

Optimal Political Career Dynamics*

Avidit Acharya[†]
Stanford

Elliot Lipnowski[‡]
Columbia

João Ramos[§]
USC

November 8, 2022

Abstract

We examine politicians’ career dynamics generated by political accountability, characterizing voter-optimal equilibrium play under repeated moral hazard. When moral hazard binds, equilibrium play is non-stationary: Re-election prospects improve with good performance and deteriorate with bad. First-term politicians are among the most electorally vulnerable and the hardest-working, and effort and electoral vulnerability both tend to decline with tenure. These dynamics enable a detailed analysis of limited voter commitment, voluntary retirement from politics, and adverse selection with politicians’ ability and effort being complementary. Our analysis highlights how a politician’s career is shaped by voters’ evolving “goodwill” toward her.

JEL Classification Codes: D72, C73, M51.

Key words: principal-agent model; repeated moral hazard; entrenchment; accountability; replacement; adverse selection; renegotiation

*We are grateful to Vincent Anesi, Dan Barron, Dan Bernhardt, Peter Buisseret, Steve Callander, Gabriel Carroll, Isa Chavez, Ernesto Dal Bó, Lucas de Lara, John Duggan, Georgy Egorov, Jim Fearon, Dana Foarta, Germán Gieczewski, Edoardo Grillo, Justin Grimmer, Marina Halac, Navin Kartik, Chad Kendall, Cesar Martinelli, Andrea Matozzi, John Matsusaka, Adam Meirowitz, Nicola Persico, Carlo Prato, Ravideep Sethi, Takuo Sugaya, and various seminar audiences for valuable conversations. Apoorva Lal, Ellen Muir, and Toby Nowacki provided excellent research assistance.

[†]Associate Professor, Department of Political Science and Graduate School of Business, Stanford University; email: avidit@stanford.edu.

[‡]Assistant Professor, Department of Economics, Columbia University; email: e.lipnowski@columbia.edu.

[§]Assistant Professor of Finance and Business Economics, Marshall School of Business, University of Southern California; email: joao.ramos@usc.edu.

1 Introduction

Since the seminal work of Barro (1973) and Ferejohn (1986), a key premise of research in political economy and public choice has been that the relationship between voters and politicians is afflicted by moral hazard. For example, if the economy stagnates under an incumbent, is it because of her mismanagement or for reasons outside her control? If the activities of politicians are hard to monitor, and the mapping from their choices to public outcomes is noisy, then how should voters hold politicians accountable?

Although these questions lie at the heart of political economy and public choice, little is known about how voters should best incentivize politicians when they are engaged in a potentially long-term relationship. Past work (surveyed thoroughly by Duggan and Martinelli (2017) and Ashworth (2012)) has either looked at settings with short time horizons or restricted attention to stationary incentives, even though optimal incentives typically condition on a politician's broader history of performance in office. From a positive standpoint, there may be good reasons to restrict attention to stationary incentives. The simplicity of these equilibria, for example, may capture the challenges that a large electorate faces in coordinating a precise calibration of politician incentives. But restricting the analysis to stationary incentives prevents us from being able to study how much citizen welfare is potentially lost as a result. The normative question of what optimal incentives look like, and how voters and politicians behave under these incentives, has been left open.

In this paper, we re-examine the canonical Ferejohn (1986) setting of repeated moral hazard. The model that we consider features a representative voter and a set of politicians, one of whom is in office each period. In each period of the model we look at, the politician in office privately observes a productivity level that takes a binary value (high or low). She then chooses her effort level from a bounded range. Effort and productivity together determine performance. Given a low productivity shock, failure is certain; otherwise, performance increases linearly with effort. Crucially, the voter cannot tell whether a failure was due to low productivity or due to low effort. Finally, the voter chooses whether to retain the politician or replace her with a new one from the pool.¹

¹For tractability, we depart from the original Ferejohn (1986) setting in two ways. First, we assume that the productivity shock is binary rather than continuous. Second, we assume that effort is bounded above, and the politician has a constant marginal cost of effort, whereas Ferejohn (1986) assumes that effort is unbounded with increasing marginal cost. We see our model as a limiting case of his—the marginal cost of effort being constant up to the bound and infinite above it. Nevertheless, we explain at the end of Section 3.1 how the main qualitative results of our analysis carry over if we assume that effort is in fact unbounded and has increasing marginal cost (while retaining the assumption of binary shocks).

We begin by characterizing the career dynamics of incumbent politicians under the voter-optimal equilibrium. We observe that if the discount factor is high, then a stationary equilibrium perfectly resolves the moral hazard problem, achieving the voter’s first-best payoff. But for lower levels of the discount factor, we show that the voter-optimal equilibrium is necessarily non-stationary. In each term of office, the politician faces an equilibrium-prescribed level of effort. She either succeeds—i.e., delivers an outcome consistent with high productivity and this effort level—or she fails. The consequences of success and failure depend on the continuation value with which she started her present term. Her continuation value improves with every success and declines with every failure, with some caveats that we will elaborate upon as follows.

On the path of play, the politician’s continuation value ranges in an interval from lowest to highest. Every politician starts her career in office at the lowest possible value in this range. A politician with the highest possible value will be forever re-elected and, therefore, always exerts zero effort: She is “fully entrenched.” Success results in certain re-election for all other values as well, and the politician starts her next term with a higher value. Failure, on the other hand, results either in removal from office or retention at a lower value. For a politician who started her current term at the lowest possible value, failure results in certain removal. For one who started her current term with a value that is higher than that value but below a certain threshold, failure results in the voter removing her from office with a probability that is decreasing in her value. If retained, she will start her next term at the lowest possible value, as if restarting her career from scratch. For a politician who started the current term at a value above the aforementioned threshold, failure results in the voter re-electing her, but at a lower continuation value.

In addition to the threshold mentioned above, there is another important equilibrium threshold. If the politician starts a term with a value that is above this threshold, then success in that term results in complete entrenchment—that is, indefinite job security. With enough consecutive successes, her value can indeed grow to exceed this key threshold. The voter-optimal equilibrium path of play is unique, and the necessary number of consecutive successes to become completely entrenched is finitely bounded, so in any voter-optimal equilibrium, some politician is eventually completely entrenched. This feature implies that one cannot conclude from the fact that a politician becomes entrenched that accountability was not at work.

The voter-optimal equilibrium grants the voter a high degree of commitment power, as it assumes continuation play constitutes a suboptimal equilibrium from many histories. A common rationale for stationary equilibria is that they home in on circumstances—germane

to political economy—in which such commitment is not possible (see, e.g., Duggan and Martinelli (2017)). We examine this view by limiting voter commitment according to (an adaptation of) the notion of renegotiation-proofness proposed by Pearce (1987) for games with repeated structure. In short, the premise is that voters will renegotiate away from anticipated low-payoff continuation play if an alternative equilibrium exists that generates a high payoff uniformly across all histories. We establish that, for some parameter values, stationary equilibria are not renegotiation-proof in this sense. In fact, the main qualitative features of voter-optimal equilibria described above carry over to all renegotiation-proof equilibria as well, except that no politician is ever completely entrenched on the path of play.² Hence, far from being ruled out by a collective commitment concern, richer dynamic equilibria are sometimes explicitly selected by it.

We also enrich the model in two directions, showing the versatility of the canonical accountability model (shed of its traditional restriction to stationary play) and generating new economic insights as we add features to the original framework. First, we allow politicians in office to voluntarily retire if they encounter a sufficiently attractive outside option. Optimal career dynamics in such a model are again similar to that of the baseline model, but the possibility of a sufficiently attractive outside offer precludes the possibility of politician entrenchment, as the incumbent will eventually voluntarily retire. In addition, we find that opportunities for politicians to work outside politics can be both helpful and harmful to voters’ interests. On the one hand, if politicians expect to voluntarily exit from politics soon, this makes it harder for the voter to provide dynamic incentives to the politician with the promise of future rewards. On the other hand, if sitting politicians are more likely than those who have been removed from office to receive these lucrative opportunities, then they have an additional benefit from office, the cost of which is not borne by the voter. Since a politician can stay in office solely at the pleasure of the voter, this gives her an additional incentive to serve voters’ interests. Thus, private-sector opportunities generate a countervailing economic force that can, in fact, aid accountability.

Second, we consider a setting with both moral hazard and adverse selection in which politicians can be either “good” (high ability) or “bad” (low ability) types; effort and ability are complements; and bad types always fail. Ashworth et al. (2017) find that whether the voter faces a tradeoff between sanctioning poor performance and selecting good types depends on whether effort and ability are complements or substitutes. If they are complements, then providing incentives serves to distinguish good types from bad types, so that

²In particular, this result implies that for the purpose of limiting entrenchment, imposing term limits of any length is also suboptimal.

no such tradeoff exists. But in the dynamic setting, we show that this tradeoff can resurface, albeit in a different form, even under such complementarity. We first show that the voter cycles through politicians until he discovers one that is good, and equilibrium play from that point on is similar to that of the baseline model.³ However, if the good type is rare in the pool of politicians, then removing a politician who is known to be good comes with the cost of having to wait through a potentially long spell of low performance. Since incentivizing more work from the politician entails lowering her initial continuation value (and, consequently, firing her sooner in expectation), the voter’s selection problem makes it costlier for him to provide future incentives. Therefore, an intertemporal tradeoff emerges between sanctioning and selection, even when effort and ability are complements.

Related literature— As we mentioned above, models in which politicians are unrestricted in the number of terms they can serve typically focus on stationary equilibria, following the early work of Barro (1973), Ferejohn (1986), Banks and Sundaram (1990, 1993) and others. A notable exception is Schwabe (2009), who allows politicians to serve an unrestricted number of terms but restricts attention to a class of equilibria that are suboptimal for the voter. In recent work, Anesi and Buisseret (2020) and Kartik and Van Weelden (2019) also allow for an unrestricted number of terms, studying models with both adverse selection and moral hazard. Anesi and Buisseret (2020), however, restrict attention to the case of minimal discounting, whereas Kartik and Van Weelden (2019), who look at settings both with and without term limits, focus on a class of Markovian equilibria in the latter.⁴ Our paper closely relates to this prior work, but focuses on long-run career dynamics under optimal political accountability when stationary equilibria are not optimal.

Our work also connects to prior work on political careers. Like much of the above literature, past work on political careers also looks mainly at models in which politicians serve only a restricted number of terms, or the relationship between voters and politicians is stationary. Ashworth (2005), for example, studies a three-period career concerns model in which politicians allocate effort across different activities, focusing on how career stage determines this allocation. Mattozzi and Merlo (2008) study a two-period model in which politicians decide whether to enter/remain in politics or work in the private sector.

Finally, our paper relates to prior work on repeated moral hazard. It is most closely related to work studying principal-optimal equilibria in settings in which the principal is

³Since bad types always fail, any success perfectly certifies a politician as a good type.

⁴In a setting in which the agents/politicians have heterogeneous expertise in assessing risky decisions, Aghion and Jackson (2016) show tenure contracts are approximately optimal for a patient principal/voter, and explore the value of term limits when the voter cannot commit.

unable to finely adjust the agent’s compensation. This includes the work on delegation by Lipnowski and Ramos (2020), Li et al. (2017) and Guo and Hörner (2020), as well as other studies in political economy, including those on indirect control and war by Padró i Miquel and Yared (2012) and Yared (2010) and a recent paper by Foarta and Sugaya (2019), who look at principal-optimal equilibria in an intervention game. In these contributions, as in our work, the standard recursive toolbox developed by Spear and Srivastava (1987) and Abreu et al. (1990) facilitates the analysis of how the future terms of a relationship can substitute for monetary incentives. These techniques have been applied much more broadly in a variety of settings—for example, in related work by Fong and Li (2017), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007), Thomas and Worrall (1990) and Atkeson and Lucas (1992), which all share with our work the feature that utility is imperfectly transferable due to limited liability or risk aversion.

In our model, the principal has only the simple choice of whether to fire or retain the agent—the voter’s so-called “blunt tool.” This choice—i.e., the provision of job security—is the main incentivizing device in many other settings besides ours, where the remuneration of workers is fixed and does not vary with performance through options, bonuses, commissions, and the like. In particular, another relevant application of our setup is to the accountability of public employees (government bureaucrats) whose salaries are fixed. To the extent that their bosses can provide incentives, it is through the provision of job security, e.g., the threat of being transferred to undesirable postings, which is relevant even in cases where it is difficult to outright fire these public sector employees.

In addition to being a realistic feature of many agency relationships, the possibility of replacement furnishes additional structure by endogenizing the principal’s outside option. This additional structure shapes substantive features of the agent’s career in our model—e.g., ensuring agents are fired for a first-term failure in our baseline model, and ensuring that the principal benefits from voluntary agent retirement. Moreover, this structure renders certain rich relational phenomena (renegotiation by principals and the simultaneity of moral hazard and adverse selection) especially tractable to analyze.

2 Political Accountability under Moral Hazard

2.1 Baseline Model

The players in our model are a representative voter and an infinite collection of politicians. Time is discrete, with an infinite horizon and indexed by $t \in \{0, 1, 2, \dots\}$. Each period

starts with a politician in office. One of the politicians begins in office at date 0. A productivity shock θ_t is drawn from $\{0, 1\}$ independently across periods and privately seen by the politician in office. The high productivity shock, $\theta_t = 1$, has probability $\mu \in (0, 1)$ and the low productivity shock, $\theta_t = 0$, has probability $1 - \mu$. After seeing the productivity level, the politician chooses how much effort to exert, $a_t \in [0, 1]$.

The voter next observes only her own payoff, the product $y_t = \theta_t a_t$, and publicly decides whether to retain the politician or remove her from office. Slightly abusing terminology, we will refer to y_t as the *outcome* of period t and will often drop the time subscript when the period is clear. We denote by ρ_t the probability with which the politician is re-elected. If the voter removes the politician, she is replaced by someone from the pool and never re-enters office in the future.⁵ In each period t , the politician in office accrues a flow payoff $1 - ca_t$, where $c > 0$ is the marginal cost of effort.⁶ All politicians that are not in office get 0. All individuals discount the future using a common discount factor, $\delta \in (0, 1)$. In each period $t > 0$, we will use the term *incumbent* to refer to the politician in office in the previous period; for period $t = 0$, the incumbent is the initial politician in office. We express continuation values in terms of average flow payoffs. We look at perfect public equilibria of this model that are optimal from the voter’s perspective.⁷

Observe that the voter and politicians never move simultaneously; any deviation by the voter is publicly observable; and an equilibrium exists that gives the voter a zero payoff (e.g., the “always-shirk, always-replace” equilibrium). Thus, it follows that voter optimal equilibria are behaviorally equivalent to play under the voter-optimal strategy profile among those that respect politicians’ incentives.⁸ We will, therefore, let an incentive-compatible

⁵The latter is without loss for voter utility: for any equilibrium in which the politician is re-elected after some delay, some weakly better equilibrium exists in which delayed re-election never occurs.

⁶We have thus fixed the office benefit to 1 in addition to fixing the maximum effort bound to also be 1, but this is without loss of generality. A setting in which the politician’s office benefit is $B > 0$, marginal cost of effort is $C > 0$, and maximum politician effort in a period is $\bar{a} > 0$ is equivalent to ours if we divide politician utility by B , voter utility by \bar{a} , and politician effort by \bar{a} —leading to a specification of our model with $c = \bar{a}C/B$. Comparative statics of various equilibrium quantities with respect to B , C and \bar{a} can be taken after accounting for these rescalings. For example, the voter’s highest equilibrium value is increasing in \bar{a} , increasing in B , and decreasing in C .

⁷See Definitions 5.2 and 5.3 in Fudenberg and Tirole (1991) for the relevant definition of a perfect public equilibrium.

⁸This equivalence follows directly when the voter uses a pure strategy. Voter incentives might still matter in principle, as the voter must be indifferent when called upon to mix. However, as our analysis establishes that such indifference will automatically be satisfied in the commitment solution, this added constraint imposes no further limits on attainable voter outcomes in this model. In the environment of Section 5 (with adverse selection), this equivalence requires a public randomization device because the voter would strictly prefer to retain the politician when called upon to mix.

(IC) policy be a strategy profile such that no politician has an incentive to deviate from her prescribed conditional effort choice at any history.

We now observe that a special kind of IC policy (and, hence, equilibrium) is without loss of optimality from the voter's perspective. Specifically, each politician uses a pure strategy; each public history h has some effort $a[h]$ such that the politician in office exerts effort zero, given a bad shock in the history- h period, and effort $a[h]$, given a good shock in the history- h period; and continuation play is identical for all outcomes $y \geq a[h]$ and identical again for all outcomes $y < a[h]$. We will refer to $a[h]$ as the equilibrium-prescribed level of effort at history h ; and we will refer to any outcome $y \geq a[h]$ as a *success* and to any outcome $y < a[h]$ as a *failure*. For such a policy, it is apparent that a politician's incentives will be satisfied as long as she always willingly chooses the equilibrium-prescribed effort over an effort of zero, following a high shock. Given this form, from now on, we will abuse notation and interpret a politician's choice at public history h to be her choice at history h conditional on a high-productivity shock in that period.

To see that the above-described policies are without loss of optimality, first note that a politician can always use a pure strategy when the voter is committed: If she were mixing over multiple pure strategies, choosing a voter's favorite among them would offer a weak improvement.⁹ Moreover, working is strictly dominated for a politician when productivity is low, since it leads to an inferior payoff in the stage game with no effect on the public outcome of the dynamic game. Thus, when productivity is low, the politician in office chooses not to exert any effort. Let $a[h]$ denote the effort that the politician chooses following a high-productivity shock, and define success and failure as in the previous paragraph. We now consider a modification to the voter's strategy in which he responds to every failure as if the outcome were 0 and every success as if it were $a[h]$. This policy clearly generates the same voter payoff as the original policy, so all that remains is to verify that it is IC. To see that it is, note that under the modified voter behavior, a politician always prefers to choose 0 over $a \in (0, a[h])$ and $a[h]$ over $a > a[h]$. Thus, the modified policy is IC as long as the politician weakly prefers $a[h]$ to zero following a high shock. As these two choices produce outcomes that are treated in the same way as under the original policy (which was assumed to be IC), the modified policy is IC as well. Hence, the given form of equilibrium is without loss in terms of voter payoffs. Throughout the paper, we focus only on such policies.

⁹A more formal argument is as follows. At a given history and following a given current shock, each effort level \tilde{a} generates some continuation payoff $\pi(\tilde{a}) \in [0, 1]$ for the voter, and this function $\pi(\cdot)$ is measurable given that strategies are. If the incumbent were to mix according to measure $\eta \in \Delta[0, 1]$, then the set $\{a \in [0, 1] : \pi(a) \geq \int \pi(\tilde{a}) d\eta(\tilde{a})\}$ would have positive η -probability, and so some choice from this set would maintain incentive compatibility and weakly raise the voter's payoff.

Finally, note that the voter's highest feasible payoff is μ . His first-best payoff is equal to $\mu \min\{1, 1/c\}$, which obtains if and only if every politician in office exerts the highest effort satisfying participation ($a = \min\{1, 1/c\}$) in every high-productivity period.

2.2 Stationary Incentives

Without loss of optimality for the voter, voter behavior at any history is a step function in current outcome y that changes only at the prescribed effort level at that history. Thus, a stationary equilibrium is parametrized by a stationary effort prescription a and a stationary pair of firing probabilities ψ_0 and ψ_a , such that the politician in office is removed with probability ψ_0 following $y < a$, and with probability ψ_a following $y \geq a$. The politician finds it optimal to exert effort $a_t = a\theta_t$ in all periods t , and the voter's equilibrium payoff, μa , is increasing in the prescribed effort level.

In addition, observe that either $\psi_0 = 1$ or $\psi_a = 0$ can be used without loss of optimality. Indeed, if both ψ_0 and ψ_a were interior, then raising ψ_0 while lowering ψ_a , so that the on-path firing rate remains constant, would preserve the politician's incentive to choose a rather than deviate to a period of shirking. Moreover, even if $\psi_0 = 1$, then lowering ψ_a would strengthen the politician's incentives by raising the value of working without changing the value of shirking. Therefore, we may take $\psi_a = 0$ without loss of optimality. The politician is, thus, retained for outcome a and fired with probability ψ_0 for outcome zero. Hereafter, we will drop the subscript and let $\psi = \psi_0$.

Facing a stationary firing rule, the politician's best response is stationary. She either works for a value satisfying $u_W = (1 - \delta)(1 - c\mu a) + \delta[1 - \psi(1 - \mu)]u_W$ or shirks for a value satisfying $u_S = (1 - \delta) + \delta(1 - \psi)u_S$. Rearranging yields

$$\frac{u_W}{u_S} = (1 - c\mu a) \frac{(1 - \delta) + \delta\psi}{(1 - \delta) + \delta\psi(1 - \mu)}.$$

Work is IC for the politician if and only if $u_W/u_S \geq 1$. As the ratio is increasing in ψ , one optimally sets $\psi = 1$. In this case, $u_W/u_S \geq 1$ if and only if $ca \leq \delta$. Thus, the optimal a is the smaller of 1 and δ/c . The key findings are as follows.

Proposition 0. *In a voter-optimal stationary equilibrium, for a politician in office in period t , we have:*

1. *The politician chooses effort $a_t = \theta_t a^*$, where $a^* = \min\{1, \delta/c\}$.*
2. *The politician is fired if she fails ($y_t < a^*$) and retained if she succeeds ($y_t \geq a^*$).*

The expected equilibrium payoff to the voter is $\mu a^* = \mu \min\{1, \delta/c\}$. In particular, a stationary equilibrium exists that achieves the voter’s first-best payoff if and only if $\delta \geq c$.

Thus far, we have replicated the main features of the Ferejohn (1986) analysis, which also focuses on optimal stationary equilibria and shows that they can be characterized by a cutoff retention rule. However, despite the environment being stationary, voter-optimal equilibria need not be stationary when stationary equilibria do not completely resolve the moral hazard problem—that is, for values of the discount factor lower than c .

3 Career Dynamics

3.1 Non-stationary Incentives

In the previous section, we showed that if $\delta \geq c$, a stationary equilibrium achieves the voter’s first-best payoff, so in this case, some stationary equilibrium is obviously voter-optimal. We now look at the case in which $\delta < c$ and study political career dynamics under voter-optimal equilibrium play.

We take a recursive approach à la Spear and Srivastava (1987). Note that there is an IC policy generating continuation value u for the politician in office if and only if $u \in [1 - \delta, 1]$. The politician can guarantee herself a payoff of at least $1 - \delta$ by shirking indefinitely, and a maximum flow payoff of 1 can be obtained by accruing the benefits of office without exerting effort. Any value in the interval can be generated by randomizing between two IC extremes—immediate replacement after the first term and unconditional entrenchment.

Given the parameters δ and c , for any politician utility u , prescribed effort a , and outcome y , it will be convenient for our main results to define the quantity

$$v_u(a, y) := \begin{cases} \frac{1}{\delta}[u - (1 - \delta)(1 - ca)] & \text{if } y \geq a \\ \frac{1}{\delta}[u - (1 - \delta)] & \text{if } y < a. \end{cases}$$

Simple algebra shows that, if a politician’s current continuation value is u , her equilibrium-prescribed effort is a , and she is (conditional on a high shock) indifferent between choosing said effort and shirking, then $v_u(a, \cdot)$ describes her average continuation value from the next period as a function of the current period’s public outcome.

Proposition 1. *Suppose that $\delta < c$. In a voter-optimal equilibrium, every first-term politician takes office with continuation value $u = 1 - \delta$; and for any politician who starts any period t with continuation value $u \in [1 - \delta, 1]$, we have:*

1. The politician chooses $a_t = \theta_t a_u$, where $a_u = \min \left\{ 1, \frac{1-u}{(1-\delta)c} \right\}$.
2. Given the outcome y in period t , if $v_u(a_u, y) \geq 1 - \delta$, then the politician is retained with probability 1 at continuation value $v_u(a_u, y)$ starting from the next period. If $v_u(a_u, y) < 1 - \delta$, she is retained with probability $v_u(a_u, y)/(1 - \delta)$ at continuation value $1 - \delta$ from the next period.

Moreover, every voter-optimal equilibrium generates the same distribution over the possible paths of play.

In short, the play described in Proposition 1 has (given the politician's current continuation value) the politician working as hard as can be made incentive compatible—and hence being made indifferent between working and shirking—and the voter replacing the politician as little as possible.¹⁰

We now make a series of observations about this voter-optimal equilibrium. On the equilibrium path, the voter prescribes an effort level a_u to the politician, which the politician chooses when productivity is high. When productivity is low, the politician optimally chooses not to exert any effort. Note that if $u = 1$, the equilibrium-prescribed effort level is $a_u = 0$, so the politician cannot fail. We say that a politician with continuation value $u = 1$ is *completely entrenched*: She remains in office forever, independent of all current and subsequent outcomes, and, therefore, never exerts any effort.

Next, note that a politician who is successful is re-elected for sure, with a continuation value higher than that of her previous period. Failure, however, results either in removal from office or retention at a continuation value lower than the previous period's. In particular, if $u = 1 - \delta$ (which is the lowest possible continuation value for a politician currently in office, and the one at which every politician starts her career), then failure results in certain removal. If $u \in (1 - \delta, 1 - \delta^2)$, then failure results in the politician being probabilistically retained, and if she is retained, then she starts the next term at value $1 - \delta$, as if starting her career from scratch. If $u \geq 1 - \delta^2$ and she fails, then she is retained for sure at a value that is lower than the value with which she started the term, but higher than that at which she started her career.

Importantly, as a politician's continuation value climbs with an increasingly long run of successes, it is possible for her to become completely entrenched on the path of play. In particular, whenever the politician's continuation value is above $1 - (1 - \delta)c$, being

¹⁰It is also worth noting that although stationary equilibria become optimal at $\delta = c$, the model's predictions are not discontinuous at this point. For example, the probability of a given politician becoming completely entrenched converges to zero as $\delta < c$ converges to c .

successful in the current term results in complete entrenchment. Indeed, if $u > 1 - (1 - \delta)c$, then the prescribed effort is $a_u < 1$, and, hence, $v_u = 1$ if the politician succeeds. A related observation is that all politicians who are more than one successful term away from complete entrenchment are prescribed the maximum feasible effort level of $a = 1$; and if a politician is only one successful term away from complete entrenchment, then she is prescribed an effort level that is either strictly less than 1 because u exceeds $1 - (1 - \delta)c$ or is equal to 1 because u exactly equals $1 - (1 - \delta)c$. The corollary below shows that with a long enough run of successes, the politician's continuation value can, indeed, cross this threshold.

Corollary 1. *Suppose that $\delta < c$. Then, the sequence of cutoffs $\{\frac{\delta}{1-\delta^k}\}_{k=1}^\infty$ is such that, if $\frac{\delta}{1-\delta^{k+1}} < c \leq \frac{\delta}{1-\delta^k}$, then in the voter-optimal equilibrium, any politician who has $k \in \mathbb{N}$ consecutive successes will be completely entrenched if successful in the following period. If $c > \frac{\delta}{1-\delta}$, any politician that produces a single success in office will immediately become completely entrenched.*

Proof. A straightforward recursive calculation shows that if the politician starts with continuation value $1 - \delta$, then, following k successes, her continuation value will be

$$u_k = 1 - \frac{1}{\delta^{k-1}} + \left(\frac{1}{\delta^k} - 1\right) c.$$

Thus, $u_k \geq 1 - (1 - \delta)c$ if and only if $c \geq \delta/(1 - \delta^{k+1})$. Therefore, if $\frac{\delta}{1-\delta^{k+1}} < c \leq \frac{\delta}{1-\delta^k}$, then any politician with k past successes will be completely entrenched after her next success. If $c > \delta/(1 - \delta)$, then every politician will be completely entrenched after her first success. ■

Because the sequence of cutoffs in the corollary converges to δ , for any value of $c > \delta$, some politician will eventually achieve enough consecutive successes to become completely entrenched. This result implies that the optimal way for the voter to incentivize effort in the early periods is to provide the politician with the promise of entrenchment. The promise is credible given that the voter can be punished for breaking the promise—for instance, if the relationship stipulates that the continuation equilibrium has every subsequent politician exerting zero effort and being replaced with certainty.

The corollary puts an upper bound on the number of consecutive successes needed for the politician to become completely entrenched. But, in fact, a politician may become entrenched with fewer consecutive successes if she has a sequence of successes following a failure that resulted in her being re-elected with a continuation value larger than $1 - \delta$. As mentioned above, this case arises if the politician's continuation value has climbed above




Figure 1: Simulated paths of continuation values for parameter values $\mu = 0.8$, $c = 0.9$, and $\delta = 0.85$. The important thresholds for continuation value are, therefore, $1 - \delta = 0.15$, $1 - \delta^2 = 0.2775$ and $1 - (1 - \delta)c = 0.865$.

$1 - \delta^2$. That said, such an outcome requires that the politician has not yet become completely entrenched prior to this point. In particular, it is possible that the entrenchment threshold $1 - (1 - \delta)c$ mentioned above is smaller than $1 - \delta^2$, in which case the politician could be completely entrenched on the path of play without ever having experienced a term with perfect short-term job security.

Figure 1 makes our observations concrete by depicting some simulated paths of the politician’s continuation value for a numerical example. The figure shows that some politicians have longer careers than others. Some become completely entrenched more quickly; others take longer; and still others may be removed from office before they achieve full job security. In the early stages of a politician’s career, failure is more prone to result in the politician being replaced or the continuation value being “reset” to $1 - \delta$. Failure later on is less likely to result in removal or a reset since later-term career continuation values tend to be higher. A politician is more likely to serve another term the more successful she has been in the past, but even politicians who start their careers with a good run may have short careers if they are unlucky enough to experience a sequence of bad shocks. Hence, while job security tends to grow over a politician’s career in office, it need not grow monotonically

on the path of play. Similarly, a politician’s effort tends to decline over her career in office, but need not grow monotonically along every path.¹¹

Finally, let us emphasize that our observations about the path of play are not merely features of a special voter-optimal equilibrium characterized in Proposition 1. Indeed, these observations are strengthened by the fact (established in the proposition) that all voter-optimal equilibria are equivalent in terms of the paths of play that they generate. We prove this in the appendix by observing that the voter’s optimal continuation payoff, as a function of the politician’s value u , is either affine or strictly concave on $[1 - \delta, 1]$. This feature implies that the equilibrium actions solved for in Proposition 1 above are, in fact, uniquely optimal.

General Costs. Our main results transport to a more general version of our model with weakly increasing marginal effort cost (and effort allowed to be bounded or unbounded). In particular, this specification includes the case of unbounded effort with strictly convex costs, which is a step closer to the original Ferejohn (1986) formulation. For example, each history still has only one relevant IC constraint for the politician in office, which continues to hold with equality; and the politician’s continuation value following the realization of the voter’s utility within the period is still one of two quantities that we can continue to call the values after “success” and “failure.” The politician’s continuation value will still weakly rise with every success and weakly fall with every failure. The voter’s decision to retain or replace the politician will continue to reflect the same dependence on the politician’s continuation value, and the politician’s equilibrium effort will still be weakly decreasing in her continuation value. Stationary equilibria remain suboptimal whenever moral hazard binds: If effort is bounded, then a threshold discount factor exists such that stationary equilibria will achieve the first-best voter payoff above this threshold, and stationary equilibria will be strictly suboptimal below it; and if effort is unbounded, stationary equilibria will always be strictly suboptimal. While the general insights carry over, we opt to focus on the simpler setting in which effort is bounded and its marginal cost is constant because this case further enables an exact closed-form characterization of politician effort at each continuation value, and hence of the voter-optimal equilibrium.

¹¹Whereas average effort is highest for newly elected politicians and converges to its lowest level as tenure in office converges to infinity, even average effort may not decline monotonically over time, because time is discrete in our model. For example, when $\delta/(1 - \delta^2) < c < \delta/(1 - \delta)$ and μ is very small, one can verify that the average third-term politician exerts higher effort than the average second-term politician.

3.2 Limiting Voter Commitment

The results of the previous section show how the voter incentivizes work from an incumbent politician (in the early and middle parts of the politician's career) through the delayed reward of entrenchment. But when a politician is completely entrenched, the voter can expect to derive no subsequent value from the politician. While such a history generates the best possible continuation play for the incumbent politician, it is the worst possible continuation play for the voter. Can the voter commit to tolerating this low value once a politician has created sufficient past value? Though Proposition 1 describes a perfect equilibrium of the game (and so, interpreted literally, requires no commitment power), a valid concern is that voters might want to somehow coordinate on a new equilibrium when facing such a dismal future.

In this section, we limit the voter's commitment power by imposing a threshold, π , such that his equilibrium payoff cannot fall below π at any history. This approach is reminiscent of the form of renegotiation-proofness introduced by Pearce (1987). In Proposition 2, we consider exogenous thresholds below which the voter's payoff cannot fall at any equilibrium history. One can think of the *extent* of the voter's commitment power as being measured by the magnitude of π : The higher π is, the less ability the voter has to tolerate unfavorable outcomes ex post. In line with Pearce (1987), we focus in Corollary 2 on the highest (so in this sense, endogenous) threshold that can possibly be sustained, essentially allowing the voter to renegotiate away from the path whenever such renegotiation is itself credible.¹²

Given an exogenous limit to the voter's commitment, what is the form of constrained-optimal play? The next proposition answers this question, showing that the sole effect is to cap the politician's achievable continuation value.

Proposition 2. *Suppose that $\delta < c$, and fix a payoff lower bound $\pi > 0$. Some $\hat{u} \in [1 - \delta, 1)$ exists such that, in any equilibrium that is optimal for the voter, subject to the constraint that the voter's continuation payoff after every history is at least π : (i) the continuation value at any history of any politician in office lies in $[1 - \delta, \hat{u}]$; (ii) every politician who enters office starts with continuation value $1 - \delta$; and (iii) if a politician starts a period t with continuation value u , we have:*

1. *The politician chooses $a_t = \theta_t a_{u, \hat{u}}$, where $a_{u, \hat{u}} = \min\{1, \frac{(1-\delta)(1-u) + \delta(\hat{u}-u)}{(1-\delta)c}\}$.*

¹²While our approach is inspired by Pearce (1987), one important difference is that Pearce (1987) chiefly focuses on strongly symmetric equilibria of a symmetric game, while we apply analogous reasoning to a principal-agent setting under the hypothesis that the principal can unilaterally initiate renegotiation. We believe that this approach to renegotiation could be appropriate in other dynamic principal-agent settings.

2. Given the outcome y in period t , if $v_u(a_{u,\hat{u}}, y) \geq 1 - \delta$, then the politician is retained with probability 1 at continuation value $v_u(a_{u,\hat{u}}, y)$ starting from the next period. If $v_u(a_{u,\hat{u}}, y) < 1 - \delta$, then she is retained with probability $v_u(a_{u,\hat{u}}, y)/(1 - \delta)$ at continuation value $1 - \delta$ from the next period.

This payoff cap \hat{u} is smaller for higher π .

An important implication of the proposition is that the optimal way to “protect” the voter from suffering low payoffs is to directly preclude entrenchment—that is, to place a permanent upper bound on the degree of job security a politician can enjoy, whatever her history is in office. For any level of voter security π , the proposition says that an “entrenchment limit” \hat{u} exists that could assure the voter a value of π from any history, and moreover that a more stringent voter security level necessitates a more stringent limit to politician entrenchment. If a politician ever achieves the highest possible continuation value of \hat{u} , then she chooses a low effort of $(1 - \hat{u})/c$ in that period, maintains the same continuation value if she succeeds, and has her continuation value fall to the strictly lower value of $[\hat{u} - (1 - \delta)]/\delta$ if she fails. Apart from this modification, the proposition shows that the dynamics of a politician’s career are qualitatively similar to those of the baseline model.

Another implication of the proposition is that, for the purposes of limiting politician entrenchment, it is suboptimal to impose a fixed term limit of any length. This suboptimality is obvious in the case of a one-term limit, as a politician facing no dynamic incentives at all would always shirk. It is also intuitive in the case of a two-term limit. Indeed, if moral hazard is binding (that is, if $\delta < c$), then a first-term politician exerts the exact same effort level as under the best stationary equilibrium for the voter, being re-elected if she attains at least that outcome and replaced otherwise.¹³ Therefore, the best equilibrium with a two-term limit gives her a payoff of $1/(1 + \delta)$ times that of the best stationary equilibrium, and it provides a worst-case voter utility of $\delta/(1 + \delta)$ times that of the best stationary equilibrium. In a sense, the two equilibria feature the same degree of entrenchment, as all politicians in office in either of these equilibria always have a continuation value of $1 - \delta$. Consequently, imposing a two-term limit also does not permit dynamics to outperform a stationary equilibrium, for either the best-history voter value or the worst-history voter value. Moreover, given the proposition, the suboptimality of imposing a fixed term limit extends to term limits of any fixed number of terms.

¹³Let π^* be the voter-optimal equilibrium payoff for the voter if politicians face a two-term limit. Since every second-term politician shirks, the voter’s continuation payoff from retaining a politician is $\delta\pi^* < \pi^*$. Therefore, the voter optimally fires the politician for failure and makes IC bind, just as in the stationary equilibrium—and so (maximizing first-term effort subject to these two constraints) a first-term politician chooses the same effort level as in the best stationary equilibrium.

The key idea behind Proposition 2 is that the voter might renegotiate away from an equilibrium when continuation play gives him a low payoff, say $\pi_L < \bar{\pi}$. However, this renegotiation will be credible only if the proposed new equilibrium is not vulnerable to the same sort of renegotiation. Internal consistency requires that, at any possible future history, the new equilibrium continuation payoff that the voter faces is itself at least π_L because, otherwise, the voter would similarly wish to renegotiate. This internal consistency can be taken a step further: If an equilibrium exists that guarantees a continuation payoff of $\pi_H > \pi_L$ at every possible future history, then how can the equilibrium that guarantees only π_L be credible? If a history from which the voter's continuation is near π_L might again occur, would the voter not renegotiate away from this new equilibrium? And would the politicians not anticipate this? Following Pearce (1987), only one level of worst-case continuation payoff to the voter, $\bar{\pi}$, satisfies this internal consistency property: the highest feasible lower bound $\bar{\pi}$. Intuitively, a higher bound is unattainable, whereas a lower bound would not be immune to renegotiation.

One may wonder whether choosing $\bar{\pi}$ to be the endogenous payoff lower bound ultimately selects (or at least allows for) stationary equilibria. The following corollary shows that, at least for some parameter values, this is not the case.

Corollary 2. *If $1/(2 - \mu) < \delta < c$, then some $\bar{\pi}$ exists such that some equilibrium gives the voter a payoff of at least $\bar{\pi}$ after every history, while every stationary equilibrium gives the voter a payoff strictly below $\bar{\pi}$. In particular, an equilibrium of the form described in Proposition 2 exists such that voters are strictly better off after every history than under any stationary equilibrium.*

The corollary shows that allowing the voter to renegotiate away from bad equilibria (that is, limiting voter commitment) can actually *preclude* stationary play. Following Pearce's justification that renegotiation itself must be credible (which selects $\bar{\pi}$ from Proposition 2 to be as large as possible), Corollary 2 shows that equilibria featuring richer dynamics can be "more credible" for voters than stationary equilibria. Note that the assumption in the proposition that $\delta < c$ simply rules out the case in which the voter can achieve his first-best payoff. The condition $1/(2 - \mu) < \delta$ says that, despite binding moral hazard, the future is still sufficiently valuable.

Put differently, the conclusion of Corollary 2 leads us to question the conventional view that stationary equilibria describe the natural outcome when voters cannot commit in the political accountability model. The result shows that this interpretation is misleading. After all, what exactly is the voter's "inability to commit"? We have argued (following Pearce)

that it is the common understanding that the voter may, in some contingencies, desire to renegotiate the promises he has made to the politician. But from where does this desire to renegotiate arise? Presumably, it is activated by the voter facing a pessimistic enough future. Lack of commitment, then, is the voter’s inability or unwillingness to tolerate low future payoffs, which is exactly what the Pearce bound formalizes. Lack of commitment, importantly, is *not* the voter’s inability or unwillingness to condition on past history.

3.3 Political Careers in Practice

The existing literature and conventional wisdom point to a variety of factors that shape political career dynamics, including the returns to seniority in the party and the legislature;¹⁴ the increases in productivity and skill that come with job experience;¹⁵ the changes in individual, constituent and ruling party ideology;¹⁶ and the creation of a loyal voting coalition among constituents, which potentially grows more stable over time.¹⁷ Our model, on the other hand, stresses a different factor that potentially shapes political careers: the evolving dynamics of optimal political incentives under moral hazard.

In our model, politicians who enjoy long spells of good luck deliver good outcomes to their constituents and are rewarded with the promise of job security. Others are expelled from office. In the political lifecycle, as in many other settings, idiosyncratic shocks such as those in our model can confound the empirical link between political career dynamics and factors such as seniority, learning-by-doing, political consolidation, and so on.

How well do the dynamics of our model correspond to the stylized facts about political careers? Are they consistent with these facts, or do they provide countervailing pressures that work against the forces that generate these patterns?

Data collected by Hibbing (1991) show that while incumbents do very well in general, first-term members of Congress are the most electorally vulnerable. Re-election rates among those seeking a new term are generally increasing in tenure but are relatively stable after the first term. From Hibbing’s analysis, it is also apparent that political entrenchment is not an unusual or even atypical phenomenon. In Figure 2, we supplement Hibbing’s data on re-election rates among U.S. House members for the period 1946-84 with data from two other

¹⁴See, for example, McKelvey and Riezman (1992) and Muthoo and Shepsle (2014) for theoretical arguments, and evidence in Hall and Shepsle (2014), Cirone et al. (2019), and Epstein et al. (1997).

¹⁵See Padró i Miquel and Snyder (2006) for evidence that legislative effectiveness increases with tenure.

¹⁶This argument comes from the conventional view that when a politician becomes out of step with her party or with the views of her constituents, she is more likely to retire from politics; see, e.g., Jacobson (2015) and Ansolabehere and Snyder (2002).

¹⁷The idea that incumbents can consolidate voters into loyal voting blocs over the course of their tenure in office goes back to Fenno (1978).




Figure 2: Re-election rates among those seeking re-election in three countries. Data for the U.S. are House member elections from 1946-84 from Hibbing (1991). Data for Norway are from Fiva and Smith (2017) and cover the 16 parliamentary elections from 1953-2013 and politicians elected in 1953 or later. Data for the U.K. are from Eggers and Spirling (2014) and cover the six parliamentary elections from 1979-2001 and politicians elected in 1979 or later.

countries: Norway and the U.K. While the electoral and party systems in these countries differ, the pattern of re-election rates over terms in office is largely similar, suggesting that the pattern uncovered by Hibbing for the U.S. is probably not entirely driven by the specifics of the American electoral and party systems.¹⁸

Given that these patterns are consistent with explanations other than dynamic moral hazard (some of which are listed above), the main lessons of our analysis for the empirical study of accountability ought to be stated carefully. The lessons are twofold.

First, before concluding that the observed career patterns are driven by any set of factors, empirical studies of the determinants of political careers should take into account the potential for non-stationarity in the accountability relationship between politicians and their voters, and recognize that these patterns could be generated partly by the equilibrium dynamics of political accountability under moral hazard. More specifically, a researcher who does not take into account the dynamics of political accountability is implicitly making

¹⁸Norway uses a proportional representation system with closed lists. The U.K. uses simple plurality.

a nontrivial identifying assumption when interpreting the data on political careers—an assumption that should be made explicit and defended, at the very least.

Second, we cannot draw broad normative conclusions about whether or not political accountability is at work in a particular context from the fact that politicians have a lifecycle. Even under the normative benchmark of maximal accountability, politicians have evolving careers, from their first term in office, when they are most electorally vulnerable, to later terms, when they may have much greater job security. For example, although it may be tempting to infer an accountability failure from the observation that some politicians appear to have become entrenched, entrenchment alone does not provide sufficient evidence for accountability failure.

Overall, the evolving dynamics of political accountability should not be ignored when interpreting the data on political careers.

4 Voluntary Retirement

4.1 Opportunities Outside of Politics

Many politicians voluntarily retire from office and find lucrative positions outside of politics, such as in the private sector, as lobbyists, advisers, industry executives, and board members of major corporations and organizations (see, e.g., Hall and Van Houweling (1995)). Suppose that we enrich the baseline model by assuming that in each period, a politician receives with some probability the opportunity to leave her career in politics and work elsewhere. How does this possibility affect the politician’s optimal incentives and career dynamics?

Specifically, consider a modified model in which a politician randomly receives outside offers—giving her utility $w > 0$ —and can, thus, decide to voluntarily retire from politics. The rate of arrival of these offers depends on whether or not the politician is in office. An in-office politician receives an offer with probability $p \in (0, 1)$, while an out-of-office politician receives an offer with probability $\underline{p} \in [0, p)$. Without loss of generality, we normalize \underline{p} to equal zero—or, equivalently, interpret politician payoffs as being payoffs net of her continuation value after being replaced. In addition, we restrict attention to the interesting case of $w \in [1 - \delta, 1]$. If $w < 1 - \delta$, every offer is turned down, and the ability to work outside of politics has no effect on political career dynamics. If $w > 1$, every offer is accepted, and while the arrival of these offers does affect a politician’s continuation value, her career dynamics are the same as in a version of the baseline model with a lower discount factor, δ , and a higher benefit from holding office (rather than the normalized benefit of 1

in our baseline model). Finally, it will be convenient to give the players access to a public randomization device, also known as “sunspots.”¹⁹

Each period t proceeds as follows. First, the realization of the public random variable is revealed. The politician in office then either privately receives an outside offer of w with probability p , or no offer with probability $1 - p$. If the politician has received an outside offer, she decides whether to accept or reject it. If she accepts the offer, a new politician arrives and potentially receives an outside offer to accept or reject. After some politician in office either does not receive or does not accept an outside offer, she receives her office benefit of 1, observes the state θ_t , and makes an effort choice a_t at a private cost of ca_t . Then, the outcome $y_t = \theta_t a_t$ is publicly observed. Finally, the voter decides whether to re-elect the incumbent or to replace her.²⁰

Proposition 3. *Suppose that $\delta < c$ and $1 - \delta \leq w \leq 1$. In a voter-optimal equilibrium, every first-term politician starts in office with continuation value $(1 - p)(1 - \delta) + pw$, and some $v^* \in \{w, 1\}$ and $u^* \in [w, 1]$ exist such that the politician’s continuation value at the start of every period lies in the interval $[(1 - p)(1 - \delta) + pw, v^*]$, and for any politician who starts any period t with continuation value \tilde{u} , we have:*

1. *If $\tilde{u} \leq w$, then she accepts an outside offer if it arrives, and her continuation value after the outside-offer stage (if she does not accept one) is $u = \frac{\tilde{u} - pw}{1 - p}$. If $\tilde{u} > u^*$, then she does not accept the outside offer even if one arrives, and her continuation value after the outside-offer stage is $u = w$. If $\tilde{u} \in (w, u^*]$, then with probability $\frac{u^* - \tilde{u}}{u^* - w}$ (determined by the sunspot), she accepts an outside offer if it arrives, and her continuation value after the outside-offer stage (if she does not accept one) is $u = w$; with probability $\frac{\tilde{u} - w}{u^* - w}$, she does not accept the outside offer even if it arrives, and her continuation value after the outside-offer stage is $u = u^*$.*
2. *Given the politician’s continuation value u after the outside-offer stage, she chooses $a_t = \theta_t a_{u, v^*}$, where $a_{u, v^*} := \min \left\{ 1, \frac{(1 - \delta) + \delta v^* - u}{(1 - \delta)c} \right\}$.*
3. *Given the outcome y in period t , if $v_u(a_{u, v^*}, y) \geq (1 - p)(1 - \delta) + pw$, then the politician is retained with probability 1 at continuation value $v_u(a_{u, v^*}, y)$ starting from*

¹⁹This assumption is useful technically, but we do not view it as substantive. That said, if one wished to take it literally in our setting, the voter and politician could react to commonly observed global events that are unrelated to the productivity shocks that drive the relationship’s moral hazard problem. Wolfers (2002), for example, finds that voters in gubernatorial elections condition their re-election decisions on economic fluctuations that are unrelated to gubernatorial actions.

²⁰Various assumptions here are irrelevant. For example, the results are substantively identical if the voter observes the politician’s offer; if the sunspot realizes at the end of the period; if a newly elected politician cannot condition on a sunspot; or if a newly elected politician cannot receive an outside offer.

the next period. If $v_u(a_{u,v^*}, y) < (1-p)(1-\delta) + pw$, she is retained with probability $v_u(a_{u,v^*}, y)/[(1-p)(1-\delta) + pw]$ at continuation value $(1-p)(1-\delta) + pw$ from the next period.

The proposition characterizes a voter-optimal equilibrium up to two details. The first unspecified feature is whether the highest admitted continuation value for a politician is w or 1. Second, if it is the latter, one must specify the length of the interval (w, u^*) in which the politician mixes between taking and leaving the outside offer.

When the highest possible continuation value for the politician is 1, effort decisions are the same as in the baseline model, conditional on the politician's continuation value. The politician and the voter use the public randomization device only to coordinate on continuation play in the interval above w in which the politician mixes between taking and leaving the outside offer. In this case, note that the voter has the discrete benefit of having the politician leave office by taking the outside offer and starting afresh with a new politician to whom he owes nothing.

When the politician's highest possible continuation value is w , effort is reined in (just as it is in Proposition 2 to limit entrenchment) to keep the politician's value from exceeding w . This way, the voter benefits from the politician taking the outside offer whenever it arrives, again due to the benefit of getting to start with a new politician who enters office at the lowest possible continuation value on the equilibrium path. Here, because the politician's continuation value is always below w , the public randomization device is never used.

Finally, note that when $v^* = w$, politicians never become entrenched, accepting every outside offer that arrives on the path of play. They will vacuously do so when $w = 1$, but when is this the case more generally? The next corollary establishes that a sufficient condition for this property is that $w < 1$ is high enough.

Corollary 3. *Suppose that $\delta < c$. If $w < 1$ is sufficiently high, then a voter-optimal equilibrium exists such that every politician has a hazard rate of at least p of leaving office in any given period. In particular, no politician is ever fully entrenched.*

Holding the politician to a maximal continuation value of w rather than 1 entails a cost of reducing the short-term effort that can be sustained, with the benefit of allowing the voter a fresh start with a new politician rather than retaining an entrenched one. To see why $w \approx 1$ should be sufficient, then, note that the incentive cost is very small in this case, while the benefit is bounded away from zero.

In sum, the model here sheds light on two countervailing economic forces, generating a cost and a benefit to the voter when public office generates attractive outside offers (i.e.,

when p and w are high). The cost is that the effective discount factor is only $\delta(1 - p) < \delta$ whenever the politician’s continuation value is below w because she will leave office with a hazard rate p . Intuitively, if the politician is less likely to be beholden to the voter for long, the voter has limited scope for providing dynamic incentives. The benefit is that outside offers give the voter a cost-effective way to reward a successful politician. Indeed, the voter provides incentives in the current period by meeting failure with future punishment and meeting success with future rewards. But in our baseline model, conferring this reward upon the politician is costly to the voter—having the incumbent shirk in office rather than starting anew with another politician. If, instead, the incumbent takes a lucrative outside offer, the voter outsources an otherwise costly reward to the private sector.

4.2 Retirements in Practice

In our model of voluntary retirement, electoral insecurity (as measured by the probability that the voter will replace the incumbent politician) is decreasing in the politician’s continuation value from staying in politics. In recent work, Egerod (2021) shows that when indicators of job prospects outside of politics improve, sitting politicians (senate and House members in the U.S.) become more likely to retire from politics. In past research, Hall and Van Houweling (1995) note that politicians that are more electorally insecure are more likely to retire (see, also, Groseclose and Krehbiel, 1994). Diermeier et al. (2005) estimate that the effect of wages on decisions to run for re-election go in the predicted direction, but the effect size is overall quite modest. In follow-up work, Keane and Merlo (2010) unpack these estimates to find that decreasing politicians’ wages induces retirement from the most-skilled politicians and from those who were relatively young when first elected (but not those who value legislative accomplishment the most).²¹ These findings lend some support to our predictions, but more work is needed to relate these features to accountability-induced political career dynamics.²²

²¹In other related work, Besley (2004) finds that higher-paid U.S. governors are better representatives of their voters, as measured by “policy congruence.” Caselli and Morelli (2004) and Messner and Polborn (2004) examine the dual question of how the wages of office affect entry into politics, and, more recently, Hall (2019) finds that devaluing political office affects entry decisions, increasing the share of more-extreme politicians among those that run for office.

²²An additional question that our model leaves unexplored is how politicians trading favors with private interest can undermine accountability. In our model, the actions of a politician in office are unrelated to her current chances of receiving outside opportunities to work, but an important feature of interest group politics is this precise connection. Our model isolates the *ceteris paribus* effects of outside opportunities for politicians, holding this connection fixed. The economic forces that we highlight would still be present (but operate alongside potentially new forces) if, for example, we made p depend on a_t .

5 Moral Hazard and Adverse Selection

5.1 Effort and Ability

We now enrich the baseline model by assuming that each politician has a perfectly persistent skill level $\omega \in \{0, 1\}$, which only she knows as of her first term in office. Skill is independent across politicians and takes the high value of 1 with probability $q \in (0, 1)$. A politician's payoffs are as before, but the outcome in each period t produced by a type ω office-holder is $y_t = \omega a_t \theta_t$. Thus, low-skilled politicians are not able to produce successes. An important consequence of this assumption is that skill and effort are complements.

As in the previous section, it will be convenient here to assume that the voter and politicians have access to a public correlation device (see footnote 8). Specifically, we assume that a uniformly distributed random variable realizes after the politician's effort decision but before the voter's re-election decision.

A starting observation is that stationary play is again suboptimal if $\delta < c$, as in the baseline model. In particular, two potential sources of non-stationarity exist under a voter-optimal equilibrium—dynamic moral hazard, as in our baseline model, and a learning problem that is itself inherently dynamic.

Another observation is that learning in this model features, in the terminology of the strategic experimentation literature, “perfectly revealing good news.” Under this special structure, the politician, after her first term in office, is either known to be the high-skilled type (due to a past success) or is less likely to be a high-skilled type than anyone from the pool of challengers. A feature of the baseline model that carries over to this model is, therefore, that if a politician fails in her first term, then she is replaced.

Lemma 1. *In a voter-optimal equilibrium, any politician who fails in her first term is removed from office.*

The observation implies that the methods that we used to characterize the voter-optimal equilibrium in the baseline setup can be adopted here, as well. Other than in a politician's first term in office, she is known to be good. Therefore, the learning problem need not make an appearance in our recursive analysis.

Proposition 4. *Suppose that $\delta < c$. Some $u^s \in [1 - \delta, 1]$ exists such that in the voter-optimal equilibrium:*

1. *If period t is her first term in office, a politician of type ω chooses $a_t = \omega \theta_t a_*$, where $a_* := \min\{1, \frac{\delta}{(1-\delta)c}\}$.*

2. If a politician fails in her first term in office ($y_0 < a_*$) then she is replaced; and if she succeeds ($y_0 \geq a_*$), then she is retained with continuation value u^s . Thus, her continuation value at the start of her career is $u_0 = (1 - \delta)(1 - \mu c a_*) + \mu \delta u^s$ if she is of high skill and $1 - \delta$ if she is of low skill.
3. If a politician who has previously produced a success starts a term in office with continuation value $u \in [1 - \delta, 1]$, then as in the baseline model, in period t :
 - (a) The politician chooses $a_t = \theta a_u$ when the state is θ , where $a_u := \min\{1, \frac{1-u}{(1-\delta)c}\}$.
 - (b) Given the outcome y in period t , if $v_u(a_u, y) \geq 1 - \delta$, then the politician is retained with probability 1 at continuation value $v_u(a_u, y)$ from the next period. If $v_u(a_u, y) < 1 - \delta$, then the politician is retained with probability $v_u(a_u, y)/(1 - \delta)$ (determined by the sunspot) at continuation value $1 - \delta$.

The proposition implies that as long as the voter does not know the type of the politician in office, he cycles through politicians until he discovers one that is good. Since bad types always fail, this discovery happens the moment a politician succeeds. After that point, the good politician's career is similar to that in the baseline model, with one important difference: It may be optimal for the voter to start the politician's career at a continuation value that is higher than $1 - \delta$ —or, equivalently, to provide a newly elected good politician strict incentives to work. The reason for this possibility is an effective replacement cost. When the voter boots the incumbent out of office, he gets a flow payoff of zero until he draws a skilled type (and high productivity). Since this may take some time, removing a politician that is known to be good is costly. Based on this cost, the following corollary provides a sufficient condition for the good type's starting continuation value to exceed $1 - \delta$, that is, for a newly elected politician to retain some rents.

Corollary 4. *If $\delta < c$ is close enough to c , and μ and q are small enough, then in the voter-optimal equilibrium a newly elected politician has continuation value strictly above $1 - \delta$ and, given a high productivity shock, has strict incentives to work.*

The corollary considers which continuation value for the politician best serves the voter. If that value is low, the politician will likely soon be fired. If it is high, she will likely soon be entrenched. The voter-optimal continuation value for the politician optimally resolves this tradeoff. In particular, the corollary provides sufficient conditions for this optimal value u^s to be strictly greater than $c(1 - \delta)/\delta$ —that is, higher than the continuation value of a politician who, in the previous period, was to be fired for a failure, was made indifferent between maximum and minimum effort, and was successful.

To obtain these sufficient conditions, the corollary considers the case in which q and μ are low. For low values of q , the cost of firing the politician is high, as it is unlikely that the voter will discover a high-skilled politician soon after removing the incumbent. For low values of μ , the risk of having to soon reward the politician with entrenchment (with the associated costs to the voter that this entails) is also low, as the likelihood of a success—even if the politician is of high quality—is low. Therefore, at least in the range where the politician exerts full effort, the voter prefers the politician’s continuation value to be as high as possible, as this provides more periods of productive cooperation.

Finally, if δ is close to c , then the range in which the politician exerts full effort will have $c(1 - \delta)/\delta$ in its interior. In this case, the voter can provide a skilled politician with strict incentives to work in the first term and still incentivize maximal effort in the next. Hence, when the cost of firing is high and risk of entrenchment is low, the politician’s “one-success” continuation value is lower than the voter would like the politician to have. The upshot is that under the conditions of Corollary 4, the working strategy is strictly incentive-compatible for a high-skilled first-term politician in a voter-optimal equilibrium.

5.2 The Sanctioning-Selection Tradeoff

In the standard two-period accountability model, when effort and ability are complements and bad types always fail (as we have assumed here), Ashworth et al. (2017) show that no tension exists between selecting good types and sanctioning poor performance. In the case of the two-period model, if the voter considers the good type to be better than the bad type even when the good type shirks, then the voter has an incentive to learn the politician’s type even if neither type works in the second period. By incentivizing first-period effort from the good type, the voter is simultaneously able to separate the two types.

This result of Ashworth et al. (2017) contrasts with the case studied by Fearon (1999) and Duggan and Martinelli (2015) in which effort and ability are substitutes. For example, suppose that the good type is a commitment type that always works, and the bad type is a strategic type that works only in response to incentives. If the voter succeeds in incentivizing work from the bad type in the first period, he makes the bad type look more like the good type in the first period, undermining his ability to screen the two types. Thus, when effort and ability are substitutes, there is a tradeoff between selecting good types and sanctioning bad performance.

However, our dynamic setting exhibits a novel tradeoff between sanctioning poor performance and selecting good types—that can arise even though effort and ability are complements. In particular, the fact that the voter has a selection problem at present means

that he has less latitude to provide incentives in the *future*. The reason is precisely the replacement cost that we mentioned earlier. If the voter replaces a politician who is known to be good, then it may take a while before he draws another good type. This makes the voter reluctant to replace a politician he knows to be good. The more reluctant the voter is to replace a known good type, the higher is the good type politician’s continuation value. And the higher this continuation value, the less the politician tends to work in the future. This intertemporal tradeoff between selection and sanctioning is a novel economic phenomenon that arises in the multi-period model but is absent from the two-period model.

6 Conclusion

What do optimal political career dynamics look like in the classical political accountability model with moral hazard?

Our answer to this question highlights the importance of the politician’s continuation value as a state variable in the accountability model. We can think of this quantity as a measure of the voter’s goodwill toward the politician. When the politician’s continuation value is high, the voter has (credibly) promised the politician many future terms in office with little effort. When it is low, the politician is either unlikely to survive in office or expected to work hard in future periods (or both). A voter’s goodwill toward a politician rises with good performance and declines with bad performance. The evolution of goodwill affects the politician’s career, even absent other forces (exogenous and, thus, held constant in our model) that shape political careers.²³

We have shown that these key features of our analysis remain as we add other realistic features to the classic Ferejohn (1986) model of political accountability, enabling novel analysis of their interaction with moral hazard. In particular, similarly rich dynamics arise also under the possibility of renegotiation by voters, voluntary strategic retirement by politicians, or adverse selection with politicians’ effort and ability being complementary—and the precise form of optimal career dynamics is altered in interpretable ways. Overall, these results demonstrate the versatility of the canonical accountability model when we move beyond the traditional focus on stationary incentives.

Finally, although our contribution is primarily a normative one (as we stated at the onset), we argued that the political career dynamics that appear in the voter-optimal equilib-

²³Although our baseline model also exhibits complete politician entrenchment on the path of play, we do not emphasize this result in summing up our contribution because it is not robust to limiting voter commitment or allowing for voluntary retirement. On the other hand, the two features that we do highlight in this paragraph do hold throughout the variants that we study.

rium are consistent with what we see in practice—arguably more consistent than stationary equilibrium, in which career dynamics are absent. Our purpose is not to “validate” a single definitive theory, since these empirical facts are also consistent with other explanations. Nevertheless, our paper also makes a modest positive contribution, because, at the very least, it raises new challenges to the empirical study of accountability. In particular, it implies that empirical research is now burdened with the task of showing (with evidence) or assuming (with justified identifying assumptions) that political careers in practice are *not* driven by the dynamics of optimal political accountability under moral hazard.

References

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–63.
- AGHION, P. AND M. O. JACKSON (2016): “Inducing leaders to take risky decisions: dismissal, tenure, and term limits,” *American Economic Journal: Microeconomics*, 8, 1–38.
- ANESI, V. AND P. BUISSERET (2020): “Making elections work: Accountability with selection and control,” Tech. rep., Mimeo, University of Chicago.
- ANSOLABEHERE, S. AND J. M. SNYDER (2002): “The incumbency advantage in US elections: An analysis of state and federal offices, 1942–2000,” *Election law journal*, 1, 315–338.
- ASHWORTH, S. (2005): “Reputational dynamics and political careers,” *Journal of Law, Economics, and Organization*, 21, 441–466.
- (2012): “Electoral accountability: recent theoretical and empirical work,” *Annual Review of Political Science*, 15, 183–201.
- ASHWORTH, S., E. BUENO DE MESQUITA, AND A. FRIEDENBERG (2017): “Accountability and information in elections,” *American Economic Journal: Microeconomics*, 9, 95–138.
- ATKESON, A. AND R. E. LUCAS (1992): “On efficient distribution with private information,” *The Review of Economic Studies*, 59, 427–453.
- BANKS, J. S. AND R. K. SUNDARAM (1990): “Incumbents, challengers, and bandits: Bayesian learning in a dynamic choice model,” Tech. rep., University of Rochester-Center for Economic Research (RCER).

- (1993): “Adverse selection and moral hazard in a repeated elections model,” *ch*, 12, 295–311.
- BARRO, R. J. (1973): “The control of politicians: an economic model,” *Public choice*, 19–42.
- BESLEY, T. (2004): “Paying politicians: theory and evidence,” *Journal of the European Economic Association*, 2, 193–215.
- CASELLI, F. AND M. MORELLI (2004): “Bad politicians,” *Journal of Public Economics*, 88, 759–782.
- CIRONE, A., G. W. COX, AND J. H. FIVA (2019): “Seniority-based Nominations and Political Careers,” Tech. rep., Technical report. mimeo, BI Norwegian Business School.
- CLEMENTI, G. AND H. HOPENHAYN (2006): “A Theory of Financing Constraints and Firm Dynamics,” *Quarterly Journal of Economics*, 121, 1131–1175.
- DEMARZO, P. M. AND M. J. FISHMAN (2007): “Optimal long-term financial contracting,” *Review of Financial Studies*, 20, 2079–2128.
- DIERMEIER, D., M. KEANE, AND A. MERLO (2005): “A political economy model of congressional careers,” *American Economic Review*, 95, 347–373.
- DUGGAN, J. AND C. MARTINELLI (2015): “Electoral accountability and responsive democracy,” .
- (2017): “The political economy of dynamic elections: Accountability, commitment, and responsiveness,” *Journal of Economic Literature*, 55, 916–84.
- EGEROD, B. C. (2021): “The lure of the private sector: career prospects affect selection out of Congress,” *Political Science Research and Methods*, 1–17.
- EGGERS, A. C. AND A. SPIRLING (2014): “Ministerial responsiveness in westminster systems: Institutional choices and House of Commons debate, 1832–1915,” *American Journal of Political Science*, 58, 873–887.
- EPSTEIN, D., D. BRADY, S. KAWATO, AND S. O’HALLORAN (1997): “A comparative approach to legislative organization: Careerism and seniority in the United States and Japan,” *American Journal of Political Science*, 965–998.

- FEARON, J. D. (1999): “Electoral accountability and the control of politicians: selecting good types versus sanctioning poor performance,” *Democracy, accountability, and representation*, 55, 61.
- FENNO, R. F. (1978): *Home style: House members in their districts*, Pearson College Division.
- FEREJOHN, J. (1986): “Incumbent performance and electoral control,” *Public choice*, 50, 5–25.
- FIVA, J. H. AND D. M. SMITH (2017): “Norwegian parliamentary elections, 1906–2013: representation and turnout across four electoral systems,” *West European Politics*, 40, 1373–1391.
- FOARTA, D. AND T. SUGAYA (2019): “Wait-and-See or Step in? Dynamics of Interventions,” .
- FONG, Y.-F. AND J. LI (2017): “Relational contracts, limited liability, and employment dynamics,” *Journal of Economic Theory*, 169, 270–293.
- FUDENBERG, D. AND J. TIROLE (1991): “Game theory, 1991,” *Cambridge, Massachusetts*, 393, 80.
- GROSECLOSE, T. AND K. KREHBIEL (1994): “Golden parachutes, rubber checks, and strategic retirements from the 102d House,” *American Journal of Political Science*, 75–99.
- GUO, Y. AND J. HÖRNER (2020): “Dynamic allocation without money,” .
- HALL, A. B. (2019): *Who Wants to Run?: How the Devaluing of Political Office Drives Polarization*, University of Chicago Press.
- HALL, A. B. AND K. A. SHEPSLE (2014): “The changing value of seniority in the us house: conditional party government revised,” *The journal of Politics*, 76, 98–113.
- HALL, R. L. AND R. P. VAN HOUWELING (1995): “Avarice and ambition in Congress: Representatives’ decisions to run or retire from the US House,” *American Political Science Review*, 89, 121–136.
- HIBBING, J. R. (1991): *Congressional careers: Contours of life in the US House of Representatives*, UNC Press Books.

- JACOBSON, G. C. (2015): “It’s nothing personal: The decline of the incumbency advantage in US House elections,” *The Journal of Politics*, 77, 861–873.
- KARTIK, N. AND R. VAN WEELDEN (2019): “Reputation Effects and Incumbency (Dis) Advantage,” *Quarterly Journal of Political Science*, 14, 131–157.
- KEANE, M. P. AND A. MERLO (2010): “Money, political ambition, and the career decisions of politicians,” *American Economic Journal: Microeconomics*, 2, 186–215.
- LI, J., N. MATOUSCHEK, AND M. POWELL (2017): “Power dynamics in organizations,” *American Economic Journal: Microeconomics*, 9, 217–241.
- LIPNOWSKI, E. AND J. RAMOS (2020): “Repeated delegation,” *Journal of Economic Theory*, 105040.
- MATTOZZI, A. AND A. MERLO (2008): “Political careers or career politicians?” *Journal of Public Economics*, 92, 597–608.
- MCKELVEY, R. D. AND R. RIEZMAN (1992): “Seniority in legislatures,” *American Political Science Review*, 951–965.
- MESSNER, M. AND M. K. POLBORN (2004): “Paying politicians,” *Journal of Public Economics*, 88, 2423–2445.
- MUTHOO, A. AND K. A. SHEPSLE (2014): “Seniority and incumbency in legislatures,” *Economics & Politics*, 26, 13–37.
- PADRÓ I MIQUEL, G. AND J. M. SNYDER (2006): “Legislative effectiveness and legislative careers,” *Legislative Studies Quarterly*, 31, 347–381.
- PADRÓ I MIQUEL, G. AND P. YARED (2012): “The political economy of indirect control,” *The Quarterly Journal of Economics*, 127, 947–1015.
- PEARCE, D. G. (1987): “Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation,” Tech. rep., Cowles Foundation for Research in Economics, Yale University.
- SCHWABE, R. (2009): “Reputation and accountability in repeated elections,” in *APSA 2009 Toronto Meeting Paper*.
- SPEAR, S. E. AND S. SRIVASTAVA (1987): “On repeated moral hazard with discounting,” *The Review of Economic Studies*, 54, 599–617.

THOMAS, J. AND T. WORRALL (1990): “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem,” *Journal of Economic Theory*, 51, 367 – 390.

WOLFERS, J. (2002): *Are voters rational?: Evidence from gubernatorial elections*, Graduate School of Business, Stanford University.

YARED, P. (2010): “A dynamic theory of war and peace,” *Journal of Economic Theory*, 145, 1921–1950.

A Appendix: Proofs

A.1 Proof of Proposition 1

This proof proceeds as follows. First, we describe a Bellman equation that characterizes the voter-optimal equilibrium. Second, we simplify the Bellman equation to show that the strategy profile given in the statement of the proposition is a voter-optimal equilibrium. Third, we use properties of the associated value functions that solve the Bellman equation to establish on-path uniqueness of the voter optimal equilibrium.

For any continuation value $u \in [1 - \delta, 1]$ for the incumbent politician, let $\pi(u)$ denote the voter's optimal continuation value among all feasible contracts that give the politician a value of u . For any average continuation value $v \in [0, 1]$ for the politician to have starting in the following period, let $\tilde{\pi}(v)$ denote the voter's optimal continuation value among all feasible contracts that give the agent an average value of v . These continuation values for the voter are defined by the following Bellman equation:

$$\begin{aligned} \pi(u) = \Phi\tilde{\pi}(u) &:= \sup_{a \in [0,1], v_s, v_f \in [0,1]} (1 - \delta)\mu a + \delta [\mu\tilde{\pi}(v_s) + (1 - \mu)\tilde{\pi}(v_f)] \\ &\text{subject to } u = (1 - \delta)(1 - \mu ca) + \delta[\mu v_s + (1 - \mu)v_f] & \text{(PK)} \\ &\text{and } (1 - \delta)1 + \delta v_f \leq (1 - \delta)(1 - ca) + \delta v_s; & \text{(IC)} \\ \tilde{\pi}(v) = \tilde{\Phi}\pi(v) &:= \sup_{\rho \in [0,1], u, u_0 \in [1-\delta,1]} \rho\pi(u) + (1 - \rho)\pi(u_0) \\ &\text{subject to } \rho u = v; & \text{(P}\tilde{\text{K)}} \end{aligned}$$

where $\Phi, \tilde{\Phi}$ are operators on the functions $\tilde{\pi}$ and π , respectively. The first of the two constraints defining π above is the voter's promise-keeping (PK) constraint, saying that the politician's current continuation value is, indeed, u if she follows the prescribed action, and the continuation values following success and failure are v_s and v_f , respectively. The second is the incentive-compatibility (IC) constraint, saying that the politician would rather follow her prescribed strategy than engage in the one-time deviation of shirking today. For the latter, the politician trades off the myopic benefit of shirking following a good shock against the expected gain in future value from playing her on-path action rather than failing. In the definition of $\tilde{\pi}$, ρ is the probability that the incumbent is retained, v is next period's average continuation value for today's incumbent, u is next period's continuation value for today's incumbent conditional on being retained, and u_0 is the continuation value for a newly elected politician. No incentive-compatibility constraint appears in this definition

because the firing choice is fully under the voter's control, and he can commit. But there is the promise-keeping (PK) constraint that says that the politician's average continuation value is a combination of the zero value she gets from being fired and the positive value she gets from being retained.

Standard arguments can be used to show that an optimal policy exists; that the value it generates to the voter, as a function of the politician's continuation value, is the unique solution to the above Bellman equation; and that π and $\tilde{\pi}$ are both continuous and concave.²⁴ With these observations, we prove the proposition.

We start by observing that there is some $v_* \in [1 - \delta, 1]$ such that $\tilde{\pi}$ is affine on $[0, v_*]$ and coincides with π on $[v_*, 1]$. Indeed, the Bellman equation tells us that $\tilde{\pi}(0) = \sup_{u_0 \in [1 - \delta, 1]} \pi(u_0)$, and so, for a given $v \in [1 - \delta, 1]$, we have $\tilde{\pi}(v) = \pi(v)$ if and only if $\pi(v) \geq \frac{v}{u} \pi(u) + (1 - \frac{v}{u}) \tilde{\pi}(0)$ for every $u \in (v, 1]$. But this condition rearranges to the requirement that

$$\pi(v) \geq \tilde{\pi}(0) + v \frac{\pi(u) - \pi(v)}{u - v}, \quad \forall u \in (v, 1].$$

Since π is concave, the given quotient is weakly increasing in u , and so the inequality is more stringent the closer u is to v . Therefore, $\tilde{\pi}(v) = \pi(v)$ if and only if either $v = 1$ or $\pi(v) \geq \tilde{\pi}(0) + v\pi'(v^+)$. Now, concavity of π implies that $\pi(v) - v\pi'(v^+)$ is weakly increasing and right-continuous in $v \in [1 - \delta, 1]$,²⁵ so that the set of v for which $\tilde{\pi}(v) = \pi(v)$ is of the form $[v_*, 1]$ for some $v_* \in [1 - \delta, 1]$. For any $v \in [0, v_*]$, then, the program defining $\tilde{\pi}(v) = \tilde{\Phi}\pi(v)$ has $u = v_*$ as an optimum, so that $\tilde{\pi}$ is affine on $[0, v_*]$.

Now, concavity of $\tilde{\pi}$ tells us that v_s and v_f are chosen to make (IC) hold with equality, except, perhaps, in the case that $v_s = v_f$. If neither of these conditions holds, we can modify the contract to one of this form without loss of optimality, bringing v_s and v_f closer together, holding a fixed, and maintaining (PK), for a weakly higher value for the voter. But even in the case that $v_s = v_f$, (IC) implies that $a = 0$. Therefore, we can restrict attention to the case that (IC) holds with equality. Then, combining the (IC) equation

²⁴These arguments proceed as follows. The voter's optimal payoff function $\pi : [1 - \delta, 1] \rightarrow \mathbb{R}$ must satisfy the Bellman equation $\pi = \Phi\tilde{\Phi}\pi$. As for the operators Φ and $\tilde{\Phi}$ between the spaces of real bounded functions on $[0, 1]$ and $[1 - \delta, 1]$ (the latter being viewed as metric spaces with respect to uniform convergence), it is straightforward to see that Φ is a contraction of modulus δ , that $\tilde{\Phi}$ is a weak contraction, and that both take concave continuous functions to concave continuous functions. The limit of concave continuous functions is itself concave and continuous, so the voter's optimal value functions, π and $\tilde{\pi}$, yield a unique solution to the Bellman equation, and both are concave and continuous.

²⁵For example, these properties can be seen from the formula

$$\pi(v) - v\pi'(v^+) = \pi(1 - \delta) - (1 - \delta)\pi'((1 - \delta)^+) + \int_{1 - \delta}^v [\pi'(\tilde{v}^+) - \pi'(v^+)] d\tilde{v}.$$

with (PK) immediately gives a solution for v_s and v_f . Specifically,²⁶

$$v_s = \bar{u}_a := \frac{1}{\delta}[u - (1 - \delta)(1 - ca)] \quad \text{and} \quad v_f = \underline{u} := \frac{1}{\delta}[u - (1 - \delta)].$$

Finally, since every $u \geq 1 - \delta$ and $a \in [0, 1]$, we have $0 \leq \underline{u} \leq \bar{u}_a$. It follows that $\underline{u}, \bar{u}_a \in [0, 1]$ if and only if $\bar{u}_a \leq 1$. Summarizing these observations yields

$$\pi(u) = \Phi \tilde{\pi}(u) = \max_{a \in [0, 1]} (1 - \delta)\mu a + \delta [\mu \tilde{\pi}(\bar{u}_a) + (1 - \mu)\tilde{\pi}(\underline{u})] \quad \text{subject to } \bar{u}_a \leq 1.$$

Since \bar{u}_a is an increasing and affine function of a , and $\tilde{\pi}$ is concave, the objective in the maximization problem above is concave in a . Moreover, its derivative with respect to a is simply $(1 - \delta)\mu[1 + c\tilde{\pi}'(\bar{u}_a)]$.²⁷ Now, let $v^* = 1$ if $\tilde{\pi}'(1) \geq -1/c$, and, otherwise, let v^* be the highest $v \in [0, 1)$ such that $\tilde{\pi}'|_{[0, v]} \geq -1/c$. By definition, notice that this means that $\tilde{\pi}$ cannot be affine in a neighborhood of v^* , which, in turn, implies that $v^* \geq v_*$.

The values of u such that $\bar{u}_1 = v^*$ and $\underline{u} = v^*$ are, respectively,

$$u_L := 1 - \delta(1 - v^*) - (1 - \delta)c \quad \text{and} \quad u_R := 1 - \delta(1 - v^*).$$

In addition, the value of a such that $\bar{u}_a = v^*$ is $a = [(1 - \delta) + \delta v^* - u] / [(1 - \delta)c]$. Therefore, the politician's optimal action takes the form

$$a_u = \begin{cases} 1 & \text{if } u < u_L, \\ \frac{(1 - \delta) + \delta v^* - u}{(1 - \delta)c} & \text{if } u \in [u_L, u_R], \\ 0 & \text{if } u > u_R \end{cases}.$$

Substituting the optimal choice of current action into the Bellman equation and differentiating gives us

$$\pi'(u) = \begin{cases} \mu \tilde{\pi}'(\bar{u}_1) + (1 - \mu)\tilde{\pi}'(\underline{u}) & \text{if } u < u_L, \\ \mu \left(-\frac{1}{c}\right) + (1 - \mu)\tilde{\pi}'(\underline{u}) & \text{if } u \in (u_L, u_R), \\ \tilde{\pi}'(\underline{u}) & \text{if } u > u_R. \end{cases}$$

²⁶Formally, we define the functions $\bar{\cdot}_a : \mathbb{R} \rightarrow \mathbb{R}$ (for $0 \leq a \leq 1$) and $\underline{\cdot} : \mathbb{R} \rightarrow \mathbb{R}$, so that \bar{u}_a and \underline{u} are functions of u .

²⁷In general, $\tilde{\pi}$ may fail to be differentiable. But since it is concave, it has one-sided derivatives defined everywhere, and these are sufficient to characterize the optimal a via a first-order condition. Whenever we refer to a derivative in this appendix without specifying a direction, we mean that the relevant claim applies to each one-sided derivative.

If $v^* = 1$, the formula for a_u will then follow directly. Assume now, for a contradiction, that $v^* < 1$. Then, $v^* < u_R$, and every $u \in [v^*, u_R)$ has

$$\tilde{\pi}'(u) = \pi'(u) \geq \mu \frac{-1}{c} + (1 - \mu)\tilde{\pi}'(\underline{u}) \geq \frac{-1}{c},$$

where the equality follows from $u \geq v^* \geq v_*$, the first inequality follows from $u < u_R$ and concavity of π , and the second inequality follows from $\underline{u} < v^*$ (since $u < u_R$). But then, $\tilde{\pi}'|_{[0, u_R)} \geq \frac{-1}{c}$, contradicting the definition of v^* . Thus, $v^* = 1$.

Finally, we verify that $v_* = 1 - \delta$, so that the voter's optimal retention rule will be as desired. To see this, it is enough to observe that both operators $\Phi, \tilde{\Phi}$ clearly take nonincreasing concave functions to nonincreasing functions (in fact, $\tilde{\Phi}$ takes every concave function to a nonincreasing function), and that a limit of nonincreasing functions is itself nonincreasing. The contraction property then tells us that π is nonincreasing. It follows that, in the optimization defining $\tilde{\pi}(v) = \tilde{\Phi}\pi(v)$, the voter optimally takes $u = \max\{1 - \delta, v\}$ —i.e., takes $u \in [1 - \delta, 1]$ as small as possible subject to (P \tilde{K}). It also follows that the voter sets $u_0 = 1 - \delta$, starting any newly elected politician at continuation value $1 - \delta$.

Having shown that the given strategy profile is a voter-optimal equilibrium, all that remains is to establish its on-path uniqueness.

Let us establish the result for two different cases. First, suppose that $\delta \leq c/(1 + c)$, so that some voter-optimal equilibrium entails $a_u = 1$ being played at every $u \in [1 - \delta, 1]$. By direct computation, π is affine, and $\tilde{\pi}$ is constant on $[0, 1 - \delta]$ and affine on $[1 - \delta, 1]$, since the Bellman operators $\Phi, \tilde{\Phi}$ preserve these two properties. But π is affine, nonnegative, and not globally 0, and it takes value 0 at 1, implying that it is strictly decreasing. Therefore, the voter retention rule that does as little firing as possible, subject to promise keeping, is uniquely optimal. All that remains is to show that it is uniquely optimal to have a politician choose $a_{1-\delta}$ when in office at continuation value $1 - \delta$. To that end, note that any strict incentive to work would give the politician a value strictly above $1 - \delta$, as would a strictly positive continuation value from shirking. Therefore, it must be that her incentive constraint is binding and she is fired for failure. Rearranging the promise-keeping and incentive constraints, and monotonically transforming the objective, the voter-optimal effort level for a newly elected politician, therefore, solves

$$\begin{aligned} & \max_{a, v_s \in [0, 1]} (1 - \delta)\mu a + \delta [\mu \tilde{\pi}(v_s)] \\ \text{subject to } & 1 - \delta = (1 - \delta)(1 - ca) + \delta v_s. \end{aligned}$$

But the objective is affine on its one-dimensional domain and strictly positive at the maximum feasible effort level $a_{1-\delta}$, and it takes value 0 at the minimum feasible effort level of 0. It is, therefore, uniquely maximized at $a = a_{1-\delta}$, as desired.

We now turn to the $c/(1+c) < \delta < c$ case. Note that uniqueness follows if π is strictly concave on its entire domain. To see this, begin with the observation that the optimal politician action a_u that we solved for is unique for every $u \in [1-\delta, 1]$ since it derives from maximizing a strictly concave objective over a convex domain. Optimality requires that the constraint (IC) hold with equality because $\tilde{\pi}$ is weakly concave on its domain, and $\tilde{\pi}$ is not affine between \underline{u} and \bar{u}_{a_u} for any $u \in [1-\delta, 1)$. Finally, π is strictly decreasing because it is nonincreasing and strictly concave, so firing as infrequently as (P \tilde{K}) allows is uniquely optimal for the voter.

We now argue that π is strictly concave when $c/(1+c) < \delta < c$. Following the proof of Lemma 2 in Guo and Hörner (2020), it follows that π is either affine or strictly concave when $\delta < c$. But observe that, for $c/(1+c) < \delta < c$ and $u \in (0, 1-\delta]$ sufficiently small,

$$\begin{aligned}\pi(u) &= (1-\delta)\mu + \delta \left[(1-\mu)\tilde{\pi}\left(\frac{u-(1-\delta)}{\delta}\right) + \mu\tilde{\pi}\left(\frac{u-(1-\delta)(1-c)}{\delta}\right) \right] \\ &= (1-\delta)\mu + \delta \left[(1-\mu)\pi(1-\delta) + \mu\pi\left(\frac{u-(1-\delta)(1-c)}{\delta}\right) \right]\end{aligned}$$

and, therefore,

$$\pi'(u) = (1-\mu)\pi' \left(\frac{u-(1-\delta)(1-c)}{\delta} \right) \neq \pi' \left(\frac{u-(1-\delta)}{\delta} \right).$$

So π cannot be affine and, thus, must be strictly concave. □

A.2 Proof of Proposition 2

Let $\mathcal{U} = \mathcal{U}(\bar{\pi})$ be the set of all incumbent politician payoffs that can arise in an equilibrium in which the voter's continuation payoff is at least $\bar{\pi}$ after every history, and let $\hat{u} = \hat{u}(\bar{\pi}) \in [1-\delta, 1]$ be the supremum of \mathcal{U} . For each $u \in \mathcal{U}$, let σ_u be some equilibrium in which the voter's continuation payoff is at least $\bar{\pi}$ after every history.

The alternative equilibrium we now consider is the voter's optimal equilibrium among all those that always yield an incumbent politician value in $[1-\delta, \hat{u}]$. Essentially the exact same proof as that of Proposition 1 shows that the form described in the statement of this proposition is an optimal such contract. In particular, notice that the modified limit on $a_{u, \hat{u}}$ is chosen to ensure that $\bar{u}_{a_{u, \hat{u}}} \leq \hat{u}$. Let $\pi_{\hat{u}} : [1-\delta, \hat{u}] \rightarrow \mathbb{R}$ denote the induced value function, so that $\pi_{\hat{u}}(u)$ is the voter's optimal value over all equilibria that (i) give the incumbent a

continuation value of u ; and (ii) never give any incumbent politician a continuation value greater than \hat{u} . Again, as in the proof of Proposition 1, the function $\pi_{\hat{u}}$ is continuous and nonincreasing. That equilibrium σ_u respected the threshold π implies, given optimality of the modified equilibrium and the definition of \mathcal{U} , that $\pi_{\hat{u}}(u) \geq \pi$ for every $u \in \mathcal{U}$. But then, that $\pi_{\hat{u}}$ is continuous implies that $\pi_{\hat{u}}(\hat{u}) \geq \pi$, and that it is nonincreasing implies that $\pi_{\hat{u}} \geq \pi$, as required.

Finally, to see that $\hat{u}(\cdot)$ is strictly decreasing, assume for a contradiction that $\hat{u}' := \hat{u}(\pi - \epsilon) \leq \hat{u}(\pi)$ for some $\epsilon \in (0, \pi)$. But then, one can construct an equilibrium that gives the politician a continuation value slightly greater than \hat{u}' from the first period; gives the politician a continuation value weakly below \hat{u}' from every other history; and gives the voter a continuation value above $\pi - \epsilon$ from every history. Existence of such an equilibrium contradicts the definition of $\hat{u}' = \hat{u}(\pi - \epsilon)$. \square

A.3 Proof of Corollary 2

To show that such a π and such an equilibrium exists whenever $\frac{1}{2-\mu} < \delta < c$ (a condition that is consistent with our standing assumptions), consider equilibria of the form described in Proposition 2, as parametrized by $\hat{u} \in [1 - \delta, 1)$, when \hat{u} is low. In what follows, we make use of notation from that proposition and its proof where convenient.

Observe that $a_{u, \hat{u}} = \min\{1, \frac{(1-\delta)(1-u) + \delta(\hat{u}-u)}{(1-\delta)c}\}$ is weakly decreasing in $u \in [1 - \delta, 1]$, and we have the following two limits:

$$\begin{aligned} \lim_{\hat{u} \rightarrow 1-\delta} a_{1-\delta, \hat{u}} &= \min\left\{1, \frac{\delta}{c} + \frac{\delta}{(1-\delta)c}[\hat{u} - (1-\delta)]\right\} = \frac{\delta}{c} \\ \lim_{\hat{u} \rightarrow 1-\delta} v_{\hat{u}}(a_{\hat{u}, \hat{u}}, 0) &= \frac{\hat{u} - (1-\delta)}{\delta} = 0. \end{aligned}$$

Therefore, fixing small enough $\epsilon \in (0, \delta)$ guarantees that both $a_{1-\delta, \hat{u}}$ and $v_{\hat{u}}(a_{\hat{u}, \hat{u}}, 0)$ are strictly below 1 for all $\hat{u} \in [1 - \delta, 1 - \delta + \epsilon]$ (since $\delta < c$). In what follows, we focus on \hat{u} from this interval.

For such \hat{u} , the equilibrium described in Proposition 2 can be equivalently described as a simple “two-state automaton” strategy profile, in which a politician in office starts the period in either a “good” (for the voter) or “bad” (for the voter) state—corresponding to continuation values $1 - \delta$ and \hat{u} , respectively, for the politician in office. The equilibrium-prescribed effort is $a_G(\hat{u}) := a_{1-\delta, \hat{u}}$ in the good state and $a_B(\hat{u}) := a_{\hat{u}, \hat{u}}$ in the bad state. In either state, a success results in the politician in office being re-elected and next period beginning in the bad state. In the good state, a politician who fails is replaced with certainty;

and in the bad state, a politician who fails is re-elected into the good state with probability $v_{\hat{u}}(a_{\hat{u},\hat{u}}, 0)/(1 - \delta)$ and replaced by a new politician with complementary probability. Each newly elected politician begins her career in the good state.

In particular, in the described equilibrium, starting from either state, the voter re-elects the politician in office to the bad state following a good shock, and he faces some politician (either an incumbent or a newly elected one) following a bad shock. We can, therefore, write the voter's continuation payoffs $\pi_G(\hat{u})$ and $\pi_B(\hat{u})$ from the good and bad states, respectively:

$$\pi_i(\hat{u}) = (1 - \delta)\mu a_i(\hat{u}) + \delta [\mu\pi_B(\hat{u}) + (1 - \mu)\pi_G(\hat{u})] \text{ for } i \in \{G, B\}.$$

As $\pi_G(\hat{u}) - \pi_B(\hat{u})$ is proportional to $a_G(\hat{u}) - a_B(\hat{u}) \geq 0$, the voter's worst-case continuation value from this equilibrium is $\pi_B(\hat{u})$. But direct computation shows that

$$a_G(\hat{u}) - a_B(\hat{u}) = \frac{(1 - \delta) + \delta\hat{u} - (1 - \delta)}{(1 - \delta)c} - \frac{(1 - \delta) + \delta\hat{u} - \hat{u}}{(1 - \delta)c} = \frac{\hat{u} - (1 - \delta)}{(1 - \delta)c},$$

and

$$\pi_B(\hat{u}) = (1 - \delta)\mu a_B(\hat{u}) + \delta\pi_B(\hat{u}) + \delta(1 - \mu) [\pi_G(\hat{u}) - \pi_B(\hat{u})].$$

From this, it follows that

$$\begin{aligned} \pi_B(\hat{u}) &= \mu a_B(\hat{u}) + \delta(1 - \mu) \frac{\pi_G(\hat{u}) - \pi_B(\hat{u})}{1 - \delta} \\ &= \mu \{a_B(\hat{u}) + \delta(1 - \mu) [a_G(\hat{u}) - a_B(\hat{u})]\} \\ &= \frac{\mu}{(1 - \delta)c} \{[(1 - \delta)(1 - \hat{u}) + \delta(\hat{u} - \hat{u})] + \delta(1 - \mu) [\hat{u} - (1 - \delta)]\} \\ &= \frac{\mu}{(1 - \delta)c} \{(1 - \delta)(1 - \hat{u}) + \delta(1 - \mu) [\hat{u} - (1 - \delta)]\}. \end{aligned}$$

And, therefore,

$$\pi'_B(\hat{u}) = \frac{\mu}{(1 - \delta)c} [-(1 - \delta) + \delta(1 - \mu)] = \frac{\mu}{(1 - \delta)c} [\delta(2 - \mu) - 1] > 0.$$

Thus, the described equilibrium with $\hat{u} = 1 - \delta$ (which is exactly the best stationary equilibrium in Proposition 0) has a strictly lower worst-case voter value than the described equilibrium with $\hat{u} = 1 - \delta + \epsilon$. The corollary follows. \square

A.4 Proof of Proposition 3

The proof of Proposition 3 follows along similar lines to that of Proposition 1. We begin by defining optimal value functions for the voter as a function of the politician's continuation value. Because the stage game has more components, it is now convenient to define three value functions rather than two.

First, for any continuation value $u \in [1 - \delta, 1]$ for a politician who has either just turned down or not received an outside offer, let $\pi(u)$ denote the voter's optimal continuation value among all feasible contracts that give the politician a value of u . For any continuation value $\tilde{u} \in [pw + (1 - p)(1 - \delta), 1]$ for the politician to have after seeing the realized public randomization outcome, but before learning whether or not an outside offer realized, let $\hat{\pi}(\tilde{u})$ denote the voter's optimal continuation value among all feasible contracts that give the politician a continuation value of \tilde{u} . Note that the optimal value of the voter when a politician has continuation value \tilde{u} *before* the public randomization device realizes is, therefore, $\text{cav } \hat{\pi}(\tilde{u})$, where $\text{cav } \hat{\pi}$ is the concave envelope of $\hat{\pi}$.²⁸ Finally, after the current period's outcome y has realized, for any average continuation value $v \in [0, 1]$ for a politician to have starting in the following period, let $\tilde{\pi}(v)$ denote the voter's optimal continuation value among all feasible contracts that give the agent an average value of v .

Now, define $\tilde{u}_L := (1 - p)(1 - \delta) + pw$, and consider the following system of Bellman equations:

$$\begin{aligned} \pi(u) = \Phi\tilde{\pi}(u) &:= \sup_{a \in [0,1], v_s, v_f \in [0,1]} (1 - \delta)\mu a + \delta [\mu\tilde{\pi}(v_s) + (1 - \mu)\tilde{\pi}(v_f)] \\ &\text{subject to } u = (1 - \delta)(1 - \mu ca) + \delta[\mu v_s + (1 - \mu)v_f] & \text{(PK)} \\ &\text{and } (1 - \delta)1 + \delta v_f \leq (1 - \delta)(1 - ca) + \delta v_s; & \text{(IC)} \\ \hat{\pi}(\tilde{u}) = \hat{\Phi}(\tilde{\pi}, \hat{\pi})(v) &:= \sup_{\lambda \in [0,1], u \in [1 - \delta, 1], \tilde{u}_0 \in [\tilde{u}_L, 1]} (1 - p\lambda)\pi(u) + p\lambda \text{cav } \hat{\pi}(\tilde{u}_0) \\ &\text{subject to } \tilde{u} = (1 - p\lambda)u + p\lambda w, & \text{(P}\hat{\text{K)}} \\ &\lambda = 1 \text{ if } u < w \text{ and } \lambda = 0 \text{ if } u > w; & \text{(I}\hat{\text{C)}} \\ \tilde{\pi}(v) = \tilde{\Phi}\hat{\pi}(v) &:= \sup_{\rho \in [0,1], \tilde{u}, \tilde{u}_0 \in [\tilde{u}_L, 1]} \rho \text{cav } \hat{\pi}(\tilde{u}) + (1 - \rho)\text{cav } \hat{\pi}(\tilde{u}_0) \\ &\text{subject to } v = \rho\tilde{u}. & \text{(P}\tilde{\text{K)}} \end{aligned}$$

²⁸For an upper semicontinuous function $f : [a, b] \rightarrow \mathbb{R}$ defined on a compact interval, $\text{cav } f$ is the pointwise smallest concave function $[a, b] \rightarrow \mathbb{R}$ that is pointwise above f .

Let us explain this system. The operator Φ , describing the optimal incentive provision for the politician's effort, is exactly as in the proof of Proposition 1. The operator $\tilde{\Phi}$, describing the optimal retention rule for a given politician continuation value is essentially identical to that from Proposition 1, with two small differences. First, the possibility of an outside offer raises the minimum possible continuation value of a retained politician—it is now \tilde{u}_L rather than only $1 - \delta$. Second, because play can condition on a sunspot after a given politician is elected or re-elected, continuation values are delivered via $\text{cav } \hat{\pi}$ rather than via $\hat{\pi}$. Now, we turn to the operator $\hat{\Phi}$. The associated promise-keeping condition ($\hat{\text{PK}}$) expresses the politician's continuation value \tilde{u} as a weighted average of w and u , where w is her continuation value if she accepts an outside offer of w , and u is her continuation value in the complementary case. Letting λ denote her (chosen) probability of leaving for the private sector conditional on receiving such an offer, her total probability of receiving and accepting an offer is $p\lambda$. As a politician who contemplates leaving office is comparing value w to u , it follows that she will willingly accept [resp. reject] an outside offer if and only if $w \geq [\leq]u$, which justifies the incentive-compatibility condition ($\hat{\text{IC}}$).

We proceed in a similar fashion to the proof of Proposition 1. First, we note that standard recursive arguments show that an optimal policy exists and that the value it generates to the voter, as a function of the politician's continuation value, is the unique bounded solution to the above system of Bellman equations. Furthermore, as all three operators above take weakly decreasing functions to weakly decreasing functions and take upper semicontinuous to upper semicontinuous functions, the four functions, π , $\hat{\pi}$, $\text{cav } \hat{\pi}$, and $\tilde{\pi}$, are all weakly decreasing and upper semicontinuous. That they are weakly decreasing implies that setting $\tilde{u}_0 = \tilde{u}_L$ is optimal in the optimizations determining $\hat{\pi}(\tilde{u})$ and $\tilde{\pi}(v)$. Also, since the concave envelope of any upper semicontinuous function on a compact interval is concave and continuous, it follows that $\tilde{\Phi}$ takes any upper semicontinuous function to a concave and continuous one, so that $\tilde{\pi}$ and π are concave and continuous.

Next, because $\text{cav } \hat{\pi}$ is concave and weakly decreasing, just as in the proof of Proposition 1, it is optimal in the program determining $\tilde{\pi}(v)$ to set the re-election probability ρ as high as the promise-keeping constraint ($\tilde{\text{PK}}$) will allow. Then, without loss of optimality, the incentive-compatibility condition (IC) in the equation defining π holds with equality, as moving (v_f, v_s) closer together while satisfying ($\tilde{\text{PK}}$) will weakly improve the voter's objective by inducing a mean-preserving contraction over continuation values evaluated via the concave function $\tilde{\pi}$. Combining binding incentive compatibility with promise keeping yields exact formulas for the success and failure continuation values as a function of the action

taken— $v_f = \underline{u}$ and $v_s = \bar{u}_a$, where both quantities are given in the proof of Proposition 1—with the only constraint on the action being that $\bar{u}_a \leq 1$.

We make four additional simplifying observations. First, $\text{cav } \hat{\pi}$ necessarily agrees with $\hat{\pi}$ at \tilde{u}_L , which is an extreme point of the two functions' domain. Second, because $\text{cav } \hat{\pi} \geq \hat{\pi}$ and setting $\tilde{u}_0 = u$ ($= w$) is feasible in the program determining $\hat{\pi}(w)$, it follows that $\sup_{\tilde{u}_0 \in [\tilde{u}_L, 1]} \text{cav } \hat{\pi}(\tilde{u}_0) \geq \pi(w)$, so setting $\lambda = 1$ is optimal when $\tilde{u} = w$. Intuitively, because the voter can, at worst, continue the same play with a newly elected politician, it always weakly benefits him when the politician takes her outside offer. Third, because the promise-keeping constraint (P $\hat{\text{K}}$) requires $\tilde{u} - w$ to have the same sign as $u - w$, the incentive-compatibility constraint (I $\hat{\text{C}}$) implies that offers are accepted with certainty ($\lambda = 1$) whenever $\tilde{u} < w$ and rejected with certainty ($\lambda = 0$) whenever $\tilde{u} > w$. Fourth, applying these three observations to $\hat{\pi}(\tilde{u}_L)$ yields the equation $\hat{\pi}(\tilde{u}_L) = (1 - p)\pi(\frac{\tilde{u}_L - pw}{1 - p}) + p\hat{\pi}(\tilde{u}_L)$, which can be rearranged to $\hat{\pi}(\tilde{u}_L) = \pi(1 - \delta)$.

Collecting the above observations, together with the observation that (since $\text{cav } \hat{\pi}$ is nonincreasing) the voter optimally sets the continuation value \tilde{u} to be as small as possible subject to P $\hat{\text{K}}$, yields a simplified system of Bellman equations:

$$\begin{aligned} \pi(u) &= \sup_{a \in [0, 1]} (1 - \delta)\mu a + \delta [\mu \tilde{\pi}(\bar{u}_a) + (1 - \mu)\tilde{\pi}(u)] \\ &\text{subject to } \bar{u}_a \leq 1; \\ \hat{\pi}(\tilde{u}) &= \begin{cases} (1 - p)\pi(\frac{\tilde{u} - pw}{1 - p}) + p\pi(1 - \delta) & \text{if } \tilde{u} \leq w \\ \pi(\tilde{u}) & \text{if } \tilde{u} > w; \end{cases} \\ \tilde{\pi}(v) &= \text{cav } \hat{\pi}(\max\{v, \tilde{u}_L\}). \end{aligned}$$

We now detail the circumstances under which public randomization is used. To that end, note that π being concave and continuous implies that the restrictions $\hat{\pi}|_{[\tilde{u}_L, w]}$ and $\hat{\pi}|_{(w, 1]}$ are both concave and continuous, as well.

Let us now establish that some $u^* \in [w, 1]$ exists for which

$$\text{cav } \hat{\pi}(\tilde{u}) = \begin{cases} \hat{\pi}(\tilde{u}) & \text{if } \tilde{u} \notin (w, u^*) \\ \xi \hat{\pi}(w) + (1 - \xi)\hat{\pi}(u^*) & \text{if } \tilde{u} = \xi w + (1 - \xi)u^* \text{ for some } \xi \in [0, 1]. \end{cases}$$

It is immediate that the given functional form is a lower bound for $\text{cav } \hat{\pi}$, so our goal is to find some $u^* \in [w, 1]$ such that it forms an upper bound. To show this, one can find a slope $m \in \mathbb{R}$ such that $\pi(x) \leq \hat{\pi}(w) + m(\tilde{u} - w)$ for every $\tilde{u} \in [w, 1]$, with equality at some

$u^* \in [w, 1]$, which is sure to exist because π is continuous. Using this u^* , let us observe that the given function lies weakly above $\text{cav } \hat{\pi}$ and, hence, coincides with it. Indeed, because $\hat{\pi}(w) \geq \hat{\pi}(w^+)$ and $\hat{\pi}'(w^-) \geq \lim_{\bar{u} \rightarrow w^+} \hat{\pi}'(x)$ if $w < 1$, it is straightforward to show that this function is concave and above $\hat{\pi}$ —hence, above $\text{cav } \hat{\pi}$. The characterization of $\text{cav } \pi$ follows.

Finally, we turn to the optimal level of effort—that is, the optimal choice of a in the optimization problem determining $\pi(u)$. Because the associated objective function,

$$a \mapsto (1 - \delta)\mu a + \delta [\mu \tilde{\pi}(\bar{u}_a) + (1 - \mu)\tilde{\pi}(\underline{u})],$$

is concave and continuous, one can solve for the optimal action (which exists by continuity) via a first-order approach. Specifically, letting u_L, v^*, u_R be exactly as defined in the proof of Proposition 1 and reasoning exactly as in that proof, an optimal choice of politician action at continuation value u is

$$a_u = \begin{cases} 1 & \text{if } u < u_L, \\ \frac{(1-\delta)+\delta v^*-u}{(1-\delta)c} & \text{if } u \in [u_L, u_R], \\ 0 & \text{if } u > u_R. \end{cases}$$

Let us now observe that $v^* \in \{w, 1\}$. By definition, v^* cannot have an open interval around it on which $\tilde{\pi}$ is affine, implying that $v^* \notin (w, u^*)$. Next, from the hypothesis that $v^* \in [u^*, 1)$, we can derive a contradiction by the argument (verbatim) in the proof of Proposition 1 that $v^* = 1$ because the three functions, π , $\tilde{\pi}$, and $\hat{\pi}$, all agree on $[u^*, 1]$. Finally, we can adapt the same argument as follows to show that $v^* \geq w$. Indeed, assume, for a contradiction, that $v^* < w$. First, because $\tilde{\pi}$ is affine on $[0, \bar{u}_L]$, the definition of v^* guarantees that $v^* \geq \underline{x}$. Therefore, any $u \in [v^*, w)$ close enough to v^* has $\underline{u} < v^*$, so that

$$\tilde{\pi}'(u) = \hat{\pi}'(u) = \pi' \left(\frac{u-pw}{1-p} \right) \geq \pi'(u) \geq \mu(-1/c) + (1 - \mu)\tilde{\pi}'(\underline{u}) \geq -1/c,$$

where the second equality follows from the chain rule, and the inequalities follow from the concavity of π and $\tilde{\pi}$. However, that $\tilde{\pi}' \geq -1/c$ in a neighborhood to the right of v^* contradicts the definition of v^* .

In summary, there exist $v^* \in \{w, 1\}$ and $u^* \in [w, 1]$ such that

$$\begin{aligned}\pi(u) &= (1 - \delta)\mu a_{u,v^*} + \delta \left[\mu \tilde{\pi}(\bar{u}_{a_{u,v^*}}) + (1 - \mu)\tilde{\pi}(\underline{u}) \right], \\ \hat{\pi}(\tilde{u}) &= \begin{cases} (1 - p)\pi\left(\frac{\tilde{u} - pw}{1 - p}\right) + p\pi(1 - \delta) & \text{if } \tilde{u} \leq w \\ \pi(\tilde{u}) & \text{if } \tilde{u} > w; \end{cases} \\ \tilde{\pi}(v) &= \begin{cases} \hat{\pi}(\tilde{u}_L) & \text{if } v < \tilde{u}_L \\ \frac{u^* - v}{u^* - w}\hat{\pi}(w) + \frac{v - w}{u^* - w}\hat{\pi}(u^*) & \text{if } w < v < u^* \\ \hat{\pi}(v) & \text{otherwise,} \end{cases}\end{aligned}$$

where

$$a_{u,v^*} = \max \left\{ 0, \min \left\{ 1, \frac{(1 - \delta) + \delta v^* - u}{(1 - \delta)c} \right\} \right\}.$$

The result then follows from observing that the equilibrium described in the proposition generates these optimal values. \square

A.5 Proof of Corollary 3

Throughout this proof, we take for granted the structure of the optimal equilibrium given by Proposition 3, and we follow the notation of its proof.

To prove the result, it suffices to show that it is optimal to the voter to have $v^* = w$ rather than having $v^* = 1$, when w is close enough to 1. In this case, no politician's continuation value will ever exceed w in the voter-optimal equilibrium described by Proposition 3. To show this, it suffices to establish that, when $w < 1$ is close enough to 1, $\pi(\cdot)$ evaluated at w is higher when using the effort level implied by $v^* = w$ than when using the effort level implied by $v^* = 1$. That is, we wish to show that

$$(1 - \delta)\mu a_{w,w} + \delta \left[\mu \tilde{\pi}(\bar{w}_{a_{w,w}}) + (1 - \mu)\tilde{\pi}(\underline{w}) \right] > (1 - \delta)\mu a_w + \delta \left[\mu \tilde{\pi}(\bar{w}_{a_w}) + (1 - \mu)\tilde{\pi}(\underline{w}) \right],$$

where a_w denotes $a_{w,1}$ as defined in Proposition 3; \bar{w}_{a_w} denotes $\bar{u}|_{u=w_{a_w}}$; and \underline{w} denotes $\underline{u}|_{u=w}$, also as defined in the proof of Proposition 3.

To show the desired inequality, focus on the case in which $w \in [1 - \delta, 1)$ has $w \geq 1 - c(1 - \delta)$, so that $a_w = 1$. Using the expressions for the value functions at the end of the proof of Proposition 3, we find that the left-hand side and right-hand side of the centered

inequality above differ by

$$\begin{aligned}
\text{LHS} - \text{RHS} &= \left\{ (1 - \delta)\mu a_{w,w} + \delta \left[\mu \tilde{\pi}(\bar{w}_{a_w,w}) + (1 - \mu)\tilde{\pi}(w) \right] \right\} \\
&\quad - \left\{ (1 - \delta)\mu a_w + \delta \left[\mu \tilde{\pi}(\bar{w}_{a_w}) + (1 - \mu)\tilde{\pi}(w) \right] \right\} \\
&= \mu\delta \left[\tilde{\pi}(\bar{w}_{a_w,w}) - \tilde{\pi}(\bar{w}_{a_w}) \right] - \mu(1 - \delta)(a_w - a_{w,w}) \\
&= \mu\delta \left[\tilde{\pi}(w) - \tilde{\pi}(1) \right] - \mu(1 - \delta) \left[\frac{(1-\delta)+\delta 1-w}{(1-\delta)c} - \frac{(1-\delta)+\delta w-w}{(1-\delta)c} \right] \\
&= \mu\delta \left[\hat{\pi}(w) - 0 \right] - \mu \left[\frac{\delta 1}{c} - \frac{\delta w}{c} \right] \\
&= \mu\delta \left[(1 - p)\pi(w) + p\pi(1 - \delta) \right] - \frac{\mu\delta}{c}(1 - w) \\
&\geq \mu\delta p(1 - \delta)\mu a_{1-\delta} - \frac{\mu\delta}{c}(1 - w),
\end{aligned}$$

which converges to $\mu^2\delta p(1 - \delta)a_{1-\delta} > 0$ as $w \rightarrow 1$. Therefore, the difference is strictly positive when $w < 1$ is close enough to 1, as required. \square

A.6 Proof of Lemma 1

Consider optimal play for the high-skilled incumbent at the start of her term. As in the baseline model, we can focus on pure strategy equilibria. If the politician were to choose zero effort, then the voter's optimal value would be zero, which is not the case. Therefore, any newly elected politician chooses strictly positive effort if she is a high-skilled type. Let v^s and v^f be the politician's continuation value from the next period following success and failure, respectively. An improvement for the voter, which would relax the politician's incentive constraint, would instead be to remove the politician and put in a new one at starting value v^f following failure. Doing so relaxes the high-skilled type politician's incentive to work and improves the voter's utility by increasing the odds that the politician in office tomorrow will be a high type. The result follows by continuing the argument recursively on the histories of the game. \square

A.7 Proof of Proposition 4

For any continuation value $u \in [1 - \delta, 1]$ for an incumbent high-type politician, let $\pi_1(u)$ denote the voter's optimal continuation value (computed under the belief that the current politician is certain to be a high type) among all continuation equilibria that give the politician a value of u . For any average continuation value $v \in [0, 1]$ for the politician to have, starting in the following period, let $\tilde{\pi}_1(v)$ denote the voter's optimal continuation value among all feasible contracts that give the agent an average value of v . Finally, let π_q

denote the voter's optimal continuation value when a new politician (who has never before acted and is, therefore, a high type with probability q) is in office.

Given Lemma 1, the voter will accrue zero flow payoffs and draw a new politician each period until the first time a high type is drawn. Therefore, these continuation values for the voter are defined by the following Bellman equation:²⁹

$$\begin{aligned} \pi_q = \Psi(\tilde{\pi}_1, \pi_q) &:= \sup_{a_*, v^s \in [0,1]} (1 - \delta)q\mu a_* + \delta [(1 - q\mu)\pi_q + q\mu\tilde{\pi}_1(v^s)] \\ &\text{subject to } 1 - \delta \leq (1 - \delta)(1 - ca_*) + \delta v^s; \quad (\text{IC}_0) \\ &\pi_1(u) = \Phi\tilde{\pi}_1(u); \\ \tilde{\pi}_1(v) = \tilde{\Phi}_q(\pi_1, \pi_q)(v) &:= \sup_{\rho \in [0,1], u \in [1-\delta,1]} \rho\pi_1(u) + (1 - \rho)\pi_q \\ &\text{subject to } \rho u = v, \quad (\text{PK}) \end{aligned}$$

where Φ is defined as in the proof of Proposition 1, and $\tilde{\Phi}_q$ and Ψ are operators on (π_1, π_q) and $(\tilde{\pi}_1, \pi_q)$, respectively.

The entire proof of Proposition 1, excluding its final paragraph, applies essentially verbatim to π_1 and $\tilde{\pi}_1$. In particular: both functions are concave and continuous; some $v_* \in [1 - \delta, 1]$ exists such that $\tilde{\pi}_1|_{[0, v_*]}$ is affine and $\tilde{\pi}_1|_{[v_*, 1]} = \pi_1|_{[v_*, 1]}$; and the optimal politician effort and continuation values prescribed by $\tilde{\Phi}$ for any given value of the incumbent high-type politician are exactly those prescribed by Proposition 1. In particular, this establishes part 3 of the proposition.

In what follows, we establish parts 1 and 2, assuming, without loss of optimality, that $v_* \in [1 - \delta, 1]$ is as large as possible subject to $\tilde{\pi}_1|_{[0, v_*]}$ being affine.

For part 1, let us study the optimization problem defining Ψ , rewriting it as

$$\begin{aligned} \Psi(\tilde{\pi}_1, \pi_q) &= \delta(1 - q\mu)\pi_q + q \sup_{a_*, v^s \in [0,1]} \{(1 - \delta)\mu a_* + \delta\mu\tilde{\pi}(v^s)\} \\ &\text{subject to } v^s \geq ca_*(1 - \delta)/\delta, \end{aligned}$$

²⁹It is standard that these optimal values would be characterized by the given Bellman equation if the voter could commit. As we have observed, any observable deviation by the voter can be made unprofitable because there is an equilibrium—for instance, one in which all future politicians shirk and the voter replaces politicians with a rate independent of performance—that provides zero payoff to the voter. The only potentially profitable voter deviations, then, take place when the voter is expected to mix. Hence, with a sunspot available (which can be conditioned upon rather than having the voter privately mix), the given Bellman equation characterizes voter-optimal equilibrium payoffs.

As a continuous optimization problem with compact domain, this problem has a maximizer (a_*, v^s) . That $ca_*(1 - \delta)/\delta \leq v^s \leq 1$ implies that $a_* \leq a_{1-\delta}$ (where $u \mapsto a_u$ is as defined in part 3 of the proposition statement) for any such maximizer. To show that $a_* \geq a_{1-\delta}$ in some solution to this program, we separately consider two cases. On the one hand, if $v^s > ca_*(1 - \delta)/\delta$ for such a maximizer, then it must be that $a_* = 1$, for, otherwise, a_* could be slightly raised (maintaining feasibility) for a strictly higher objective. On the other hand, if $v^s = ca_*(1 - \delta)/\delta$, then (a_*, v^s, v^f) is optimal among all feasible (a, v^s, v^f) in the program defining $\Phi\tilde{\pi}_1(1 - \delta)$ in which the IC constraint binds (which, we showed in the proof of Proposition 1, is without loss of optimality). Therefore, following the proof of Proposition 1, $(a_{1-\delta}, ca_{1-\delta}(1 - \delta)/\delta)$ is optimal in the same program, which establishes part 1 of the proposition.

Now for part 2, notice that $\pi_q \leq \max_{v \in [0,1]} \pi_1(v)$ because the voter is (in a best equilibrium) better off starting with a high-type politician than with a politician of uncertain type.³⁰ We also know that $\pi_q > 0$ since, for instance, any positive-value equilibrium play from the environment without adverse selection can yield a positive (though lower) profit in the current environment with adverse selection. Therefore, it must be that $v_* < 1$. Then, for $u \in (v_*, 1)$ high enough,

$$\tilde{\pi}_1(u) = \pi_1(u) = \mu\left(\frac{1-u}{c}\right) + \delta \left[(1 - \mu)\tilde{\pi}_1\left(\frac{u-(1-\delta)}{\delta}\right) + \mu\tilde{\pi}_1(1) \right].$$

Therefore,

$$\tilde{\pi}'_1(u) = \mu\left(-\frac{1}{c}\right) + (1 - \mu)\tilde{\pi}'_1\left(\frac{u-(1-\delta)}{\delta}\right) \geq \mu\left(-\frac{1}{c}\right) + (1 - \mu)\tilde{\pi}'_1(u),$$

where the inequality follows since $\tilde{\pi}_1$ is concave. Therefore, $\tilde{\pi}'_1(u) \geq -1/c$.

Thus, we can extend $\tilde{\pi}_1$ to a concave function on $[1 - \delta, \infty)$ by letting it take value $\tilde{\pi}_1(u) = \tilde{\pi}_1(1) + (1 - u)/c$ for $u > 1$. Assume now, toward a contradiction, that the interval $(1 - \delta, v_*)$ is nonempty. Then, for all $u \in (1 - \delta, v_*)$, we have

$$\begin{aligned} \tilde{\pi}'_1(u) &\leq \pi'_1(u) = (1 - \mu)\tilde{\pi}'_1\left(\frac{u-(1-\delta)}{\delta}\right) + \mu\tilde{\pi}'_1\left(\frac{u-(1-\delta)(1-c)}{\delta}\right) \\ &= (1 - \mu)\tilde{\pi}'_1(u) + \mu\tilde{\pi}'_1\left(\frac{u-(1-\delta)(1-c)}{\delta}\right), \end{aligned}$$

³⁰Given any feasible contract for a voter facing a politician of unknown type, a voter who knows the newly elected politician to be the high type can randomize play to simulate the firing rule of a voter who is uncertain of the initial politician's type and get a weakly higher payoff along every path of play.

where the first line inequality follows from the fact that $\tilde{\pi}_1|_{[0, v_*]}$ is affine; π_1 is concave; and $\tilde{\pi}_1(v_*) = \pi_1(v_*)$; and the second line follows since, again, $\tilde{\pi}_1|_{[0, v_*]}$ is affine. Therefore,

$$\tilde{\pi}'_1(u) \leq \tilde{\pi}'_1\left(\frac{u - (1 - \delta)(1 - c)}{\delta}\right),$$

which implies that $u \geq [u - (1 - \delta)(1 - c)]/\delta$, and, thus, $u \leq 1 - c < 1 - \delta$, a contradiction. Therefore, it follows that $v_* = 1 - \delta$, establishing part 2.

A.8 Proof of Corollary 4

Our first step is to consider a modified contracting problem that differs from the original one in two ways. First, in the modified problem, the voter knows that the current politician is a high type and that all other politicians are low types; hence, he gets zero continuation value from replacement. Second, the voter gets the payoff associated with a politician's intended effort in a given period, even if the state turns out to be the low-productivity state.³¹ Moreover, we consider this modified contracting problem for all values of μ , including the value of $\mu = 0$ that our model precludes.

We recursively define the associated value functions, denoted as π_μ° and $\tilde{\pi}_\mu^\circ$ for this modified problem, parametrized by the voter's probability of having to reward, rather than punish, the politician in office. While we find the interpretations of π_μ° and $\tilde{\pi}_\mu^\circ$ above to be instructive (and arguments essentially identical to those supporting Proposition 4 would show that the optimal value functions from this modified contracting problem solve the recursive equations that define π_μ° and $\tilde{\pi}_\mu^\circ$, as well), these interpretations are unnecessary for establishing Corollary 4. The rest of the proof can take π_μ° and $\tilde{\pi}_\mu^\circ$ as purely mathematical objects defined by the given recursive equations below. Given such functions, we show that $\tilde{\pi}_0^\circ$ is maximized to the right of the continuation value $c(1 - \delta)/\delta$. When q and μ are small, then, we show that $\tilde{\pi}_1$ is also maximized to the right of the same continuation value because this function is well-approximated by $\tilde{\pi}_0^\circ$.

Suppose that $\delta \in (0, c)$ is close enough to c to ensure that $1 - (1 - \delta)c > c(1 - \delta)/\delta$. For any $\mu \in [0, 1)$, let the bounded functions $\pi_\mu^\circ : [1 - \delta, 1] \rightarrow \mathbb{R}$ and $\tilde{\pi}_\mu^\circ : [0, 1] \rightarrow \mathbb{R}$ be defined

³¹Instead of the second modification, one could interpret this contracting problem as one with the voter's payoff scaled up by a factor of $1/\mu$ whenever $\mu > 0$. This scaling is strategically irrelevant.

by the recursive equations

$$\begin{aligned}\tilde{\pi}_\mu^\circ(v) &= \frac{v}{\max\{v, 1 - \delta\}} \pi_\mu^\circ(\max\{v, 1 - \delta\}) \\ \pi_\mu^\circ(u) &= (1 - \delta)a_u + \delta \left[(1 - \mu)\tilde{\pi}_\mu^\circ\left(\frac{u - (1 - \delta)}{\delta}\right) + \mu\tilde{\pi}_\mu^\circ\left(\frac{u - (1 - \delta)(1 - ca_u)}{\delta}\right) \right],\end{aligned}$$

where a_u is as defined in the proposition's statement. By routine use of the contraction mapping theorem, several results are straightforward to establish. First, there is a unique pair of functions $(\pi_\mu^\circ, \tilde{\pi}_\mu^\circ)$ that solves this system, and π_μ° and $\tilde{\pi}_\mu^\circ$ are both continuous and concave functions. Next, the mapping $\mu \mapsto \tilde{\pi}_\mu^\circ$ is a continuous mapping with respect to the supremum-norm on the space of continuous functions. More substantively, whenever $\mu > 0$, we have $0 \leq \pi_q \leq q\mu$ and $\mu\tilde{\pi}_\mu^\circ \leq \tilde{\pi}_1 \leq \mu\tilde{\pi}_\mu^\circ + \pi_q$. Finally, focusing on $\mu = 0$, we have $\tilde{\pi}_0^\circ(v) = v$ for $v \in [0, 1 - (1 - \delta)c]$.

Using the above observations, we can now establish the corollary. Given the form of Ψ , we must show that, when $q, \mu \in (0, 1)$ are both sufficiently low, we have

$$\begin{aligned}\bar{u}_1 \notin \arg \max_{v^s \in [0, 1]} (1 - \delta)q\mu + \delta [(1 - q\mu)\pi_q + q\mu\tilde{\pi}_1(v^s)] \\ \text{subject to } 1 - \delta \leq (1 - \delta)(1 - c) + \delta v^s.\end{aligned}$$

Transforming the objective and simplifying the constraint, the goal is to show that

$$\bar{u}_1 \notin \arg \max_{v^s \in [0, 1]} \tilde{\pi}_1(v^s)/\mu \quad \text{subject to } (1 - \delta)c/\delta \leq v^s.$$

But this follows from the fact that since $v^s = 1 - (1 - \delta)c > (1 - \delta)c/\delta = \bar{u}_1$, we have

$$\frac{1}{\mu}\tilde{\pi}_1(v^s) - \frac{1}{\mu}\tilde{\pi}_1(\bar{u}_1) \geq \tilde{\pi}_\mu^\circ(v^s) - [\tilde{\pi}_\mu^\circ(\bar{u}_1) + \pi_q/\mu] \geq [\tilde{\pi}_\mu^\circ(v^s) - \tilde{\pi}_\mu^\circ(\bar{u}_1)] - q,$$

and the expression on the very right converges to $\tilde{\pi}_0^\circ(v^s) - \tilde{\pi}_0^\circ(\bar{u}_1) > 0$ as $q, \mu \rightarrow 0$. \square