when the voter is uninformed, posturing entails a large welfare loss. In fact, in equilibrium both types of politician disregard private information in policy-making. It follows that if the voter cares about decision-making in office A but is uninformed about optimal policy, she would do strictly better if she could commit not to use politician's policy choices to decide whether or not to appoint the politician to office B. This result echoes the insights regarding the "wrong type" of transparency identified by Prat (2005).

The welfare loss incurred in the pooling equilibrium is especially severe when the politician is highly able. On the other hand, the loss cannot be too big, for the politician can always deviate and play the no career concerns strategy $\hat{\sigma}_\theta$. This is most clearly illustrated in the case of an expert politician, who can observe the state perfectly; i.e., $\theta = 1$. First, note that because the informed expert doesn't make mistakes, his payoff from following the no career concerns strategy at any point t is simply $\hat{W}_{\theta=1}(t) = 0$. Thus, he will be willing to keep playing the pooling strategy in period t as long as the per period benefit from pooling (the career concern δ) is larger than the per period expected cost of pooling, given beliefs $\Pr(\omega_t = 1|h_{t+1}(s);\theta) = \beta_\theta \in \{0,1\}$. Note that the best pooling equilibrium for politician is to pool on $\tilde{d}^{PL} = (1,1,\ldots,1)$ if $\overline{\mu} > b$ and to pool in $\tilde{d}^{PL} = (0,0,\ldots,0)$ if $\overline{\mu} < b$. Suppose for concreteness that $\overline{\mu} > b$, so that $\tilde{d}^{PL} = (1,1,\ldots,1)$. Then the per period expected cost of pooling, given beliefs β_θ is (see Lemma 3 in the Appendix)

$$(1-\delta)b\left\{\left(\frac{T-t}{T}\right)(1-\overline{\mu}) + \frac{1}{T}(\overline{\mu}-\beta_\theta)\left[\frac{1-\kappa^{T-t+1}}{1+\kappa}\right]\right\},$$

where $\kappa \equiv \gamma_0 + \gamma_1 - 1$, and as before, $\overline{\mu} = \frac{1-\gamma_0}{2-\gamma_0-\gamma_1}$. Note that the cost of pooling is larger for $\beta_\theta = 0$ and $t = 1$. Thus, there exists an equilibrium with pooling if and only if

$$\delta \geq \frac{\left(\frac{T-1}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]}{1/b + \left(\frac{T-1}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]} \equiv \underline{\delta}(1),$$

The cost of pooling and the threshold $\underline{\delta}(1)$ are larger (i) the longer is politician's tenure in office A, (ii) the larger is politician's bias b to follow an active policy $d_t = 1$ in any given period, (iii) the smaller is the long-run average probability that an active policy is optimal $\overline{\mu}$, which is increasing in γ_1 and decreasing in γ_0, and (iv) the fastest past information depreciates (the smaller is $\gamma_0 + \gamma_1$).

4.3 Career Concerns with (Partial) Transparency

We have shown that when the voter is completely uninformed, posturing dominates policy-making, in the sense that politicians chooses a single sequence of policies ignoring all information about the optimal policy process, independently of their ability. In contrast to this negative result, we show that endowing the voter with a noisy observation of the optimal policy can be enough to break posturing. In particular, we show that even with an arbitrarily large noise about the outcomes of politician's decisions in office, when the voter can observe enough decisions she can discriminate between good and bad decision makers with high probability. This in turn allows the existence of an equilibrium with no posturing, in which each type of politician follows the no career concern strategy $\hat{\sigma}_\theta$.

The proof of this result builds on the fact that when the expert politician chooses policy with the no career concerns strategy $\hat{\sigma}_e$, the voter can use her private information to identify almost surely decision records that are generated by an expert politician.[10] Because the inference of the voter relies on information that is not known to the politician, and the amateur is in fact of lower ability than the expert, the amateur cannot mimic the expert independently of how he sets policy. And because in equilibrium politicians will play according to $\hat{\sigma}$, the voter will be almost sure about politician's type almost always; i.e., $\Pr\left(\Pr(\theta = e|h_T(d))|e, \hat{\sigma}_e\right) \to_p 1$ and $\Pr\left(\Pr(\theta = a|h_T(d))|a, \hat{\sigma}_a\right) \to_p 1$. These conditions do not imply that the voter will always be able to appoint the expert and dismiss the amateur politician. There are some histories for which the voter will make mistakes. However for long enough policy records, these histories have a very small probability.

In the proof we make use of the following refinement, which simplifies the inferences of the voter following unexpected decision records off the equilibrium path.

A Refinement. With probability $1 - \xi$ (for ξ close to zero), the politician is fully rational. Conditional on being rational, he is an amateur with probability π, and an expert with probability $1 - \pi$. With probability ξ the politician is a behavioral type. There is a large number K of behavioral types. A behavioral type $k = \{1, \ldots, K\}$ plays d_1 with probability $1/2$, and then $\Pr(d_t = 1|d_{t-1} = 1) = \Pr(d_t = 0|d_{t-1} = 0) = 1/2 + k/2K$. Conditional on the politician being behavioral, he is of type k with probability r_k.

We are now ready to state formally our result.

[10] This result does not follow from the results in Kalai and Lehrer (1993) (Theorem 1), for in our case the decision-maker's type is not given only by her ability, but also by the signals he observed. As a result, the type space grows with T.

Theorem 2 *For any precision $q > 1/2$ of the voter's information, any career concern $\delta \in (0,1)$ of the politician, and any $\varepsilon \in (0,1)$, there is a $\overline{T}(\delta, q, \varepsilon)$ such that if $T > \overline{T}(\delta, q, \varepsilon)$, then there is a PBE in which (i) both types of politicians "ignore" career concerns and choose policies in office A according to $\hat{\sigma}$, and (ii) the voter appoints an "expert" politician but not an "amateur" politician at task B with a probability of at least $1 - \varepsilon$.*

Theorem 2 highlights the primary importance of transparency on equilibrium outcomes. With no transparency about the outcomes of politician's decisions in office, low ability politicians disregard useful information in order to mimic the behavior of expert politicians and decision records become uninformative, jeopardizing selection efforts by the voter. In contrast, with even an arbitrarily small amount of transparency about the outcomes of politician's decisions in office, when politician's tenure is long enough (i) the voter appoints good types and fires bad types with probability close to one, and (ii) politician chooses policy as if he had no career concerns (implying efficiency when he has the same preferences as the voter). Furthermore, the more transparent is the environment (i.e., the larger is q), the shorter is the evaluation period $\overline{T}(\delta, q)$ needed to break posturing. (This is a straightforward corollary of Theorem 2, and follows intuitively from considering the limiting case of q approaching 1.)

The role of transparency here is indirect. While a very noisy observation of the optimal policy process would not be enough for the voter to choose the right policy in the absence of delegation, it is enough to sort out the agency relation when the evaluation period is long enough. This suggests that an unsophisticated media can greatly improve the quality of political choices taken in office, and may help voters to appoint high ability decision makers to office.

4.3.1 Ability and Policy Outcomes

Theorem 2 shows that if tenure is long enough, in equilibrium the politician chooses policy according to $\hat{\sigma}$, even after taking into account career concerns. This allows us to understand more deeply the connection between policy outcomes and the ability of the decision maker. In particular, a natural question in this setting is whether observing frequent policy reversals tells us something about the ability of the decision-maker. Does flip-flopping on policy choices reflect poor decision-making, and consistent track records virtue?

The characterization of $\hat{\sigma}$ in (2) pins down the policy chosen by the politician as a *contingent plan*. To say more about the expected *outcomes*, we need to investigate further the properties of the t-period prior process.

To do this it is useful to write p_t^θ recursively, exploiting the Markovian structure of the state process. Note that $p_t^\theta = \sum_{\omega_{t-1}} \Pr(\omega_t = 1|\omega_{t-1}) \Pr(\omega_{t-1}|h_t(s);\theta)$, which can be written as

$$p_t^\theta = \beta_\theta(s_{t-1}, p_{t-1}^\theta)(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0).$$

Substituting $\beta_\theta(s_{t-1}, p_{t-1}^\theta)$ from (1), and letting $w(p_t^\theta, \theta) \equiv \theta p_t^\theta + (1-\theta)(1-p_t^\theta)$, we obtain:

$$p_{t+1}^\theta = \begin{cases} \frac{\theta p_t^\theta}{w(p_t^\theta,\theta)}(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0) & \text{if } s_t = 1 \\ \frac{(1-\theta) p_t^\theta}{1-w(p_t^\theta,\theta)}(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0) & \text{if } s_t = 0. \end{cases} \tag{3}$$

It follows that the politician's beliefs are given by the stochastic process

$$p_{t+1}^\theta = \begin{cases} \frac{\theta p_t^\theta}{w(p_t^\theta,\theta)}(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0) & \text{w.p. } w(p_t^\theta, \theta) \\ \frac{(1-\theta) p_t^\theta}{1-w(p_t^\theta,\theta)}(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0) & \text{w.p. } 1 - w(p_t^\theta, \theta). \end{cases} \tag{4}$$

Hence, both the size of the jump from p_t^θ to p_{t+1}^θ and the transition probabilities $w(p_t^\theta, \theta) = \theta p_t^\theta + (1-\theta)(1-p_t^\theta)$ and $1 - w(p_t^\theta, \theta)$ are functions of the state of the process in period t. Note moreover that while the probability of moving away from the mean is higher than the probability of moving towards the center (by the memory of the process), the size of the jump after updating against the current belief is larger than one resulting from reinforcing the current belief (by the concavity of learning). However, expression (4) implies that

$$E[p_{t+1}^\theta | p_t^\theta] = (1 - \gamma_0) + (\gamma_1 + \gamma_0 - 1) p_t^\theta. \tag{5}$$

Hence $E[p^\theta_{t+1}|p^\theta_t] < p^\theta_t$ if and only if $p^\theta_t > (1-\gamma_0)/(2-(\gamma_0+\gamma_1)) \equiv \overline{\mu} \equiv \lim_{t\to\infty} \Pr(\omega_t = 1)$. Thus, politician's prior beliefs fluctuate around the long term probability $\overline{\mu}$ that $\omega_t = 1$. Figure 1 illustrates one possible realization of this process.

[Figure 1 about here]

As equation (5) shows, the expected value of p^θ_{t+1} conditional on p^θ_t does not depend on the politician's ability, θ. This property does not extend to other moments of the distribution. In particular, using (4) we can compute the conditional variance,

$$V(p^\theta_{t+1}|p^\theta_t) = \frac{(\gamma_1+\gamma_0-1)^2}{x(\theta)\,(1-x(\theta))}[p^\theta_t(1-p^\theta_t)]^2,$$

where $x(\theta) \equiv \theta p^\theta_t + (1-\theta)(1-p^\theta_t)$. From this it follows that

$$\frac{\partial V(p^\theta_{t+1}|p^\theta_t)}{\partial\theta} = \frac{(\gamma_1+\gamma_0-1)^2}{[x(\theta)(1-x(\theta))]^2}[p^\theta_t(1-p^\theta_t)(2p^\theta_t-1)]^2 > 0.$$

Thus, better politicians have more volatile beliefs (and more so the slowest past information depreciates; i.e., the larger are γ_0, γ_1). Now, in order to characterize behavior, we are interested in the variation in p^θ_t *relative to* $P(\theta,b) \equiv (\underline{p}(\theta,b), \overline{p}(\theta,b))$. And as we have shown before (see Remark 2) $P(\theta,b)$ increases faster with θ than the support of the politician's t-period prior. As a result, sufficiently high ability politicians always have beliefs contained within $P(\theta,b)$. Thus if the politician is sufficiently good relative to the persistence of the optimal policy process, all of his decisions are based only on current period information. In particular, there exists a threshold $\hat{\theta}(\gamma_0,\gamma_1)$, increasing in γ_0 and γ_1, such that $p^\theta_t \in P(\theta,b)$ for all $t = 1,\ldots,T$ whenever $\theta > \hat{\theta}(\gamma_0,\gamma_1)$.

Whenever the politicians' abilities are sufficiently high, experts' policy records are more *consistent* than those of amateur politicians. We establish this under two alternative notions of consistency. First, we define the (long-run) flip-flop rate out of a $j = \{0,1\}$ decision of a type θ politician given strategy σ_θ as $\phi_j(\theta, \sigma_\theta) \equiv \lim_{t\to\infty} \Pr(d_{t+1} \neq j | d_t = j; \sigma_\theta, \theta)$. This is the long-run probability that politician changes policy in any given period. We say that a type θ politician is more likely to flip-flop from a j decision than a type θ' politician given strategies $\sigma_\theta, \sigma_{\theta'}$ if $\phi_j(\theta, \sigma_\theta) > \phi_j(\theta', \sigma_{\theta'})$. Second, we say that a type θ politician is more consistent on active (passive) decisions than a type θ' politician given $\sigma_\theta, \sigma_{\theta'}$ if for $j = 1$ ($j = 0$) and for all $\ell' > \ell > 1$

$$\frac{\Gamma^j_{\ell'}(\theta, \sigma_\theta)}{\Gamma^j_{\ell'}(\theta', \sigma_{\theta'})} > \frac{\Gamma^j_{\ell}(\theta, \sigma_\theta)}{\Gamma^j_{\ell}(\theta', \sigma_{\theta'})} > 1,$$

where $\Gamma^j_{\ell'}(\theta, \sigma_\theta) \equiv \lim_{t\to\infty} \Pr(d_{t+\ell'} = \ldots = d_{t+1} = d_t = j; \theta, \sigma_\theta)$. Thus, a type θ politician is more consistent on active decisions than a type θ' politician if for any $\ell' > 1$, θ is more likely to generate a chain of active decisions of length ℓ' than θ', and if longer chains are increasingly likely to have been generated by type θ vis a vis θ'. Consistency is therefore a more expansive concept than flip-flopping.

In Proposition 1 we show that when all politicians' abilities are sufficiently high, so that past information depreciates fast enough (Lemma 2), the flip-flop rate out of the most likely optimal policy of both experts and amateurs is decreasing in the persistence of that policy, and increasing in the persistence of the less likely optimal policy. Furthermore, we show that experts are less likely than amateurs to flip-flop out of the long-run most likely optimal policy.

Proposition 1 *Suppose that $e > a > \hat{\theta}(\gamma_0, \gamma_1)$. If politicians have no-career concerns, (i) experts are less likely to flip-flop from a j decision than amateurs whenever $\gamma_j > \gamma_{-j}$; (ii) the flip-flop rate out of j decisions is decreasing in γ_j and increasing in γ_{-j} whenever $\gamma_j > \gamma_{-j}$.*

Using our more expansive notion of consistency, we can also show that under the same conditions of Proposition 1, experts are also more *consistent* than amateurs. We also show, however, that under alternative conditions – and in particular if the difference in competence between the expert and the amateur is sufficiently large – we can reverse this result, so that the amateur is more consistent than the expert.

Thus in general it is not possible to rank experts and amateurs based on consistency only. The next proposition formalizes our findings.

Proposition 2 *Fix $\gamma_1 \geq \gamma_0$.*[11] *Suppose that in equilibrium, politicians choose policy according to the no career concerns strategy $\hat{\sigma}$. Then (i) if $e > a > \hat{\theta}(\gamma_0, \gamma_1)$, experts are more consistent than amateurs on active decisions. On the other hand, (ii) if $b < \overline{\mu}$, there exist $\overline{\theta}(\gamma_0, \gamma_1) < 1$ and $\underline{\theta}(\gamma_0, \gamma_1) \in (1/2, \overline{\theta})$ both increasing in γ_1 and decreasing in γ_0, such that if $a < \underline{\theta} < \overline{\theta} < e$, the amateur is more consistent than then expert on active decisions.*

When policy choices are informative, if all politicians are sufficiently competent and/or information depreciates at a fast rate, flip-flopping on policy choices reflects poor decision-making skills. In this case we would expect politicians with consistent policy records to get reelected, and/or climb the ladder of political offices. On the other hand, this ranking does not hold in all situations. For example, if information does not depreciate fast enough for the worst possible politician, consistency can reflect poor decision-making skills.

4.4 Two Polar Cases: Full and No Information Depreciation

In the paper, we focused on the case of partial depreciation of past information. Since $\Pr(\omega_{t+1} = \omega_t) \in (1/2, 1)$, knowledge of the optimal policy in the current period is an informative but imperfect signal about what would be the optimal policy in the future. In this section, we discuss briefly the polar cases of full information depreciation, i.e. $\gamma_1 = \gamma_0 = 1/2$, and no information depreciation, i.e. $\gamma_1 = \gamma_0 = 1$. While in the former case past information is completely useless for current decisions, in the latter case the optimal decision is the same in all periods.

Full Information Depreciation. The case of full information depreciation falls entirely within the analysis in the paper. Theorems 1 and 2 apply unchanged, and so does the characterization of the no-career concerns strategy $\hat{\sigma}$ in Section 4.1, although

[11] A similar statement expressed in terms of passive decisions holds when $\gamma_0 > \gamma_1$.

in a simplified form. Note that in this case by definition $p_t^\theta = \Pr(\omega_t = 1|s_1, \ldots, s_{t-1}) = \Pr(\omega_t = 1)$, which approaches $1/2$ as t gets large. Thus, for large t, a politician with no career concerns and bias equal to b follows the t period information if

$$\underline{p}(\theta, b) \equiv \frac{(1-\theta)b}{(1-\theta)b + \theta(1-b)} < 1/2 < \frac{\theta b}{\theta b + (1-\theta)(1-b)} \equiv \overline{p}(\theta, b),$$

and he always takes active decisions ($d_t = 1$) when $\theta/(1-\theta) < (1-b)/b$, and always takes passive decisions ($d_t = 0$) when $\theta/(1-\theta) < b/(1-b)$. Clearly, an unbiased politician ($b = 1/2$) always bases his t-period decision on t-period information.

One Truth: No Information Depreciation. The case of no information depreciation is conceptually different from the case in which $\Pr(\omega_{t+1} = \omega_t) < 1$, both in the characterization of the no-career concerns strategy $\hat{\sigma}$ and in the strategic considerations behind Theorems 1 and 2, which do not apply in this environment.

Consider first the no-career concerns strategy. The fundamental difference with the analysis thus far is that in the case of full persistence there is an "unchanging truth" that the politician gradually learns about, becoming more and more informed about it as time goes by. Formally, in this case the politician receives independent signals from the *same* Bernoulli distribution, with success probability θ when $\omega = 1$ and $1 - \theta$ when $\omega = 0$. Thus, letting $m(k, t)$ denote the event in which exactly k out of t signals are equal to one, the politician's belief that $\omega = 1$ after observing $m(k, t)$ is

$$\beta_\theta(m(k,t)) = \left[1 + \frac{1-p_1}{p_1}\left(\frac{\theta}{1-\theta}\right)^{t\left(1-\frac{2k}{t}\right)}\right]^{-1}.$$

Now suppose that $\omega = 1$. Then $k/t = \left(\sum_{\ell=1}^{t} s_\ell\right)/t \to_p \theta > 1/2$. And with $k/t > 1/2$, $\beta_\theta(m(k,t)) \to 1$ for large t. Thus $\beta_\theta(m(\cdot,t)) \to_p 1$. Similarly, $\beta_\theta(m(\cdot,t)) \to_p 0$ when $\omega = 0$. That is, with full persistence the politician gradually learns the state of the world, and after sufficiently many observations knows the state almost perfectly almost always. The same of course is true of p_t^θ, which in this case is exactly equal to β_{t-1}^θ. If follows that if $b \in (0,1)$, there is a $\tilde{t}$ such that if $t > \tilde{t}$, then $d_t = 1$ almost surely when $\omega = 1$, and $d_t = 0$ almost surely when $\omega = 0$.

The second difference with the benchmark case is in terms of the strategic interactions, and the ability of the voter to distinguish between high ability and low ability politicians. Suppose for example that both types of politician are of relatively high ability, and play the no career concerns strategy $\hat{\sigma}$. Then with high probability both expert and amateur would converge within a few periods to a decision to which they will stick for the duration of their tenure in office. This illustrates the basic difference with the benchmark case of partial information depreciation. When the optimal policy is static, the voter cannot rely on long-run tests to discriminate among types, and must rely instead only on the initial "learning" period, and on how fast the politician converges to a decision.

5 Conclusion

Constant development is the law of life. This is true in the policy-making realm as well, where the optimal policy in any given period evolves in response to a changing environment. In this paper we tackled the joint problem of a career minded politician, who is learning about the changing environment in order to choose policy, and a voter, who is learning about the ability of the politician in order to decide whether to keep/promote him or dismiss him altogether. A key notion in this setting is how the actions of the politician are interpreted by the voter. Here we considered two alternative informational environments. In the first case, the voter can observe the choices of the politician, but cannot observe the process of optimal policies (the voter is uninformed). In the second one, the voter can observe the optimal policy in each period with noise. The level of transparency is parametrized by the precision of the voter's signal about the optimal policy.

Our results point to a relevant strategic effect of transparency and the media. When the voter is uninformed, and the politician puts enough weight on career concerns, in equilibrium policy decisions are independent of – and uninformative about – the ability of the politician. Moreover, posturing completely dominates policy-making, as both high and low ability politicians choose the same given sequence of policies, disregarding all private information about the optimal policy process.

However, with an arbitrarily small level of transparency, if the voter can observe a sufficiently large number of policy decisions, (i) the politician sets policy as if he had no career concerns, and (ii) the voter appoints high ability politicians and dismisses low ability politicians with probability close to one. Thus, while a very noisy observation of the optimal policy process would not be enough for the voter to choose the right action in the absence of delegation, this limited information is enough to sort out the agency relation and possibly lead to large gains in efficiency.

A key assumption that we maintained throughout the paper is that the voter only decides whether to reelect/appoint the politician to office or not, and cannot commit to punishing the politician for having chosen certain policies come election time. We believe that this coarseness of the set of instruments at the disposition of the voter is indeed one of the main characteristics of a political setting. In other environments, however, the principal has a wider array of instruments at her disposal.

The question that naturally emerges with a larger contracting space is what contract best balances selection and incentives considerations for the principal. A similar problem arises even within the political setting if we take the perspective of constitutional design: what institution best serves the voters' interests? A relatively simple but attractive version of this problem arises if we allow the voter to replace the politician at any given $t \leq T$. Analyzing this problem is also interesting from a theoretical perspective. On the one hand, the voter has now an additional instrument with which she can both discipline acting politicians and replace bad politicians early. On the other hand, the threat of being fired at any period increases the incentives of incompetent politicians to mimic high ability types, slowing down the rate of learning.

Another interesting extension of the model is to consider the case in which the politician's bias is private information. In the paper we focused on a situation in which the voter is uninformed about the ability of the politician, but not about his ideology. In this setting we obtained the result that consistent behavior can be indicative of high ability. When this is the case, politicians with consistent records would get reappointed (or reelected) and politicians who flip-flop thrown out of office. In some applications, however, it is reasonable to assume that not only the ability but also the bias of the politician is imperfectly observed by the voter. This would be the case for example in the case of judges or CEOs. In these cases, a consistent behavior in office may signal extremism. Understanding when the voter will be able to disentangle bias and ability in a changing environment is a natural next step in this research agenda.

6 Appendix A: Proofs

Proof of Remark 1. We show that $d_{t-1} = 1$ and $d_t = 0$ imply $s_t = 0$ (the other case is symmetric). Note that $d_{t-1} = 1$ implies either $p^\theta_{t-1} \geq \overline{p}(\theta, b)$ or $p^\theta_{t-1} \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$ and $s_{t-1} = 1$. Similarly, $d_t = 0$ implies that either $p^\theta_t \leq \underline{p}(\theta, b)$ or $p^\theta_t \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$ and $s_t = 0$. We then have four possibilities: (i) $p^\theta_{t-1} \geq \overline{p}(\theta, b)$ and $p^\theta_t \leq \underline{p}(\theta, b)$, (ii) $p^\theta_{t-1} \geq \overline{p}(\theta, b)$, $p^\theta_t \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$ and $s_t = 0$, (iii) $p^\theta_{t-1} \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$, $s_{t-1} = 1$ and $p^\theta_t \geq \underline{p}(\theta, b)$, or (iv) $p^\theta_{t-1} \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$, $s_{t-1} = 1$, $p^\theta_t \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$, and $s_t = 0$. If (ii) or (iv), we are done. Case (iii) is impossible, because $p^\theta_t > p^\theta_{t-1}$ when $s_{t-1} = 1$ (see equation (3)). Furthermore, for case (i) to be true, it must be that $s_t = 0$ and therefore $d_t = s_t$. ■

Proof of Remark 2. First, note that for all $\theta \in (1/2, 1)$, $p^\theta_t < \Pr(\omega_t = 1|\omega_{t-1} = 1) = \gamma_1$, and $1 - p^\theta_t < \Pr(\omega_t = 0|\omega_{t-1} = 0) = \gamma_0$, so that $p^\theta_t > 1 - \gamma_0$. Note that as long as $b \in (0, 1)$, $\overline{p}(\theta, b)$ is a strictly increasing continuous function of θ with $\overline{p}(1, b) = 1$, and $\underline{p}(\theta, b)$ is a strictly decreasing continuous function of θ with $\underline{p}(1, b) = 0$. Then there exists a $\theta_H < 1$ such that $\overline{p}(\theta, b) > \gamma_1$ whenever $\theta > \theta_H$, and $\theta_L < 1$ such that $\underline{p}(\theta, b) < 1 - \gamma_0$ whenever $\theta > \theta_L$. Thus $p^\theta_t \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$ for all $\theta > \max\{\theta_L, \theta_H\}$. In fact, it is easy to compute $\theta_H = \gamma_1(1-b)/(\gamma_1(1-b) + b(1-\gamma_1))$, increasing in γ_1 and decreasing in b, and $\theta_L = \gamma_0 b/(\gamma_0 b + (1-\gamma_0)(1-b))$, increasing in γ_0 and increasing in b. ■

Proof of Lemma 1. Let $D(\sigma)$ denote the set of records that have positive probability given σ; i.e., $D(\sigma) \equiv \{d : f(h_T(d)|\sigma_\theta) > 0 \text{ for some } \theta\}$. First we show that there exists a $\underline{\delta} \in (0, 1)$ such that if $\delta > \underline{\delta}_1$, then in any equilibrium politician's decisions are uninformative; i.e., $\beta^P_\sigma(h_T(d); \emptyset) \equiv \Pr(\theta = a|h_T(d); \sigma, \emptyset) = \pi$ for all $h_T(d) \in D(\sigma)$, and the voter appoints the politician to office B with positive probability.

Take $\delta \in (0, 1)$ given, and suppose that it is not true that $\beta^P_\sigma(h_T(d); \emptyset) = \pi$ for all $h_T(d) \in D(\sigma)$ in any equilibrium. Then there exist records $h^+_T(d) \in D(\sigma)$ and $h^-_T(d) \in D(\sigma)$ such that $\beta^P_\sigma(h^-_T(d); \emptyset) > \pi$ and $\beta^P_\sigma(h^+_T(d); \sigma) < \pi$ in equilibrium. Hence, the set of records $D^-(\sigma) \equiv \{h_T(d) \in D(\sigma) : y(h_T(d)) = 0\}$ and $D^+(\sigma) \equiv$

$\{h_T(d) \in D(\sigma) : y(h_T(d)) = 1\}$ are nonempty. Furthermore, since $\beta^P_\sigma(h^-_T(d);\emptyset) > \pi$, it must be that $\Pr(h^-_T(d)|e;\sigma,\emptyset) < \Pr(h^-_T(d)|a;\sigma,\emptyset)$. Thus in particular the amateur plays $h^-_T(d)$ with positive probability.[12]

Say that $h_t(d)$ is *compatible* with the set $D^+(\sigma)$ at time t (alternatively, with $D^-(\sigma)$) if $\exists\ d^0 \in D^+(\sigma)$ (in $D^-(\sigma)$) such that $h_t(d^0) = h_t(d)$. Define the function $m(\cdot;\sigma)$ mapping decision histories to $\{0,1\}$ by $m(h_t(d);\sigma) = 1$ if $h_t(d)$ is compatible with the set $D^+(\sigma)$, and $m(h_t(d);\sigma) = 0$ if $h_t(d)$ is compatible with the set $D^-(\sigma)$. Now, recall that the politician's payoff is

$$V_t(c_t(d);h_t(d),h_{t+1}(s),\theta) = \delta T y(c_t(d);h_t(d)) + (1-\delta)\sum_{m=t}^{T} E\left[u(d_m,\omega_m)|h_{t+1}(s),\theta\right]$$

and, for $t = 1,\ldots,T$, define the value $W^\theta_t(m,\beta)$ as follows

$$W^\theta_T(1,\beta) = \max_{d_T}\{\delta T m(h_T(d),d_T) + (1-\delta)E[u(d_T,\omega_T)|\beta_T = \beta]\}$$

and

$$W^\theta_T(0,\beta) = (1-\delta)\max_{d_T} E[u(d_T,\omega_T)|\beta_T = \beta],$$

where we write β_t instead of $\beta_t(h_{t+1}(s))$ for simplicity. Then, for any $t = 0,\ldots,T-1$, let

$$\begin{aligned} W^\theta_t(0,\beta) &= (1-\delta)\max_{d_t} E[u(d_t,\omega_t)|\beta_t = \beta;\theta] + \sum_{s_{t+1}} \Pr(s_{t+1}|\beta_t = \beta) W^\theta_{t+1}(0,\beta'(\beta,s_{t+1})) \\ &= (1-\delta)E\left[\sum_{n=t}^{T}\max_{d_n} E[u(d_n,\omega_n)|\beta_n;\theta]|\beta_t = \beta;\theta\right] \end{aligned} \tag{6}$$

and

$$W^\theta_t(1,\beta) = \max_{d_t}\left\{\begin{array}{c} (1-\delta)E[u(d_t,\omega_t)|\beta_t = \beta;\theta]+ \\ m(h_t(d),d_t;\sigma)\sum_{s_{t+1}} W^\theta_{t+1}(1,\beta'(\beta,s_{t+1}))\Pr(s_{t+1}|\beta_t = \beta;\theta)+ \\ (1-m(h_t(d),d_t);\sigma)\sum_{s_{t+1}} W^\theta_{t+1}(0,\beta'(\beta,s_{t+1}))\Pr(s_{t+1}|\beta_t = \beta;\theta) \end{array}\right\}$$

[12] It follows immediately then that if $\delta \geq 1/2$, this cannot be possible, for in this case getting hired dominates any possible payoff gain from better decision-making for all agent types, and thus no type of agent can play $h^-_T(d)$ in equilibrium. This is a contradiction, since $h^-_T(d) \in D(\sigma)$ by hypothesis.

Since $D^-(\sigma)$ is nonempty by hypothesis, and d^- is played with positive probability by the amateur, for equilibrium it must be that there exists some β that is consistent with a possible realization of signals such that

$$\left\{ \begin{array}{c} (1-\delta)\max_{d_t} E[u(d_t,\omega_t)|\beta_t=\beta;\theta=a] \\ + \\ \sum_{s_{t+1}} W^{\theta}_{t+1}(0,\beta'(\beta,s_{t+1}))\Pr(s_{t+1}|\beta_t=\beta;\theta=a) \end{array} \right\} > \left\{ \begin{array}{c} (1-\delta)\min_{d_t} E[u(d_t,\omega_t)|\beta_t=\beta;\theta=a] \\ + \\ \sum_{s_{t+1}} W^{\theta}_{t+1}(1,\beta'(\beta,s_{t+1}))\Pr(s_{t+1}|\beta_t=\beta;\theta=a) \end{array} \right\}$$

Note that the difference between the LHS and the RHS is maximized for $\beta=\overline{\beta}$ if $b\geq 1/2$ and for $\beta=\underline{\beta}$ if $b\leq 1/2$. Fix $b\geq 1/2$ without loss of generality. Then this cannot be an equilibrium if

$$b+E\left[\sum_{n=t}^{T}\max_{d_n} E[u(d_n,\omega_n)|\beta_n;a]|\beta_t=\overline{\beta};a\right] < \frac{1}{(1-\delta)}E[W^{\theta}_{t+1}(1,\beta'(\overline{\beta},s_{t+1}))|\overline{\beta};a]$$

While the LHS is constant in δ, the RHS is a continuous increasing function of δ. It is strictly smaller than LHS for $\delta\to 0$ and strictly larger than LHS for $\delta\to 1$. Then for any σ there exists a unique value of δ, say $\underline{\delta}(\delta)\in(0,1)$, such that if $\delta>\underline{\delta}(\sigma)$, then in any equilibrium σ, $\beta(d;\sigma)\equiv\Pr(\theta=a|d,\sigma)=\pi$ for all $d\in D(\sigma)$. Then define $\underline{\delta}=\max_{\sigma}\underline{\delta}(\sigma)$.

Now consider $\delta>\underline{\delta}$. Our previous argument establishes that the sets $D^-(\sigma)\equiv\{h_T(d)\in D(\sigma):y(h_T(d))=0\}$ and $D^+(\sigma)\equiv\{h_T(d)\in D(\sigma):y(h_T(d))=1\}$ cannot both be nonempty. Then either $\varphi(h_T(d))>0$ for all $h_T(d)\in D(\sigma)$, or $\varphi(h_T(d))<1$ for all $h_T(d)\in D(\sigma)$. Moreover, in this case there must exist a set of decision histories of the politician for which the voter appoints the politician with positive probability. Otherwise $y(h_T(d))=0$ for all $h_T(d)\in D(\sigma)$, the politician doesn't have a career concern, and best responds by playing $\hat{\sigma}$. But this contradicts $\beta(d;\sigma)\equiv\Pr(\theta=a|d,\sigma)=\pi$ for all $h_T(d)\in D(\sigma)$. ∎

Proof of Lemma 2. Suppose first that if indifferent, the voter appoints the politician with probability one. First, note that $\Pr(h_T(d)|\sigma_e)=\Pr(h_T(d)|\sigma_a)$ for all $h_T(d)\in D(\sigma)$ implies that $\Pr(h_t(d)|\sigma_e)=\Pr(h_t(d)|\sigma_a)$ for all t, for all $h_t(d)$ consistent with some $h_T(d)\in D(\sigma)$. So suppose (towards a contradiction) that there is a t and a signal realization $\tilde{h}_t(s)$ such that $\sigma_e(\tilde{h}_t(s),s_t=1)\neq\sigma_e(\tilde{h}_t(s),s_t=0)$. The choice strategy $\sigma_{e,\{1,2,\ldots,t\}}(\cdot)$ induces a probability distribution $f(\cdot|\sigma_e)$ over records $h_t(d)$. Then we can compute $\Pr(h_{t+1}(d)|\sigma_e) = \Pr(h_t(d)|\sigma_e)\Pr(d_t|\sigma_e, h_t(d))$. Since $\Pr(h_t(d)|\sigma_e) = \Pr(h_t(d)|\sigma_a)$ for all t and $h_t(d)$ consistent with some $h_T(d) \in D(\sigma)$, we must have that

$$\Pr(h_t(d)|\sigma_e)\Pr(d_t|\sigma_e, h_t(d)) = \Pr(h_t(d)|\sigma_a)\Pr(d_t|\sigma_a, h_t(d))$$

and $\Pr(h_t(d)|\sigma_e) = \Pr(h_t(d)|\sigma_a)$, so that we must have $\Pr(d_t|\sigma_e, h_t(d)) = \Pr(d_t|\sigma_a, h_t(d))$ as well. Now, since $\sigma_e(\tilde{h}_t(s), s_t = 1) \neq \sigma_e(\tilde{h}_t(s), s_t = 0)$, then $\Pr(d_t|\sigma_e, h_t(d)) \in (0,1)$, say $\Pr(d_t = 1|\sigma_e, h_t(d)) = \alpha \in (0,1)$. Inducing $\Pr(d_t|\sigma_a, h_t(d)) = \alpha$ will generically entail mixing by the amateur at either $s_t = 1$ or $s_t = 0$ (or both). But recall that the no career concerns strategy $\hat{\sigma}_a$ generically involves playing $d_t = 1$ or $d_t = 0$ after every history $h_{t+1}(s)$. Thus the amateur politician will generically not be indifferent and will prefer to deviate and play $\hat{d}(h_{t+1}(s))$. Therefore it must be that for all t and all $h_t(s)$, $\sigma_e(h_t(s), s_t = 1) = \sigma_e(h_t(s), s_t = 0)$. But this is perfect pooling on one record.

The same result holds if we allow the possibility that the voter mixes, using different conditional appointment probabilities for different decision histories of the politician. As before, suppose that there is a t and a signal realization $\tilde{h}_t(s)$ such that $\sigma_e(\tilde{h}_t(s), s_t = 1) \neq \sigma_e(\tilde{h}_t(s), s_t = 0)$. Then the amateur must be mixing over $d_t = 1$ and $d_t = 0$ in decision t. We established before that if the voter appoints the politician with probability one when indifferent, the amateur politician has incentives to deviate, and play the decision consistent with his no-career concerns strategy $\hat{\sigma}_a$, say $\hat{d}_t$. So it must be that there are (at least) two classes of records, say D_1 and D_2, with associated probabilities of appointment ϕ_1 and $\phi_2 \neq \phi_1$. Now suppose without loss of generality that the decision history up to period t was consistent with D_1, and that $\hat{\sigma}(h_{t+1}(s)) = 1$. If both $d_t = 1$ and $d_t = 0$ are consistent with D_1, then the amateur politician does not have incentives to mix (this is just as in the case in which the voter always appoints when indifferent). So it must be that $(h_t(d), d_t = 1)$ is consistent with D_1 and $(h_t(d), d_t = 0)$ is consistent with D_2, and that ϕ_1 and ϕ_2 are such that the amateur politician is indifferent between $d_t = 1$ and $d_t = 0$ at t. But this is not possible. Suppose in fact that ϕ_1 and ϕ_2 are such that given a history of signals $h_{t+1}(s)$, the amateur politician is indifferent between $d_t = 1$ and $d_t = 0$ at t. Then the amateur will not be willing to mix for a history of signals $h'_t(s) = h_t(s)$ for all $t \neq i$ and $s'_i = 1 - s'_i$. We conclude that it must be that for all t and all $h_t(s)$, $\sigma_e(h_t(s), s_t = 1) = \sigma_e(h_t(s), s_t = 0)$. ■

Lemma 3 *Consider an expert politician that can observe the state perfectly, $\theta = 1$, and suppose $\overline{\mu} > b$. The expected payoff from following the no career concerns strategy at any point t is $\hat{W}_{\theta=1}(t) = 0$. The (average) expected payoff from following the pooling equilibrium strategy from period t on, given beliefs $\Pr(\omega_t = 1|h_{t+1}(s);\theta) = \beta_\theta$ is*

$$W^{\theta}_{PL}(t, \beta_\theta) = \delta - (1-\delta)b\left\{\left(\frac{T-t}{T}\right)(1-\overline{\mu}) + \frac{1}{T}(\overline{\mu} - \beta_\theta)\left[\frac{1-\kappa^{T-t+1}}{1+\kappa}\right]\right\},$$

where $\kappa = \gamma_0 + \gamma_1 - 1$, and $\overline{\mu} = \frac{1-\gamma_0}{2-\gamma_0-\gamma_1}$. Thus, there exists an equilibrium with pooling if and only if

$$\delta \geq \frac{\left(\frac{T-1}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]}{1/b + \left(\frac{T-1}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]} \equiv \underline{\delta}(1),$$

Moreover, for any $\theta \in (1/2, 1]$, $\underline{\delta}(\theta) \geq \tilde{\delta}$, where

$$\tilde{\delta} \equiv \frac{\left(\frac{T-1}{T}\right)(b-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]}{1/b + \left(\frac{T-1}{T}\right)(b-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right]}$$

Proof of Lemma 3. The best pooling equilibrium for the politician is to pool in

$$\tilde{d}^{PL} = \begin{cases} (1,1,\ldots,1) & \text{if } \overline{\mu} > b \\ (0,0,\ldots,0) & \text{if } \overline{\mu} < b. \end{cases}$$

So suppose for concreteness that $\overline{\mu} > b$, so that $\tilde{d}^{PL} = (1,1,\ldots,1)$. In period t, politician has a belief $\Pr(\omega_t = 1|h_{t+1}(s);\theta)$. Thus, if in period t, politician has a belief β_θ, his equilibrium payoff in the best pooling equilibrium is

$$
\begin{aligned}
W^{\theta}_{PL}(t,\beta_\theta) &= \delta T + (1-\delta)\sum_{m=t}^{T} E[u(1,\omega_m)|h_{t+1}(s),\theta] \\
&= \delta T - (1-\delta)\sum_{m=t}^{T} \Pr(\omega_m = 0|h_{t+1}(s),\theta)b.
\end{aligned} \tag{7}
$$

Now, iterating,

$$
\Pr(\omega_{t+n} = 1|h_{t+1}(s),\theta) = \beta_\theta(\gamma_1+\gamma_0-1)^n + (1-\gamma_0)\left[\frac{1-(\gamma_1+\gamma_0-1)^s}{2-\gamma_1-\gamma_0}\right]
$$

Then

$$
W^{\theta}_{PL}(t,\beta_\theta) = \delta T - (1-\delta)\sum_{m=t}^{T}\left\{1-\beta_\theta(\gamma_1+\gamma_0-1)^{m-t} - (1-\gamma_0)\left[\frac{1-(\gamma_1+\gamma_0-1)^{m-t}}{2-\gamma_1-\gamma_0}\right]\right\} \tag{8}
$$

Now consider the payoff for a type θ politician of deviating and playing $\hat{\sigma}_\theta$. This is

$$
\hat{W}_\theta(t,\beta_\theta) = -(1-\delta)\sum_{m=t}^{T}\sum_{h_{m+1}(s)} \Pr(h_{m+1}(s)|h_{t+1}(s)) \left[\begin{array}{c} \Pr(\omega_m = 1|h_{m+1}(s),\theta)(1-\hat{\sigma}_{\theta,m}(h_{m+1}(s)))(1-b) \\ \Pr(\omega_m = 0|h_{m+1}(s),\theta)\hat{\sigma}_{\theta,m}(h_{m+1}(s))b \end{array}\right] \tag{9}
$$

Note that for $\theta = 1$, $\hat{W}_{\theta=1}(t,\beta_\theta) = 0$. Then for a $\theta = 1$ not to deviate from the pooling eq., (write $\kappa = \gamma_1+\gamma_0-1$, and recall that $\overline{\mu} = \frac{1-\gamma_0}{2-\gamma_1-\gamma_0}$)

$$
\begin{aligned}
\delta &\geq (1-\delta)b\frac{1}{T}\sum_{m=t}^{T}\left\{1-\beta_\theta\kappa^{m-t} - (1-\gamma_0)\left[\frac{1-\kappa^{m-t}}{1+\kappa}\right]\right\} \\
&= (1-\delta)b\frac{1}{T}\left\{(T-t)(1-\overline{\mu}) + (\overline{\mu}-\beta_\theta)\sum_{m=t}^{T}\kappa^{m-t}\right\} \\
&= (1-\delta)b\left\{\left(\frac{T-t}{T}\right)(1-\overline{\mu}) + \frac{1}{T}(\overline{\mu}-\beta_\theta)\left[\frac{1-\kappa^{T-t+1}}{1+\kappa}\right]\right\}.
\end{aligned} \tag{10}
$$

Since this has to hold for all β_θ (for all possible histories $h_t(s)$), then

$$
\delta \geq (1-\delta)b\left\{\left(\frac{T-t}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^{T-t+1}}{1+\kappa}\right]\right\}.
$$

And since this has to hold for all $t = 1, \ldots, T-1$, then in particular for $t = 1$ (where it binds), so

$$\delta \geq (1-\delta)b \left\{ \left(\frac{T-1}{T}\right)(1-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right] \right\}.$$

Now consider an arbitrary $\theta \in (1/2, 1]$. Recall that

$$\hat{\sigma}_{\theta,m}(h_{m+1}(s)) = \begin{cases} 1 & \text{if } \Pr(\omega_m = 1 | h_{m+1}(s), \theta) \geq b \\ 0 & \text{if } \Pr(\omega_m = 1 | h_{m+1}(s), \theta) < b. \end{cases}$$

Call the term in brackets in (9) Z. Then when $\Pr(\omega_m = 1 | h_{m+1}(s), \theta) \geq b$, $Z = \Pr(\omega_m = 0 | h_{m+1}(s), \theta) b < (1-b)b$, and when $\Pr(\omega_m = 1 | h_{m+1}(s), \theta) < b$, $Z = \Pr(\omega_m = 1 | h_{m+1}(s), \theta)(1-b) < (1-b)b$. Thus $Z < (1-b)b$. But then

$$\hat{W}_\theta(t, \beta_\theta) > -(1-\delta)b(1-b)(T-t)$$

Then cannot have pooling for whatever expert this is if this is larger than the pooling payoff , so here we need

$$\delta > (1-\delta)b \left\{ \left(\frac{T-t}{T}\right)(b-\overline{\mu}) + \frac{1}{T}(\overline{\mu} - \beta_\theta)\left[\frac{1-\kappa^{T-t+1}}{1+\kappa}\right] \right\},$$

and since this has to hold for all (feasible) β_θ (for all feasible histories $h_t(s)$) then so with $\beta_\theta = 0$ (approx), and $t = 1$, which gives

$$\delta \geq (1-\delta)b \left\{ \left(\frac{T-1}{T}\right)(b-\overline{\mu}) + \frac{1}{T}\overline{\mu}\left[\frac{1-\kappa^T}{1+\kappa}\right] \right\},$$

■

Proof of Theorem 2. Consider first the case $e \geq \hat{\theta}(\gamma_0, \gamma_1)$. Here $p_t^e \in (\underline{p}, \overline{p})$ for all t, and thus according to $\hat{\sigma}_e$, $d_t^e = s_t$ for all t. It follows, in particular, that in equilibrium d_t^e is independent of the history of signals up to period t, $h_t(s)$.

Suppose that the Principal could observe the realization of the state, $h_{T+1}(\omega) = \{\omega_1, \ldots, \omega_T\}$. Let $\mathcal{D}_\omega \equiv \{d_t : \omega_t = \omega\}$, with $|\mathcal{D}_1| = U$, and $|\mathcal{D}_0| = B$. Then relabel these so that the first element of $\mathcal{S}_\omega$ is d_1, the second d_2, and so on. Then the observations $d_t \in \mathcal{D}_1$ are i.i.d. with $\Pr(d_t = 1 | d_t \in \mathcal{D}_1) = \theta$, and the observations $d_t \in$

$\mathcal{D}_0$ are i.i.d. with $\Pr(d_t = 1|d_t \in \mathcal{D}_0) = 1 - \theta$. Thus by the LLN, $\frac{1}{U}\sum_{d_t \in \mathcal{D}_1} d_t \to_p \theta$, and $\frac{1}{B}\sum_{d_t \in \mathcal{D}_0} d_t \to_p 1 - \theta$.

Now define the sets $\hat{\mathcal{D}}_0$ and $\hat{\mathcal{D}}_1$ by $\hat{\mathcal{D}}_\omega \equiv \{d_t : z_t = \omega\}$. Intuitively, $\hat{\mathcal{D}}_\omega$ is the set of observations in periods t that are classified as consistent with a realization $\omega_t = \omega$ given only the information in z_t. Call the set of observations $d_t \in \hat{\mathcal{D}}_\omega$ that are correctly classified $C_\omega \equiv \{d_t \in \hat{\mathcal{D}}_\omega : z_t = \omega = \omega_t\}$, and the set of observations $d_t \in \hat{\mathcal{D}}_\omega$ that are incorrectly classified $I_\omega \equiv \{d_t \in \hat{\mathcal{D}}_\omega : z_t = \omega \neq \omega_t\}$. Now, focusing on $\hat{\mathcal{D}}_1$,

$$\frac{1}{|\hat{\mathcal{D}}_1|}\sum_{d_t \in \hat{\mathcal{D}}_1} d_t = \frac{|C_1|}{|\hat{\mathcal{D}}_1|}\left[\frac{1}{|C_1|}\sum_{d_t \in C_1} d_t\right] + \frac{|I_1|}{|\hat{\mathcal{S}}_1|}\left[\frac{1}{I_1}\sum_{d_t \in I_1} d_t\right]$$

As before, here $\frac{1}{|C_1|}\sum_{d_t \in C_1} d_t \to_p \theta$, and $\frac{1}{|I_1|}\sum_{d_t \in I_1} d_t \to_p 1 - \theta$. Moreover, $\lim_{t\to\infty} \Pr(z_t = 1|\omega_t = 1)\Pr(\omega_t = 1) = \mu q$, so that in the long-run distribution the z's are draws from almost identical distributions. Then $|C_1|/|\hat{\mathcal{D}}_1| \to_p q$. Thus $\frac{1}{|\hat{\mathcal{D}}_1|}\sum_{d_t \in \hat{\mathcal{D}}_1} d_t \to_p q\theta + (1-q)(1-\theta) \equiv \phi(\theta)$. Similarly, $\frac{1}{|\hat{\mathcal{D}}_0|}\sum_{d_t \in \hat{\mathcal{D}}_0} d_t \to_p \phi(\theta)$. In other words, if the record $h_T(d)$ is generated by an agent of type θ with choice strategy $\hat{\sigma}_\theta$, then for any $\varepsilon > 0$ and $\eta > 0$ there is a T^* such that if $T > T^*$, then $\forall \omega \in \{0, 1\}$.

$$\Pr\left(\phi(\theta) - \varepsilon \leq \frac{1}{|\hat{\mathcal{D}}_\omega|}\sum_{d_t \in \hat{\mathcal{D}}_\omega} d_t \leq \phi(\theta) + \varepsilon |(\theta, \hat{\sigma}_\theta)\right) > 1 - \eta. \qquad (11)$$

Define for any $\varepsilon > 0$ the sets $A_e(\varepsilon) \equiv (\phi(e) - \varepsilon, \phi(e) + \varepsilon)$ and $A_a(\varepsilon) \equiv (\phi(a) - \varepsilon, \phi(a) + \varepsilon)$, and let $\mathcal{A}_\theta(\varepsilon) \equiv \left\{h_T(d) : \frac{1}{|\hat{\mathcal{D}}_\omega|}\sum_{d_t \in \hat{\mathcal{D}}_\omega} d_t \in A_\theta(\varepsilon)\ \forall \omega \in \{0, 1\}\right\}$. It follows from our previous argument that for any positive ε, η there is a T^* such that if $T > T^*$, then $\Pr(\mathcal{A}_\theta(\varepsilon)|\theta, \hat{\sigma}_\theta) > 1 - \eta$. Moreover, for small enough ε, these sets are disjoint. Furthermore, for small ε, $\lim_{T\to\infty} \Pr(\mathcal{A}_\theta(\varepsilon)|B_k) = 0$.

So consider a record $h_T(d) \in \mathcal{A}_e(\varepsilon)$. Then, writing $m(d) \equiv \Pr(h_T(d)|\mathcal{A}_e(\varepsilon))$ for simplicity

$$\Pr(\theta = e|h_T(d)) > \frac{m(d)(1-\eta)(1-\pi)\lambda}{m(d)[(1-\eta)(1-\pi) + \eta\pi]\lambda + (1-\lambda)\sum_k \Pr(h_T(d)|\theta = B_k)} \equiv 1 - \alpha,$$

which can be made arbitrarily close to 1 for large enough T. (The inequality comes from (11), and the following: we know the amateur is going to be in $\mathcal{A}_a(\varepsilon)$

with probability at least $1-\eta$, so he will be everywhere else with probability at most η, and in particular, will be in $\mathcal{A}_e(\varepsilon)$ with probability strictly less than η.) Now, because for $T > T^*$ we have $\Pr\left(d : \frac{1}{|\hat{\mathcal{D}}_\omega|}\sum_{d_t \in \hat{\mathcal{D}}_\omega} d_t \in A_\theta | \theta, \hat{\sigma}_\theta\right) > 1-\eta$, it follows that when the decision record is generated by an expert following choice strategy $\hat{\sigma}_e$, then for any $\alpha > 0$, $\eta > 0$ there exists $\overline{T}$ such that if $T > \overline{T}$ then

$$\Pr\left[\Pr(\theta = e | h_T(d)) > 1-\alpha | e, \hat{\sigma}_e\right] > 1-\eta,$$

or equivalently, $\Pr\left(\Pr(\theta = e|h_T(d))|e, \hat{\sigma}_e\right) \rightarrow_p 1$.

By the same reasoning, if $h_T(d) \in \mathcal{A}_a(\varepsilon)$, then $\Pr\left(\Pr(\theta = a|h_T(d))|a, \hat{\sigma}_a\right) \rightarrow_p 1$. Thus in equilibrium with large T, the voter appoints the expert politician and doesn't appoint the amateur politician with probability close to one.

Next consider a deviation by an agent of type θ to a choice strategy $\sigma_\theta \neq \hat{\sigma}_\theta$. Note that any deviation from $\hat{\sigma}_\theta$ is costly at the decision-making stage. Thus, if it is to be profitable, it must induce a higher probability of getting hired. Now, by definition, $\hat{\sigma}_\theta$ is the optimal choice strategy with no career concerns, and therefore for any $\sigma_\theta \neq \hat{\sigma}_\theta$ and $\omega \in \{0,1\}$ is must be that $\Pr(d_t = \omega | \omega_t = \omega, h_t(d), h_t(s), \hat{\sigma}_\theta) = \Pr(s_t = \omega | \omega_t = \omega, \hat{\sigma}_\theta) \geq \Pr(d_t = \omega | \omega_t = \omega, h_t(s), h_t(d), \sigma_\theta)$ for all $t, h_t(s), h_t(d)$, with strict inequality for some $t, h_t(s), h_t(d)$. Therefore, for $\theta = a, e$,

$$\Pr\left(h_T(d) \notin \mathcal{A}_\theta(\varepsilon) | \theta, \sigma_\theta\right) > \Pr\left(h_T(d) \notin \mathcal{A}_\theta(\varepsilon) | \theta, \hat{\sigma}_\theta\right),$$

and (for $\theta = a$)

$$\Pr\left(h_T(d) \notin \mathcal{A}_e(\varepsilon) | a, \sigma_a\right) > \Pr\left(h_T(d) \notin \mathcal{A}_e(\varepsilon) | a, \hat{\sigma}_a\right).$$

But because for large enough T $\Pr\left(d : \frac{1}{|\hat{\mathcal{D}}_1|}\sum_{d_t \in \hat{\mathcal{D}}_1} d_t \in A_\theta | \theta, \hat{\sigma}_\theta\right) > 1-\eta$ for η arbitrarily small, it follows that for any record $h_T(d)$ such that $\frac{1}{|\hat{\mathcal{D}}_1|}\sum_{d_t \in \hat{\mathcal{D}}_1} d_t \notin A_\theta$ $\Pr\left(d | \theta, \hat{\sigma}_\theta\right) < \eta$ for η arbitrarily small, and therefore that for any such record $h_T(d)$

$$\Pr\left(\Pr(\theta \in B | h_T(d)) | a, \sigma_a\right) \rightarrow_p 1,$$

where as before $B = \cup_k B_k$.

Because deviations to $\sigma_a \neq \hat{\sigma}_a$ are costly in terms of task A, and put higher probability on these decision histories $h_T(d) \notin A_e(\varepsilon) \cup A_a(\varepsilon)$ and do not increase the probability of decision histories in $A_e(\varepsilon)$, the amateur doesn't have a profitable deviation. Differently to the amateur, the expert could feasibly deviate to a σ_e generating decision histories consistent with $A_a(\varepsilon)$ with high probability, but this only decreases his payoffs.

We now drop the assumption of $e \geq \hat{\theta}(\gamma_0, \gamma_1)$, and consider the general case. When e is relatively large, $p_t^e \in (\underline{p}(e,b), \overline{p}(e,b))$ for all t, and thus according to $\hat{\sigma}_e$, $d_t^e = s_t$ for all t. As a result,

$$\frac{1}{|C_1|} \sum_{d_t \in C_1} d_t \to_p \theta, \quad \text{and} \quad \frac{1}{|C_0|} \sum_{d_t \in C_0} d_t \to_p 1 - \theta.$$

In general, however, for both the expert and the amateur, $p_t^\theta \notin (\underline{p}(e,b), \overline{p}(e,b))$, and thus according to $\hat{\sigma}_e$, it is possible that $d_t^e = 1$ when $s_t = 0$ (if $p_t^e > \overline{p}(e,b)$), or $d_t^e = 0$ when $s_t = 1$ (if $p_t^e < \overline{p}(e,b)$). This implies that the probability of $d_t = 1$ for $(\theta, \hat{\sigma}_\theta)$ now depends on the realization of p_t^θ, which in general will be different for different decisions, depending on the realization of $h_T(\omega)$ and then $h_T(s)$. As we show, however (see (i) below), it is still the case that $\mu(e) \equiv \lim_{t\to\infty} \Pr(d_t^e = 1|\omega_t = 1, \hat{\sigma}_e) > \lim_{t\to\infty} \Pr(d_t^a = 1|\omega_t = 1, \hat{\sigma}_a) \equiv \mu(a) > 1/2$. Furthermore, while the observations d_t are drawn from distributions that are not i.i.d., these differences average out in the long run distribution (this is conceptually equivalent to drawing a random sample from C_1), so that it is still the case that $1/|C_1| \sum_{d_t \in C_1} d_t \to_p \mu(\theta)$. We argue this more formally below, in (ii). With these amenpoliticianents, the previous proof then extends to the general case.

We begin by showing (i) that $\mu(e) \equiv \lim_{t\to\infty} \Pr(d_t^e = 1|\omega_t = 1, \hat{\sigma}_e) > \lim_{t\to\infty} \Pr(d_t^a = 1|\omega_t = 1, \hat{\sigma}_a) \equiv \mu(a) > 1/2$. To do this, consider any realization of the process ω_t, and define the random variable

$$X_t = \begin{cases} 1 & \text{if } d_t^e - d_t^a \geq 0 \\ 0 & \text{if } d_t^e - d_t^a < 0. \end{cases}$$

Given a particular history of states, equation (3) describes the evolution of p_t^θ for each type of agent, and the random variable X_t records the (random) events in which the decision of the expert would differ from the decision of the amateur. We want to show that $\mathbb{E}(X_t|\omega_t = 1, \hat{\sigma}_\theta) > 1/2$ for large t. If the latter inequality is true, it follows that in the long period $\Pr(d_t^e \geq d_t^a|\omega_t = 1, \hat{\sigma}_e, \hat{\sigma}_a) > 1/2$ and i) must follow, i.e., when $\omega_t = 1$, the expert choose $d_t = 1$ with higher probability than the amateur. We proceed under the assumption that t is large, so that for all s close to t, $\Pr(\omega_s = 1)$ is close to $(1-\gamma_0)/(2-(\gamma_1+\gamma_0))$. Let $m_t^\theta \equiv \{p_t^\theta : p_t^\theta \in (\underline{p}(\theta, b), \overline{p}(\theta, b))\}$, $u_t^\theta \equiv \{p_t^\theta : p_t^\theta > \overline{p}(\theta, b)\}$, and $l_t^\theta \equiv \{p_t^\theta : p_t^\theta < \underline{p}(\theta, b)\}$. Then we have that

$$E(X_t|\omega_t = 1, \hat{\sigma}_\theta) = \Pr(u_t^e|\omega_t = 1) + \Pr(m_t^e, u_t^a|\omega_t = 1)e + \Pr(m_t^e, m_t^a|\omega_t = 1)(1-(1-e)a) +$$

$$+ \Pr(m_t^e, l_t^a|\omega_t = 1) + \Pr(l_t^e, m_t^a|\omega_t = 1)(1-a) + \Pr(l_t^e, l_t^a|\omega_t = 1).$$

Since $1-(1-e)a > e$, it follows that

$$\Pr(m_t^e, u_t^a|\omega_t = 1)e + \Pr(m_t^e, m_t^a|\omega_t = 1)(1-(1-e)a) + \Pr(m_t^e, l_t^a|\omega_t = 1) > \Pr(m_t^e|\omega_t = 1)e,$$

and therefore

$$\begin{aligned} E(X_t|\omega_t = 1, \hat{\sigma}) &> \Pr(u_t^e|\omega_t = 1) + \Pr(m_t^e|\omega_t = 1)e + \Pr(l_t^e, m_t^a|\omega_t = 1)(1-a) + \Pr(l_t^e, l_t^a|\omega_t = 1) \\ &> \Pr(u_t^e|\omega_t = 1) + \Pr(m_t^e|\omega_t = 1)e + (\Pr(l_t^e|\omega_t = 1) - \Pr(l_t^e, u_t^a|\omega_t = 1))(1-a) \\ &> \Pr(u_t^e|\omega_t = 1) + \Pr(m_t^e|\omega_t = 1)e + \Pr(l_t^e|\omega_t = 1)(1 - \Pr(u_t^a|\omega_t = 1))(1-a), \end{aligned}$$

where the second inequality follows from $\Pr(l_t^e, m_t^a|\omega_t = 1)(1-a) + \Pr(l_t^e, l_t^a|\omega_t = 1) > (\Pr(l_t^e|\omega_t = 1) - \Pr(l_t^e, u_t^a|\omega_t = 1))(1-a)$, and in the last inequality we use the fact that since $p_t^e(\cdot|\omega_t = 1)$ and $p_t^a(\cdot|\omega_t = 1)$ are affiliated, $\Pr(l_t^e, u_t^a|\omega_t = 1) < \Pr(l_t^e|\omega_t = 1)\Pr(u_t^a|\omega_t = 1)$. But now

$$\Pr(u_t^e|\omega_t = 1) + \Pr(m_t^e|\omega_t = 1)e + \Pr(l_t^e|\omega_t = 1)(1 - \Pr(u_t^a|\omega_t = 1))(1-a) >$$

$$\Pr(u_t^e|\omega_t = 1)\frac{1}{2} + \Pr(m_t^e|\omega_t = 1)e + \Pr(l_t^e|\omega_t = 1)\left(\frac{1}{2} + (1 - \Pr(u_t^a|\omega_t = 1))(1-a)\right)$$

where we used the fact that $\Pr(u_t^e|\omega_t = 1)$ is larger than $\Pr(l_t^e|\omega_t = 1)$. But since the last inequality above is a convex combination of numbers bigger than 1/2, it must be the case that $E(X_t|\omega_t = 1, \hat{\sigma}) > 1/2$.

Now consider (ii). Note that according to $\hat{\sigma}_\theta$, $d_t^\theta = 1$ iff either $p_t^\theta > \overline{p}(\theta, b)$ or $p_t^\theta \in (\underline{p}(\theta, b), \overline{p}(\theta, b))$ and $s_t = 1$. Now, we can show (see (3), (4)) that

$$p^\theta_{t+1} = \begin{cases} \frac{\theta p^\theta_t}{\theta p^\theta_t + (1-\theta)(1-p^\theta_t)}(\gamma_1+\gamma_0-1)+(1-\gamma_0) & \text{w.p. } \theta p^\theta_t + (1-\theta)(1-p^\theta_t) \\ \frac{(1-\theta)p^\theta_t}{1-[\theta p^\theta_t + (1-\theta)(1-p^\theta_t)]}(\gamma_1+\gamma_0-1)+(1-\gamma_0) & \text{w.p. } 1-[\theta p^\theta_t + (1-\theta)(1-p^\theta_t)]. \end{cases}$$

Then

$$E[p^\theta_{t+1}|p^\theta_t] = (1-\gamma_0) + (\gamma_1+\gamma_0-1)p^\theta_t.$$

Hence it follows that $E[p^\theta_{t+1}|p^\theta_t] < p^\theta_t$ if and only if $p^\theta_t > \frac{1-\gamma_0}{2-(\gamma_0+\gamma_1)} \equiv \overline{\mu} \equiv \lim_{t\to\infty}\Pr(\omega_t = 1)$. Thus the politician's beliefs fluctuate around the long term probability that $\omega_t = 1$, $\overline{\mu}$. This implies that the process p^θ_t is a martingale if and only if $\theta = 1/2$. However, when $\theta > 1/2$, the process $y^\theta_t \equiv p^\theta_t - \overline{\mu}$ satisfies $E[y^\theta_{t+1}|y^\theta_t] < y^\theta_t$ for all t, and is thus a supermartingale. To see this, note that

$$y^\theta_{t+1} = \begin{cases} \left(\frac{\theta\left(y^\theta_t+\overline{\mu}\right)}{\theta\left(y^\theta_t+\overline{\mu}\right)+(1-\theta)\left(1-\left(y^\theta_t+\overline{\mu}\right)\right)} - \overline{\mu}\right)(\gamma_1+\gamma_0-1) & \text{w.p. } \theta\left(y^\theta_t+\overline{\mu}\right)+(1-\theta)\left(1-\left(y^\theta_t+\overline{\mu}\right)\right) \\ \left(\frac{(1-\theta)\left(y^\theta_t+\overline{\mu}\right)}{1-\left[\theta\left(y^\theta_t+\overline{\mu}\right)+(1-\theta)\left(1-\left(y^\theta_t+\overline{\mu}\right)\right)\right]} - \overline{\mu}\right)(\gamma_1+\gamma_0-1) & \text{otherwise.} \end{cases}$$

It follows that

$$E[y^\theta_{t+1}|y^\theta_t] = \left(y^\theta_t+\overline{\mu}\right)(\gamma_1+\gamma_0-1) - \overline{\mu}(\gamma_1+\gamma_0-1) = (\gamma_1+\gamma_0-1)y^\theta_t < y^\theta_t.$$

Since $E[|y^\theta_t|] \in [l(\theta;\gamma_1,\gamma_0)-\overline{\mu}, u(\theta;\gamma_1,\gamma_0)-\overline{\mu}]$ is always bounded, by Doob's Forward Convergence Theorem, y^θ_t converges almost surely to a random variable Z. It follows that p^θ_t converges to a random variable X with c.d.f. F. We can then apply the strong law of large numbers to X to conclude that the empirical distribution function $\hat{F}_t(a) = \frac{1}{T}\sum_{t=1}^T 1\{p^\theta_t < a\}$ converges to $F(a)$ as $T\to\infty$ almost surely, for every value of a. This then implies that

$$\frac{1}{|C_1|}\sum_{d_t\in C_1} d_t \to_p [1-F(\theta)] + [F(\theta)-F(1-\theta)]\theta = \mu(\theta).$$

This completes the proof. ■

Proof of Proposition 1. Note that because $\theta > \hat{\theta}(\gamma_0, \gamma_1)$, $\phi_1(\theta) \equiv \phi_1(\theta, \hat{\sigma}_\theta) = \lim_{t\to\infty} \Pr(s_{t+1} = 0|s_t = 1)$. Now,

$$\begin{aligned}
\Pr(s_{t+1} = 1|s_t = 1) &= \sum_{\omega_{t+1}=0,1} \Pr(s_{t+1} = 1|\omega_{t+1}) \Pr(\omega_{t+1}|s_t = 1) \\
&= \sum_{\omega_{t+1}=0,1} \Pr(s_{t+1} = 1|\omega_{t+1}) \left[\sum_{\omega_t=0,1} \Pr(\omega_{t+1}|\omega_t) \Pr(\omega_t|s_t = 1)\right] \\
&= [\theta(1-\gamma_0) + (1-\theta)\gamma_0] + \left(\frac{\theta\mu_t}{\theta\mu_t + (1-\theta)(1-\mu_t)}\right)(\gamma_0 + \gamma_1 - 1)(2\theta - 1),
\end{aligned} \tag{12}$$

where $\mu_t = \Pr(\omega_t = 1)$. Recall $\lim_{t\to} \mu_t = (1-\gamma_0)/(2-(\gamma_0+\gamma_1)) \equiv \overline{\mu}$. Let $q \equiv \theta\overline{\mu}/(\theta\overline{\mu} + (1-\theta)(1-\overline{\mu})) = \theta(1-\gamma_0)/(\theta(1-\gamma_0) + (1-\theta)(1-\gamma_1)) \in (0, 1)$. Then

$$\phi_1(\theta) = 1 - [\theta(1-\gamma_0) + (1-\theta)\gamma_0] - q(\gamma_0 + \gamma_1 - 1)(2\theta - 1)$$

Note that $q|_{\theta=1/2} = \overline{\mu} > 1/2$ iff $\gamma_1 > \gamma_0$ and

$$\frac{\partial q}{\partial \theta} = \frac{(1-\gamma_0)(1-\gamma_1)}{(\theta(1-\gamma_0) + (1-\theta)(1-\gamma_1))^2} = q^2 \frac{1-\gamma_1}{\theta^2(1-\gamma_0)} = \frac{q(1-q)}{\theta(1-\theta)} > 0.$$

Therefore $q > 1/2$ whenever $\gamma_1 > \gamma_0$. Furthermore,

$$\frac{\partial q}{\partial \gamma_1} = \frac{\theta(1-\theta)(1-\gamma_0)}{(\theta(1-\gamma_0) + (1-\theta)(1-\gamma_1))^2} = q^2 \frac{1-\theta}{\theta(1-\gamma_0)} = \frac{q(1-q)}{1-\gamma_1} > 0,$$

and

$$\frac{\partial q}{\partial \gamma_0} = -\frac{\theta(1-\theta)(1-\gamma_1)}{(\theta(1-\gamma_0) + (1-\theta)(1-\gamma_1))^2} = -q^2 \frac{(1-\theta)(1-\gamma_1)}{\theta(1-\gamma_0)^2} = -\frac{q(1-q)}{1-\gamma_0} < 0.$$

Using these results we can conclude that

$$\frac{\partial \phi_1(\theta)}{\partial \gamma_1} = -(2\theta-1)\left((\gamma_0 + \gamma_1 - 1)\frac{\partial q}{\partial \gamma_1} + q\right) = -\frac{q(2\theta-1)}{1-\gamma_1}(\gamma_0(1-q) + q(1-\gamma_1)) < 0,$$

and

$$\frac{\partial \phi_1(\theta)}{\partial \gamma_0} = -(2\theta-1)\left(-(1-q) + (\gamma_0 + \gamma_1 - 1)\frac{\partial q}{\partial \gamma_0}\right) = +\frac{(1-q)(2\theta-1)}{1-\gamma_0}((1-\gamma_0)(1-q) + q\gamma_1) > 0$$

This establishes part (i). For part (ii), note that

$$\frac{\partial \phi_1(\theta)}{\partial \theta} = (2\gamma_0 - 1) - (\gamma_0 + \gamma_1 - 1)\left(2q + (2\theta - 1)\frac{\partial q}{\partial \theta}\right),$$

and notice that when $\gamma_1 \geq \gamma_0$ we have that

$$\frac{\partial \phi_1(\theta)}{\partial \theta} < 2\gamma_0 - 1 - \gamma_0 - \gamma_1 + 1 = \gamma_0 - \gamma_1 \leq 0,$$

where the first inequality follows from the fact that $\gamma_1 \geq \gamma_0$ implies $q > 1/2$. ■

Proof of Proposition 2. Part (i). Suppose without loss of generality that $\gamma_1 > \gamma_0$. Then we want to show that better politicians are more consistent on active decisions. Because $e > a > \hat{\theta}(\gamma_0, \gamma_1)$, then for all types θ, $\hat{\sigma}_\theta$ implies $d_t = s_t$. Then

$$\Gamma^j_{\ell'}(\theta, \sigma_\theta) = \lim_{t\to\infty} \prod_{\ell=1}^{\ell'} \Pr(s_{t+\ell} = 1 | s_t = \ldots = s_{t+\ell-1} = 1) \Pr(s_t = 1).$$

Furthermore, $\lim_{t\to\infty} \Pr(s_t = 1) = \overline{\mu}\theta + (1 - \overline{\mu})(1 - \theta)$, and $C_{\ell'} \equiv \lim_{t\to\infty} \Pr(s_{t+\ell'} = 1 | s_t = \ldots = s_{t+\ell'-1} = 1)$ is

$$\begin{aligned} C_{\ell'} &= \lim_{t\to\infty} \sum_{\omega_{t+\ell'}} \Pr(s_{t+\ell'} = 1 | \omega_{t+\ell'}) \Pr(\omega_{t+\ell'} | s_t = \ldots = s_{t+\ell'-1} = 1) \\ &= (1 - \theta) + (2\theta - 1) \lim_{t\to\infty} \Pr(\omega_{t+\ell'} = 1 | s_t = \ldots = s_{t+\ell'-1} = 1). \end{aligned} \tag{13}$$

From (3), $x_{\ell'} \equiv \lim_{t\to\infty} \Pr(\omega_{t+\ell'} = 1 | s_t = \ldots = s_{t+\ell'-1} = 1)$ is given

$$x_{t+1} = \frac{\theta x_t}{\theta x_t + (1 - \theta)(1 - x_t)}(\gamma_1 + \gamma_0 - 1) + (1 - \gamma_0), \tag{14}$$

evaluated at $t + 1 = \ell'$, with initial condition $x_t = \overline{\mu}$. Now, from (13),

$$\frac{\partial C_{\ell'}}{\partial \theta} = (2x_{\ell'} - 1) + (2\theta - 1)\frac{\partial x_{\ell'}}{\partial \theta}$$

Because $x_t = \overline{\mu} > 1/2$ if and only if $\gamma_1 > \gamma_0$ and (14) defines an increasing sequence, it follows that $x_{\ell'} > 1/2$. Therefore $\partial C_{\ell'}/\partial \theta > 0$ whenever $\partial x_{\ell'}/\partial \theta > 0$. This in turn follows from the fact that (i) $x_t = \overline{\mu}$ is increasing in θ when $\gamma_1 > \gamma_0$ and that (ii) each step in (14) is increasing in θ. Since each term $C_{\ell'}$ is increasing in θ and

$\lim_{t\to\infty}\Pr(s_t = 1) = \overline{\mu}\theta + (1-\overline{\mu})(1-\theta)$ is increasing in θ when $\gamma_1 > \gamma_0$, it follows that

$$\frac{\Gamma^j_{\ell'}(\theta,\sigma_\theta)}{\Gamma^j_{\ell'}(\theta',\sigma_{\theta'})} > \frac{\Gamma^j_{\ell}(\theta,\sigma_\theta)}{\Gamma^j_{\ell}(\theta',\sigma_{\theta'})} > 1$$

for all $\ell' > \ell > 1$.

Part (ii). First, note that p^θ_t is bounded above by a number $u(\theta) \in (1/2, 1)$, and bounded below by $l(\theta) \in (0, 1/2)$. To see this, note that the process p^θ_t must be below the deterministic sequence given by (3) with $s_t = 1$ for all t, and above the deterministic sequence given by (3) with $s_t = 0$ for all t. First consider the upper sequence. Because this upper sequence is increasing and bounded, it converges. To compute the limit, solve

$$p^\theta_{t+1} = \frac{\theta p^\theta_t}{\theta p^\theta_t + (1-\theta)(1-p^\theta_t)}(\gamma_1+\gamma_0-1) + (1-\gamma_0)$$

for $p^\theta_{t+1} = p^\theta_t = u(\theta;\gamma_0,\gamma_1)$. Note that $u(\theta;\gamma_0,\gamma_1)$ solves

$$W(u) \equiv u(\theta u + (1-\theta)(1-u)) - \theta u(\gamma_1+\gamma_0-1) - (1-\gamma_0)(\theta u + (1-\theta)(1-u)) = 0$$

Since $W(u)$ is convex, $W(0) = -(1-\gamma_0)(1-\theta) < 0$ and

$$W\left(\frac{1-\gamma_0}{2-(\gamma_1+\gamma_0)}\right) = -\frac{(1-\gamma_0)(1-\gamma_1)(\gamma_1+\gamma_0-1)(2\theta-1)}{(2-(\gamma_1+\gamma_0))^2} < 0,$$
$$W(\gamma_1) = (1-\theta)(1-\gamma_1)(\gamma_1+\gamma_0-1) > 0,$$

and $\gamma_1 > (1-\gamma_0)/(2-(\gamma_1+\gamma_0))$, it follows that $u(\theta;\gamma_0,\gamma_1) \in ((1-\gamma_0)/(2-(\gamma_1+\gamma_0)), \gamma_1)$ exists and it is unique. Furthermore, when $\gamma_1 \geq \gamma_0$, we have that

$$\frac{\partial u(\theta;\gamma_0,\gamma_1)}{\partial\theta} = -\frac{2u^2 - u(2+\gamma_1-\gamma_0) + (1-\gamma_0)}{2\theta u + (1-\theta)(1-2u) - \theta(\gamma_1+\gamma_0-1) - (1-\gamma_0)(2\theta-1)} > 0.$$

To sign the derivative, notice that the denominator is positive since it is the derivative of $W(u)$ evaluated at $u(\theta;\gamma_0,\gamma_1)$, i.e., where $W(u)$ is increasing. As for the numerator, when $\gamma_1 \geq \gamma_0$ it is always increasing in u, and hence $2u^2 - u(2+\gamma_1-\gamma_0) + (1-\gamma_0) < 2\gamma_1^2 - \gamma_1(2+\gamma_1-\gamma_0) + (1-\gamma_0) = -(\gamma_1+\gamma_0-1)(1-\gamma_0) < 0$. Furthermore,

$$\frac{\partial^2 u(\theta;\gamma_0,\gamma_1)}{\partial^2\theta} = -\frac{\frac{\partial^2 W(u)}{\partial^2\theta}\left(\frac{\partial W(u)}{\partial u}\right)^2 - 2\frac{\partial^2 W(u)}{\partial\theta\partial u}\frac{\partial W(u)}{\partial u}\frac{\partial W(u)}{\partial\theta} + \frac{\partial^2 W(u)}{\partial^2 u}\left(\frac{\partial W(u)}{\partial\theta}\right)^2}{\left(\frac{\partial W(u)}{\partial u}\right)^3} =$$

$$\frac{\partial^2 u(\theta;\gamma_0,\gamma_1)}{\partial^2\theta} = -2\frac{-(4u-(2+\gamma_1-\gamma_0))\frac{\partial W(u)}{\partial u}+(2\theta-1)\frac{\partial W(u)}{\partial \theta}}{(2\theta u+(1-\theta)(1-2u)-\theta(\gamma_1+\gamma_0-1)-(1-\gamma_0)(2\theta-1))^3}\frac{\partial W(u)}{\partial\theta} < 0.$$

To sign the derivative, notice that as we show above $\partial W(u)/\partial u$ is positive, the denominator is positive, $\partial W(u)/\partial\theta$ is negative, and

$$4u-(2+\gamma_1-\gamma_0) > 4\frac{1-\gamma_0}{2-(\gamma_1+\gamma_0)}-(2+\gamma_1-\gamma_0) = \frac{\gamma_1^2-\gamma_0^2}{2-(\gamma_1+\gamma_0)} > 0 \text{ iff } \gamma_1 > \gamma_0.$$

We also have that

$$\frac{\partial u(\theta;\gamma_0,\gamma_1)}{\partial\gamma_1} = -\frac{-\theta u}{2\theta u+(1-\theta)(1-2u)-\theta(\gamma_1+\gamma_0-1)-(1-\gamma_0)(2\theta-1)} > 0,$$

and

$$\frac{\partial u(\theta;\gamma_0,\gamma_1)}{\partial\gamma_0} = -\frac{(1-\theta)(1-u)}{2\theta u+(1-\theta)(1-2u)-\theta(\gamma_1+\gamma_0-1)-(1-\gamma_0)(2\theta-1)} < 0.$$

Recall that $\overline{p}(1/2,b)=b$, $\overline{p}(1,b)=1$, and notice that $\overline{p}(\theta,b)$ is always increasing in θ, and convex (concave) in θ if and only if $b<1/2 (b>1/2)$. Since $1>u(1;\cdot)=\gamma_1>(1-\gamma_0)/(2-(\gamma_1+\gamma_0))=u(1/2;\cdot)$, it follows that if $u(1/2;\cdot)=\overline{\mu}>b$ then $u(\theta;\gamma_0,\gamma_1)$ and $\overline{p}(\theta,b)$ must cross an odd number of times. By letting $\overline{\theta}(\gamma_0,\gamma_1)$ the largest value of θ such that $u(\theta;\gamma_0,\gamma_1)=\overline{p}(\theta,b)$, it follows that if $\theta>\overline{\theta}(\gamma_0,\gamma_1)$, $p_t^\theta<\overline{p}(\theta,b)$ for all t. Notice that $\overline{\theta}(\gamma_0,\gamma_1)$ is increasing in γ_1 and decreasing in γ_0. Furthermore, if $b\leq 1/2$, $\overline{\theta}(\gamma_0,\gamma_1)$ must be unique.

Consider now the lower sequence. Because this lower sequence is decreasing and bounded, it converges. To compute the limit, solve

$$p_{t+1}^\theta = \frac{(1-\theta)p_t^\theta}{(1-\theta)p_t^\theta+\theta(1-p_t^\theta)}(\gamma_1+\gamma_0-1)+(1-\gamma_0)$$

for $p_{t+1}^\theta=p_t^\theta=l(\theta;\gamma_0,\gamma_1)$. Note that $l(\theta;\gamma_0,\gamma_1)$ solves

$$K(l)\equiv l((1-\theta)l+\theta(1-l))-(1-\theta)l(\gamma_1+\gamma_0-1)-(1-\gamma_0)((1-\theta)l+\theta(1-l))=0$$

Since $K(u)$ is concave, $K(1)=(1-\theta)(1-\gamma_1)>0$ and

$$K(1-\gamma_0) = -(1-\theta)(1-\gamma_0)(\gamma_1+\gamma_0-1) < 0,$$

$$K\left(\frac{1-\gamma_0}{2-(\gamma_1+\gamma_0)}\right) = \frac{(1-\gamma_0)(1-\gamma_1)(\gamma_1+\gamma_0-1)(2\theta-1)}{(2-(\gamma_1+\gamma_0))^2} > 0,$$

and $(1-\gamma_0)/(2-(\gamma_1+\gamma_0)) > 1-\gamma_0$, it follows that $l(\theta;\gamma_0,\gamma_1) \in (1-\gamma_0,(1-\gamma_0)/(2-(\gamma_1+\gamma_0)))$ exists and it is unique. Furthermore, when $\gamma_1 \geq \gamma_0$, we have that

$$\frac{\partial l(\theta;\gamma_0,\gamma_1)}{\partial \theta} = -\frac{-2l^2 + l(2+\gamma_1-\gamma_0) - (1-\gamma_0)}{2(1-\theta)l + \theta(1-2l) - (1-\theta)(\gamma_1+\gamma_0-1) + (1-\gamma_0)(2\theta-1)} < 0.$$

To sign the derivative, notice that the denominator is positive since it is the derivative of $K(l)$ evaluated at $l(\theta;\gamma_0,\gamma_1)$, i.e., where $K(l)$ is increasing. As for the numerator, when $\gamma_1 \geq \gamma_0$ it is easy to check that is concave in l, positive at the boundaries, and hence always positive. Since $l(\theta;\gamma_0,\gamma_1)$ is decreasing in θ, $l(1/2;\gamma_0,\gamma_1) = u(1/2;\gamma_0,\gamma_1) = \overline{\mu}$, when $b < \overline{\mu}$, there exist a unique $\underline{\theta}(\gamma_0,\gamma_1)$ such that if $\theta < \underline{\theta}(\gamma_0,\gamma_1)$, $l(\theta;\gamma_0,\gamma_1) > \overline{p}(\theta,b)$. But this implies that whenever $a < \underline{\theta} < \overline{\theta} < e$ the amateur is more consistent than then expert on j decisions. Since

$$\frac{\partial l(\theta;\gamma_0,\gamma_1)}{\partial \gamma_1} = -\frac{-(1-\theta)l}{2(1-\theta)l + \theta(1-2l) - (1-\theta)(\gamma_1+\gamma_0-1) + (1-\gamma_0)(2\theta-1)} > 0,$$

and

$$\frac{\partial l(\theta;\gamma_0,\gamma_1)}{\partial \gamma_0} = -\frac{\theta(1-l)}{2(1-\theta)l + \theta(1-2l) - (1-\theta)(\gamma_1+\gamma_0-1) + (1-\gamma_0)(2\theta-1)} < 0,$$

it also follows that $\underline{\theta}(\gamma_0,\gamma_1)$ is increasing in γ_1 and decreasing in γ_0. ■

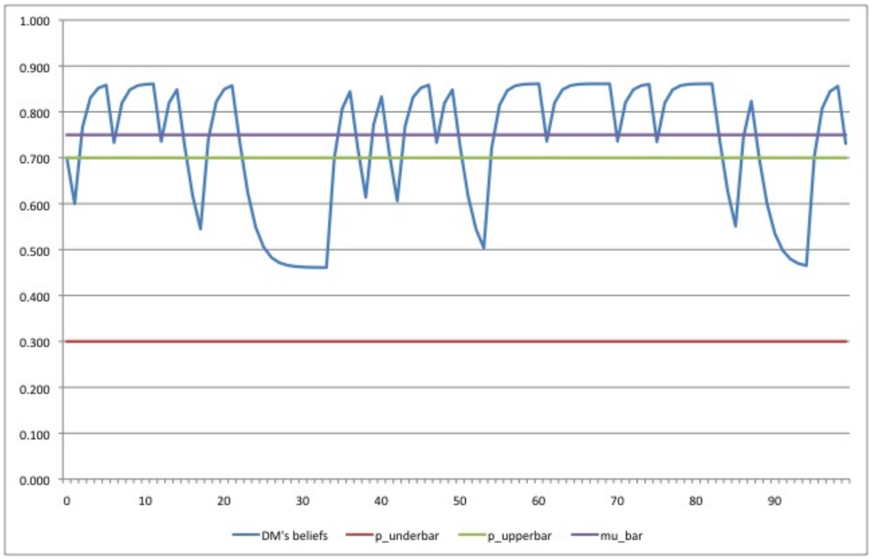

Figure 1: A possible path of the politician's t-period prior beliefs, together with $\overline{\mu}$ (upper purple line), $\overline{p}$ (mid green line) and $\underline{p}$ (lower red line). In this example, $\gamma_0 = 0.7$, $\gamma_1 = 0.9$, and $b = 1/2$.

References

ALESINA, A., AND G. TABELLINI (2007): "Bureaucrats or Politicians? Part I: A Single Policy Task," *The American Economic Review*, 97(1), 169–179.

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

BANKS, J., AND R. SUNDARAM (1998): "Optimal Retention in Agency Problems," *Journal of Economic Theory*, 82(2), 293–323.

BARRO, R. (1973): "The Control of Politicians: an Economic Model," *Public Choice*, 14(1), 19–42.

BENABOU, R., AND G. LAROQUE (1992): "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*, 107(3), 921–958.

BERNHARDT, D. (1995): "Strategic Promotion and Compensation," *Review of Economic Studies*, 62(2), 315–339.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

CANES-WRONE, B., M. HERRON, AND K. SHOTTS (2001): "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, 45(3), 532–550.

CANES-WRONE, B., AND K. SHOTTS (2007): "When Do Elections Encourage Ideological Rigidity?," *American Political Science Review*, 101(2), 273–288.

CRAWFORD, V. P., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

FEREJOHN, J. (1986): "Incumbent Performance and Electoral Control," *Public Choice*, 50, 5–25.

FOX, J., AND R. VAN-WEELDEN (2011): "Costly Transparency," Forthcoming, *Journal of Public Economics.*

KALAI, E., AND E. LEHRER (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61(5), 1019–1045.

LEVY, G. (2005): "Careerist Judges and the Appeals Process," *The Rand Journal of Economics*, 36(2), 275–297.

Li, W. (2007): "Changing One's Mind When the Facts Change: Incentives of Experts and the Design of Reporting Protocols," *The Review of Economic Studies*, 74(4), 1175–1194.

Maskin, E., and J. Tirole (2004): "The politician and the judge: Accountability in government," *American Economic Review*, 94, 1034–1054.

Meyer, M. A. (1991): "Learning from Coarse Information: Biased Contests and Career Profiles," *The Review of Economic Studies*, 58(1), 15–41.

Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

Moscarini, G., M. Ottaviani, and L. Smith (1998): "Social Learning in a Changing World," *Economic Theory*, 11(3), 657–665.

Olszewski, W., and A. Sandroni (2008): "Manipulability of Future-Independent Tests," *Econometrica*, 76(6), 1437–1466.

——— (2009): *Annals of Statistics*37(2), 1013–1039.

Ottaviani, M., and P. Sorensen (2006a): "Professional Advice," *Journal of Economic Theory*, 126(1), 120–142.

——— (2006b): "Reputational Cheap Talk," *Rand Journal of Economics*, 37(1), 155–175.

Prat, A. (2005): "The Wrong Kind of Transparency," *American Economic Review*, 95(3), 862–877.

Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104(6), 1105–1134.

Sobel, J. (1985): "A Theory of Credibility," *Review of Economic Studies*, 52(4), 557–573.