

Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits

Philippe Aghion and Matthew O. Jackson *

June 2014

Abstract

In this paper we analyze the problem of whether and/or when to replace a leader (agent) when no monetary rewards are available, and it is the leader's competence rather than effort that is being evaluated. The only decisions that the leader takes over time are whether to undertake risky but potentially high payoff projects, the choice of which can reveal the leader's competency. If the value of foregone projects are observed, then the probability that a leader is replaced is bell-shaped and saw-toothed over time. If the value of foregone projects are not observed, and the leader's competency is only indirectly inferrable through the success or failure of projects that the leader undertakes, then the incentives of the leader depend on the replacement strategy. If the principal can commit to a replacement strategy in advance, then we show that (approximately) optimal mechanisms either involve a probationary period and then indefinite tenure, or else a random dismissal strategy. If instead commitment is impossible, and for instance voters regularly choose whether to replace the leader, then there are poor incentives and inefficiently low payoffs, even below that of simply replacing the leader in every period. Incentives can be improved via term limits.

Keywords: Information, tenure, dismissal, replacement, elections, term limits, leadership, discretion

JEL Classification Codes: D72, D82, D86, M12, C72

*Aghion is at the Department of Economics, Harvard University, Cambridge, Massachusetts USA, and is a member of CIFAR. Jackson is at the Department of Economics, Stanford University, Stanford, California 94305-6072 USA, and is an external faculty member at the Santa Fe Institute and a member of CIFAR. Emails: paghion@fas.harvard.edu and jacksonm@stanford.edu. We gratefully acknowledge financial support from ARO MURI award No. W911NF-12-1-0509. We thank Daron Acemoglu, Tim Besley, Renee Bowen, Andres Drenik, Stephen Nei, Alessandro Pavan, Torsten Persson, Larry Samuelson, Guido Tabellini, Jean Tirole, and seminar participants at CIFAR, Northwestern, Stanford, INSEAD and the IIES at Stockholm University, for helpful comments and suggestions.

If General McClellan isn't going to use his army, I'd like to borrow it for a time. *Abraham Lincoln, Jan 10, 1862, before relieving George B. McClellan of command (for the first time).*

1 Introduction

How can we motivate individuals to take risky decisions in situations where there is limited scope for using adjustable monetary compensation and the only lever providing incentives is a threat of dismissal? Examples of such situations include the relationship between political leaders and voters, as well as many employees in systems with fixed wages (including academics, bureaucrats, and some government employees), and even wealthy top level managers for whom monetary incentives are secondary to reputation or power within a position.

These examples are not only such that firing/replacement is the main incentive device, but also that the incentive issue is not so much one of inducing more effort as much as one of inducing the agent to use her discretion to take risky decisions. Thus Abraham Lincoln's main concern with General McClellan was not so much that McClellan was not putting enough effort, but rather the fact that McClellan was missing opportunities to act, meanwhile the Confederate Army was dangerously close to Washington. Nor was the lack of (higher) monetary incentives the reason for McClellan's reluctance to act. Eventually, Lincoln replaced McClellan in order to prompt action.

Similarly, the primary concern in motivating a high-ranking politician is not getting them to work long hours, as they tend to already be driven to do so, but instead to motivate them to make the "right" decisions. In many settings, the consequences of some decisions can reveal how competent the decision-maker is at identifying good decisions from bad ones, and this leads to an important but relatively under-studied set of incentive issues.¹ In this paper, we investigate performance incentives in such environments, providing new insights into contracts such as tenure and term limits, as well as the timing of replacement decisions, and the effects on the resulting behavior of decision makers, abstracting from standard moral-hazard effort considerations, and instead focusing on incentives regarding which decision to make.

Our model is one where a principal can hire a new agent, henceforth called the "leader", in any period at some fixed cost. The leader can be either competent or incompetent, and whether or not she is competent is initially unknown to her and to the principal. Moreover, the leader has discretion over a choice of actions. She can choose to take a "conservative" action which yields a sure payoff (normalized to zero) but does not reveal anything about her level of competence. Or she can take a "risky" action, which could lead to a positive

¹There are some important papers that consider competence and incentives, with the closest in terms of the basic issue of motivating a decision maker being that of Scharfstein and Stein (1990), and some of the literature that followed. However, they consider herding behavior among fund managers and so the ultimate context and analysis is quite different from ours here.

or negative payoff. An incompetent leader receives uninformative signals about whether the risky action is likely to lead to positive or negative payoffs, whereas a competent leader receives informative signals about whether the risky action is more likely to lead to positive or negative payoffs. Each day there is a new state of nature, new information, and a new choice of action to be made, thus over time a leader's competence can be learned by tracking the payoffs on the days on which she took the risky action. Based on that information, the principal decides each period whether to keep or replace the current leader.

We focus attention on settings in which the leader does not respond to monetary incentives and only draws benefits from being kept on the job. In this model in which wage incentives have no bite, firing/replacement becomes the main incentive instrument, but it involves subtle trade-offs. On the one hand, not replacing a leader may result into getting stuck with an incompetent leader. On the other hand, a threat of replacing a leader may induce leaders to avoid risky decisions that reveal competence. The question then to design the sequence of replacement decisions over time in a way that incentivizes the leader to take appropriate decisions, but still results in learning and the replacement of incompetent leaders. We examine this question in three different contexts.

We first consider the case where the leader's information and the eventual state are eventually publicly observed: so even if the conservative decision is taken, the principal sees what would have happened if the risky decision had been taken. In this case the model delivers three main predictions. First, the probability that the principal replace the current leader follows either an increasing, or a decreasing, or bell-shaped sawtooth pattern. The intuition is that early on it is worth keeping the current leader on the job, even if she takes the risky action unsuccessfully, as there is an option value of acquiring more information about her level of competence; but then, as time goes on and the number of failed attempts increases, so does the probability of dismissing the current leader; but eventually, leaders that have survived longer are increasingly more likely to be competent, and therefore the dismissal probability will start to decrease over time. The bell-shaped pattern is more likely to be observed when actions are not very revealing, so that learning is slow; however, when actions are very revealing, learning is fast and therefore in this case the replacement probability is essentially decreasing over time. ²

We then move to the case where the signals are privately observed by the leader - and therefore the state is only indirectly observable through the leader's actions. If the leader takes the conservative action, nothing is learned, while if the leader takes the risky action some information is revealed. In this setting, the dismissal threat can lead to inaction, as the leader is reluctant to choose the risky action and thereby risk future dismissal. How to motivate the leader in this case?

²The bell-shaped pattern appears to be in line with evidence on the dynamics of job separation, for example as shown by Farber (1999) for the US labor market, and or more recently by Jia et al (2014) when looking at the patterns of dismissals and promotions of provincial leaders in China. In this case, we also get the surprising finding that the cumulative probabilities of replacing incompetent leaders, as well as mistakenly replacing competent leaders, are non-monotonic in the leader's actual degree of competence.

Here the results depend very much on whether the principal can commit in advance to a specific mechanism for replacing the leader, or else acts based on her beliefs in every period. We first analyze the case of commitment. We argue that there are essentially two main mechanisms that emerge as optimal ways to motivate the agent, namely the "carrot" and the "stick". The "carrot" consists in granting tenure to the agent after a finite number of periods (sometimes immediately). The "stick" consists in replacing the leader in any period with positive probability if she chooses the safe action in the previous period, while at the same time not committing to any kind of tenure arrangement. The cost of choosing the carrot, i.e. the tenure strategy, is that the principal might get stuck forever with an incompetent leader. The cost of choosing the stick, i.e. the threat of dismissal strategy, is that you may end up firing a competent leader for taking the safe action, even when it was the right thing to do. We show that the stick strategy dominates when the replacement cost is sufficiently low, otherwise the carrot strategy dominates. More generally, we show that it is optimal to use the stick (dismissal) strategy in earlier periods and the carrot (tenure) strategy in later periods, and that the lower the discount rate the later should tenure occur.

Finally, we consider the case in which the principal cannot commit in advance to keeping or replacing the leader in future periods. A leading example of such a situation is that of an elective democracy, where voters (the principal) choose whether or not to replace the leader after each period, without any precommitment. We abstract from voter bias issues to concentrate on voters' decisions of whether to replace the leader as a function of the information voters receive over time about the leader's level of competence. A main result in that case is that every (Markov Perfect) equilibrium results in a negative net payoff from hiring a new leader, in particular a payoff that is lower than simply replacing the leader in every period. The reason is that an incumbent leader's best response in the no-commitment case involves taking the safe action when they are faced with reelection probabilities: indeed, doing so maintains voters' belief about the leader's competence at its initial level; this, together with the fact that replacing the current leader by a new leader is costly but leads to the same probability of competence, implies that the current leader will not be fired. On the other hand, if current leaders are sure to be replaced then they are free to take risky decisions. We also prove that something similar is true of a class of non-Markov equilibria: no matter how successful the leader has been in the past the leader eventually stops taking risky actions and is replaced. We then show that an optimal term limit can substantially improve the leader's incentives to take the risky action. The optimal term limit in turn depends on the speed of learning and on the cost of replacement.

The paper relates to several strands of literature. First, there is the literature on job matching and turnover (see Farber (1999) for a survey of that literature). Most closely related to our paper in that literature is Jovanovic (1979). Jovanovic develops a continuous time model of job matching between an employer and an employee, where the quality of the match is initially unknown to all but progressively revealed over time through observing successive output realizations. Jovanovic assumes that the worker responds to monetary incentives and is paid her expected output each period. Under this assumption, the relevant

decision lies with the worker, and it to choose between staying on the job or quitting. The main findings are that the expected wage increases with job tenure and that the probability of quitting conditional upon having remained on the job follows a bell-shaped pattern over time.

We contribute to this first literature by first characterizing the equilibrium dismissal pattern in discrete time under symmetric information for various parameter configurations, seeing more precisely when the bell-shaped pattern emerges. But more importantly, we extend the analysis by moving from symmetric to asymmetric information situations.

Our paper also relates to a literature on tenure and "up-or-out" contracts (e.g see Kahn and Huberman (1988), Carmichael (1988), Waldman (1990), or Burdett and Coles (2003)). A main argument in this literature is that tenure promotion serves as a commitment device for the principal not to underreport the agent's value ex post (if the agent's value was truly low then the up-and-out contract allows the principal to simply fire the worker), which in turn preserves the worker's ex ante incentives to provide effort. In Carmichael (1988), granting academic tenure to current faculty helps ensure that good potential candidates to become new faculty are not dismissed by current faculty because the former would represent a threat for the latter. We provide a completely different perspective on tenure based on three important features that such contracts provide: incentives to take risky actions during the probationary period, ability to sort competent from incompetent, and incentives to take appropriate actions during the tenure phase.

The paper also contributes to a political economy literature on term limits, and the more general literature on career concerns. In particular, Besley and Case (1995) (see also Alt, Bueno de Mesquita and Rose (2011)) build on Holmstrom (1982)'s career concerns model, and more specifically on Banks and Sundaram's (1993) model and analysis, to argue and verify empirically that allowing for reelection improves political leaders' incentives.³ These models are built upon standard signalling structures: agents have costs of effort that are decreasing in their type (competency). The equilibria that they focus upon (there can be many) are such that more competent agents have incentives to work harder given the greater marginal payoff to their effort. Agents are retained as long as they are successful and fired once they fail. In that context, offering longer term limits (more chances for reelection) increases incentives to put in effort - especially in early periods.⁴

³For other related signalling-based career concerns models, see Dewatripont, Jewitt, and Tirole (1999 a and b) and Alesina and Tabellini (2007, 2008). More generally, a literature on reputations (Scharfstein and Stein (1990), Allen and Gorton (1993), Tirole (1996), Tadelis (1999), Taylor (2000), Mailath and Samuelson (2001,2006)) couples actions and some hidden type with outcomes.

⁴More recently, Smart and Sturm (2013) developed a model where incumbent politicians may be either "public-spirited" (i.e with payoffs that coincide with voters' payoff) or "biased" towards a particular choice of action. By reducing her expected payoff from reelection, term limits reduces a "public-spirited" politician's incentives to otherwise deviate from efficient decision making in order to signal her type so as to increase her probability of reelection. Like in our model, seeking reelection induces inefficient decision making by current leaders, and term limits help overcome this problem. However, once again, our analysis does not rely on any standard kind of incentive problem (be it moral hazard or signaling) and moreover in our model politicians

We depart from these two literatures by focusing on situations where discretionary risk-taking, not effort, is the main issue at stake. Our main question is: how can one induce leaders to seize decision opportunities that are socially desirable but yet exposes them to a higher risk of losing their job as this reveals information about their competency of making choices. In that respect, a more closely related paper is Canes-Wrone, Herron, and Shotts (2001). Canes-Wrone, Herron, and Shotts (2001) develop a model where, like in ours, leaders can have high-quality or high-quality signals about the true state of nature, and worry that being found out as incompetent they will be voted out of office. In their (two-period) model, the fear of being found out as being incompetent may lead incumbent leaders to sometimes "pander" to their current voters' beliefs.⁵

A main contribution of our work, in this light, is to examine how repeated elections compare to the case of commitment and mechanisms. If a voter/principal can commit to a particular sequence of evaluations, and actors are sufficiently patient, then competent leaders can be identified, kept, and appropriately incentivized. In contrast, if the voter/principal cannot commit and repeatedly evaluates leaders, then incentives unravel and leaders inefficiently follow safe actions. Repeated elections biases leaders towards sitting on their laurels and taking safe actions: much as the fear of failure that might have led to McLellan's inaction and Lincoln's frustration in finding a general who would act in the early parts of the U.S. Civil War. In our setting, tenure contracts or term limits can alleviate fear and provide leaders with incentives to make choices that they would not make under the repeated microscope of retention.

2 Basic model

2.1 The players

An organization is operated by an "agent" whom we often refer to as a "leader", for reasons that will become clear. A "principal" decides on keeping or replacing the leader.

We focus on settings in which the leader is paid a fixed wage and so the only relevant decision is whether to keep or replace the leader.

A leader is either "competent" or "incompetent", denoted by *Comp* and *Incomp*. The prior probability that the leader is competent is $\lambda_0 \in (0, 1)$.

A given leader's type does not change over time. If a leader is replaced, then the new leader is competent with probability λ_0 . This is also the prior about the initial leader's competence with which the principal begins at the start of period 1: in other words, there is no asymmetric information ex ante about the leader's level of competence.

are not a priori biased against voters' preferences.

⁵See also Rodriguez-Barraquer and Tan (2014), who examine students' choices of fields to work on, which then can reveal their abilities. Their model is quite different from ours and focuses on herding effects in competitive settings, and not on dynamic incentives.

2.2 Time, states, and signals

Time proceeds in discrete periods $t \in \{1, 2, \dots, T\}$. On occasion, we a case in which $T = 2$ or $T = 3$ for illustrative purposes, but then look at the infinite period case $T = \infty$ for general results.

In each period a state of nature, $\omega_t \in \{X, Y\}$, is realized. States X and Y occur with equal probability, independently across periods.

In the beginning of each period t , the leader sees a signal $s_t \in \{X, Y\}$ that may be informative about the state of nature. If the leader is *competent* then the probability that the signal is equal to the state is $p > 1/2$, and thus the signal is informative. If the leader is *incompetent* then the two signals $s_t = X$ and $s_t = Y$ are equally likely in both states X and Y , and so an incompetent leader's signals are completely uninformative.

Thus,

$$\Pr(s_t = \omega_t | \text{Comp}) = p > 1/2$$

$$\Pr(s_t = \omega_t | \text{Incomp}) = 1/2.$$

2.3 Actions and replacement costs

We can think of the state being whether there some sort of “opportunity” (e.g., profitable investment) available: in state Y there is such an opportunity and in state X there is not. If the state is Y then it is best to invest: choosing action y then pays 1 to the principal, while in state X it is better not to invest and then choosing y results in a loss of value, $-v$, to the principal. Not investing (choosing action x) always leads to a payoff of 0 to the principal. Thus, the principal's payoff from the action as a function of the state in any period is:

	X	Y
x	0	0
y	$-v$	1

Given that $v > 1$, if there is no information about the state, then action x offers a higher expected payoff than action y , and so it is only competent leaders whom the principal would ever want to take the action y .

Thus, we can think of y as ‘taking action’, while x can be interpreted as not acting or sticking with a status quo. Foregoing actions is a safe alternative that might provide no information of the value of the foregone opportunities.

Throughout what follows, we presume that $p - v(1 - p) > 0$. If this were violated, even a competent leader could never lead to a positive expected payoff from action y , and so action x is the only action that should every be taken.

2.4 The order of moves

The sequence of moves within each period is as follows. At the beginning of the period the current leader sees her signal. Next, the leader takes an action x or y . Subsequently, the principal sees the the payoff and updates his beliefs about the leader's competence. We also consider a case in which the principal sees the state along with the payoff (we make the information scenarios more explicit below). At the end of the period the principal decides whether to keep the current leader or replace her by a new leader.

The system repeats itself with a new draw of signal and state at each time, but the type of any given leader remains fixed over time. Once a leader is replaced that leader never returns to the game. Replacing the leader leads to a cost of $c \geq 0$ in that period for the principal.

2.5 Payoffs

All players are expected payoff maximizers and discount time with the same discount factor, δ , such that $0 \leq \delta \leq 1$.

Payoffs are as follows: (i) A leader gets private benefits b per period that she is on the job,⁶ and (ii) The principal gets the per-period payoffs from the matrix above as a function of the action taken and the state, less any costs of replacing leaders.

2.6 Two informational scenarios

We consider two informational scenarios. In one case the principal learns the state and signal regardless of the action taken by the leader, while in the other case the principal only observes payoffs and not the state or signal.⁷

Scenario 1: Observed States, Signals, and Payoffs

The state, signal, and payoff are observed by both parties at the end of each period regardless of which action was taken in that period.

A t -period history is thus a sequence

$$h^t = (\omega_1, s_1, a_1, d_1 ; \dots ; \omega_t, s_t, a_t, d_t),$$

where ω_t is the state, s_t is the leader's signal, $a_t \in \{x, y\}$ is the action that the leader took, and $d_t \in \{K, R\}$ indicates whether the leader was *Kept* or *Replaced* at the end of the period by the principal.

Let H be the set of all finite histories.

⁶This isolates the incentive problem. We could also allow the leader to prefer to make successful decisions without changing the main content of the results, as long as the payoff from being in office was large enough.

⁷One could also consider an intermediate case in which only signals and payoffs are observed by both parties. The incentives in that setting are similar to the complete information setting, and lead to little additional insight and so we omit it.

Scenario 2: Observed Payoffs

The leader privately observes the signal and then chooses an action. The principal only observes his payoff. Thus, if the leader chooses x then the principal's payoff is 0 and the principal does not learn the true state nor anything about the leader's signal. If the leader chooses y , then the principal's payoff is either 1 or $-v$ and therefore he can infer the state from observing his payoff. Neither the leader nor principal sees the state except via inference from the realized outcome of the action.

In this scenario, the histories differ for the leader and the principal, as they observed different things.

A t -period history for the principal is a sequence

$$h_P^t = (u_1, d_1 ; \dots ; u_t, d_t),$$

where $u_t \in \{0, -v, 1\}$ is the payoff and $d_t \in \{K, R\}$ indicates whether the leader was *Kept* or *Replaced* at the end of the period by the principal.

Here the history for the leader also includes the signals she observed, but only since that leader was in office. So, for a leader that was in place since time τ , a history is then

$$h_L^t = (u_1, d_1 ; \dots, u_{\tau-1}, R ; s_\tau, u_\tau, d_\tau ; \dots ; s_t, u_t, d_t).$$

So, the latest leader has the same information as the principal about past leaders and then additional information about her own signals during her own reign from time τ through t .⁸

So, a history in this second scenario is a pair $h^t = (h_P^t, h_L^t)$ of related histories, and again we let H^t denote the set of possible histories.

We presume that a new leader begins with a prior λ_0 on her own competence, although as will become clear this is largely inconsequential in what follows.

2.7 Some useful preliminaries

2.7.1 Posteriors and updating

A useful expression is the posterior that a leader is competent if the leader is known to have had m correct signals out of n total signals given a prior of λ_0 , denoted $\Lambda(m, n, \lambda_0)$. This is:

$$\Lambda(m, n, \lambda_0) = \frac{\lambda_0 p^m (1-p)^{n-m}}{\lambda_0 p^m (1-p)^{n-m} + (1-\lambda_0)/2^n} \quad (1)$$

Note that $\Lambda(m, n, \lambda_0)$ is increasing in m and decreasing in n .

Generally, we let λ_t denote the posterior belief of the principal about the competence of the current leader after t periods of observations, which moves randomly over time depending on the leader's type and the actions and realized states. This will also be dependent upon the leader's equilibrium strategy.

⁸Note that u_t fully reveals the action a_t , so we do not need to add that to the histories.

A relevant expression is the probability that the leader's signal is correct when the leader sees a Y (or equivalently an X), conditional on a current posterior probability λ_t :

$$f(\lambda_t) = \Pr(\omega = Y | s = Y, \lambda_t) = \frac{\frac{1}{2}(\lambda_t p + (1 - \lambda_t)/2)}{\frac{1}{2}(\lambda_t p + (1 - \lambda_t)/2) + \frac{1}{2}(\lambda_t(1 - p) + (1 - \lambda_t)/2)}. \quad (2)$$

or

$$f(\lambda_t) = \lambda_t p + (1 - \lambda_t)/2.$$

A useful benchmark is the principal's belief, λ , for which the expected current payoff resulting from the leader taking action y when the leader sees signal Y is just equal to zero, namely:

$$f(\lambda_t) - (1 - f(\lambda_t))v = 0,$$

which solves into:

$$\lambda^* = \frac{v/(1 + v) - 1/2}{p - 1/2},$$

which decreases in p and increases in v .

2.7.2 The principal's expected payoff

The following lemma is direct but useful.

LEMMA 1 *Let λ_t denote the principal's belief at the end of some period t about the current leader's competence. The principal's expected payoff for the next period of the leader taking an action that matches the signal is:*

$$u(\lambda_t) = \frac{1}{2}[f(\lambda_t) - (1 - f(\lambda_t))v].$$

The lemma follows from noting that with probability 1/2 the signal is Y and that playing y when the signal is Y yields expected payoff $[f(\lambda_t) - (1 - f(\lambda_t))v]$ to the principal.⁹

3 Public information: an increasing, decreasing or bell-shaped replacement pattern?

We begin with the case where the state and signal are publicly observed after each period. The most interesting incentives issues arise in Scenario 2, thus we consider Scenario 1 as a benchmark, mainly to understand some of the basics of what a replacement profile would look like without any incentive issues.

Given this symmetric information, we examine a situation in which the leaders simply follow the signal in each period. This is easily enforced simply by using a mechanism in

⁹The two states are equally probable and the probability of a leader getting the correct signal is the same no matter the state.

which a leader who does not follow signals is fired immediately. Thus, we simply examine the principal’s optimal choice of how long to keep any given leader. This is a variation on a standard “bandit problem”, and the optimal strategy for the principal can be expressed via a simple cut-off belief such that the leader is replaced at the end of period t whenever the posterior belief on her competence is lower than the cutoff.

Here, the “typical” replacement probability is bell-shaped over time, i.e., first increasing and then eventually decreasing. The intuition is follows. Given that replacing the leader is costly, unless signals are extremely accurate it will not be optimal for the principal to replace the leader immediately. In other words, there is an initial ”honeymoon” period where the leader is not replaced as there is some initial number and fraction of failures before the leader’s competence could begin to be revealed. This is for the short run. As for the very long run, if the leader survives for a long enough time, by the law of large numbers she is very likely to be competent, in which case she is unlikely to be replaced. Hence, in the very long run, the replacement probability also becomes small. It is in the middle range where the substantial replacement probability falls, as enough information to identify competence with some confidence has accumulated. Overall, we thus expect a bell-shaped replacement probability of the incumbent leader over time.

We suppose that $u(\lambda_0) - c > 0$. This guarantees that it is better to get a new leader than to keep a leader who is thought sufficiently incompetent that the principal would rather them not even try to take action y even with a good signal.

Let $P(t)$ be the probability that a principal replaces a leader at time t (and kept the leader in all periods before t). An optimal strategy for the principal is to retain the leader as long as $\lambda_t \geq \lambda$ for some $0 < \lambda \leq \lambda_0$. We now show that the replacement probability follows a sawtooth pattern for any threshold strategy (optimal or not).

PROPOSITION 1 *Suppose that the principal starts with some prior λ_0 and retains the leader at the end of period t if and only if $\lambda_t \geq \lambda$ for some $0 < \lambda \leq \lambda_0$. There exists $t > 1$ such that $P(t) > 0$ while $P(t + 1) = 0$ and $P(t + k) > 0$ for some $k \geq 1$.*

The sawtooth pattern suggested in the proposition is illustrated in Figure 1. That $P(t)$ is strictly positive in period 7 but zero in period 8 can be explained as follows. A leader who is fired in period 7 but not before necessarily had 7 failures in a row over the first seven periods. This follows since a leader is not fired in any previous period regardless of the number of failures, including six straight failures. So consider a leader who survived until period 8. Such a leader could not have had eight failures and no success as in that case she would already been fired at the end of period 7. Nor could she have had failures in all first seven periods followed by one success in period 8 because once again she would have been fired at the end of period 7. Thus the only possibility is that she had at most six failures and at least one success over the first seven periods. But then even if she had another failure in period 8 this is better than having had seven failures in a row over the first seven periods,




Figure 1: The probability of replacing a leader (for the first time) in various periods for $p = .55$, $\lambda_0 = 1/2$ and $\lambda(c) = 1/3$.

which was the threshold for replacing the leader, and having seven failures and one success is closer to the posterior of having six failures and no successes than seven failures and no success.¹⁰ This in turn implies that a leader who survived until period 8 will not be fired in period 8. In a nutshell: *a success over the first periods buys the current leader some additional "grace period" where she is not fired.*

We can also examine just the positive probability dates as pictured in Figure 2. The fact that successes and failures come in integers leads the curves to be non-monotonic even when we look only at dates with positive probabilities.

We also see some interesting comparisons between situations where competent leaders are "barely competent" so that $p = .55$ and so hard to tell apart from incompetent leaders, compared to situations where competent leaders are "highly competent" so that $p = .95$ and very different from incompetent leaders. In the left-hand panel where competent leaders are

¹⁰Although our reasoning is particular to our discrete time setting, the sawtooth pattern will not fully disappear if we move to continuous time. In the continuous time case, after any success there will still be periods of time during which the leader is kept for sure. Again, a leader who makes it past some particular time must have had some success, and so there can be periods in which the leader is fired with positive probability followed by ones in which the leader is not fired at all.

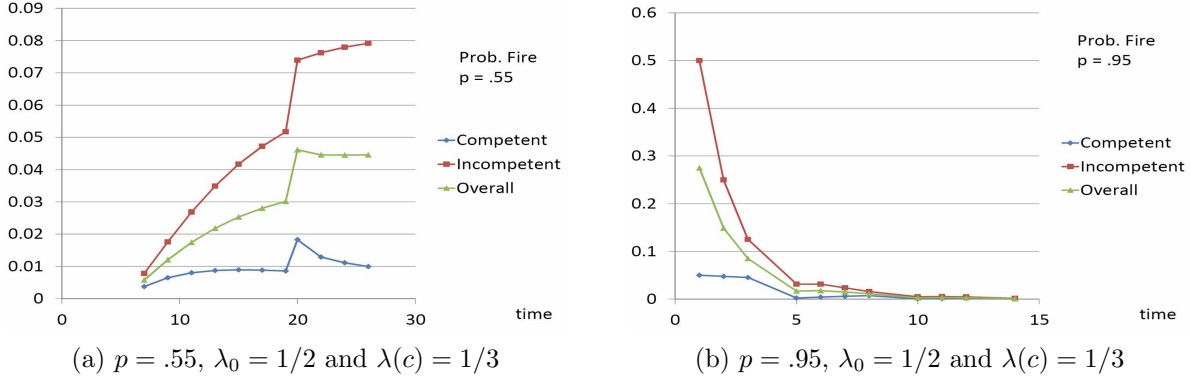


Figure 2: The probability of replacing a leader (for the first time) in various periods, *just the dates positive probabilities*.

barely competent, it takes time to sort out leaders, and so the probability of replacement is growing over time. Also, the probability of making a “false-positive” or type I error (replacing a competent leader) conditional on making a replacement starts out at roughly 1/2 and then drops over time reaching about 1/8 by period 25. Notice also that the probability of replacing a leader in any given period is quite small, less than .1 in all of the first 25 periods. Also, the first period where any replacement occurs is not even until period 7. In contrast, when competent leaders are “highly competent” then they are much easier to distinguish from incompetent ones, and replacements begin in period 1, and have a much higher probability (.5 for incompetent leaders in the first period). Moreover, the probabilities in that case are decreasing over time. The relative probability a type I error conditional on making a replacement starts out at 1/10 and then actually increases for a few periods.¹¹

4 Privately informed leaders under full commitment: tenure mechanisms

We now turn to our main concern: the case where the principal only observes his payoff and does not directly see the signals. In particular, he learns nothing if the leader chooses action x , but can infer the state from observing his payoff if the leader chooses y . We examine two cases in order: a first in which the principal can commit to specific evaluation times and decisions conditional on histories. Effectively, this becomes a mechanism design problem. The second case, is one in which the principal cannot commit to specific evaluation patterns, but instead can replace the leader at any time. The two different scenarios have different applications and their contrast provides some of our central insights.

In either scenario, without proper incentives the leader will prefer to always choose action x in order to avoid the risk of being replaced. In particular, under the mechanism analyzed

¹¹The cumulative probabilities of replacing incompetent leaders, as well as mistakenly replacing competent leaders, also exhibit some interesting patterns as pictured in Figure 3 in the appendix.

in the previous section, for which the leader is replaced if the posterior belief falls below some threshold, will no longer be optimal. Indeed, given such a mechanism, by choosing x all periods the leader can guarantee not being replaced, and so any equilibrium must lead to the leader not being replaced at any point.

To induce the leader to take the risk of choosing y when the true state of nature is Y , there are a variety of mechanisms that can be used. In particular, the principal can randomly dismiss the leader if he chooses action x : this we refer to as the stick. Or the principal can provide incentives by simply guaranteeing to keep the leader, we refer to as the carrot.

We restrict attention to the case where the leader does not know her type, although this has little impact on the results.

4.1 The two period case

To gain intuition on the commitment case, we begin by looking at the two period case, in which we can fully characterize the optimal mechanism. In this scenario, immediate tenure dominates if the replacement cost is sufficiently large as immediate tenure avoids replacement altogether; while the stick of random replacement after safe actions dominates if the replacement cost is small, as it avoids being stuck with a leader who turns out to be incompetent. Let us examine this in more detail.

We presume that $u(\Lambda(0, 1, \lambda_0)) > 0$, so even after a failure in the first period, there is no question of whether or not the principal would like the leader to follow the signal, so the only issue is providing proper incentives.

More formally, a mechanism is characterized by two parameters: π be the probability of retention if y is chosen and the leader fails in period 1; and q the probability of retention if x is chosen in period 1. It is clear that any optimal mechanism (i.e., which maximizes the principal's expected payoff subject to incentive compatibility of the leader) involves keeping the leader if the leader chooses y and is successful.

To induce the leader to choose the risky action y requires that the following incentive compatibility constraint to be satisfied:

$$q \leq f(\lambda_0) + \pi(1 - f(\lambda_0)),$$

where the right-hand side is the overall probability that the leader is retained at the end of period 1 if she chooses action y in that period (this is equal to the success probability $f(\lambda_0)$ times 1 plus the failure probability $1 - f(\lambda_0)$ times π). Given the cost of replacement, it is straightforward to see that this constraint is binding in equilibrium, and thus

$$q = f(\lambda_0) + \pi(1 - f(\lambda_0)). \tag{3}$$

The overall ex ante expected profit of the principal is then

$$U(\lambda_0, \pi, q) = \frac{1}{2} \left[\begin{aligned} & f(\lambda_0) [1 + \delta u(\Lambda(1, 1, \lambda_0))] + \\ & (1 - f(\lambda_0)) [-v + \pi \delta u(\Lambda(0, 1, \lambda_0)) + (1 - \pi)(-c + \delta u(\lambda_0))] \end{aligned} \right] \\ + \frac{1}{2} [-(1 - q)c + \delta u(\lambda_0)],$$

Substituting from (3), maximizing the principal's expected payoff with respect to π and q is equivalent to maximizing:

$$(1 - f(\lambda_0)) [\pi \delta u(\Lambda(0, 1, \lambda_0)) + (1 - \pi)(-c + \delta u(\lambda_0))] + [f(\lambda_0) + \pi(1 - f(\lambda_0))] c.$$

If

$$c > \frac{\delta}{2}[u(\lambda_0) - u(\Lambda(0, 1, \lambda_0))] = \bar{c}(\lambda_0, p, v), \quad (4)$$

then it is better to set $\pi = q = 1$, and otherwise it is better to set $\pi = 0$ and $q = f(\lambda_0)$.

The above inequality (4) compares between the cost of replacing the leader and the potential gain of having a more competent leader, and leads to the following propositions.

PROPOSITION 2 *There exists a cut-off value $\bar{c}(\lambda_0, p, v)$ such that the optimal mechanism for the principal is: (i) if $c > \bar{c}(\lambda_0, p, v)$, then "use the carrot" - grant immediate tenure and retain the leader regardless of the outcome the leader; (ii) if $c < \bar{c}(\lambda_0, p, v)$, then "use the stick" - fire the leader if she takes the risky action and fails in the first period, keep the leader if she takes the risky action and succeeds in the first period, and keep the leader with probability $q = f(\lambda_0)$ if she chooses x in the first period.*

Higher costs clearly favor non-replacement and so the tenure mechanism instead of the random dismissal mechanism. We also note how the optimal mechanism varies with other parameters.

PROPOSITION 3 *The cutoff value is given by*

$$\bar{c}(\lambda_0, p, v) = \frac{\delta(1+v)}{4} \left(\frac{\lambda_0(1-\lambda_0)(p-\frac{1}{2})^2}{\lambda_0(1-p) + (1-\lambda_0)/2} \right).$$

Thus, the random dismissal mechanism is optimal for a wider set of costs as p increases and as v increases. The set of cost values for which it is optimal is initially increasing in λ_0 and eventually decreasing in λ_0 and so non-monotone in the prior probability of competence.

A higher accuracy of signals of competent leaders, and a higher v both increase the relative value of having a competent leader compared to an incompetent one – this increases the relative payoff from the dismissal mechanism compared to the instant tenure mechanism. The comparative statics in λ_0 are not monotone. If λ_0 is near 0 there is no value in replacing the leader as the replacement is likely to be incompetent. If λ_0 is close to 1, then the leader is likely to be competent and there is no reason to replace regardless of the first period outcome. It is in intermediate cases in which it becomes worthwhile to replace the leader.

Proposition 3 follows directly from the fact that

$$\frac{\delta}{2}[u(\lambda_0) - u(\Lambda(0, 1, \lambda_0))] = \frac{\delta(1+v)}{4} \left(\frac{\lambda_0(1-\lambda_0)(p-\frac{1}{2})^2}{\lambda_0(1-p) + (1-\lambda_0)/2} \right).$$

Note that this two-period case does not capture all of the aspects of a tenure contract – as it is essentially a guaranteed contract - there is no decision made after seeing some output from the agent. The comparison between contracts is that of a carrot (guaranteed employment) with a stick (firing for either decision) in terms of motivating the leader. In order to get richer tenure possibilities, we move to three periods.

4.2 The three period case

Moving from two to three periods introduces the possibility that non-immediate tenure dominates immediate tenure or no tenure for suitable parameter values.

In this case we now see more flexibility emerging. If costs of replacement are sufficiently high, then it makes sense to keep the leader in place forever. If costs are low, then it makes sense to constantly evaluate the leader, with some threat of random replacement for a choice of x in order to maintain incentives to choose action according to signal. For intermediate costs, it can become optimal to conditionally evaluate the leader: if the leader performs well in early periods then the leader is tenured and kept in later periods without any fear of replacement.

More formally, we now compare the following three mechanisms.

- (i) A *Random Retention Mechanism* whereby at the end of each period the leader is replaced if she takes action y and fails, and is randomly replaced with positive probability if she takes the safe action x .
- (ii) A *Probationary Tenure Mechanism* whereby the leader is kept for sure after the first period, and then is kept after the second period if: she took action y and was successful in the first period, if she took action x in the first period and y in the second period and was successful, is kept with probability $f(\lambda_0)$ if she took action x in both the first and second periods. The leader is fired after the second period in all other cases.¹²
- (iii) An *Immediate Tenure Mechanism* whereby the leader is never replaced regardless of her actions and the outcomes.

These three mechanisms do not comprise the full space of mechanisms, as there are some hybrids. But the full exploration of the space yields little additional insights, and so for the purpose of exposition we compare on these three, which can each be optimal for some parameter values.¹³

We presume that $u(\Lambda(0, 2, \lambda_0)) > 0$, so that even after two failures it is better to have the leader follow signals than to stop taking the risky action, as that simplifies some of the

¹²Thus, the leader is tenured immediately after being successful on the first attempt; and also with some random probability if both choices in the first two periods were x , and is replaced otherwise, but only replaced after the second period. She cannot be replaced immediately after failure in the first period, or that will distort incentives. The random probability of tenuring after two x 's keeps the leader from being forced to take action y regardless of signal in the second period.

¹³More formally, let x denote an x choice, 1 denote a successful y attempt and 0 a failed y attempt and then p_1, p_0, p_x be the corresponding retention probabilities after the first period. So $p_x = .7$ indicates that if x was chosen in the first period then the leader is fired with probability .3 and retained with probability .7. Let p_{x0} (resp. p_{x1}) denote the probability of retention in the second period if an x was played in the first period and then a y was played and failed (resp. succeeded) in the second period. Let p_{00} (resp. p_{01}) denote the probability of retention in the second period if y was chosen in the first period and failed and then a y was played and failed (resp. succeeded) in the second period. And let p_{10} (resp. p_{11}) denote the probability of retention in the second period if a y was successfully played in the first period and then a y was played

calculations.¹⁴ We also consider the case in which $\delta = 1$ to simplify calculations, as with the finite horizon, the discounting case adds little insight and the calculations already involve many subcases.

The following proposition follows from comparison of payoffs of the principal.

PROPOSITION 4 *There exist cut-off values $c^{IP} > 0$, $c^{PR} > 0$, and $c^{IR} > 0$ of the replacement cost, such that:*

- *the immediate tenure mechanism leads to higher payoffs for the principal than the probationary tenure mechanism if $c > c^{IP}$, with the reverse if $c < c^{IP}$,*
- *the probationary tenure mechanism leads to higher payoffs for the principal than the random retention mechanism if $c > c^{PR}$, with the reverse if $c < c^{PR}$, and*
- *the immediate tenure mechanism leads to higher payoffs for the principal than the random retention mechanism if $c > c^{IR}$, with the reverse if $c < c^{IR}$.*

Thus, for high costs, the immediate tenure mechanism is optimal and for low costs the random retention mechanism is optimal. For some parameter values (combinations of p, v, λ_0), $c^{IP} > c^{PR}$, in which case the probationary tenure mechanism is optimal for intermediate costs.

Intuitively, when the replacement cost c is very large, then it is optimal to never replace the leader and then immediate tenure is optimal. When the replacement cost is very small, it is optimal to re-evaluate the leader in each period, but then random dismissal following action x must be used as otherwise the leader would never take action y since it could lead to failure and dismissal. Tenuring the leader after a success in the first two periods, but firing after a failure (after the second period), emerges as an optimal solution for intermediate values of the replacement cost - as then it is worthwhile to replace a leader who has failed, but to provide incentives by guaranteed employment to a leader who has demonstrated sufficient competence via success.

4.3 The infinite horizon case

Moving to the infinite-period case, we say that a principal uses a *tenure mechanism* if:

and failed (resp. succeeded) in the second period. So, a contract is a specification of

$$p_1, p_0, p_x, p_{11}, p_{10}, p_{1x}, p_{01}, p_{00}, p_{0x}, p_{x1}, p_{x0}, p_{xx}.$$

There are various incentive constraints tying these together, but the full variations of potential mechanisms extend beyond the three considered for our illustration.

¹⁴More generally, under the immediate tenure mechanism, it could be optimal to incentivize the leader to only take action x in some circumstances, but that case adds little insight to the analysis.

- The principal sets a date τ and a nonnegative integer M and a fraction f .
- The principal commits to:
 - keep the leader forever if the leader chooses y *exactly* M times¹⁵ and is successful at least a fraction f of the time by date τ , and
 - replace the leader otherwise.

Note that we could also define a class of mechanisms which require the principal's posterior be at least some λ by time τ . These are not exactly equivalent, but would provide a result similar to the one that we now state.¹⁶

PROPOSITION 5 *For any $\varepsilon > 0$ there exists $\bar{\delta} < 1$ and a tenure mechanism such that if $\delta \geq \bar{\delta}$ and principal employs that mechanism, then in all equilibria of the mechanism (in which a leader follows signals once tenured)¹⁷ the principal's discounted stream of expected utilities is at least $(1 - \varepsilon)$ times the utility the principal would have from having a competent leader in all periods who always chooses based on the signal.*

The implication is that there exists a tenure mechanism that does two things:

- ends up firing all incompetent leaders and eventually keeping a competent leader with an arbitrarily high probability, and
- giving the leader the correct incentives to choose actions consistent with signals in subsequent periods.

Intuitively, with long horizons and sufficient patience, the principal can prolong the test period during which the leader is assessed before deciding whether to tenure her or not, without damaging the overall long run expected utility. This allows him to make sure that incompetent leaders are fired while competent leaders are kept with arbitrarily large probability. The advantage of probationary tenure over the random replacement mechanism is that the former saves on replacement costs, while still providing good incentives for risk-taking both during the test period and then during the post-tenure period.

¹⁵If the leader allows a choice of *at least* M choices of y , this motivates an incompetent leader to try y more than M times which is costly for the principal.

¹⁶Working only off of the posterior can lead an incompetent leader to try action y excessively following failures in order to try to boost the posterior, which is both costly for the principal and can lead to higher chances of incorrectly passing incompetent leaders. This is related to the literature on “calibration” and the literature that followed (e.g., Dawid (1982), Foster and Vohra (1998), Sandroni, Smorodinsky, and Vohra (2003), Dekel and Feinberg (2006)). Our more general replacement strategy avoids the negative results from that literature (that any checking rule in some class can be passed) by allowing restrictions on the strategies and enriching the way in which outcomes can be checked. Here, competent and incompetent agents can be statistically distinguished with high accuracy.

¹⁷It would be enough to have the leader have lexicographic preferences for following signals, after maximizing total benefits. The proposition is also true if a tenured leader only follows signals as long as that leads to positive expected payoffs, but stops if it is eventually revealed that she is likely enough to be incompetent enough to lead to negative expected continuation values. So, one could have the leader value the principal's payoff lexicographically.

4.4 Lessons from the commitment case

The take away from our analysis of commitment by the principal is that in the case in which the state and signal are privately observed by the leader but yet the principal can commit in advance to a dynamic replacement strategy contingent upon past payoff history: (i) when the replacement cost is sufficiently low, the principal will use the stick: fire the leader with positive probability at the end of some periods if the leader either chose the safe action or chose the risky action but failed in that period;¹⁸ (ii) when the replacement cost is high, he will grant immediate or early tenure to the leader; (iii) when the replacement cost is intermediate, and there is sufficient patience it is (always at least approximately) optimal to resort to conditional tenure: during a trial period the leader’s performance is evaluated and the principal decides whether to keep or replace the leader, and then the leader is kept forever if thought to be competent with sufficient probability.

5 Privately informed leader without commitment

The mechanisms that we have considered up until now presume that the principal can fully commit to firing a leader in specific situations. In this section we consider the case where the principal cannot commit to a mechanism and can always decide whether or not to keep the current leader. The principal makes the choice that maximizes his expected future discounted stream of payoffs at each such time.

In this case, we refer to the principal as a “voter” as a central application of this scenarios is to voting settings. Here, we need only one voter since we only consider the leader’s competence and abstract from partisan policies.¹⁹

5.1 A benchmark

A useful benchmark is that of the one-period term limit institution whereby the current leader is replaced every period no matter her past achievement. The payoff to this benchmark is

$$V^1 = \frac{u(\lambda_0) - \delta c}{1 - \delta}. \quad (5)$$

Note that this is a highly inefficient benchmark as successful leaders are replaced immediately and so there is no opportunity to take advantage of any learning about a leader’s competence. Beyond the lack of learning and optimal retention, the cost of replacing a leader is incurred in every period. The only advantage to such a system is that the leader has an incentive to choose according to signal since the leader will be replaced regardless of the outcome.

¹⁸In the fully optimal mechanism, these firing probabilities, both for failures and safe actions, can depend in complex ways on the history and current beliefs.

¹⁹See Canes-Wrone, Herron, and Shotts (2001) and Kartik and McAfee (2007) analyses of how a quality/competence dimension can interact with a policy dimension.

5.2 The inefficiency of unlimited democracy

We begin by an analysis of general forms of equilibria, and then specialize to Markovian equilibria for further results.

5.2.1 Definitions and Preliminary Results

Strategies are defined for the leader, $\sigma_L : H_L \rightarrow \Delta\{x, y\}$ and the principal, $\sigma_P : H_P \rightarrow \delta\{K, R\}$, as a function of possible histories, where H_L and H_P are the sets of all finite histories for the leader and principal, respectively.

We consider perfect Bayesian equilibria of the game.

Let $V_\sigma(h_P^t, \lambda_{tP})$ denote the value function to the principal in the continuation of a perfect Bayesian equilibrium σ if the leader is kept after a history h_P^t and a corresponding posterior λ_{tP} .

It is possible for the beliefs of the voter and leader to differ, but this can only happen when the leader takes action y conditional upon an X signal. If the leader only takes action y in cases in which a Y signal has been observed (but might also be taking action x in such cases), then the beliefs will not differ. Thus, in most equilibrium settings the beliefs will be the same, but it is conceivable that they will differ after some histories.

For some, but not all, results we consider Markov strategies that condition only upon the beliefs. In this case, the principal can only condition upon his beliefs, while the leader knows both beliefs, and so can condition upon both her beliefs and the principal's.

So, Markov strategies for the principal are functions $\sigma_P : [0, 1] \rightarrow \Delta(\{K, R\})$; while Markov strategies for the leader are functions $\sigma_L : [0, 1]^2 \times \{X, Y\} \rightarrow \Delta(x, y)$, with $\sigma = (\sigma_P(\lambda_{tP}), \sigma_L(\lambda_{tP}, \lambda_{tL}, S_t))$ representing a generic strategy.²⁰

We first note that there always exists an equilibrium and in fact there always exists a Markov perfect equilibrium.

LEMMA 2 *There exists a (Markov perfect) equilibrium in which a leader is never voted out of office and all leaders choose x in all situations if elected.*

The proof of Lemma 2 is obvious and so omitted.

5.2.2 Non-Markovian equilibria

There are many equilibria of the game, but they can be sorted into two broad classes as we now show.

First, there is a class of equilibria in which, through threats of future leaders always choosing x , some sort of current behavior is enforced. Effectively, this allows the principal

²⁰The principal could be inferring something from the time period about the leader's beliefs in some out of equilibrium conditions. So, the Markov assumption is actually that the strategy conditions only upon the posterior belief, but we allow that belief to be derived from the full history. So, to be careful, a Markov strategy σ is actually a function of h^t , but depends on h^t only through the corresponding induced $\lambda_{tP}, \lambda_{tL}$.

an arbitrary degree of commitment. The basic idea is that all leaders must act as they would under some mechanism as if that mechanism was committed to, and if a leader deviates from how he should act under that mechanism then all leaders resort to playing the equilibrium in which x is chosen forever by all leaders regardless of history. This can enforce many (but not all) mechanisms as if the leader could commit to the mechanism. The reason that it cannot enforce arbitrary mechanisms is that it can only motivate actions by current leaders that provide nonnegative continuation values to each leader at each point on the equilibrium path.

PROPOSITION 6 *For any $\varepsilon > 0$ there exists $\bar{\delta} < 1$ and a tenure mechanism (τ, M, f) as described above, such that if $\delta \geq \bar{\delta}$ there is an equilibrium in which the principal and leaders act as if the tenure mechanism was in place, and if any player ever deviates from behavior that would be consistent with the tenure mechanism then all players in the future resort to the equilibrium described in Lemma 2.*

The proof of Proposition 6 is straightforward, and so omitted.²¹ Essentially, it points out that arbitrarily efficient equilibria exist, via threats of reversion to a bad equilibrium. This is a way of enforcing commitment by the principal.

On the one hand, such equilibria are somewhat natural: the principal states that he will find a certain pattern over time, but if the principal deviates from that then he is no longer trusted and leaders revert to the equilibrium of never taking action y , which is self-enforcing. On the other hand, this sort of construction relies heavily upon the ability of future leaders to see the full play of the game, and thus would be ruled out for example if the model were such that leaders do not see or pay attention to the history of payoffs before their appointment.²² Also, if there are new voters over time, and current voters are not punished for the actions of past generations, then such constructions would be precluded.

More generally, there is a fundamental difficulty with motivating the current leader that stems from the fact that in any situation where the leader is guaranteed not to be replaced, the leader has an incentive to “sit on her laurels” and stop taking the risky action, and this leads to a contradiction.

²¹The play is as follows, which can be directly checked to be an equilibrium. On the equilibrium path, the principal’s actions are as in a tenure mechanism. The leader’s actions are such that any leader who is tenured then takes action y if and only if there is a Y signal and the current belief $\lambda_{Lt} = \lambda_{Pt}$ is such that $u(\lambda_{Lt}) > 0$. Leaders who are not tenured try y the first M times they get a Y signal and otherwise choose action x , unless $\tau - t = M - m$ where t is the current period and m is the cumulative number of y actions taken by the leader to date, in which case they take action y for each remaining period. If the principal ever deviates from prescribed play, then all players play as described in Lemma 2. If a tenured leader ever deviates from prescribed by play, then all players play as described in Lemma 2 in the continuation. If a non-tenured leader deviates from prescribed play, then the play continues as in the tenure mechanism (note that the only detectable deviations would be such that the leader takes action y too many times, or too few times, to reach exactly M after τ periods, in which case the leader will already be fired after the τ periods). In any other situation, all players play as described in Lemma 2.

²²Refinements that impose some sort of renegotiation-proofness can also rule out such constructions, but can be difficult to work with in such settings.

In particular in cases where there is no reversion to a 0 equilibrium, then the approximate efficiency achieved under Proposition 6 is precluded, as we now show.

Let us say that an equilibrium σ is *has non-trivial continuations* if a new leader always provides a net positive expected continuation payoff of at least some $\gamma > 0$ to the principal.²³

THEOREM 1 *Consider any perfect Bayesian equilibrium, σ that has nontrivial continuations. For any λ that is reached on the equilibrium path (no matter how high), there is some (h_P^t, λ_{tP}) reached with positive probability such that $\lambda_{tP} = \lambda$, $V_\sigma(h_P^t, \lambda_{tP}) \leq V_\sigma(\emptyset, \lambda_{0P}) - c$, and the leader is replaced with positive probability conditional upon that history.*

Theorem 1 implies that regardless of how high the posterior is that the leader is competent, the leader cannot avoid being eventually replaced with positive probability - provided future leaders provide some positive payoffs. This underscores the fundamental difficulty in motivating any leader with *any* history to take the risky action y in the absence of term limits.

The intuition is as follows. If conditional upon some history, the leader was never replaced for any continuation that had the same belief, then by taking action x indefinitely, the leader would be guaranteed to never be fired. Thus, the leader cannot expect to be fired in any continuation that lies on the equilibrium path. Then in the (nonzero probability) case in which she is incompetent, the leader must eventually stop taking the risky action in order to be kept and not replaced. However, this leads to a continuation value of 0 for the principal, and since this is less than the value of a new leader she must be replaced at that time, which is a contradiction. Thus, there must be some histories with the original belief under which the leader is replaced with positive probability.

We point out an important corollary of Theorem 1.

Let us say that an equilibrium is weakly efficient if there exists some $\bar{\lambda}$ such that for any $(h^t, \lambda_{tL}, \lambda_{tP})$, if $\lambda_{tP} > \bar{\lambda}$, then σ is such that the leader follows signals.

Thus, weak efficiency is a minimal sort of requirement in terms of having the leader take appropriate actions. It merely requires that there is some high enough belief, such that whenever beliefs of the leader's competency are above that high enough level, the leader follows signals. It does not require that the leader always feel compelled or secure in following signals, only that this happen at least while beliefs are close enough to 1 that the leader is competent. As the following corollary states, there is no weakly efficient equilibrium.

COROLLARY 1 *There do not exist any weakly efficient perfect Bayesian equilibria that have non-trivial continuations.*

The corollary is derived from Theorem 1 as follows. If an equilibrium is weakly efficient, then there are some high enough levels of λ_{tP} , for which the leader always follows the signal. This then implies that the leader will not be fired by the principal in the next period,

²³That is, $V_\sigma(h_P^t, \lambda_{0P}) - c \geq \gamma$ for some $\gamma > 0$ whenever h_P^t has a new leader at the end of period t ($h_P^t = (\dots, R)$).

conditional upon such histories.²⁴ Thus, the leader will not be replaced if she takes x an arbitrary number of times in a row since beliefs do not change and she is known to be following signals. This contradicts the Theorem.

5.2.3 Markov Perfect equilibria

Theorem 1 shows that motivating a leader to follow signals when believed to be competent with arbitrarily high probability is impossible. That does not necessarily imply that there is not some complicated mixing so that the leader follows the signal often. Although the full set of nontrivial equilibria are difficult to characterize, we can obtain strong bounds and show a strong form of inefficiency for a subclass of them: Markov perfect equilibria.

Let $V_\sigma(\lambda_{tP})$ denote the expected discounted payoff of the principal/voter conditional upon starting a period with a belief that the leader is competent with probability λ_{tP} and given Markov strategies σ .

THEOREM 2 *The entire present discounted value to the principal of any Markov perfect equilibrium, σ , is no more than the one-period cost of replacement: $V_\sigma(\lambda_{0P}) \leq c$. Thus, if the value of constant replacement exceeds the cost of replacement ($V^1 > c$), then the value of any Markov perfect equilibrium is worse than simply replacing the leader in every period.*

Theorem 2 makes a strong statement about the inefficiency of *all* Markov perfect equilibria in a democracy. In particular, if $V^1 > c$, then the ex ante discounted expected value $V_\sigma(\lambda_0)$ of unlimited democracy (as well as the value in any possible continuation, as shown in the proof) is strictly less than the value of the benchmark one-term institution whereby the current leader is replaced in every period.

A basic intuition or heuristic proof for this result (see the appendix for the full proof) is as follows. First, given the cost of replacement, any leader who has the same probability of being competent as a replacement must be kept with probability one. This then implies that a new leader can keep taking action x and never be replaced. Thus, for the leader to take continue to take action y , it must be that she has the assurance of never being replaced thereafter no matter the sequence of successes and failures that ensues. However, with positive probability the leader is incompetent, in which case if she continues to take action y , by the law of large numbers the principal's belief about her ability will end up dropping to a point where she will have to stop taking action y in order to avoid being

²⁴The careful argument is more subtle that it appears, as the principal can expect the posterior to drop in the next period with some probability, and if he knows that he will replace the leader in the next period with some probability, he may also be willing to replace her today. The argument that for high enough λ_{tP} he will keep her for at least one period is as follows. Consider any number of periods $T > 0$. If the leader follows signals, then there must then exist some level $\lambda'_{tP} > \bar{\lambda}$ which is reached with positive probability on the equilibrium path (simply by hitting a sequence of successes given that the leader follows the signals), and such that starting from that level of at least λ'_{tP} , beliefs are guaranteed to stay above $\bar{\lambda}$ for at least T periods. Thus, the principal is guaranteed at least $u(\bar{\lambda})$ for at least T periods. By choosing T large enough, this exceeds the value of replacing the leader, and so the leader will be kept.

replaced (as otherwise she result in a negative utility for the principal). However, at that point the continuation value of keeping her on the job is simply zero (as she must take action x forever after) whereas the continuation value of hiring a new leader is $V_\sigma(\lambda_{0P}) - c$. For the current leader not to be fired in equilibrium even at that point, it must be the case that $V_\sigma(\lambda_{0P}) \leq c$.

5.3 The costs and benefits of longer term limits

Given that unlimited democracy may be incapable of providing proper incentives for choosing risky actions indefinitely, we now investigate whether term limits can provide better incentives. This is a limited form of commitment, where the principal/voter cannot commit to any particular decision in each period, but can commit to an absolute limit on re-elections.

For this analysis, we consider the case in which $V^1 > c$, so that there is some benefit in having the leader make appropriate choices, even if that involves frequent replacement. Under this condition, we know that at a minimum, setting a term limit of one period, already improves over the unlimited equilibrium. We compare short versus longer term limits.

5.3.1 Choice between one-period and two-period term limits

On the one hand longer term limits lower the leader's incentives since the leader has more to lose in terms of foregone future private benefits if she chooses the risky action early and is consequently voted out of office in case she failed: in other words, the longer the term limit, the more one has to wait before the leader takes the risky action. On the other hand, a longer term limits lowers the aggregate inter-temporal replacement cost. The higher the replacement cost, the longer the optimal term limit should be.

To get some first intuition of how this trade-off plays out, suppose first that only one-period and two-period term limits are available and also suppose that $V^1 > c$. Theorem 2 implies that some term limit dominates unlimited terms, and in particular dominated by one-period term limits. Let V^2 denote the expected utility to the principal of replacing a leader every two periods and having leaders follow signals in all periods.

The following cases may occur:

Case 1: The replacement cost c is sufficiently large so that:

$$V^2 - c \leq u(\Lambda(0, 1, \lambda_0)) + \delta(V^2 - c). \quad (6)$$

In this case, under a two-period term limit it is always optimal to keep the current leader no matter her performance on action y in her first period of employment. In this case, the leader has proper incentives to take action y in both periods, as she does not fear replacement.

This lead to an equilibrium payoff to the principal under a two-period term limit equal to

$$V^2 = \frac{u(\lambda_0)}{1 - \delta} - \frac{\delta^2 c}{1 - \delta^2}. \quad (7)$$

This equation allows us to reexpress condition (6) as

$$c > \tilde{c}(\lambda_0, p, v) = (1 + \delta)[u(\lambda_0) - u(\Lambda(0, 1, \lambda_0))]$$

or, using our analysis in Section 3.1 above:

$$c > \tilde{c}(\lambda_0, p, v) = \frac{(1 + \delta)(1 + v)}{2} \left(\frac{\lambda_0(1 - \lambda_0)(p - \frac{1}{2})^2}{\lambda_0(1 - p) + (1 - \lambda_0)/2} \right).$$

Equation (7) also implies that

$$V^2 > V^1 = \frac{u(\lambda_0) - \delta c}{1 - \delta}.$$

Hence in this case a two-period term limit dominates a one-period term limit whenever condition (6) is satisfied.

Case 2: The replacement cost is sufficiently low that under a two-period term limit it would be optimal to dismiss the current leader if she took action y and failed in her first period of employment. Without commitment, the principal cannot commit to randomly firing the leader if x is chosen in the first period, and then that leads the leader to no longer want to choose y in the first period for fear of failure. In this case a one-period term limit becomes optimal.

5.3.2 Longer Term Limits

The above logic extends to term limits of greater lengths.

Let V^T be the present expected discounted value starting at λ_0 of having a leader follow the signal in all T periods and replacing the leader every T -th period, but not sooner. This is given by

$$V^T = \frac{u(\lambda_0)}{1 - \delta} - \frac{\delta^T c}{1 - \delta^T}.$$

Let T^* be the largest number of periods T for which it is better to keep the current leader after $T^* - 1$ periods even if she failed in all periods rather than replace her, but not after T^* . That is,

$$V^{T^*-1} - c \leq u(\Lambda(0, T^* - 1, \lambda_0)) + \delta(V^{T^*-1} - c),$$

but

$$V^{T^*} - c > u(\Lambda(0, T^*, \lambda_0)) + \delta(V^{T^*} - c),$$

where

PROPOSITION 7 *The term limit that maximizes the principal's ex ante expected discounted utility is at least T^* .*

The proof is straightforward, and so we just point out the obvious extension from the case of $T = 2$. Just as in the two-period example, the T^* -period term limit provides the leader with correct incentives, and leads to higher expected utility ex ante than any shorter term limit, as the per period cost is lower but the expected utility in every period is the same irrespective of the term limit (provided it is not more than T^*).

Once one moves beyond the T^* expressed above, the equilibria become more complicated, as they can involve mixing or non-signal following in some periods by leaders with some probability. Whether the payoffs to such equilibria actually exceed the T^* -period term limit depends on the circumstances, but it could be optimal to have a term limit longer than T^* .

Note that in term limit with a limit not exceeding T^* , the leader is never replaced except by the limit. Under the optimal term limit, which could exceed the T^* discussed above, there would still be a strong incumbent advantage.

5.4 Wrapping-up

In this section we have shown that if the signal of the leader is private information and the principal only observes payoff realizations *and* dismissals are decided at the end of each period, then the equilibria that have nontrivial continuations lead to poor incentives for risk-taking by leaders. This contrasts with the commitment case, in which tenure contracts achieve (approximate) efficiency.

6 Summary and discussion

We have analyzed the problem of how to motivate individuals who have discretion over decision making in situations in which firing/replacement threats are the main motivation. There are many settings where this sort of decision making rather than effort really is the main source of hidden action: many high level management, political and judiciary positions (e.g., supreme court justices). In contrast to models where effort rather than discretion interacts with competence, we find that contracts involving commitment to specific periods of evaluation (tenure contracts) or without commitment but then with limited horizons (term limits) are optimal. We have seen that in this world a marked contrast between the case in which the principal can commit to keep the leader from cases in which he cannot.

Undoubtedly, there are settings in which both discretion and effort are both at play, and given that we are finding different results from the effort-based literature, it makes sense to examine how the two interact when effort does not satisfy the usual signaling assumptions.

The analysis in this paper can be extended in several other potentially interesting directions. Two such extensions are straightforward. First, what happens if job separation also occurs exogenously with positive probability? Second, what about if a new leader inherits bad outcomes from previous leaders, which impacts on the probability of failure this period?

On the former (exogenous job separation): suppose that no matter the past history of

success and failures, in each period there is some probability the current leader leaves the job for some exogenous reason. Effectively, this does acts like a decrease in the discount factor for the leader, and a random cost and restart for the principal. In the case of the tenure result, it means that the patience of the principal must now also include a low probability of exogenous separation, otherwise a random retention mechanism will now dominate and a long tenure evaluation period becomes obsolete since leaders are never expected to stay long enough to be properly evaluated.

Now moving to the case where past failures affect current failure: in our model so far we have assumed that the occurrence of successes and failures is uncorrelated across successive leaders. But suppose instead that the occurrence of a failure last period increases the likelihood of a failure this period (even) when action y is chosen in state Y . This naturally occurs in various settings such as coaching or managing (where a past coach or manager's athletes still comprise the team, and it may take some time to evaluate the new coach/manager's ability to choose good athletes), as well as politician's who inherit various economic or political crises (as well as successes) from previous administrations. This has an interesting effect as it can provide a lag before any real updating can take place. Thus, with complete or incomplete information, it can lead to some time periods where the leader is kept for sure, before we return to a setting where the decisions lead to informative signals, and then to the analysis as we have conducted it here.

Other extensions are equally interesting to investigate. One would be to introduce the possibility of promotions rather than just hiring and firing. A conjecture is that as long as there is a limited set of possible promotions, many of the underlying insights would still remain, but that requires a careful analysis.

Another extension would be to allow for multiple leaders or to multiple principals. Having several leaders and one principal would allow the principal to observe the outcome from the challenger's actions over time, not just the outcome from the incumbent's action. A conjecture is that under the current setting this should enhance the leaders' motivation to take risky actions, as it results in a sort of contest. On the other hand, having several principals compete for a limited number of leaders might increase the likelihood of early tenure, with probationary periods that would be shorter than what would be the optimal mechanism from any single principal's perspective. This emerges from the fact that once a leader begins to have success, a principal can entice that leader away from other principals by offering early tenure. More generally, we conjecture that having more leaders than principals should lead to excessive risk-taking by current leaders, whereas having more principals than leaders should lead to earlier tenure - in either case leading to increased inefficiencies.

Another extension is to analyze the case where the principal and the leader do not discount the future at the same rate. Here, we conjecture that the higher the leader's discount factor, the higher the probability of retention under the random retention mechanism, and also the larger the range of replacement costs for which tenure dominates random retention

as this will increase the leader's expected discounted return from currently following signal.²⁵

One could analyze the implications, in the no-commitment case, of having elections being held every X periods instead of being held every period, where $X > 1$. However, the basic inefficiency issues underlying Theorems 1 and 2 would still hold in that case.

Finally, the model delivers empirical predictions, in particular on the efficiency of tenure or term limit arrangements, which should be confronted with the evidence on the performance of firms and countries with differing contractual practices, corporate charters or political constitutions. These and other extensions await further research.

References

- [1] Alesina, A., and Tabellini, G. (2007) "Bureaucrats or politicians? Part I: a single policy task," *American Economic Review*, 97, 169–179.
- [2] Alesina, A., and Tabellini, G. (2008) "Bureaucrats or politicians? Part II: Multiple policy tasks," *Journal of Public Economics*, 92, 426–447.
- [3] Allen, F. and Gorton, G. (1993) "Churning Bubbles" *Review of Economic Studies*, 60:4, 813-836.
- [4] Alt, J, Bueno de Mesquita, E, and Rose, S (2011) "Disentangling Accountability and Competence in Elections: Evidence from US Term Limits. *Journal of Politics*, 73, 171-186.
- [5] Banks, J, and Sundaram, R (1998), "Optimal Retention in Agency Problems", *Journal of Economic Theory*, 82, 293-323.
- [6] Besley, T (2006), *Principled Agents? The Political Economy of Good Government*, Oxford: Oxford University Press.
- [7] Besley, T, and Case, A (1995a), "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits", *Quarterly Journal of Economics*, 110, 769-98.
- [8] Besley, T, and Case, A (1995b), "Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition." *American Economic Review*, 85, 25-45.

²⁵Recall that early tenure in the two-period case dominates whenever

$$c > \bar{c}(\lambda_0, p, v, \eta) = \frac{\delta}{2}[u(\lambda_0, \eta) - u(\Lambda(0, 1, \lambda_0), \eta)],$$

whereas two-period term limits dominate one- period term limits whenever:

$$c > \tilde{c}(\lambda_0, p, v) = (1 + \delta)[u(\lambda_0) - u(\Lambda(0, 1, \lambda_0))].$$

- [9] Besley, T, and Case, A (2003), "Political Institutions and Policy Choices: Evidence from the United States", *Journal of Economic Literature*, 41, 7-73.
- [10] Burdett, K, and Coles, M (2003), "Equilibrium Wage-Tenure Contracts", *Econometrica*, 71, 1377-1404.
- [11] Canes-Wrone, B, and Shotts, K (2007), "When Do Elections Encourage Ideological Rigidity?", *American Political Science Review*, 101, 273-288.
- [12] Canes-Wrone, B, Herron, M, and Shotts, K (2001), "Leadership and Pandering: A Theory of Executive Policymaking", *American Journal of Political Science*, 45, 532-50.
- [13] Carmichael, L (1988), "Incentives in Academics: Why Is There Tenure?", *Journal of Political Economy*, 96, 453-472.
- [14] Dawid, A. P. (1982), "The Well Calibrated Bayesian," *Journal of the American Statistical Association*, 77, 605-613.
- [15] Dekel, E. and Feinberg, Y. (2006) "Non-Bayesian testing of a stochastic prediction," *Review of Economic Studies* 73:4, 893 - 906.
- [16] Dewatripont, M, Jewitt, I, and Tirole, J (1999a), "The Economics of Career Concerns, Part I: Comparing Information Structures", *Review of Economic Studies*, 66, 183-198
- [17] Dewatripont, M, Jewitt, I, and Tirole, J (1999a), "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies", *Review of Economic Studies*, 66, 199-217
- [18] Farber, H (1999), "Mobility and Stability: The Dynamics of Job Changes In Labor Markets", in *Handbook of Labor Economics*, Ashenfelter, O, and Card, D (Eds) , Elsevier: North-Holland
- [19] Fearon, J (1999), "Electoral Accountability and the Control of Politicians" In *Democracy, Accountability and Representation*, Przeworski, A, Stokes, S, and Manin, B (Eds) New York: Cambridge University Press, 55-97.
- [20] Ferejohn, J (1986), "Incumbent Performance and Electoral Control", *Public Choice* 50, 5-26.
- [21] Ferraz, C, and Finan, F (2009), "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments", *NBER Working Paper No 14937*.
- [22] Foster, D. P., and Vohra, R. (1998) "Asymptotic Calibration," *Biometrika*, 85:2, 379 - 390.
- [23] Gordon, S, Huber, G, and Landa, D (2007), "Challenger Entry and Voter Learning", *American Political Science Review*, 101, 303-20.

- [24] Harrington, J (1993), "Economic Policy, Economic Performance and Elections," *American Economic Review*, LXXXIII, 27-42.
- [25] Holmstrom, B (1982) "Managerial Incentive Problems: A Dynamic Perspective," in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics.
- [26] Jackson, M, and Lazear, E (1991) "Stock, Options, and Deferred Compensation," *Research in Labor Economics*, 2, 41-62.
- [27] Jia, R, Kudamatsu, M, and Seim, D (2014), "Political Selection in China: Complementary Role of Connections and Performance", Working Paper, IIES, Stockholm University.
- [28] Jovanovic, B (1979), "Job Matching and the Theory of Turnover", *Journal of Political economy*, 87, 972-990.
- [29] Kahn, C, and Huberman, G (1988), "Two-Sided Uncertainty and Up-or-Out Contracts," *Journal of Labor Economics*, 6, 423-444.
- [30] Kartik, N. and McAfee, R.P. (2007) "Signaling Character in Electoral Competition," *American Economic Review*, 97(3), 852-870.
- [31] Mailath, G. and L. Samuelson (2001), "Who Wants a Good Reputation?" *Review of Economic Studies* 68(2): 415-441.
- [32] Mailath, G.J., and Samuelson, L. (2006), *Repeated Games and Reputations*. Oxford University Press.
- [33] Mattozzi, A, and Merlo, A (2008), "Political Careers or Career Politicians?" *Journal of Public Economics*, 92 , 597-608.
- [34] Padro i Miquel, G, and Snyder, J (2006), "Legislative Effectiveness and Legislative Careers", *Legislative Studies Quarterly*, 31, 347-81.
- [35] Rodriguez-Barraquer, T, and Tan, X (2014), "A Model of Competitive Signaling," Unpublished Working Paper, Universitat Autònoma de Barcelona.
- [36] Sandroni, A., Smorodinsky, R. and Vohra, R. (2003) "Calibration with many checking rules," *Mathematics of Operations Research*, 28:1, 141 - 153.
- [37] Scharfstein, D, and Stein, J (1990) ",Herd Behavior and Investment", *American Economic Review*, 80, 465-479
- [38] Smart, M, and Sturm, D. (2013), "Term Limits and Electoral Accountability", mimeo London School of Economics.

- [39] Tadelis, S. (1999), “What’s in a Name? Reputation as a Tradeable Asset,” *American Economic Review*, 89(3), 548-563
- [40] Taylor, C. (2000), “The Old-Boy Network and the Young-Gun Effect,” *International Economic Review*, 41(4), 871-891
- [41] Tirole, J. (1996), “A Theory of Collective Reputations (With Applications to the Persistence of Corruption in the Firm Quality),” *Review of Economic Studies*, 63, 1, 1-22.
- [42] Waldman, M (1990), ”Up-Or-Out Contracts: A Signaling Perspective”, *Journal of Labor Economics*, 8, 230-250.

Appendix: Proofs

Consider the case in which states are observed as well as signals. In that case, let $o_t = 1$ if the leader is correct (a “success”) and $o_t = 0$ if the leader is incorrect (a “failure”). A sufficient statistic for λ_t given λ_0 is the cumulative number of successes through time t :

$$O_t = \sum_{t' \leq t} o_{t'}.$$

As noted above, the posterior probability of having a competent leader $\lambda_t = \Lambda(O_t, t, \lambda_0)$ is increasing in the number of successes through a given time, O_t .

Let

$$O_t(\lambda_0, \lambda) = \{O \in \mathbb{N} : \Lambda(O, t, \lambda_0) < \lambda\}$$

denote the set of possible cumulative successes through time t for which the posterior falls below the threshold λ . Given that $\Lambda(O, t, \lambda_0)$ is increasing in O , there exists

$$\underline{O}_t(\lambda_0, \lambda) = \max O_t(\lambda_0, \lambda).$$

Thus, $O \in O_t(\lambda_0, \lambda)$ if and only if $O \leq \underline{O}_t(\lambda_0, \lambda)$ and $O \in \mathbb{N}$.

Proof of Proposition 1:

The proof makes use of the following two lemmas.

LEMMA 3 *Suppose that the principal starts with some prior λ_0 and continues to retain the leader as long as $\lambda_t \geq \lambda$ for some $0 < \lambda \leq \lambda_0$. Then $\underline{O}_t(\lambda_0, \lambda)$ is the unique value of O_t for which the principal fires the leader at period t and not before t (if $P(t) > 0$).*

Proof of Lemma 3:

If $\underline{O}_t(\lambda_0, \lambda) = 0$, then this holds since this is the maximum value out of $O_t(\lambda_0, \lambda)$ and since there are no lower values, it must be unique. So, suppose $\underline{O}_t(\lambda_0, \lambda) > 0$, and also that

to the contrary $o = \underline{Q}_t(\lambda_0, \lambda) - k$, with $k > 0$, is another value for which the follower stops at period t and not before t . Note that since $k > 0$

$$\Lambda(o, t - 1, \lambda_0) < \Lambda(o + k, t, \lambda_0).$$

Since $\Lambda(o + k, t, \lambda_0) < \lambda$ (noting that $o + k = \underline{Q}_t(\lambda_0, \lambda)$), it follows that $\Lambda(o, t - 1, \lambda_0) < \lambda$. This implies that $O_{t-1} \leq o$ would lead to the follower stopping at time $t - 1$ (or possibly before), and so o could not lead to a first stopping at time t since $O_t = o$ implies $O_{t-1} \leq o$ which would lead to stopping by $t - 1$.

LEMMA 4 *Suppose that the follower starts with some prior λ_0 and continues to retain the leader as long as $\lambda_t \geq \lambda$ for some $0 < \lambda \leq \lambda_0$. Let $P(t') > 0$ and $P(t) > 0$. Then $\underline{Q}_{t'}(\lambda_0, \lambda) > \underline{Q}_t(\lambda_0, \lambda)$ whenever $t' > t$, and $\underline{Q}_{t'}(\lambda_0, \lambda) = 1 + \underline{Q}_t(\lambda_0, \lambda)$ if $t' = t + 1$.*

Proof of Lemma 4:

First, note that $\underline{Q}_{t'}(\lambda_0, \lambda) \geq \underline{Q}_t(\lambda_0, \lambda)$. To see this, suppose to the contrary that $o = \underline{Q}_{t'}(\lambda_0, \lambda) < \underline{Q}_t(\lambda_0, \lambda)$. However, if $O_{t'} = o$ then $O_t \leq o < \underline{Q}_t(\lambda_0, \lambda)$ which implies that the follower would have stopped by t .

Next let us show that $\underline{Q}_{t'}(\lambda_0, \lambda) \neq \underline{Q}_t(\lambda_0, \lambda)$. This follows since $O_t \leq O_{t'}$ and so then the principal would have fired the leader by t if the two were the same. Next, let us show that $\underline{Q}_{t+1}(\lambda_0, \lambda) = \underline{Q}_t(\lambda_0, \lambda) + 1$. Note that $\underline{Q}_t(\lambda_0, \lambda) \geq \underline{Q}_{t+1}(\lambda_0, \lambda) - 1$ since

$$\Lambda(o, t, \lambda_0) < \Lambda(o + 1, t + 1, \lambda_0),$$

for any o . So if $\Lambda(o + 1, t + 1, \lambda_0) < \lambda$, it follows that $\Lambda(o, t, \lambda_0) < \lambda$. ■

To show the proposition, let \underline{t} be the first t for which $P(t) > 0$.²⁶

We now invoke the two lemmas above. They imply that if $P(t) > 0$ for all $t > \underline{t}$, then it would have to be that

$$\underline{Q}_t = \underline{Q}_{\underline{t}} + t - \underline{t}.$$

This would imply that $\underline{Q}_t/t \rightarrow 1$, and so $\lambda_t \rightarrow 1$, which contradicts the fact that the leader would be followed. Thus, there is some finite t for which $P(t) = 0$. Note also, that $P(t) > 0$ infinitely often: it is possible that the leader is correct for any arbitrary initial number of periods and then incorrect thereafter for any given number of periods. Thus, for any τ and $\varepsilon > 0$ it is possible to have $\lambda_t > \lambda_0$ for all $t < \tau$, and then $\lambda_t < \varepsilon$ for large enough $t > \tau$. This implies that $P(t) > 0$ for some $t > \tau$. Since this is true for any τ , it is true infinitely often. ■

Proof of Proposition 4: We begin by characterizing the principal's payoff from each of the mechanisms.

Recall that $\delta = 1$ and that $u(\Lambda(0, 2, \lambda_0)) > 0$, so that it is optimal for a leader to follow signal in any period regardless of history.

²⁶There exists such a t since a long enough string 0's will lead λ_t into any given neighborhood of 0. Any long enough string of 0's has positive probability.

The principal's payoff from the immediate tenure mechanism It follows directly that the ex ante expected utility from the immediate tenure mechanism is

$$U^{IT} = 3u(\lambda_0).$$

The principal's payoff from the probationary tenure mechanism Recall that tenure is offered with probability $f(\lambda_0)$ in the case of two x actions in order to satisfy the incentive compatibility condition to induce the leader to follow the signal in the first two periods. Then, careful calculation (considering each possible realization of signals and actions) shows that the ex ante expected utility from the probationary tenure mechanism is

$$U^{PT} = \frac{1}{2} \left[\begin{array}{c} f(\lambda_0)(1 + 2u(\Lambda(1, 1, \lambda_0))) \\ +(1 - f(\lambda_0))(-v + u(\Lambda(0, 1, \lambda_0)) - c + u(\lambda_0)) \end{array} \right] \\ + \frac{1}{4} [f(\lambda_0)(1 + u(\Lambda(1, 1, \lambda_0))) + (1 - f(\lambda_0))(-v - c + u(\lambda_0))] + \frac{1}{4} [u(\lambda_0) - c(1 - f(\lambda_0))].$$

The principal's payoff from the random retention mechanism The leader is kept after a success and fired after a failure, so to characterize the payoffs from this mechanism, we first need to solve for the various retention probabilities following x actions in both the first and second period.

There are two different versions of this mechanism. In one case $p_x = 1$ and $p_{xx} = f(\lambda_0) - 2 + 2f(\Lambda(1, 1, \lambda_0)) = g(\lambda_0)$; and in the other case $p_x = f(\Lambda(1, 1, \lambda_0))$ and $p_{xx} = f(\lambda_0)$.

To see this, note that the last period incentive compatibility (IC) constraint implies that

$$f(\lambda_0) \geq p_{xx} \geq 1 - f(\lambda_0).$$

Also,

$$f(\Lambda(1, 1, \lambda_0)) \geq p_{1x} \geq 1 - f(\Lambda(1, 1, \lambda_0)).$$

Next there is a first period incentive constraint. If the leader takes action y conditional on a Y signal, then the expected payoff is

$$b[1 + f(\lambda_0) + \frac{1}{2}(f(\Lambda(1, 1, \lambda_0)) + p_{1x})].$$

If the leader instead chooses x in the first period (and acts according to signal in the second period) then the payoff is

$$b[1 + p_x + \frac{1}{2}(f(\lambda_0) + p_{xx})].$$

Setting $p_{1x} = f(\Lambda(1, 1, \lambda_0))$ is optimal to guarantee that the leader will want to take action y when seeing a Y signal.

Then the first period constraint becomes

$$f(\lambda_0) + f(\Lambda(1, 1, \lambda_0)) \geq p_x + \frac{1}{2}(f(\lambda_0) + p_{xx}).$$

A solution to this involves pushing at least one of p_x or p_{xx} to its maximum constrained value (or else it does not matter): So, one possibility is to set $p_x = 1 = p_x^1$ and then

$$f(\lambda_0) + f(\Lambda(1, 1, \lambda_0)) = 1 + \frac{1}{2}(f(\lambda_0) + p_{xx})$$

or

$$p_{xx} = f(\lambda_0) - 2 + 2f(\Lambda(1, 1, \lambda_0)).$$

The other possibility is to set $p_{xx} = f(\lambda_0)$ and then

$$f(\lambda_0) + f(\Lambda(1, 1, \lambda_0)) = p_x + f(\lambda_0)$$

and so

$$p_x = f(\Lambda(1, 1, \lambda_0)).$$

Both of these lead to the same behavior by the leader, *and no difference in the posterior beliefs regarding any retained leader since it only conditions on sequences of x behaviors* and so the only difference comes in expected replacement costs. It is straightforward to check that the overall replacement costs are identical, and so the two mechanisms lead to identical payoffs.

Thus, we focus on the mechanism with $p_x = 1$ and $p_{xx} = f(\lambda_0) - 2 + 2f(\Lambda(1, 1, \lambda_0))$ as that leads to the easiest comparison with the other two mechanisms. Also, $p_{1x} = f(\Lambda(1, 1, \lambda_0))$.

Given these three probabilities, we then have the following expression for the random retention mechanism (the first part is what if the principal gets if there is failure in the first period, the second part is what he gets if there is success in the first period; the third part is what the principal get if x is chosen in the first period):

$$\begin{aligned} U^{RR} &= \frac{1}{2}(1 - f(\lambda_0)) \left[-v - c + u(\lambda_0) + \frac{1}{2} [f(\lambda_0) (u(\Lambda(1, 1, \lambda_0)) + u(\lambda_0)) + 2(1 - f(\lambda_0)) (u(\lambda_0) - c)] \right] \\ &+ \frac{1}{2} f(\lambda_0) \left[1 + u(\Lambda(1, 1, \lambda_0)) + \frac{1}{2} [f(\Lambda(1, 1, \lambda_0)) (u(\Lambda(2, 2, \lambda_0)) + u(\Lambda(1, 1, \lambda_0))) \right. \\ &\quad \left. + (1 - f(\Lambda(1, 1, \lambda_0))) (-v - 2c + 2u(\lambda_0))] \right] \\ &+ \frac{1}{2} \left[0 + \frac{1}{2} [f(\lambda_0) (1 + u(\Lambda(1, 1, \lambda_0))) + (1 - f(\lambda_0)) (-v - c + u(\lambda_0))] \right. \\ &\quad \left. + \frac{1}{2} [0 + u(\lambda_0) - (1 - f(\lambda_0) + 2 - 2f(\Lambda(1, 1, \lambda_0)))c] \right]. \end{aligned}$$

We now prove the result by direct comparison of the expressions.

A) *Comparison between Immediate Tenure (IT) and Probationary Tenure (PT):*

The difference between Immediate Tenure (IT) and Probationary Tenure (PT) is that PT involves firing the current leader after period 2 either if she chooses x twice in a row or if she chooses y in period 1 and fails or if she choose x in the first period followed by a failed y in the second period.

Thus PT dominates IT if and only if:²⁷

$$\begin{aligned}
& \frac{1}{2}(1 - f(\lambda_0))[-c + u(\lambda_0)] \\
& + \frac{1}{4}(1 - f(\lambda_0))[-c + u(\lambda_0)] \\
& + \frac{1}{4}\{(1 - f(\lambda_0))[-c + u(\lambda_0)] \\
& + f(\lambda_0)u(\lambda_0)\} \\
> & \frac{1}{2}(1 - f(\lambda_0))u(\Lambda(0, 1, \lambda_0)) \\
& + \frac{1}{4}(1 - f(\lambda_0))u(\Lambda(0, 1, \lambda_0)) \\
& + \frac{1}{4}u(\lambda_0)
\end{aligned}$$

or equivalently

$$c < \frac{3}{4} [u(\lambda_0) - u(\Lambda(0, 1, \lambda_0))] = c^{IP}.$$

B) *Comparison between Probationary Tenure (PT) and Random Retention (RR):*

These two mechanisms lead to the same outcomes whenever there are two successes, or an x followed by either a success or a failure; but differ in all other cases.

Thus PT will dominate RR if and only if:

$$\begin{aligned}
& \frac{1}{2}(1 - f(\lambda_0)) [u(\Lambda(0, 1, \lambda_0)) + u(\lambda_0) - c] + \frac{1}{2}f(\lambda_0) [2u(\Lambda(1, 1, \lambda_0))] \\
& + \frac{1}{4} [u(\lambda_0) - (1 - f(\lambda_0))c] \\
> & \frac{1}{2}(1 - f(\lambda_0)) \left[u(\lambda_0) - c + \frac{1}{2} [f(\lambda_0) (u(\Lambda(1, 1, \lambda_0)) + u(\lambda_0)) + 2(1 - f(\lambda_0)) (u(\lambda_0) - c)] \right] \\
& + \frac{1}{2}f(\lambda_0) \left[u(\Lambda(1, 1, \lambda_0)) + \frac{1}{2} [f(\Lambda(1, 1, \lambda_0)) (u(\Lambda(2, 2, \lambda_0)) + u(\Lambda(1, 1, \lambda_0)))] \right. \\
& \left. + \frac{1}{2} [(1 - f(\Lambda(1, 1, \lambda_0))) 2(u(\lambda_0) - c)] \right] . \\
& + \frac{1}{4} [u(\lambda_0) - (1 - f(\lambda_0) + 2 - 2f(\Lambda(1, 1, \lambda_0)))c]
\end{aligned}$$

which boils down to

$$c > c^{PR},$$

²⁷This ignores the exact indifference. In the left hand side of this inequality, the first term corresponds to the principal's payoff under PT conditional upon action y being chosen and failed upon in the first period, the second term corresponds to the principal's payoff conditional upon choosing x in the first period followed by a failed y in the second period, and the third term corresponds to the principal's payoff conditional upon x being chosen twice in a row; the right hand side of the inequality is the sum of the corresponding principal's payoffs under IT.

where

$$\underline{c} = \frac{f(\lambda_0)E + (1 - f(\lambda_0))F}{2(1 - f(\lambda_0))^2 + 2f(\lambda_0)(1 - f(\Lambda(1, 1, \lambda_0))) + 2(1 - f(\Lambda(1, 1, \lambda_0)))},$$

where

$$E = f(\Lambda(1, 1, \lambda_0))[u(\Lambda(2, 2, \lambda_0)) + u(\Lambda(1, 1, \lambda_0))] \\ + (1 - f(\Lambda(1, 1, \lambda_0)))2u(\lambda_0) - 2u(\Lambda(1, 1, \lambda_0))$$

and

$$F = f(\lambda_0)[u(\Lambda(1, 1, \lambda_0)) + u(\lambda_0)] \\ + (1 - f(\lambda_0))2u(\lambda_0) - 2u(\Lambda(0, 1, \lambda_0))$$

Note that

$$u(\lambda_0) = f(\lambda_0)u(\Lambda(1, 1, \lambda_0)) + (1 - f(\lambda_0))u(\Lambda(0, 1, \lambda_0)),$$

and so

$$u(\lambda_0) - u(\Lambda(0, 1, \lambda_0)) = f(\lambda_0)[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]. \quad (8)$$

This then helps us simplify E and F , so that

$$E = f(\Lambda(1, 1, \lambda_0))[u(\Lambda(2, 2, \lambda_0)) - u(\Lambda(1, 1, \lambda_0))] \\ - 2(1 - f(\lambda_0))(1 - f(\Lambda(1, 1, \lambda_0)))[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]$$

and

$$F = f(\lambda_0)(3 - f(\lambda_0))[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]$$

Therefore,

$$c^{PR} = f(\lambda_0)[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))] \\ \cdot \frac{f(\Lambda(1, 1, \lambda_0)) \frac{[u(\Lambda(2, 2, \lambda_0)) - u(\Lambda(1, 1, \lambda_0))]}{[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]} + (1 - f(\lambda_0))(1 - f(\lambda_0) + 2f(\Lambda(1, 1, \lambda_0)))}{2(1 - f(\lambda_0))^2 + 2(1 + f(\lambda_0))(1 - f(\Lambda(1, 1, \lambda_0)))}.$$

C) *Comparison between Immediate Tenure (IT) and Random Retention (RR):*

The existence of the cost threshold c^{IR} follows along similar lines. Here is sufficient to note that U^{RR} is decreasing in c and U^{IT} is independent of c , and that $U^{RR} > U^{IT}$ for $c = 0$.

To complete the proof we need only verify that there exist parameter values for which $c^{IP} > c^{PR}$.

Note that by (8) we can rewrite

$$c^{IP} = f(\lambda_0) \frac{3}{4} [u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))].$$

Thus, to show that $c^{IP} > c^{PR}$ it is sufficient to show that there exist parameter values for which

$$\frac{3}{4} > \frac{f(\Lambda(1, 1, \lambda_0)) \frac{[u(\Lambda(2, 2, \lambda_0)) - u(\Lambda(1, 1, \lambda_0))]}{[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]} + (1 - f(\lambda_0))(1 - f(\lambda_0) + 2f(\Lambda(1, 1, \lambda_0)))}{2(1 - f(\lambda_0))^2 + 2(1 + f(\lambda_0))(1 - f(\Lambda(1, 1, \lambda_0)))}. \quad (9)$$

It is useful to consider the limit case where $p = 1$, as the calculations simplify dramatically and it is then easy to see that the results hold in a neighborhood of that case. When $p = 1$ (and $\lambda_0 \in (0, 1)$):

$$\Lambda(0, 1, \lambda_0) = 0, \quad \Lambda(1, 1, \lambda_0) = \frac{2\lambda_0}{\lambda_0 + 1}, \quad \Lambda(2, 2, \lambda_0) = \frac{4\lambda_0}{3\lambda_0 + 1},$$

and

$$f(\Lambda(0, 1, \lambda_0)) = \frac{1}{2}, \quad f(\Lambda(1, 1, \lambda_0)) = \frac{\lambda_0}{\lambda_0 + 1} + \frac{1}{2}, \quad f(\Lambda(2, 2, \lambda_0)) = \frac{2\lambda_0}{3\lambda_0 + 1} + \frac{1}{2}.$$

Then noting that

$$\frac{[u(\Lambda(2, 2, \lambda_0)) - u(\Lambda(1, 1, \lambda_0))]}{[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]} = \frac{(f(\Lambda(2, 2, \lambda_0)) - f(\Lambda(1, 1, \lambda_0)))(1 + v)}{(f(\Lambda(1, 1, \lambda_0)) - f(\Lambda(0, 1, \lambda_0)))(1 + v)},$$

it follows that

$$\frac{[u(\Lambda(2, 2, \lambda_0)) - u(\Lambda(1, 1, \lambda_0))]}{[u(\Lambda(1, 1, \lambda_0)) - u(\Lambda(0, 1, \lambda_0))]} = \frac{f(\Lambda(2, 2, \lambda_0)) - f(\Lambda(1, 1, \lambda_0))}{f(\Lambda(1, 1, \lambda_0)) - f(\Lambda(0, 1, \lambda_0))} = \frac{1 - \lambda_0}{1 + 3\lambda_0}. \quad (10)$$

Next, in the limit in which $\lambda_0 \rightarrow 0$, we find that all of the $f(\cdot)$ s tend to $1/2$, and that the right hand side of (10) tends to 1. Thus, after some algebra, the right hand side of (9) tends to $5/8$. Therefore, for parameters near $p = 1$ and $\lambda_0 = 0$, (9) is satisfied. This completes the proof of the proposition. ■

Proof of Proposition 5:

Choose f halfway between $1/2$ and p , so $f = \frac{1}{4} + \frac{p}{2}$.

Also choose $M = \tau/2$. Let us show that if τ is large enough to satisfy the above.

The proposition is established via the following claims.

1. For any $\varepsilon' \in (0, 1/2)$ there exists a large τ such that a competent leader will have at least a fraction of f successes out of any number of more than $\tau/2$ choices of y from Y signals with probability at least $1 - \varepsilon'$ and an incompetent will have a fraction of f successes under y choices from $\tau/2$ choices of y (regardless of signals) with probability less than ε' .

2. For any λ_0 and $\varepsilon'' \in (0, 1)$ there exists some number K such that the probability of at having at least one competent leader out of K tries is at least $1 - \varepsilon''$.

2. follows directly from setting K such that $(1 - \lambda_0)^K < \varepsilon''$.

1. follows from the law of large numbers. Choose τ large enough so that the probability of a fraction of f successes out of $\tau/2$ tries is no more than $1 - \varepsilon'$ when the coin has a probability of $1/2$ but is more than $1 - \varepsilon'$ when the coin has a probability p .

By following a strategy of choosing according to signals, a competent leader will successfully pass the tenure with a probability that is above $1 - 2\varepsilon'$ for large enough τ . Regardless of the strategy followed an incompetent leader will pass with probability less than $1 - \varepsilon'$.

Let us now argue that in all equilibria competent leaders will pass the test with probability at least $1 - 2\varepsilon'$ and the incompetent leaders will fail the test with probability at least $1 - \varepsilon'$.

If the leader is competent (and is aware of this), then she can pass the test with

By 2 applied to $\varepsilon'' = \varepsilon/4$, we can find K such that within K draws of leaders the probability of having a competent leader is at least $1 - \varepsilon/4$.

Find ε' such that $(1 - \varepsilon'^K > 1 - \varepsilon/4$. It then follows from 1, using 2 to set K as above, that there exists τ such that within K periods there is a probability of at least $(1 - \varepsilon/4)^2 > (1 - \varepsilon/2)$ of having a competent leader pass the test and not having any incompetent leader pass the test. This follows since there is a probability of at least $(1 - \varepsilon/4)$ of having at least one competent leader in the first K draws. Then under 1, the probability that all incompetent leader in the first K draws fail and any competent leader in the first K draws (if that many are needed) passes the test is at least $(1 - \varepsilon'^K$. By the selection of ε' , this later probability is at least $1 - \varepsilon/4$.

Thus, for any ε we can find τ and K such that there is a probability of at least $1 - \varepsilon/2$ that a competent leader will pass the before time $K\tau$ and such that the probability of having an incompetent leader pass the test is less than $\varepsilon/2$.

Set $\bar{\delta}$ such that $\bar{\delta}^{K\tau} > 1 - \varepsilon/2$.

Thus, for any $\varepsilon \in (0, 1/2)$ we can find a tenure mechanism such that there is a probability of at least $1 - \varepsilon/2$ that a competent leader will pass the test (an no incompetent leader will) within $K\tau$ periods, and then lead to the truthful and competent payoffs for all time after $K\tau$. For $\delta > \bar{\delta}$ this leads to a total payoff of at least $1 - \varepsilon/2$ of the payoff as if truthful and competent payoffs were obtained in all periods. Thus, the ex ante expected discounted sum of payoffs from this mechanism are at least $(1 - \varepsilon/2)^2$ of the fully competent and truthful payoffs, which establishes the proposition. ■

Proof of Theorem 1:

Suppose the contrary of the proposition. Let h_P^t, λ_{tP} be reached with positive probability and suppose that a leader is never fired in any continuation h_P^t, λ_{tP} that has the same posterior λ_{tP} that is reached with positive probability. This implies that the leader is never fired in any continuation that is reached with positive probability on the equilibrium path, since the leader can guarantee a payoff of b in perpetuity by simply choosing x forever (as any finite sequences of x choices occurs with positive probability and leaves the principal's belief unchanged). Next, note that with some positive probability the leader is incompetent, conditional on the principal's observations (or, actually, under any other observations of history to date). Let us then consider the case in which the leader is actually incompetent.

Consider any $\varepsilon > 0$, and let A be the set of all sequences of possible payoffs (for which all finite truncations have positive probability in equilibrium) for the principal that have the limsup of the fraction of successes divided by failures not exceed $1 + \varepsilon$. Then, if the

leader is incompetent, the probability that the continuation sequence lies in A is 1 regardless of the strategy. It follows from Kolmogorov's 0-1 law, the beliefs of the principal on the event A (which is a tail event) converge to 1. For small enough $\varepsilon > 0$ this implies that the continuation value of the principal must lie below γ with positive probability (where γ is as defined in the statement of the theorem). This is a contradiction, since the principal will then replace the leader. ■

Proof of Theorem 2: First note that $V_\sigma^D(\lambda_0) > V_\sigma^D(\lambda_0) - c$. This implies that $P(\lambda_0) = 1$, as keeping a leader with λ_0 is better than paying the cost of replacement and getting a leader with the same probability of competence who would follow the same strategy. This implies if a leader takes action x in perpetuity the principal's belief about her competence will remain equal to λ_0 and the leader will be kept in perpetuity. This implies that since the leader has a strategy that guarantees being kept forever response to a Markovian strategy of the voter, the leader must be kept with probability 1 under the leader's best response since the leader values only being in office. Note that since $\lambda_0 < 1$, there is a positive probability that the leader is incompetent (namely, $1 - \lambda_0$). For an incompetent leader, either there is a positive probability that the leader hits some λ_t after which the leader no longer takes action Y , and then $V(\lambda_t) = 0$, or else there is a positive probability one that the leader continues to take action Y an infinite number of times, in which case λ_t converges to 0 by the law of large numbers. In that second case, it must be that $V(\lambda_t) \leq 0$ for some λ_t reached with positive probability on the path, since an incompetent leader has a negative expected payoff from taking action Y . Thus, in either case, there is a positive probability of hitting some λ_t such that $V(\lambda_t) \leq 0$. Yet, as we argued above the leader is never replaced. This implies that $0 \geq V(\lambda_t) \geq V_\sigma^D(\lambda_0) - c$, which implies that $V_\sigma^D(\lambda_0) \leq c$. ■

Cumulative Firing Probabilities in the Complete Information Case

In Figure 3 we see that the cumulative probabilities are not ordered with respect to how competent a competent leader is. Although the replacement curves in the case of incompetent leader are ordered by the p 's (see Panel 3a), those for competent leaders are not ordered monotonically (see Panel 3b). The highest probability of replacing a good leader starts out highest for the case of $p = .8$ and lowest for $p = .55$, with $p = .95$ and $p = .65$ in between. The two cases of $.65$ and $.55$ eventually overtake the others, but start out lower because information is slow to accumulate in those cases and so replacement probabilities are low initially.

By the law of large numbers, incompetent leaders will be recognized with a probability 1 (so $\lambda_t \rightarrow 0$ with probability 1 in the case of an incompetent leader), and so eventually any incompetent leader will be replaced as long as the cost of replacement is not so high that one would not replace an incompetent leader even if known to be incompetent. There are





Figure 3: The cumulative probability of replacing a leader (for the first time).

two main differences for what happens as a function of the cost of replacement: how quickly incompetent leaders are replaced, and with what probability competent leaders are replaced. These present a trade-off.