Received: 30 January 2023 | Revised: 31 January 2024 | Accepted: 2 February 2024

DOI: 10.1111/ajps.12860

AJPS AMERICAN JOURNAL of POLITICAL SCIENCE

ARTICLE

Political accountability under moral hazard

Avidit Acharya[1] | **Elliot Lipnowski**[2] | **João Ramos**[3]

[1]Department of Political Science, Stanford University, Stanford, California, USA

[2]Department of Economics, Columbia University, New York, New York, USA

[3]Marshall School of Business, University of Southern California, Los Angeles, California, USA

Correspondence
Avidit Acharya, Professor of Political Science, Encina Hall West #100, Stanford University, Stanford, CA 94305, USA.
Email: avidit.acharya@gmail.com

Abstract
Viewing the relationship between politicians and voters as a principal–agent interaction afflicted by moral hazard, we examine how political careers are shaped by the incentives that voters provide incumbents to work in the public interest. When moral hazard binds, the optimal way for voters to hold politicians accountable is to provide re-election incentives that evolve dynamically over their careers in office. Under these incentives, first-term politicians are among the most electorally vulnerable and the hardest-working; politician effort rises with electoral vulnerability; electoral security increases following good performance and decreases following bad performance; and both effort and electoral vulnerability tend to decline with tenure. In extensions, we study limited voter commitment, voluntary retirement from politics, and adverse selection.

A longstanding view in the political accountability literature that originated with the seminal works of Barro (1973) and Ferejohn (1986) is that the relationship between voters and politicians is a principal–agent relationship that is afflicted by moral hazard. For instance, if the economy stagnates under an incumbent officeholder—if inflation rages or if unemployment soars—voters cannot easily tell whether it is because of mismanagement by the politician, or factors outside her control. Given this uncertainty, what is the best way for voters to hold politicians accountable if their relationship is a potentially long-term one, and the only lever of control voters have is the decision of whether to re-elect or replace the incumbent officeholder?

Prior work on political accountability (surveyed by Ashworth, 2012, and Duggan & Martinelli, 2017) has focused mainly on settings in which the politician and voter are short-lived, or in which the politician is term-limited; or, it imposes the ad hoc restriction that voters condition their vote choice only on outcomes observed during the incumbent's present term in office. From a positive standpoint, there may be good reasons to impose these restrictions. But in this paper we are concerned with the normative question of what re-election rules are optimal for voters when they are constrained by nothing other than their own credibility.

The model that we consider is very similar to that of Ferejohn (1986). It features a representative voter and a set of politicians, one of whom is in office each period of time. In every period, the politician in office privately observes a shock that takes a binary value (good/high or bad/low) and is drawn independently each period from an identical distribution. When the shock is good, the sitting politician has control over the outcome and voter welfare is determined by the politician's action in office. When the shock is bad, the politician has no control over the outcome, and voter utility is guaranteed to be low no matter what the politician does.

After observing the shock, the politician chooses her action, which represents the extent to which she serves the public interest at the expense of her private interest. (Thus, a key feature of our model, as in any accountability setting, is the lack of perfect congruence between the voter and the politician.) For simplicity, and following the broader accountability literature, we refer to the politician's action as *effort*, though it could represent the choice of any policy that favors the voter more while favoring the politician less.

Politician effort, together with the shock, determines performance. Under the bad shock, the politician cannot deliver for the voter, and so there is no reason not to act in her private interest, that is, put in no effort. Under the good shock, performance

increases linearly with effort. Crucially, the voter cannot tell whether a bad outcome was due to a bad shock 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]

In our model, a re-election rule that considers only the politician's performance in her present term in office is not always optimal. Instead, when the misalignment of interests between the voter and politician is high (as measured by the marginal cost of effort), the optimal re-election rule has the voter conditioning re-election on a broader history of the politician's performance in office. In particular, under the optimal rule, politicians have careers, with voters extracting more service out of early-term politicians via the reward of future job security. A politician's electoral security evolves stochastically and tends to grow with time in office. With certainty, some politician eventually becomes fully entrenched in office.

We also examine three variants of our model. First, we suppose that the voter will credibly renegotiate away from any continuation play that delivers him a low payoff. Similar dynamics arise in this case as in the baseline model, except that no politician ever becomes fully entrenched. Second, we look at the case in which politicians receive stochastic outside job offers while in office. We show that voters may prefer to have the politician take these offers, thus outsourcing past effort's reward to the private sector. Third, we add adverse selection to the model by assuming that politicians have types (good or bad) and type is complementary to effort. We show that a novel dynamic trade-off arises between sanctioning bad performance and selecting good types. This trade-off is absent from existing two-period models. Together, these results reveal the versatility and tractability of the canonical accountability model when we move beyond the traditional focus on stationary incentives.

Existing work on political accountability does not reveal much in terms of long-run career dynamics, as it focuses on stationary re-election rules or short time horizons. For example, Ashworth (2005) studies a three-period career concerns model in which politicians allocate effort across different activities, focusing on how career stage determines this allocation; and 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. Yet, prior literature has documented the fact that politicians do have evolving careers. Hibbing (1991), for example, shows that while incumbents do very well in general, first-term members of Congress are the most electorally vulnerable and re-election rates among those seeking a new term is generally increasing in tenure. A variety of factors shape the political lifecycle, including the returns to seniority in the party and the legislature (McKelvey & Riezman, 1992; Muthoo & Shepsle, 2014), increases in productivity and skill that come with job experience (Padró i Miquel & Snyder, 2006), changes in individual, constituent, and ruling party ideology (Ansolabehere & Snyder, 2002; Jacobson, 2015), and the creation of a loyal voting coalition among constituents, which potentially grows more stable over time (Fenno, 1978). Our model shows that if voters were to implement the normative benchmark, then the unobserved occurrence of idiosyncratic shocks could confound the link between political career dynamics and factors such as seniority, learning-by-doing, political consolidation, and so on.

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. An 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 while 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.[2] Our paper differs from this work by focusing on long-run career dynamics under optimal accountability when stationary equilibria are suboptimal.

Our paper also 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 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 (2021), 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

[1] For tractability, we depart from the Ferejohn (1986) model in two ways. First, we assume the shock is binary rather than continuous. Second, we assume effort is bounded above, and the politician has a constant marginal cost of effort, whereas Ferejohn (1986) assumes effort is unbounded with increasing marginal cost. We explain at the end of the "Career Dynamics" section how the main qualitative results of our analysis carry over if we assume effort is unbounded and/or has increasing marginal cost (while retaining the assumption of binary shocks).

[2] In 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.

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.

A distinguishing aspect of our work is that our repeated principal–agent interaction has all three of the following features: the principal (i) cannot fine-tune the agent's monetary rewards (e.g., options, bonuses, commissions), (ii) can choose to replace or retain the agent at the end of each interaction (the voter's so-called "blunt tool"), and (iii) incurs no cost to replace the agent. In addition to being a realistic feature of the political agency relationship, these features together furnish additional structure by endogenizing the principal's outside option. This allows us to further enrich the political agency model, which we demonstrate by studying settings in which the voter can renegotiate with the politician, the politician can voluntarily retire from politics, and adverse selection is also present in the agency relationship.

BASELINE MODEL

The players in the model are a representative voter and an infinite collection of politicians. Time is discrete, with an infinite horizon and indexed by $t \in \{0, 1, 2, ...\}$. Each period starts with a politician in office. One of the politicians begins in office at date 0. In each period, a political state variable θ_t is drawn from $\{0, 1\}$ independently across periods and privately seen by the officeholding politician. In state $\theta_t = 1$, the officeholder has the ability to serve the voter's interests, while in state $\theta = 0$ she does not. Hereafter, we refer to $\theta_t = 1$ as the good/high shock and $\theta_t = 0$ as the bad/low shock. We assume that in each period the good/high shock has probability $\mu \in (0, 1)$ and the bad/low shock has probability $1 - \mu$.

After seeing the shock, the politician in office chooses how much to serve the public interest, $a_t \in [0, 1]$. High levels of a_t means that she serves the public interest at the expense of her private interest, while low levels mean the opposite. Following the broader accountability literature, we refer to a_t simply as "effort," but it should be understood to mean effort at serving the public interest.

Finally in each period t, the voter 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.[3]

In each period t, the politician in office accrues a flow payoff $1 - ca_t$, where $c > 0$ is the marginal cost of effort, which measures the degree of conflict of interest between politicians and voters.[4] All politicians not in office get 0. All individuals discount the future using a common discount factor, $\delta \in (0, 1)$. In periods $t > 0$, 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.[5]

Observe that as far as on-path behavior is concerned, voter-optimal equilibrium is equivalent to the strategically simpler setting in which the voter can commit to a strategy. That is, even though our solution concept is voter-optimal equilibrium (taking voter incentives into account), studying this simpler setting with commitment yields the same predictions about the dynamics of political accountability.

To see this equivalence, considering any strategy profile σ that respects politician incentives and makes the voter indifferent whenever mixing, we can modify it to an alternative strategy profile with the same on-path behavior that additionally respects the voter's incentives. Although there are typically many such equilibria, one example is the strategy profile that agrees with σ at all of its on-path histories, and has every politician shirking and being immediately replaced at all of its off-path histories. At any history at which the voter is expected to fire or re-elect the politician, doing so will be incentive compatible (IC) because behaving differently will lead to zero effort in the future.[6] So the only relevant voter incentive constraint is that he be indifferent between firing and re-electing whenever expected to mix. We will ignore

[3] The 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 in that equilibrium.

[4] We 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.

[5] See Definitions 5.2 and 5.3 in Fudenberg and Tirole (1991) for the relevant definition of a perfect public equilibrium, which we refine by focusing on voter-optimal equilibrium.

[6] An alternative construction is one in which every politician behaves in the same way as the politician who was "supposed to be" in office. This equilibrium strikes us as more plausible, but is more cumbersome to describe formally.

this constraint in our analysis for our baseline model—it will be satisfied by the voter-optimal re-election rule described in Propositions 1 and 2.[7] We thus let an IC policy refer to a strategy profile under which no politician has an incentive to deviate from her prescribed conditional effort choice at any history.

Next, observe that a special kind of IC policy 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 and effort $a[h]$ given a good shock; 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 good 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 good 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—all of which would then be best responses—choosing a voter's favorite (or near-favorite) among them would offer a weak improvement. Moreover, working is never optimal for a politician under the bad shock, since it leads to an inferior payoff in the stage game with no effect on the public outcome of the dynamic game. So when the shock is bad, the politician in office chooses not to exert any effort. Now, let $a[h]$ denote the effort that the politician chooses following a good shock, and define success and failure as in the previous paragraph. 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 generates the same voter payoff as the original policy because it generates the same path of play. So, all that remains is to verify that it is IC.

To see that it is IC, note that under the modified voter behavior, a politician always prefers to choose 0 over $a \in (0, a[h])$ and $a[h]$ over $a > a[h]$. Therefore, 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.

The voter's first-best payoff is μ, which obtains if and only if every politician in office exerts maximum effort in every high-shock period.

Before building on the above observations to study voter-optimal policies, we use them to study the benchmark class of stationary equilibria. The following proposition reports the main features of optimal stationary equilibria and shows that they are characterized by a cutoff retention rule.

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

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 $\mu a^* = \mu \min\{1, \delta/c\}$. Thus, a stationary equilibrium exists that achieves the voter's first-best payoff if and only if $\delta \geq c$.

The proofs of all formal results in this paper appear in the Online Appendix.

CAREER DYNAMICS

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. Before proceeding, we note that we have not assumed that $c < 1$; therefore, the analysis of this case also concerns the case of high δ (arbitrarily close to 1) if the marginal cost of effort is $c \geq 1$.

We take a recursive approach à la Spear and Srivastava (1987). More specifically, we think of the voter as solving a dynamic programming problem whose relevant state variable is the in-office politician's continuation value from the current period onward. We start by noting that the politician's continuation value u starting from any period must lie in the interval $[1-\delta, 1]$. She can guarantee herself a continuation payoff of $1-\delta$ by shirking indefinitely, and can attain a payoff of at most 1 by accruing the benefits of office in every period without ever exerting any effort. The voter can generate any continuation payoff for the politician in the interval $[1-\delta, 1]$ by randomizing between two extremes—immediate replacement after the current term and the offer of unconditional full job security.

[7] In our extension with adverse selection, firing will have an (endogenous) cost, so we will augment the model with a public randomization device to allow the voter to optimally use a pure strategy.

Recall (as noted in the previous section) all successes are treated equally, and all failures treated equally, at a given history. Thus, the optimal reelection rule takes the politician's continuation value $u \in [1-\delta, 1]$ at the start of the period to two possible continuation values at the start of the subsequent period, the value v_s associated with the politician succeeding to deliver the voter's equilibrium prescribed payoff in the current period, and the value v_f associated with the politician failing to do so.

To characterize the law of motion of politician continuation values, two observations of the voter's optimal value function are useful. First, the function is concave: Variation in the politician's continuation value is costly to the voter. Second, the function is decreasing: Firing the current politician weakly benefits the voter since the newly elected politician could be offered the same continuation equilibrium as the in-office one. This property tells us that a newly elected politician begins her career at continuation value $1-\delta$.

The law of motion of politician continuation values then satisfies two constraints. The first is the voter's promise-keeping (PK) constraint, which says that a politician's starting continuation value is the value that the voter must give her:

$$u = \mu[(1-\delta)(1-ca) + \delta v_s] + (1-\mu)[(1-\delta)1 + \delta v_f]. \tag{PK}$$

The other constraint is the politician's IC constraint, which says that the politician must have an incentive to put in the equilibrium prescribed effort a when the shock is good:

$$(1-\delta)(1-ca) + \delta v_s \geq (1-\delta)1 + \delta v_f. \tag{IC}$$

For the voter to be optimizing, this constraint in fact holds with equality—because concavity of the optimal value function means superfluous variation of the politician's continuation value is costly to the voter. Thus, the (PK) constraint and the (IC) constraint holding with equality together give us two equations in two unknowns, namely, v_f and v_s. Solving these equations gives us the exact form for the law of motion of politician continuation values:

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

We now specify the politician's equilibrium prescribed effort a as a function of her continuation value u at the start of the period; denote this a_u. We show this effort is as high as the politician's incentives (and the promised continuation value of u) will allow: It is the maximum possible effort level of 1 if the continuation value is low enough, and otherwise makes the politician indifferent when success is rewarded with perfect future job security.[8] That is, since v_s cannot exceed 1, when u is close enough to 1, we must have the equilibrium prescribed effort be $a_u = (1-u)/[(1-\delta)c]$.

Finally, because the politician's continuation value conditional on being retained cannot be below $1-\delta$, if u is low enough, then following failure by the politician we must have the voter mix between retaining the politician at continuation value $1-\delta$ from the next period on, and firing her to give her a payoff of 0. To guarantee that the politician's continuation value following failure is indeed v_f, this mixing probability must be $v_f/(1-\delta)$ for small enough values of u. Moreover, that the value function is decreasing lets us show the voter should only fire at this minimum feasible rate.

We now have our main result for the baseline model, which summarizes the observations above about voter-optimal accountability.

Proposition 1. *Suppose $\delta < c$. 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.

The behavior described in Proposition 1 has the politician working as hard as can be made IC (given the politician's current continuation value)—and hence being made indifferent between working and shirking—and the voter replacing the politician as little as possible.[9]

[8] This feature is the only feature of Proposition 1, or its proof, that relies on our assumptions of bounded effort with linear costs. In the general case (outlined under the "General costs" heading at the end of this section), a closed form is inaccessible, but one can use concavity of the voter's optimal to show that a_u is weakly decreasing in u.

[9] Although stationary equilibria become optimal at $\delta = c < 1$, 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.

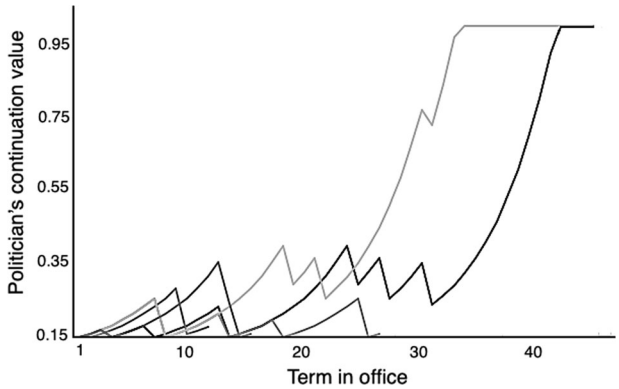

FIGURE 1 Simulated continuation value paths. *Note*: The figure shows examples of simulated paths of politician continuation values for the following parameter values: $\mu = 0.8$, $c = 0.9$, and $\delta = 0.85$. The important thresholds are: $1 - \delta = 0.15$, $1 - \delta^2 = 0.2775$, and $1 - (1 - \delta)c = 0.865$.

As a politician's continuation value climbs with a stream of successes, it is possible for her to earn full job security on the path of play. Whenever her continuation value is above $1 - (1 - \delta)c$, being successful in the current term results in complete entrenchment. The corollary below shows that with a long enough run of successes, the politician's continuation value can cross this threshold.

Corollary 1. *Suppose $\delta < c$, and let $k \in \mathbb{N}$ be smallest such that $\frac{\delta}{1-\delta^k} \leq c$. Then, any politician who has k consecutive successes will be completely entrenched.*

We now make a series of observations. First, because the sequence of cost cutoffs in the corollary converges to δ, for any value of $c > \delta$, some politician will eventually achieve enough consecutive successes to become entrenched in office.

Second, 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 $1 - \delta$.

Third, to give a political interpretation to the conditions stated in the corollary, note that since the conditions do not rely on μ, we can interpret δ as reflecting both the frequency of elections and how voters and politicians trade off the future against the present. Given this interpretation, the corollary implies that the minimum time to entrenchment is shorter if elections are less frequent, if politics is more short-termist, and if the conflict of interest c between the voter and politicians is greater.

Figure 1 makes our observations concrete by depicting some simulated paths of the politician's continuation value for a numerical example. It shows that some politicians have longer careers than others. Some become completely entrenched more quickly; others take longer; and still others are 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 their careers being reset (i.e., the continuation value falling to $1 - \delta$). Failure later on in a politician's career is less likely to result in removal or a reset. 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. Therefore, while job security tends to grow over a politician's career in office, it need not grow monotonically. Similarly, a politician's effort tends to decline over her career in office, but also need not decline monotonically.[10]

We emphasize that our observations about the path of play are not merely features of a special voter-optimal equilibrium characterized in Proposition 1. These observations are strengthened by the uniqueness result stated in the proposition: All voter-optimal equilibria are equivalent in terms of the paths of play that they generate. We prove this in Online Appendix Section A.2 (pp. 3–8).

Finally, in Online Appendix Section B (pp. 23–24) we numerically compute the voter's optimal equilibrium payoff for different parameter values, and compare it to his payoff in the best stationary equilibrium and his first-best payoff. The computations show that there are significant welfare gains from removing the restriction to stationary equilibria.

General costs. The qualitative structure of the voter's preferred equilibrium transports to a more general version of our model with weakly increasing marginal effort cost, and effort allowed to be bounded or unbounded. This specification includes the case of unbounded effort with strictly convex costs, exactly as in the Ferejohn (1986) formulation. In this more general specification, each history still has only one relevant IC constraint for the politician in office, which continues to hold with equality in the voter's preferred equilibrium; 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 rise with success and fall with failure. The voter's decision to

[10] Whereas 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.

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 decreasing in her continuation value. All of these observations can be proven by following the exact same proof, mutatis mutandis, as that of Proposition 1.

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. As the following three sections of the paper show, having such a closed-form solution facilitates various detailed extensions that let us further explore the role of dynamics in political careers.

The other aspect of our model that differs from Ferejohn (1986) is the shock distribution: We assume binary shocks whereas Ferejohn (1986) has continuous shocks. This feature of our model is an important tractability assumption.

RENEGOTIATION-PROOF EQUILIBRIA

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 full job security. But when a politician becomes completely entrenched, the voter can expect to derive no subsequent value from her. Can the voter commit to tolerating this low 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.

We now limit the voter's ability to commit to tolerating these low future payoffs by imposing a threshold, $\underline{\pi}$, such that his equilibrium payoff cannot fall below $\underline{\pi}$ at any history. This approach is reminiscent of the concept 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 $\underline{\pi}$: The higher $\underline{\pi}$ is, the less ability the voter has to tolerate unfavorable outcomes ex post. After documenting play that will hold for any such threshold $\underline{\pi}$, we then turn (in Corollary 2) to endogenizing $\underline{\pi}$. In line with Pearce (1987), we focus on the highest threshold that can possibly be sustained, essentially allowing the voter to renegotiate away from the path whenever such renegotiation is itself credible.[11]

Given an exogenous limit to the voter's ability to commit, the next proposition shows that the sole effect of this limit on voter-optimal play is to cap the politician's achievable continuation value.

Proposition 2. *Suppose $\delta < c$, and fix a payoff lower bound $\underline{\pi} > 0$. Some $\hat{u} \in [1-\delta, 1)$ exists that is decreasing in $\underline{\pi}$ 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 $\underline{\pi}$: (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 the following:*

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

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 $\underline{\pi}$, the proposition says that an "entrenchment limit" $\hat{u}$ exists that could assure the voter a value of $\underline{\pi}$ 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 dynamics of a politician's career are qualitatively similar to those of the baseline model.[12]

The main justification for a lower bound $\underline{\pi}$ on continuation payoffs is that the voter might renegotiate

[11] While our approach is inspired by Pearce (1987), one key 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 assuming that the principal can unilaterally initiate renegotiation.

[12] Another implication that follows immediately from the proposition is that, for the purposes of limiting politician entrenchment, it is suboptimal to impose a fixed term limit of any length.

away from an equilibrium when continuation play gives him a low payoff, say $\pi_L < \underset{\sim}{\pi}$. 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 consistency requirement can be taken a step further by asking for a sort of external consistency: If an equilibrium exists that guarantees a continuation payoff of $\pi_H > \pi_L$ at every possible future history, then the equilibrium that guarantees only π_L not 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, $\underset{\sim}{\pi}$, satisfies both consistency properties: the highest feasible lower bound $\underset{\sim}{\pi}$. Intuitively, a higher bound is unattainable, whereas a lower bound would not be immune to renegotiation.

One may wonder whether choosing $\underset{\sim}{\pi}$ 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 $\underset{\sim}{\pi}$ exists such that some equilibrium gives the voter a payoff of at least $\underset{\sim}{\pi}$ after every history, while every stationary equilibrium gives the voter a payoff strictly below $\underset{\sim}{\pi}$. In particular, an equilibrium of the form described in Proposition 2 exists such that the voter is strictly better off after every history than under any stationary equilibrium.*

The corollary shows that allowing the voter to renegotiate away from bad equilibria (i.e., limiting voter commitment) can actually *preclude* stationary play. Following Pearce's justification that renegotiation itself must be credible (which selects $\underset{\sim}{\pi}$ 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.

VOLUNTARY RETIREMENT

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. For example, in recent work, Egerod (2021) finds that when indicators of job prospects outside of politics improve, sitting politicians become more likely to retire from politics, while Hall and Van Houweling (1995) note that politicians that are more electorally insecure are more likely to retire from politics (see, also, Diermeier et al., 2005; Groseclose & Krehbiel, 1994; Keane & Merlo, 2010).

Suppose we enrich the baseline model to allow for voluntary retirement by assuming that in each period, a politician receives with some probability the opportunity to leave her career in politics and work elsewhere. In what follows, we study how this possibility affects the politician's optimal incentives and career dynamics.

Specifically, consider a modified model in which the in-office 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. The 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."[13]

Each period t proceeds as follows. First, the sunspot 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.[14]

Proposition 3. *Suppose $\delta < c$ and $1-\delta \le w \le 1$. 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 the following:*

1. *If $\tilde{u} \le w$, 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 $a_{u,v^*} := \min\{1, \frac{(1-\delta)+\delta v^*-u}{(1-\delta)c}\}$.*
3. *Given the outcome y in period t, if $v_u(a_{u,v^*}, y) \ge (1-p)(1-\delta) + pw$, then the politician is retained with probability 1 at continuation value $v_u(a_{u,v^*}, y)$ starting from 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 it is available.

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 (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 a more detailed calculation is required in the general case. The next corollary establishes that a sufficient condition for this property is that $w < 1$ is high enough.

Corollary 3. *Suppose $\delta < c$. 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 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

[13] This assumption is useful technically, but we do not view it as substantive. 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.

[14] Various 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 immediately receive the outside opportunity.

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. 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.

ADDING ADVERSE SELECTION

We now enrich the baseline model by assuming that each politician has a perfectly persistent type $\omega \in \{0,1\}$, which only she knows as of her first term in office. Types are independent across politicians and take 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-type politicians do not produce any value to voters no matter how strongly they are incentivized to do so. We thus refer to type $\omega = 1$ as the good/high type and type $\omega = 0$ as the bad/low type.

As in the previous section, it will be convenient here to assume the voter and politicians have access to a public correlation device. Specifically, we assume a randomization device 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 nonstationarity 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 what the strategic experimentation literature calls "perfectly revealing good news." Under this special structure, the politician, after her first term in office, is either known to be a good type (due to a past success) or is less likely to be the good type than an untried challenger. A feature of the baseline model that carries over to this model is, therefore, that if a politician fails in her first term, she is immediately replaced.

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

The lemma 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 $\delta < c$. 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 that in the model the voter faces a replacement cost. When he boots the incumbent out of office, he gets a flow payoff of zero until he draws a skilled type (and the high shock). 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 good 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 trade-off. 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. These conditions entail q and μ being low. For low values of q, the cost of firing the politician is high, as it is unlikely that the voter will discover a good politician soon after removing the incumbent. For low values of μ, the risk of having to soon reward the politician with full job security is also low, as the likelihood of a success—even if the politician is good—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. 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 good-type politician with strict incentives to work in the first term and still incentivize maximal effort in the next. Therefore, 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 her to have.

In the standard two-period accountability model, when type and effort 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. 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.[15]

However, our dynamic setting exhibits a novel trade-off between sanctioning and selection that can arise even though effort and ability are complements. 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 him reluctant to replace a known good type. The more reluctant he is, the higher is the good type's continuation value; and the higher this continuation value, the less the politician tends to work in the future. This intertemporal trade-off between selection and sanctioning is a novel political phenomenon that arises in the multiperiod model but is absent from the two-period model.

[15] Their observation contrasts with the case studied by Fearon (1999) and Duggan and Martinelli (2020) in which effort and type are substitutes. For example, suppose 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, he makes the bad type look more like the good type, undermining his ability to screen them. Thus, when effort and type are substitutes, there is a trade-off between selecting good types and sanctioning bad performance.

CONCLUSION

The restriction to stationary strategies in infinite-horizon models of political accountability precludes an understanding of how optimal re-election rules shape the careers of politicians. Looking at a model of dynamic moral hazard, we have clarified the link between these re-election rules and the career paths of politicians, showing that optimal re-election rules produce rich career dynamics when the conflict of interest between voters and politicians is large enough. The key observation that comes out of our analysis is that politicians tend to work less in the voter's interest while becoming electorally more secure. An important takeaway from this observation is that one cannot conclude from the fact that a politician becomes entrenched that the voter has not been holding her accountable.

More specifically, 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. 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.

ACKNOWLEDGMENTS

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, Toby Nowacki, and Sam Wycherley provided excellent research assistance. A previous draft of this paper was circulated with the title "Optimal Political Career Dynamics."

REFERENCES

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58(5): 1041–63.

Aghion, Philippe, and Matthew O. Jackson. 2016. "Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits." *American Economic Journal: Microeconomics* 8(3): 1–38.

Anesi, Vincent, and Peter Buisseret. 2020. "Making Elections Work: Accountability with Selection and Control." Technical report Mimeo, University of Chicago.

Ansolabehere, Stephen, and James M. Snyder. 2002. "The Incumbency Advantage in US Elections: An Analysis of State and Federal Offices, 1942–2000." *Election Law Journal* 1(3): 315–38.

Ashworth, Scott. 2005. "Reputational Dynamics and Political Careers." *Journal of Law, Economics, and Organization* 21(2): 441–66.

Ashworth, Scott. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work." *Annual Review of Political Science* 15: 183–201.

Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg. 2017. "Accountability and Information in Elections." *American Economic Journal: Microeconomics* 9(2): 95–138.

Atkeson, Andrew, and Robert E. Lucas. 1992. "On Efficient Distribution with Private Information." *The Review of Economic Studies* 59(3): 427–53.

Banks, Jeffrey S., and Rangarajan K. Sundaram. 1990. "Incumbents, Challengers, and Bandits: Bayesian Learning in a Dynamic Choice Model." Technical Report University of Rochester-Center for Economic Research, RCER.

Banks, Jeffrey S., and Rangarajan K. Sundaram. 1993. "Adverse Selection and Moral Hazard in a Repeated Elections Model." *University of Rochester-Center for Economic Research (RCER)* 12: 295–311.

Barro, Robert J. 1973. "The Control of Politicians: An Economic Model." *Public Choice* 14(1): 19–42.

Clementi, Gian Luca, and Hugo Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *Quarterly Journal of Economics* 121(1): 1131–75.

DeMarzo, Peter M., and Michael J. Fishman. 2007. "Optimal Long-term Financial Contracting." *Review of Financial Studies* 20(6): 2079–128.

Diermeier, Daniel, Michael Keane, and Antonio Merlo. 2005. "A Political Economy Model of Congressional Careers." *American Economic Review* 95(1): 347–73.

Duggan, John, and César Martinelli. 2017. "The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness." *Journal of Economic Literature* 55(3): 916–84.

Duggan, John, and César Martinelli. 2020. "Electoral Accountability and Responsive Democracy." *The Economic Journal* 130(627): 675–715.

Egerod, Benjamin C. K. 2021. "The Lure of the Private Sector: Career Prospects Affect Selection Out of Congress." *Political Science Research and Methods* 10(4): 722–38.

Fearon, James D. 1999. "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance." In *Democracy, Accountability, and Representation*, edited by Bernard Manin, Adam Przeworski and Susan Stokes, 55–97. Cambridge: Cambrige University Press.

Fenno, Richard F. 1978. *Home Style: House Members in Their Districts*. Boston, MA: Little, Brown, and Co.

Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50(1-3): 5–25.

Foarta, Dana, and Takuo Sugaya. 2021. "Wait-and-See or Step In? Dynamics of Interventions." *American Economic Journal: Microeconomics* 13(1): 399–425.

Fong, Yuk-fai, and Jin Li. 2017. "Relational Contracts, Limited Liability, and Employment Dynamics." *Journal of Economic Theory* 169(C): 270–93.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: The MIT Press.

Groseclose, Timothy, and Keith Krehbiel. 1994. "Golden Parachutes, Rubber Checks, and Strategic Retirements from the 102nd House." *American Journal of Political Science* 38(1): 75–99.

Guo, Yingni, and Johannes Hörner. 2020. "Dynamic Allocation without Money." TSE Working Paper.

Hall, Richard L., and Robert P. Van Houweling. 1995. "Avarice and Ambition in Congress: Representatives' Decisions to Run or Retire from the US House." *American Political Science Review* 89(1): 121–36.

Hibbing, John R. 1991. *Congressional Careers: Contours of Life in the US House of Representatives*. Chapel Hill, NC: UNC Press Books.

Jacobson, Gary C. 2015. "It's Nothing Personal: The Decline of the Incumbency Advantage in US House Elections." *Journal of Politics* 77(3): 861–73.

Kartik, Navin, and Richard Van Weelden. 2019. "Reputation Effects and Incumbency (Dis) Advantage." *Quarterly Journal of Political Science* 14(2): 131–57.

Keane, Michael P., and Antonio Merlo. 2010. "Money, Political Ambition, and the Career Decisions of Politicians." *American Economic Journal: Microeconomics* 2(3): 186–215.

Li, Jin, Niko Matouschek, and Michael Powell. 2017. "Power Dynamics in Organizations." *American Economic Journal: Microeconomics* 9(1): 217–41.

Lipnowski, Elliot, and Joao Ramos. 2020. "Repeated Delegation." *Journal of Economic Theory* 188(C): 105040.

Mattozzi, Andrea, and Antonio Merlo. 2008. "Political Careers or Career Politicians?" *Journal of Public Economics* 92(3-4): 597–608.

McKelvey, Richard D., and Raymond Riezman. 1992. "Seniority in Legislatures." *American Political Science Review* 86(4): 951–65.

Muthoo, Abhinay, and Kenneth A. Shepsle. 2014. "Seniority and Incumbency in Legislatures." *Economics & Politics* 26(1): 13–37.

Padró i Miquel, Gerard, and James M. Snyder. 2006. "Legislative Effectiveness and Legislative Careers." *Legislative Studies Quarterly* 31(3): 347–81.

Padró i Miquel, Gerard, and Pierre Yared. 2012. "The Political Economy of Indirect Control." *Quarterly Journal of Economics* 127(2): 947–1015.

Pearce, David G. 1987. "Renegotiation-proof Equilibria: Collective Rationality and Intertemporal Cooperation." Technical report Cowles Foundation for Research in Economics, Yale University.

Schwabe, Rainier. 2009. "Reputation and Accountability in Repeated Elections." In *APSA 2009 Toronto Meeting Paper*.

Spear, Stephen E., and Sanjay Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *Review of Economic Studies* 54(4): 599–617.

Thomas, Jonathan, and Tim Worrall. 1990. "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem." *Journal of Economic Theory* 51(2): 367–90.

Yared, Pierre. 2010. "A Dynamic Theory of War and Peace." *Journal of Economic Theory* 145(5): 1921–50.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Acharya, Avidit, Elliot Lipnowski, and João Ramos. 2024. "Political accountability under moral hazard." *American Journal of Political Science* 1–12. https://doi.org/10.1111/ajps.12860

Supplemental Online Appendix for "Political Accountability Under Moral Hazard"

Avidit Acharya* Stanford
Elliot Lipnowski† Columbia
João Ramos‡ USC

January 31, 2024

Contents

*Professor of Political Science, Stanford University; email: avidit@stanford.edu.
†Assistant Professor 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.

A Proofs

A.1 Proof of Proposition 0

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 parameterized by a stationary effort prescription a and a stationary pair of firing probabilities ψ_0 and ψ_a, such that the officeholder is removed with probability ψ_0 following $y < a$, and with probability ψ_a following $y \geq a$.[1] 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.

In addition, observe that either $\psi_0 = 1$ or $\psi_a = 0$ can be set without loss of optimality. If both ψ_0 and ψ_a were interior, then raising ψ_0 while lowering ψ_a so that the average 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 0. 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 proposition follows.

[1] Following the literature, we use the term "stationary" to describe play that may condition on occurrences within the current period, but not on those from past periods or on calendar time.

A.2 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:

$$\pi(u) = \Phi\tilde{\pi}(u) := \sup_{a\in[0,1],\ v_s,v_f\in[0,1]} (1-\delta)\mu a + \delta\left[\mu\tilde{\pi}(v_s) + (1-\mu)\tilde{\pi}(v_f)\right]$$

$$\text{subject to} \quad u = (1-\delta)(1-\mu ca) + \delta[\mu v_s + (1-\mu)v_f] \tag{PK}$$

$$\text{and} \quad (1-\delta)1 + \delta v_f \leq (1-\delta)(1-ca) + \delta v_s; \tag{IC}$$

$$\tilde{\pi}(v) = \tilde{\Phi}\pi(v) := \sup_{\rho\in[0,1],\ u,u_0\in[1-\delta,1]} \rho\pi(u) + (1-\rho)\pi(u_0)$$

$$\text{subject to} \quad \rho u = v; \tag{$\tilde{\text{PK}}$}$$

where Φ, $\tilde{\Phi}$ 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 $\tilde{\pi}$, the quantity ρ is the probability that the incumbent is retained, v is next period's average continu-

ation 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 (as we explained in the main text) we can characterize voter-optimal equilibria by studying the fictitious environment in which the voter can commit. But there is the promise-keeping ($\tilde{\text{PK}}$) 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.[2] 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 $\pi(v) \geq \frac{v}{u}\pi(u) + \left(1 - \frac{v}{u}\right)\tilde{\pi}(0)$ for every $u \in (v, 1]$. But this condition rearranges to the requirement that

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

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

[2]These arguments proceed as follows. The voter's optimal payoff function $\pi : [1-\delta, 1] \to \mathbb{R}$ must satisfy the Bellman equation $\pi = \Phi\tilde{\Phi}\pi$, and is bounded because the stage game has bounded payoffs. 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 (appealing to the Blackwell monotonicity conditions) 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.

weakly increasing and right-continuous in $v \in [1-\delta, 1)$,[3] 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 with (PK) immediately gives a solution for v_s and v_f. Specifically,[4]

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

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

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

Since $\overline{u}_a$ 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 $(1-\delta)\mu[1 + c\tilde{\pi}'(\overline{u}_a)]$.[5] 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,

[3]For example, these properties can be seen from the formula

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

[4]Formally, we define the functions $\overline{\cdot}_a : \mathbb{R} \to \mathbb{R}$ (for $0 \leq a \leq 1$) and $\underline{\cdot} : \mathbb{R} \to \mathbb{R}$, so that $\overline{u}_a$ and $\underline{u}$ are functions of u.

[5]In 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.

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 $\overline{u}_1 = v^*$ 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 $\overline{u}_a = v^*$ 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

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

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 sets $u_0 = 1 - \delta$, starting any newly elected politician at continuation value $1 - \delta$. It also follows that the voter optimally takes $u = \max\{1-\delta, v\}$—i.e., takes $u \in [1-\delta, 1]$ as small as possible subject to ($\tilde{\text{PK}}$). To see the latter, observe that $v_* = 1 - \delta$ if and only if $\tilde{\pi}(1 - \delta) = \pi(1 - \delta)$, which we have noted earlier in the proof is equivalent to having $\pi(v) \geq \tilde{\pi}(0) + v\pi'(v^+)$ for $v = 1 - \delta$. But this inequality reduces to $\pi(1 - \delta) \geq \pi(1 - \delta) + (1 - \delta)\pi'(1 - \delta)$, which holds because π is nonincreasing.

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

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

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 $\overline{u}_{a_u}$ for any $u \in [1-\delta, 1)$. Finally, π is strictly decreasing because it is nonincreasing and strictly concave, so firing as infrequently as ($\tilde{\text{PK}}$) 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,

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

and, therefore,

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

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

A.3 Proof of Corollary 1

A straightforward induction argument shows that if the politician starts with continuation value $1-\delta$, then following k successes (using the law of motion in Proposition 1 following success) her continuation value will be the minimum of 1 and

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

Therefore, the politician will be completely entrenched if and only if $u_k \geq 1$, or equivalently $c \geq \delta/(1-\delta^k)$. Finally, because the success continuation value is weakly increasing in the current continuation value, it follows by induction that any politician with k consecutive successes becomes completely entrenched. □

A.4 Proof of Proposition 2

Let $\mathcal{U} = \mathcal{U}(\underset{\sim}{\pi})$ be the set of all incumbent politician payoffs that can arise in an equilibrium in which the voter's continuation payoff is at least $\underset{\sim}{\pi}$ after every history, and let $\hat{u} = \hat{u}(\underset{\sim}{\pi}) \in [1-\delta, 1]$ 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 $\underset{\sim}{\pi}$ 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 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 $\overline{u}_{a_{u,\hat{u}}} \leq \hat{u}$. Let $\pi_{\hat{u}} : [1-\delta, \hat{u}] \to \mathbb{R}$ 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 $\underset{\sim}{\pi}$ implies, given optimality of the modified equilibrium and the definition of $\mathcal{U}$, that $\pi_{\hat{u}}(u) \geq \underset{\sim}{\pi}$ for every $u \in \mathcal{U}$. That $\pi_{\hat{u}}$ is continuous implies that $\pi_{\hat{u}}(\hat{u}) \geq \underset{\sim}{\pi}$, and that it is nonincreasing implies that $\pi_{\hat{u}} \geq \underset{\sim}{\pi}$, as required.

Finally, to see that $\hat{u}(\cdot)$ is strictly decreasing, assume for a contradiction that $\hat{u}' := \hat{u}(\underset{\sim}{\pi} - \epsilon) \leq \hat{u}(\underset{\sim}{\pi})$ for some $\epsilon \in (0, \underset{\sim}{\pi})$. 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 $\underset{\sim}{\pi} - \epsilon$ from every history. Existence of such an equilibrium contradicts the definition of $\hat{u}' = \hat{u}(\underset{\sim}{\pi} - \epsilon)$. $\square$

A.5 Proof of Corollary 2

To show that such a $\underset{\sim}{\pi}$ 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 parameterized 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:

$$\lim_{\hat{u}\to 1-\delta} a_{1-\delta,\hat{u}} = \min\{1, \frac{\delta}{c} + \frac{\delta}{(1-\delta)c}[\hat{u} - (1-\delta)]\} = \frac{\delta}{c}$$

$$\lim_{\hat{u}\to 1-\delta} v_{\hat{u}}(a_{\hat{u},\hat{u}}, 0) = \frac{\hat{u} - (1-\delta)}{\delta} = 0.$$

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:

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

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

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

From this, it follows that

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

And, therefore,

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

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. □

A.6 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, $\operatorname{cav}\hat{\pi}(\tilde{u})$, where $\operatorname{cav}\hat{\pi}$ is the concave envelope of $\hat{\pi}$.[6] 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:

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

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

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

condition on a sunspot after a given politician is elected or re-elected, continuation values are delivered via $\operatorname{cav}\hat{\pi}$ 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, $\hat{\pi}$, $\operatorname{cav}\hat{\pi}$, 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 $\operatorname{cav}\hat{\pi}$ 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, $\operatorname{cav}\hat{\pi}$ necessarily agrees with $\hat{\pi}$ at $\tilde{u}_L$, which is an extreme point of the two functions' domain. Second, because $\operatorname{cav}\hat{\pi} \geq \hat{\pi}$ and setting $\tilde{u}_0 = u$ $(= w)$ is feasible in the program determining $\hat{\pi}(w)$, it follows that $\sup_{\tilde{u}_0 \in [\tilde{u}_L, 1]} \operatorname{cav}\hat{\pi}(\tilde{u}_0) \geq \pi(w)$, 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 ($\hat{\text{PK}}$) requires $\tilde{u} - w$ to have the same sign as $u - w$, the incentive-compatibility constraint ($\hat{\text{IC}}$) 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 $\operatorname{cav}\hat{\pi}$ is nonincreasing) the voter optimally sets the continuation value $\tilde{u}$ to be as small as possible subject to $\tilde{\text{PK}}$, yields a simplified system of Bellman equations:

$$\pi(u) = \sup_{a \in [0,1]} (1-\delta)\mu a + \delta\left[\mu\tilde{\pi}(\bar{u}_a) + (1-\mu)\tilde{\pi}(\underline{u})\right]$$

$$\text{subject to} \quad \bar{u}_a \leq 1;$$

$$\hat{\pi}(\tilde{u}) = \begin{cases} (1-p)\pi(\frac{\tilde{u}-pw}{1-p}) + p\pi(1-\delta) & \text{if } \tilde{u} \leq w \\ \pi(\tilde{u}) & \text{if } \tilde{u} > w; \end{cases}$$

$$\tilde{\pi}(v) = \operatorname{cav}\hat{\pi}(\max\{v, \tilde{u}_L\}).$$

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

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

It is immediate that the given functional form is a lower bound for $\operatorname{cav}\hat{\pi}$, 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 $\operatorname{cav}\hat{\pi}$ and, hence, coincides with it. Indeed, because $\hat{\pi}(w) \geq \hat{\pi}(w^+)$ and $\hat{\pi}'(w^-) \geq \lim_{\tilde{u}\to w^+} \hat{\pi}'(x)$ if $w < 1$, it is straightforward to show that this function is concave and above $\hat{\pi}$—hence, above $\operatorname{cav}\hat{\pi}$. The characterization of $\operatorname{cav}\pi$ 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,

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

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 $[0, \tilde{u}_L]$, 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

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

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

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

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. □

A.7 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

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

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

A.8 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 tomorrow will be a high type. The result follows by continuing the argument recursively on the histories of the game. □

A.9 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:[7]

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

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 proof of Proposition 1, up to the paragraph establishing $v_* = 1-\delta$, 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 $\tilde{\pi}_1|_{[v_*,1]} = \pi_1|_{[v_*,1]}$;

[7] It is standard that these optimal values are 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—e.g., 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.

and the optimal politician effort and continuation values prescribed by Φ 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

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

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.[8] 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,

[8]Given 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.

it must be that $v_* < 1$. Then, for $u \in (v_*, 1)$ high enough,

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

Therefore,

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

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

$$\begin{aligned}\tilde{\pi}_1'(u) \leq \pi_1'(u) &= (1-\mu)\tilde{\pi}_1'\left(\tfrac{u-(1-\delta)}{\delta}\right) + \mu\tilde{\pi}_1'\left(\tfrac{u-(1-\delta)(1-c)}{\delta}\right) \\ &= (1-\mu)\tilde{\pi}_1'(u) + \mu\tilde{\pi}_1'\left(\tfrac{u-(1-\delta)(1-c)}{\delta}\right),\end{aligned}$$

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,

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

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.10 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 shock turns

out to be low.[9] 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, parameterized 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 $\pi_\mu^\circ : [1-\delta, 1] \to \mathbb{R}$ and $\tilde{\pi}_\mu^\circ : [0, 1] \to \mathbb{R}$ be defined by the recursive equations

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

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 $v \in [0, 1-(1-\delta)c]$.

[9] Instead 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.

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

$$\bar{u}_1 \notin \underset{v^s \in [0,1]}{\arg\max}\ (1-\delta)q\mu + \delta\left[(1-q\mu)\pi_q + q\mu\tilde{\pi}_1(v^s)\right]$$
$$\text{subject to} \quad 1-\delta \leq (1-\delta)(1-c) + \delta v^s.$$

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

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

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

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

and the expression on the very right converges to $\tilde{\pi}^\circ_0(v^s) - \tilde{\pi}^\circ_0(\bar{u}_1) > 0$ as $q, \mu \to 0$. $\square$

B Voter Payoff Gains

We now quantify the welfare gains the voter can obtain once we depart from the stationary equilibrium benchmark. To that end, we focus on the ex-ante payoff of the voter in the voter-optimal equilibrium described in the paper and compare it to the ex-ante voter payoff in the best stationary equilibrium. Numerically computing the voter-optimal equilibrium payoff is especially tractable given our analytical solution for the voter-optimal policy. We display the voter's payoff as a percentage of the first-best voter payoff for different payoff specifications, in Table B.1.[10]

We would like to highlight some features of these payoff calculations. First, in the case in which players are patient enough, $\delta > c$, a stationary equilibrium attains the first-best payoff. Of course, in this case, the stationary equilibrium achieves the same voter payoff as the voter-optimal equilibrium, as represented in Table B.1's green

	$c = 0.9$	$c = 1.4$	$c = 1.9$
$\delta = 0.5$	83.4%	54.9%	40.5%
$\delta = 0.7$	99.5%	87.1%	70.7%
$\delta = 0.9$	100%	99.9%	96.8%

(a) Voter-optimal Equilibrium ($\mu = 0.3$)

	$c = 0.9$	$c = 1.4$	$c = 1.9$
$\delta = 0.5$	69.2%	44.6%	32.9%
$\delta = 0.7$	94.1%	68.1%	51.1%
$\delta = 0.9$	100%	96.9%	78.3%

(b) Voter-optimal Equilibrium ($\mu = 0.6$)

	$c = 0.9$	$c = 1.4$	$c = 1.9$
$\delta = 0.5$	58.5%	37.6%	27.7%
$\delta = 0.7$	82.2%	53.7%	39.6%
$\delta = 0.9$	100%	70.3%	52.0%

(c) Voter-optimal Equilibrium ($\mu = 0.9$)

	$c = 0.9$	$c = 1.4$	$c = 1.9$
$\delta = 0.5$	55.5%	35.7%	26.3%
$\delta = 0.7$	77.7%	50.0%	36.8%
$\delta = 0.9$	100%	64.3%	47.4%

(d) Best Stationary Equilibrium

Table B.1: The four tables display the ex-ante voter payoff achievable in equilibrium as a percentage of the first-best voter payoff, for different values of discount factor and politician's marginal cost of effort. The first three tables show the voter payoff achievable in the voter-optimal equilibrium, for various probabilities (μ) of the politician's effort mattering. Finally, the fourth table displays the voter payoff in the best stationary equilibrium (which is independent of μ), again as a percentage of the first-best payoff. Finally, in all four tables, we use colors to indicate the patience of the players. The cells in green indicate high patience, $\delta \geq c$; while the cells in red indicate low patience, $\delta \leq \frac{c}{1+c}$.

[10] The first-best voter payoff is defined as the voter payoff if the politician exerted full effort in every period that $\theta = 1$—the highest expected payoff the voter can obtain with any strategy profile, ignoring all incentive constraints.

cells. However, in all other cases, the two payoffs will differ (and both be strictly less than the first-best payoff of μ). For instance, for $\mu = 0.3$, $\delta = 0.9$, and $c = 1.4$, the voter optimal equilibrium delivers 99.9% of the first-best equilibrium, whereas the best stationary equilibrium delivers only 64.3%. Similarly, for $\mu = 0.6$, $\delta = 0.7$, and $c = 1.9$, the voter optimal equilibrium delivers 51.1% of the first-best equilibrium, while the best stationary equilibrium delivers only 36.8%, an increase of 39%. Second, we note that the performance of the voter-optimal equilibrium, relative to the first-best payoff, is increasing in the patience δ; decreasing in the cost c; and decreasing in μ. It is straightforward to show analytically that these three results hold generally in our model. The least obvious of these three comparative statics results is that with respect to μ; intuitively, increasing μ raises the likelihood of tenuring a politician, and has no effect on the relationship between the politician's job security and effort conditional on a good shock.

References

Guo, Yingni and Johannes Hörner. 2020. "Dynamic Allocation without Money." *TSE Working Paper* .