

Replacement and Reputation*

Navin Kartik[†]

Elliot Lipnowski[‡]

Harry Pei[§]

December 15, 2025

Abstract

Does electoral replacement ensure that officeholders eventually act in voters' interests? We study a reputational model of accountability. Voters observe incumbents' performance and decide whether to replace them. Politicians may be "good" types who always exert effort or opportunists who may shirk. We find that good long-run outcomes are always attainable, though the mechanism and its robustness depend on economic conditions. In environments conducive to incentive provision, some equilibria feature sustained effort, yet others exhibit some long-run shirking. In the complementary case, opportunists are never fully disciplined, but selection dominates: every equilibrium eventually settles on a good politician, yielding permanent effort.

*We thank Richard Van Weelden for discussions that led to this project, as well as Nageeb Ali, Sandeep Baliga, Nina Bobkova, Alessandro Bonatti, Joyee Deb, Jeffrey Ely, Mehmet Ekmekci, Drew Fudenberg, Rohit Lamba, David Levine, Qingmin Liu, Alessandro Pavan, Carlo Prato, Doron Ravid, [Refine.ink](#), João Ramos, Larry Samuelson, Andrzej Skrzypacz, Bruno Strulovici, Mike Ting, Juuso Välimäki, Alexander Wolitzky, and various audiences for helpful comments. Maxim Ventura and Cole Wittbrodt provided excellent research assistance.

[†]Department of Economics, Yale University. Email: navin.kartik@yale.edu.

[‡]Department of Economics, Yale University. Email: elliot.lipnowski@yale.edu.

[§]Department of Economics, Northwestern University. Email: harrydp@northwestern.edu.

1. Introduction

In various delegated decision-making settings, a principal’s ultimate tool is the power to replace an agent. A firm’s board can fire a poorly performing CEO; patients can switch health-care providers after bad service; and ineffective bureaucrats or organizational leaders may be dismissed. But the most prominent context is democratic politics, where the primary instrument voters use to control officeholders is re-election. An influential body of work, starting with [Barro \(1973\)](#) and [Ferejohn \(1986\)](#), has studied the consequences of voters’ authority to replace incumbents. Such authority is essential not only to discipline politicians’ behavior, but also to select good politicians ([Fearon, 1999](#); [Besley, 2005](#)).

Indeed, much further back, [Mill \(1859\)](#) stressed the importance of the *long-run* implications of selection, writing that “The worth of a State, in the long run, is the worth of the individuals composing it.” Yet we know surprisingly little from existing formal analyses about the long-run consequences of replacement in environments with both adverse selection and moral hazard. Much of the theoretical political-accountability literature (elaborated subsequently) analyzes models with short horizons or with short term limits.¹ A fundamental question thus remains: absent other frictions, does the replacement mechanism ensure that eventually voters can identify good politicians, and/or that officeholders will act in voters’ interests?

This paper answers that question in a simple and stylized model that we view as a natural benchmark. Our model takes inspiration from not only prior work on accountability, but also the reputation literature (e.g., [Fudenberg and Levine, 1992](#)). While the model is applicable more broadly, for concreteness we filter it through a political lens.

A pool of long-lived politicians interacts with a sequence of short-lived (or myopic) voters. In each period, the current voter observes the incumbent’s past performance and decides whether to retain him or replace him with a fresh draw from the politician pool. The officeholder—old or new—then chooses, covertly, whether to exert effort (work or shirk), which maps stochastically into performance. Politicians are one of two types: a *good* type who always exerts effort, or an *opportunistic* type who values office but dislikes effort. Voters simply want effort;² but this means they instrumentally benefit from good types, and politicians have an incentive to build a reputation. Apart from the retention decision, voters

¹ This contrasts with an empirical literature on the considerable length of political careers. See, for example, structural estimates by [Diermeier, Keane, and Merlo \(2005\)](#) and a study on legislative effectiveness by [Padró-i Miquel and Snyder Jr. \(2006\)](#), both highlighting how performance and retention of officeholders evolve significantly over long tenures.

² They may also incur a cost of replacing incumbents, but we suppress that for this introduction.

have no recourse to any other incentive instruments such as transfers.

In this setting, we obtain a sharp but nuanced answer to the motivating question of long-run outcomes. On the one hand, there is always at least one equilibrium that eventually has sustained effort from incumbents.³ On the other hand, whether this good long-run outcome arises in *every* equilibrium depends on the economic environment. Intuitively, the environment is conducive to incentive provision if effort costs are low, monitoring (i.e., how performance reflects effort) is precise, and politicians are patient. Then there is an equilibrium in which even opportunistic incumbents exert effort in every period (even in the short run). But we show that in this case, there also exist equilibria in which, even in the long run, opportunistic politicians will periodically hold office and shirk. We identify a novel feature of all such equilibria that do not deliver good long-run outcomes: *replacement despite a favorable reputation*. That is, voters sometimes replace an incumbent even when his past performance indicates that he is more likely to be a good type than the average politician. Voters do so when they (correctly) anticipate that an opportunistic incumbent with the favorable reputation is sufficiently more likely to shirk than a newly-installed opportunist. The upshot is that in environments conducive to incentive provision, equilibrium selection—which can be viewed as political norms—crucially matter for outcomes, even in the long run.

Now consider the complementary case when parameters are not conducive to incentive provision: there is no equilibrium that sustains effort in every period from opportunists. Paradoxically perhaps, this turns out to guarantee effort in the long run. The mechanism operates through selection. We show that now, in every equilibrium, every opportunistic officeholder is eventually replaced, whereas eventually some good-type incumbent will be retained forever. The logic assuring long-run selection is subtle. To appreciate why, suppose good types are scarce. We establish that the value of a newly-installed politician (a voter's endogenous “outside option”) is then *arbitrarily low* across all equilibria, meaning that newly-installed officeholders almost certainly shirk. That makes voters reluctant to replace any incumbent. However, it turns out that opportunists are always driven to histories at which voters believe the incumbent is overwhelmingly likely to be an opportunist who will shirk—so much so that the low outside option is preferable. Conversely, every good type has a chance of reaching and then maintaining a high enough reputation to avoid replacement.

The dichotomy in the two previous paragraphs is summarized in our main result, Theorem

³ As elaborated in Section 2, our equilibrium notion is that of “personal symmetric” weak perfect Bayesian equilibrium. Roughly speaking, “personal” means that the interaction between every incumbent and the voters is self-contained, and “symmetric” means it is ex-ante identical across all the politicians.

rem 1: there is no equilibrium that has “full effort” (incumbents exert effort in every period) if and only if all equilibria have “eventual full effort” (after some point, incumbents exert effort in all subsequent periods). The result crystallizes a sharp tradeoff between the *possibility* of strong discipline (i.e., even in the short run) and the *guarantee* of either discipline or selection in the long run.

Beyond Electoral Accountability. Although we cast our analysis in terms of politicians and voters, our framework applies more broadly to other principal-agent settings with replacement. For example, consider the relationship between politicians and career bureaucrats. Earlier work has modeled bureaucrats as long-lived agents and politicians as short-lived principals subject to electoral turnover (Spiller and Urbiztondo, 1994; Alesina and Tabellini, 2007; Huber and Ting, 2021). Indeed, the horizon asymmetry in this relationship can be viewed as a consequence of politicians’ electoral accountability: the pressure to satisfy voters transmits a short horizon to politicians in their bureaucratic relationships. More broadly, ‘short-lived’ and ‘long-lived’ are relative, defined with respect to the counterparty in a given accountability relationship.

For this application, our results connect to two themes. First, they speak to the perils of bureaucratic politicization. In environments conducive to strong incentives, there can be bad equilibria characterized by persistent turnover and shirking. Such dynamics offer a cautionary perspective on anti-entrenchment efforts: bureaucratic replacement can become a self-fulfilling norm, destroying the reputational incentives that drive performance and resulting in a permanent cycle of mediocrity. This interpretation of our results is consistent with Akhtari, Moreira, and Trucco’s (2022) empirical evidence linking bureaucratic replacement and poor performance to politician turnover. By contrast, when incentives are limited, we show that selection is powerful in the long run. This aligns with Besley and Ghatak’s (2005) and Prendergast’s (2007) emphasis on the identification and retention of intrinsically motivated bureaucrats.

Related Literature. Our paper primarily contributes to and sits at the intersection of two literatures: electoral accountability and reputation formation.

In the accountability literature, the idea that replacement is an instrument for selection, rather than just controlling moral hazard, was popularized by Fearon (1999).⁴ Much of the literature studying selection, surveyed by Ashworth (2012) and Duggan and Martinelli (2017), uses models with either just two periods or a two-term limit; exceptions include Banks

⁴ We will not attempt to discuss the large literature on accountability with just moral hazard; see Duggan and Martinelli (2017) for a survey.

and Sundaram (1993), Duggan (2000), Bernhardt, Dubey, and Hughson (2004), and Schwabe (2009). None of these papers, nor virtually any others, have results on selection leading to voter-optimal outcomes in the long run—in any equilibrium, much less all equilibria, as we find under some conditions. The only exception we are aware of is Anesi and Buisseret (2022).⁵ They construct one equilibrium of their model that is approximately first best when each politician’s type follows an irreducible Markov chain and all agents—politicians and the voter—are almost fully patient. By contrast, we have fully-persistent politician types with fixed discounting, and short-lived (or impatient) voters. Anesi and Buisseret’s (2022) construction relies on patient voters using sophisticated punishment schemes that are themselves supported by “grim-trigger” strategies; in our setting, it is instead the very inability to discipline opportunists that guarantees eventual positive selection.

When our environment is conducive to incentive provision, our construction of an equilibrium in which incumbents shirk even in the long run uncovers a new mechanism for inefficiency. It owes to the dynamics of opportunistic incumbents’ effort over their careers, and the resulting replacement despite a favorable reputation described earlier.⁶ In Section 3.3 we explain why the mechanism is distinct from that of the inefficient equilibrium of Myerson (2006), which owes to large replacement costs. The coexistence of good and bad equilibria in our incentive-conducive environments may also be reminiscent of folk theorems in repeated complete-information games (e.g., Fudenberg and Maskin, 1986). However, we have short-lived voters who may learn about politician types (so there is incomplete information), which makes the reputation literature discussed next—with quite different prior results—the more relevant comparison.

Most models in the reputation literature have a *fixed* long-lived player interacting with

⁵ In their extension that incorporates adverse selection, Acharya, Lipnowski, and Ramos (2025) study a model with an opportunistic type and a bad type who always shirks. They show that a voter-optimal equilibrium weeds bad politicians out from office; yet, because opportunistic politicians still face moral hazard, equilibria need not yield sustained effort in the long run. We discuss adding bad types to our model in Section 4. Another paper concerned with long-run outcomes is Duggan and Forand (2025); they assume that an incumbent’s type is publicly revealed after he assumes office, similar to one section of Kartik and Van Weelden (2019) that does not have term limits but assumes types are revealed after two periods in office. Such exogenous type revelation limits (in the latter paper) or even eliminates (in the former) the reputation-building horizon. Kartik and Van Weelden (2019) focus on how their short-horizon reputation-building affects incumbency advantage. Duggan and Forand (2025) focus on how incumbents can use current policies to manipulate future states and policies, and how patient voters may nevertheless attain good outcomes.

⁶ One can contrast this equilibrium feature with the literature on managerial turnover with dynamic contracting; see, for example, Garrett and Pavan (2012) and references therein. Papers in that literature characterize ex-ante optimal retention and compensation policies for a long-lived principal with commitment power. Optimal contracts often exhibit “entrenchment” in the sense that retention standards become more lenient over time to reduce information rents. Such entrenchment also features in certain analyses of accountability, like Acharya, Lipnowski, and Ramos (2025).

a sequence of short-lived players. [Fudenberg and Levine \(1989, 1992\)](#) derive tight bounds for the long-lived player’s equilibrium payoffs as he becomes fully patient. [Cripps, Mailath, and Samuelson \(2004\)](#) show that when monitoring is imperfect, although the [Fudenberg and Levine \(1992\)](#) bounds apply *ex ante*, the long-lived player will eventually lose his reputation in the sense that his opponents will learn his type.⁷ By contrast, our model has multiple long-lived players, and our focus is on their endogenous replacement. The possibility of replacement implies that the short-lived players may not learn any given long-lived player’s type, and because the replacement probabilities are endogenous, the long-lived players’ effective discount factors are endogenous and evolving, possibly bounded away from full patience. We instead focus on outcomes from the short-lived players’ vantage.

Our replacement structure is reminiscent of [Fudenberg, Levine, and Tirole \(1987\)](#), who study a Coasian bargaining model. Similar to our paper, their seller has an “outside option”, whose value is determined endogenously, to restart the interaction afresh with a new buyer. Their analysis concerns trading dynamics that end with a sale; *inter alia*, they show how the switching option allows the seller to effectively commit to a price in some equilibria. By contrast, we study an infinitely repeated reputation model with moral hazard. We analyze how replacement motivates effort and sorts types, with a focus on characterizing when these forces assure good outcomes in the long run.

There is also prior work on reputation models with competition between long-lived players, notably [Hörner \(2002\)](#) and, more recently, [Deb and Fanning \(2025\)](#). Besides other differences (e.g., in the behavior of committed types and, in the former paper, endogenous pricing), an important distinction is the nature of their short-lived players’ outside option. In these competitive models, the consumers’ outside option is the equilibrium value of switching to a rival firm (whether a previously established firm or a new entrant), which evolves with the history and produces alternative forces to simply restarting afresh.⁸

Finally, ongoing work by [Ali \(2025\)](#) studies a model related to ours. He has only two politicians (among whom voters may rotate), and studies a class of equilibria with a “pure Markovian” structure. While [Ali \(2025\)](#) also finds conditions under which all his equilibria have eventual full effort, the mechanism does not operate through selection; to the contrary, a full-effort equilibrium exists under his conditions. [Section 4](#) elaborates.

⁷ A number of subsequent papers propose models in which reputation can be sustained in the long run. The mechanisms include competition between multiple long-lived players (e.g., [Hörner, 2002](#); [Deb and Fanning, 2025](#)), the long-lived player’s type changing over time (e.g., [Phelan, 2006](#); [Wiseman, 2008](#); [Ekmekci, Gossner, and Wilson, 2012](#)), and the short-lived players having limited memories (e.g., [Liu and Skrzypacz, 2014](#); [Pei, 2024](#)).

⁸ Naturally, considerations of when to restart afresh are not present when the long-lived player exits according to some exogenous process (e.g., [Mailath and Samuelson, 2001](#); [Tadelis, 2002](#)).

2. Model

Setting. We consider an infinite-horizon game in discrete time, with periods indexed by $t \in \mathbb{Z}_{\geq 0} \equiv \{0, 1, 2, \dots\}$. There is a pool of ex-ante identical long-lived players, whom we refer to as *politicians*. Each period, only one long-lived player, whom we call the *incumbent*, is active, or synonymously, holds office. An incumbent remains active until he is replaced, after which he is never active again. There is also an infinite sequence of short-lived players, whom we call *voters*. The voters arrive one in each period, and each voter plays the game only in the period they arrive.

In each period other than the initial period, the current voter decides whether to *keep* the previous period's incumbent or *replace* him with a new politician. In the initial period 0 (there being no prior incumbent), the voter makes no choice. (We will typically suppress the qualifier “current” hereafter.) Then, the incumbent chooses an action $a \in A := \{0, 1\}$. We refer to $a = 0$ as *shirking* and $a = 1$ as, interchangeably, *working* or *exerting effort*. This action stochastically generates a public signal s drawn from a finite set S according to a monitoring structure $f : A \rightarrow \Delta S$. Let $f_a(s)$ denote the probability of signal s conditional on action $a \in A$, and let $f_{\tilde{a}}(s) := \tilde{a}f_1(s) + (1 - \tilde{a})f_0(s)$ denote the aggregate probability of signal s if the incumbent exerts effort with probability $\tilde{a} \in [0, 1]$.

When a politician is active, his stage-game payoff from action a is $u_P(a)$. Once a politician is replaced, he is never again active and receives a fixed payoff, which we normalize to 0. The voter's payoff from action a is $u_V(a)$, and they further incur an additive cost $c \geq 0$ if they replace the incumbent.⁹ We assume that the voter's preferred action is different from the incumbent politician's (myopically) favorite action, and that a politician's least favorite outcome is to be replaced. Without further loss of generality, we normalize payoffs so that $u_V(a) = a$ and $u_P(a) = 1 - \kappa a$ for some $\kappa \in (0, 1)$. We also assume the monitoring structure is informative, that is, $f_0(\cdot) \neq f_1(\cdot)$. Finally, assume no signal perfectly reveals the politician to be shirking, that is, $f_1(s) > 0$ for every $s \in S$.

Following Besley (2006) and Besley and Smart (2007), among others, we take seriously that some politicians are intrinsically motivated. With probability $\pi_0 \in (0, 1)$, a given politician is a *good type* who always plays $a = 1$ when active. With complementary probability $1 - \pi_0$, he is an *opportunistic type* who maximizes the expectation of his discounted average payoff $(1 - \delta) \sum_{t=\tau_0}^{\tau_1-1} \delta^{t-\tau_0} u_P(a_t)$, where a_t is his period- t action, τ_0 is the period that he becomes active, τ_1 is the period that he is replaced (which we label as ∞ if he is never replaced),

⁹ We allow for a positive replacement cost to clarify that our analysis is not driven by excessive voter indifference, e.g., being indifferent to replacement when they expect the same incumbent behavior after retention or replacement. We will be interested in the case where the replacement cost is not too large.

and $\delta \in (0, 1)$ is his discount factor. Each politician's type is his private information, and types are independent. Hence, after a voter replaces an incumbent, a newly drawn politician is the good type with probability π_0 . We assume that the replacement cost (which can be zero) is small relative to this uncertainty: $c < \min\{\pi_0, 1 - \pi_0\}$.

Interpretations. An incumbent's choice of effort can be interpreted broadly to capture various familiar conflicts of interest in political agency. Working ($a = 1$) can correspond to costly activities that benefit the electorate, such as the effective provision of public goods. Conversely, shirking ($a = 0$) can represent rent-seeking behavior, such as engaging in corruption or implementing inefficient pork-barrel projects that benefit special interests.

We model voters as short lived for two reasons. First, it is a way to capture an electorate composed of many long-lived but uncoordinated voters. Rather than a mass electorate coordinating on complex strategies guided by forward-looking considerations, we take the view that each voter simply votes for the alternative that maximizes their payoff in the current period. Such myopic behavior—which can also be viewed as an equilibrium refinement with many voters, none of whom is pivotal—is mathematically equivalent to a sequence of short-lived voters. Second, we are ruling out sophisticated forward-looking coordination schemes between politicians and voters. Ultimately, as in the canonical reputation literature, our assumption sharpens the focus on how opportunistic politicians' career concerns, rather than voters' long-term considerations, shapes outcomes. The assumption is also appropriate for other settings in which principals are transient, such as the relationship between elected officials and the bureaucracy discussed in the [Introduction](#). We return to this issue in [Section 4](#).

Strategies and equilibria. We will focus on equilibria in which players' behavior depends only on the current officeholder's past outcomes. To that end, the *(public) career history* of a given incumbent records the sequence of public signals he has generated in the past. Let $\mathcal{H} := S^{<\infty} = \bigsqcup_{t=0}^{\infty} S^t$ denote the set of career histories.¹⁰ Note that $\mathcal{H} \setminus \{\emptyset\} = \mathcal{H} \times S$. Our solution concept entails a politician strategy $\sigma_P : \mathcal{H} \rightarrow [0, 1]$, a voter strategy $\sigma_V : \mathcal{H} \setminus \{\emptyset\} \rightarrow [0, 1]$, and a voter belief map $\pi : \mathcal{H} \rightarrow [0, 1]$. Here, $\sigma_P(h)$ denotes the probability that an opportunistic incumbent works when his career history is h ; $\sigma_V(h)$ denotes the probability the current voter replaces an officeholder whose career history is $h \neq \emptyset$; and $\pi(h)$ is the probability voters assign to a politician being the good type (his *reputation*) when his career history is h . Given such a triple, it is useful to record the voter-expected incumbent efforts and the opportunistic incumbents' continuation values. Formally, suppressing their

¹⁰The full public history—the full history available to voters—consists of an incumbent career history, all past replacement decisions, and the career histories associated with all previously replaced politicians.

dependence on $(\sigma_P, \sigma_V, \pi)$ for brevity, we define $\alpha : \mathcal{H} \rightarrow [0, 1]$ via

$$\alpha(h) := \pi(h) + [1 - \pi(h)] \sigma_P(h),$$

and $v : \mathcal{H} \rightarrow [0, 1]$ as the unique solution to the recursive equation

$$v(h) = (1 - \delta) [1 - \kappa \sigma_P(h)] + \delta \sum_{s \in S} f_{\sigma_P(h)}(s) [1 - \sigma_V(h, s)] v(h, s), \quad \forall h \in \mathcal{H}.$$

We call a triple $(\sigma_P, \sigma_V, \pi)$ a **personal symmetric weak perfect Bayesian equilibrium** (hereafter **equilibrium**) if it satisfies the following three properties:

1. Any $h \in \mathcal{H}$ has

$$\sigma_P(h) \in \arg \max_{a \in [0, 1]} \left\{ (1 - \delta)(1 - \kappa a) + \delta \sum_{s \in S} f_a(s) [1 - \sigma_V(h, s)] v(h, s) \right\}.$$

2. Any $h \in \mathcal{H} \setminus \{\emptyset\}$ has $\sigma_V(h) \in \arg \max_{x \in [0, 1]} \{(1 - x)\alpha(h) + x [\alpha(\emptyset) - c]\}$.
3. We have $\pi(\emptyset) = \pi_0$, and any $h \in \mathcal{H}$ and $s \in S$ with either $h = \emptyset$ or $\sigma_V(h) < 1$ has

$$f_{\alpha(h)}(s) \pi(h, s) = \pi(h) f_1(s).$$

The first equilibrium condition is a recursive formulation of sequential rationality for incumbents. The second condition is for voters: being short-lived, they make replacement decisions that maximize the current effort they expect from an officeholder net of replacement cost. The third condition incorporates voters' Bayesian updating: it requires that (i) any newly-installed incumbent has reputation π_0 , the prior probability of a good type; and (ii) reputation at any career history (h, s) is derived from Bayes' rule applied to h so long as either $h = \emptyset$ (so s is the first signal generated by the incumbent) or the incumbent has been retained with positive probability at h .¹¹

Following Banks and Sundaram's (1993) terminology, the "personal" nomenclature in our equilibrium concept reflects that the interaction between any officeholder and the voters depends only on the sequence of signals generated by that politician—not, for example, on the calendar time at which the politician took office nor on what transpired prior to when he

¹¹ We refer to our solution concept as "weak" PBE because beliefs can be arbitrary at certain non-initial career histories that are off path—specifically, at any (h, s) with $h \neq \emptyset$ and $\sigma_V(h) = 1$. It is still stronger than textbook definitions of weak PBE that impose no belief restrictions off path (Mas-Colell, Whinston, and Green, 1995, Definition 9.C.3, p. 285). Section 3.1 elaborates on our concept vs. "full" PBE.

took office.¹² The “symmetric” adjective reflects that all opportunistic officeholders respond to their career history in the same way and are treated identically by voters. Combined, these two properties reflect the premise that the role of replacement is to start the interaction with a new politician afresh; put differently, any equilibrium has a single (endogenously-determined) “outside option”—a voter’s value from replacement. This approach is standard in the political-economy literature; besides [Banks and Sundaram \(1993\)](#), see, for example, [Duggan \(2000\)](#) or [Van Weelden \(2013\)](#). It is also tantamount to [Fudenberg, Levine, and Tirole’s \(1987\)](#) “stationarity” restriction. As we shall see, it does not preclude rich dynamics within any incumbent’s career.

We highlight that our setting is *not* a “trust game” in which a principal’s refusal to trust effectively ends the stage-game interaction. That game would trivially have a “no trust, no effort” equilibrium. By contrast, replacement in our model installs a new incumbent who makes an active effort choice. There is no “always replace, never work” equilibrium: the presence of good types implies that if opportunists were never to work, there would be learning and a voter would never replace an incumbent at a history with favorable reputation (i.e., at h such that $\pi(h) > \pi_0$).

As a preliminary step, we ensure—non-constructively—equilibrium existence despite our refinements.

Proposition 1. *An equilibrium exists.*

The proof in [Appendix A.1](#) constructs an auxiliary game whose perfect Bayesian equilibria (PBE) are equilibria of our game. The auxiliary game has only one politician. It mirrors the interaction in our game between a given politician and the voters, modifying it in two ways. First, the politician’s first-period choice is $\tilde{a} \in [0, 1]$, representing the probability with which he chooses action 1; in all other periods, the politician makes a choice from $A \equiv \{0, 1\}$. While the politician can mix in subsequent periods, we require him to not mix (over the action set $[0, 1]$) at the initial history. Second, the voter’s payoff when they replace the politician is $\pi_0 + \tilde{a}(1 - \pi_0) - c$. Intuitively, our specification of the voter’s payoff upon firing the politician captures the symmetry requirement in the original game. We require the politician not to mix in the first period to avoid voters updating beliefs as play progresses about their payoff from replacement (contrary to symmetry in the original game). Making the politician’s first-period choice continuous secures the existence of such PBE in the auxiliary game, following standard arguments à la [Kreps and Wilson \(1982\)](#) and [Fudenberg and Levine \(1983\)](#).

¹² We have also imposed that the politician conditions only on his own *public* career history, not on his own past actions. This restriction is for notational simplicity: Because the game has “product monitoring”, standard arguments show this feature is without loss in terms of outcome equivalence.

3. Results

Our goal is to study whether replacement ensures that incumbents exert effort. To that end, we introduce two criteria. Using bold font for random variables, let $\mathbf{a}_t \in \{0, 1\}$ indicate the incumbent's chosen effort in period t , which will in particular be equal to 1 if the incumbent is the good type. We say that an equilibrium **attains full effort** if $\mathbf{a}_t = 1$ almost surely for every $t \in \mathbb{Z}_{\geq 0}$, and it **attains eventual full effort** if there almost surely exists some $\tau \in \mathbb{Z}_{\geq 0}$ such that $\mathbf{a}_t = 1$ for all $t \geq \tau$. In words, an equilibrium attains full effort (FE) if incumbents always work (on the path of play), and it attains eventual full effort (EFE) if, eventually, incumbents work in every period.¹³ The distinction between FE and EFE is thus one of timing; an equilibrium with FE also has EFE, but the converse is not generally true. In addition, one can also draw a distinction between whether, for any given parameters, either property holds in *no* equilibrium, *some* equilibrium, or *all* equilibria.

A strong desideratum is that all equilibria attain FE; conversely, it would be troubling if no equilibrium attains even EFE. We will see that neither of these is possible, no matter the parameters.¹⁴ Instead, our main result is that parameters divide sharply into two cases: either there is an equilibrium that attains FE, but then there is also an equilibrium that fails even EFE; or there is no equilibrium that attains FE, but then all equilibria attain EFE. The following table summarizes.

	Effort always (FE)	Effort eventually (EFE)
Condition FEI	<i>Some</i> equilibrium	<i>Some</i> equilibrium
\neg Condition FEI	<i>No</i> equilibrium	<i>All</i> equilibria

Table 1: The equilibrium possibilities. Full Effort Incentives (FEI) is a condition on parameters.

The table's **Condition FEI** on parameters is the following:

Condition FEI (Full Effort Incentives). *Some* $v : S \rightarrow \mathbb{R}_{\geq 0}$ exists such that

$$(1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s)v(s) \geq (1 - \delta)1 + \delta \sum_{s \in S} f_0(s)v(s) \quad (\text{IC}_{\text{FEI}})$$

¹³ One might be interested in weaker variants of both FE and EFE. Not only would our positive results obviously extend, but, as elaborated in Subsections 3.2 and 3.3, so do our negative results.

¹⁴ One can show that absent good types, no equilibrium attains EFE when **Condition FEI** below fails. Even with a good type, no equilibrium attains EFE in the canonical noisy-monitoring product-choice game between a long-lived firm and short-lived consumers when there is no possibility of replacement (Cripps, Mailath, and Samuelson, 2004). In a different model, Pei (2024) presents a result in the vein of all equilibria attaining FE.

and

$$(1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s)v(s) \geq \max_{s \in S} v(s). \quad (\text{PK}_{\text{FEI}})$$

[Condition FEI](#) is independent of politicians' initial reputation π_0 and voters' cost of replacement c because the condition stems from contemplating an auxiliary game without good types and without a replacement cost. In the auxiliary game, if all incumbents always exert effort, then voter incentives are trivially satisfied. Standard recursive arguments then imply that under some conditions, there are equilibria in which appropriate voter behavior can incentivize incumbents to always exert effort. The associated incentive-compatibility and promise-keeping constraints are precisely those in [Condition FEI](#), inequalities [\(IC_{FEI}\)](#) and [\(PK_{FEI}\)](#) respectively. [Condition FEI](#) is satisfied if and only if the politicians' discount factor δ is high enough, their effort cost κ is small enough (recall we have normalized their benefits of holding office to 1), and the monitoring structure f is sufficiently informative.¹⁵ Parameters that satisfy [Condition FEI](#) thus correspond to environments that are conducive to incentive provision.

3.1. Main Result

We now present our main result.

Theorem 1. *The following are equivalent:*

1. *All equilibria attain eventual full effort;*
2. *No equilibrium attains full effort;*
3. [Condition FEI](#) fails.

An implication is that no matter the parameters, there is always at least one equilibrium that attains EFE, but also at least one equilibrium without FE.

We view the main economic lesson from the theorem as the sharp contrast summarized in [Table 1](#): environments conducive to incentive provision yield the *possibility* of FE; those

¹⁵ We emphasize that [Condition FEI](#) is a joint condition on the parameters (δ, κ, f) . So, for instance, the required discount factor is higher when the effort cost is higher or the monitoring structure is less informative in the [Blackwell \(1953\)](#) sense. As an illustration, with binary signals of precision $p \in (1/2, 1)$ (that is, $S = \{0, 1\}$ and $f_a(a) = p$), it can be verified using [Lemma 12](#)—which shows that, no matter the monitoring structure, there no loss in restricting to $v(s) \in \{0, \bar{v}\}$ for some $\bar{v} > 0$ —that [Condition FEI](#) reduces to

$$\delta \geq \frac{\kappa}{p - (1 - p)(1 - \kappa)},$$

where the right-hand side is decreasing in p and increasing in κ .

that are not conducive *guarantee* EFE. More precisely, when **Condition FEI** holds, although there are equilibria that attain FE, there are also equilibria that fail EFE. By contrast, when **Condition FEI** fails, there is no FE equilibrium, but perpetual effort is inevitable after finite time in every equilibrium. As we will elaborate, EFE is guaranteed through (eventual) selection of a good type; paradoxically, when **Condition FEI** holds, the very possibility of providing strong incentives undermines assuring good selection or assuring good long-run outcomes. To our knowledge, this is a novel tradeoff between discipline and selection.

The proof of [Theorem 1](#) is in [Appendix A.2](#). The remainder of this section discusses the intuition and forces underlying the theorem. But before that, we comment on an aspect of using (personal, symmetric) *weak* PBE as our solution concept. Condition (3) of the equilibrium definition allows for arbitrary beliefs off path immediately after a voter retains an incumbent who should have been certainly replaced. Dropping the condition's qualification that $\sigma_V(h) < 1$ would impose belief requirements off path that yield (personal, symmetric) PBE, without the “weak” moniker. [Appendix A.1](#) assures existence for this solution concept as well. Our equilibrium constructions under **Condition FEI** benefit from the off-path latitude allowed by weak PBE, as flagged in [Appendix A.2](#). At the same time, the weaker solution concept strengthens what we view as the most important conclusion of [Theorem 1](#), namely that all equilibria attain EFE when **Condition FEI** fails.

3.2. Condition FEI and Full Effort

The equivalence between **Condition FEI** and the existence of FE equilibria stems from there being no learning in an FE equilibrium. So there is an FE equilibrium if and only if there is one in an auxiliary game without commitment types. It turns out that for a small enough replacement cost, it is further equivalent to consider the auxiliary game without a replacement cost. As mentioned earlier, **Condition FEI** characterizes when there are FE equilibria in that auxiliary game; the logic follows standard recursive reasoning ([Abreu, Pearce, and Stacchetti, 1990](#)). In fact, when **Condition FEI** fails and replacement costs are sufficiently small, there is a positive lower bound (across all equilibria) on the probability with which opportunistic incumbents shirk in their first period in office; see [Corollary 2](#) in the Appendix. This implies that when **Condition FEI** fails, there is some $\varepsilon > 0$ such that no equilibrium has $\mathbb{P}(\mathbf{a}_t = 1) \geq 1 - \varepsilon$ for all $t \in \mathbb{Z}_{\geq 0}$; indeed, the inequality already fails for $t = 0$.

3.3. What Prevents Eventual Full Effort?

Our argument that there is an equilibrium without EFE when **Condition FEI** holds is constructive. Notice that when replacement and effort costs are small, and monitoring is very

precise, a newly-installed opportunistic incumbent has to work with positive probability: if not, a voter would retain an incumbent if he has a favorable reputation (i.e., a posterior π_t strictly above the prior π_0), which would incentivize newly installed incumbents to work toward such a reputation. At least when the replacement cost is zero, a newly-installed incumbent cannot work with certainty either: otherwise, voter incentives imply that any retained incumbent must be working, which implies EFE.

With that in mind, let us describe the non-EFE equilibrium we construct. For simplicity, suppose there are only two signals: a “Pass” signal and a “Fail” one, with the former having the higher likelihood ratio $f_1(s)/f_0(s)$. Opportunistic incumbents work with interior probability in their first period in office. Thereafter, if they have always generated a Fail signal, they are replaced with an interior probability and, if retained, work with an interior probability.¹⁶ At any history with only Pass signals, the incumbent is retained and works. At histories where a Pass signal is followed by a Fail signal, the incumbent is replaced. See [Figure 1](#) for an automaton depiction of the strategy profile. We defer to the [Appendix](#) the explanation of why this profile constitutes an equilibrium. Crucially, however, this equilibrium has the features that every incumbent is eventually replaced, and incumbents in their first period shirk with positive probability. Hence, the equilibrium does not attain EFE; in fact, almost surely there is shirking infinitely often, and average voter welfare (even gross of replacement costs) is below first best:

$$\limsup_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbf{a}_t \right] < 1.$$

The mechanism underlying the above equilibrium is a novel rationale for undesirable long-run outcomes in accountability models. In particular, it ought to be contrasted with [Myerson \(2006\)](#). For some parameters of his model, his Theorem 1 exhibits equilibria in which opportunistic incumbents always shirk.¹⁷ Yet incumbents are never replaced (and hence EFE does not obtain) because the cost of replacement is large relative to the probability

¹⁶ Since reputation is deceasing in the number of Fail signals, the opportunistic type’s probability of work increases to keep voter incentives fixed. Meanwhile, the voter’s replacement probability is constant to keep politician incentives fixed. Although the underlying forces are quite different, the idea that opportunistic incumbents misbehave less over time is reminiscent of the literature on the repeated prisoner’s dilemma with voluntary separation (e.g., [Ghosh and Ray, 1996](#); [Eeckhout, 2006](#)) and the “starting-small” literature (e.g., [Watson, 1999](#)).

¹⁷ [Myerson’s \(2006\)](#) model for unitary democracy is like ours, except that he assumes a single long-lived and forward-looking voter who can perfectly observe the incumbents’ effort choices, and he allows for large replacement costs. But the relevant part of his Theorem 1 also holds with short-lived voters who can only observe noisy signals, as in our model.

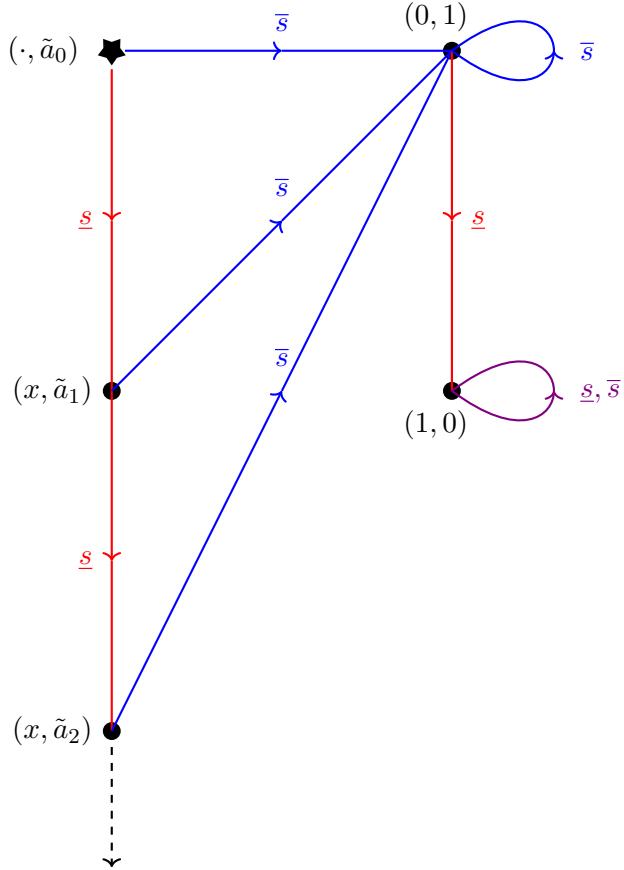


Figure 1: Under Condition **FEI**, an equilibrium that does not attain eventual full effort when, for simplicity, there are only two signals, \underline{s} and \bar{s} , with $f_1(\underline{s})/f_0(\underline{s}) < f_1(\bar{s})/f_0(\bar{s})$. Each node (black star or dot) represents an automaton state, with the initial state—corresponding to a new incumbent—indicated by the star. The vector at each node shows the corresponding strategy profile: the first entry is the probability that the incumbent is replaced (not defined at the initial state), and the second entry is the probability that an opportunistic incumbent exerts effort if retained. Each edge represents the state transition following the corresponding signal. The equilibrium satisfies $0 < \tilde{a}_0 < \tilde{a}_1 < \dots < 1$ and $x \in (0, 1]$.

of a good type. We are interested, instead, in small replacement costs. In our non-EFE equilibrium, voters replace incumbents at some histories despite a favorable reputation (i.e., when the posterior π_t strictly exceeds the prior π_0) because of the rational expectation that an opportunistic incumbent at such a history will shirk.

In fact, this phenomenon of replacing an incumbent despite a favorable reputation is a necessary feature of any equilibrium that does not attain EFE. Formally, we have:

Proposition 2. *Any equilibrium that does not attain eventual full effort entails a positive probability of an incumbent being replaced with a reputation greater than the prior.*

An intuition for the argument is as follows. As we will see shortly (Proposition 3, part 1), any opportunistic incumbent is eventually replaced. Hence, if EFE fails, then any good incumbent must also eventually be replaced—otherwise, eventually some good incumbent will stay in office forever, yielding EFE. Moreover, at least absent a replacement cost, there must be some learning in an incumbent’s first period.¹⁸ It follows that incumbents must be replaced at some histories with favorable reputation.

3.4. What Guarantees Eventual Full Effort?

We now turn to why a failure of Condition FEI implies that all equilibria attain EFE. The key is that precisely because the environment is not conducive to incentive provision, selection is now ineluctable. The following result, a key input to Theorem 1, formalizes the point.

Proposition 3. *Consider any equilibrium.*

1. *Every opportunistic officeholder is eventually replaced.*
2. *If Condition FEI fails, then every good officeholder has a positive probability of being retained forever.*

Hence, if Condition FEI fails, then some good officeholder is retained forever.

The last statement of Proposition 3 follows from the first two parts and an application of the Borel-Cantelli lemma.

The logic for the first part of the proposition is reminiscent of Cripps, Mailath, and Samuelson (2004). Here is the intuition. Since any incumbent’s reputation is a martingale, it converges almost surely. If an opportunistic incumbent is not eventually replaced, then

¹⁸ Otherwise, a new incumbent must be working, which implies FE when there is no replacement cost.

because signals statistically identify effort, eventually either (i) he must be exerting effort—behaving like the good type, so that belief updating stops—or (ii) his type must be revealed. Case (i) contradicts him not being eventually replaced, because it is only the possibility of replacement that provides effort incentives. The same logic applies if we are in case (ii) and yet the incumbent eventually still exerts effort with positive probability. Finally, if we are in case (ii) and the incumbent is certain to shirk, then it is optimal for voters to replace him.

Now consider the second part of [Proposition 3](#). When [Condition FEI](#) fails, full-effort incentives cannot be sustained; in particular, as noted earlier, a newly installed incumbent shirks with positive probability. Thus, once an incumbent’s reputation is sufficiently high, a voter strictly prefers retention to replacement, since replacement exposes her to some chance of shirking. Put differently, there is a reputation threshold above which the voter never replaces. The key step is that these high no-replacement beliefs are reached with positive probability. In fact, we show that reputation can get arbitrarily high. The logic is that if the supremum of on-path beliefs were interior, then at beliefs close to that supremum an opportunistic incumbent would have to exert effort with high probability—otherwise some signal would push the posterior above the supposed supremum, contradicting that being a supremum. But such effort is intuitively at odds with the failure of [Condition FEI](#),¹⁹ and so the supremum must equal one. Now, once beliefs enter the high region where there is no replacement, martingale arguments imply that there is a positive probability of never leaving that region. Combined with the first part of the proposition, it follows that every good-type officeholder is retained forever with positive probability when [Condition FEI](#) fails.

[Proposition 3](#) implies that when [Condition FEI](#) fails, every equilibrium asymptotically yields voters their first best welfare, even net of replacement cost. More precisely:

Corollary 1. *Suppose [Condition FEI](#) fails. In any equilibrium, the period- t voter’s expected payoff converges to first best as $t \rightarrow \infty$.*

We note that the selection logic of [Proposition 3](#) holds even if the good type were to exert effort with probability less than 1. The welfare conclusion of [Corollary 1](#) would, of course, be attenuated as that probability decreases.

Although [Proposition 3](#) shows that selection is extremely powerful in the long run in every equilibrium (when [Condition FEI](#) fails), it prompts a question about the short run. To assure that opportunists are eventually replaced, what makes replacement attractive to voters in the first place—especially when good types may be rare? One might guess that the voters’

¹⁹ This is only an intuition because [Condition FEI](#) does not actually preclude the possibility of effort with certainty at some histories. The formal argument is substantially more involved.

endogenous **outside option**—their equilibrium payoff from replacing an incumbent, gross of replacement cost—must be high for selection to operate effectively.²⁰ In fact, the following result says that the exact opposite is true when good types are scarce and replacement costs are small.

Proposition 4. *Suppose Condition FEI fails. For any $\varepsilon > 0$, there exist $\bar{c} > 0$ and $\bar{\pi}_0 \in (0, 1)$ such that, if $c \leq \bar{c}$ and $\pi_0 \leq \bar{\pi}_0$, then every equilibrium has voters' outside option below ε .*²¹

In other words, assuming Condition FEI fails, the combination of scarce good types and low replacement costs implies that voters' outside option is arbitrarily low—or equivalently, newly-installed opportunists are almost certainly shirking. So, replacement can look arbitrarily unattractive in the short run, despite culminating powerfully in the long run. While we defer the logic of Proposition 4 to the Appendix, why does a low outside option not undermine eventual selection? It is because opportunists are assured to reach histories at which voters believe the incumbent is overwhelmingly likely to be an opportunist who will shirk—so much so that even the low outside option is attractive enough to trigger replacement, given small replacement costs.

4. Discussion

This section discusses some extensions of our model.

Bad types. The possibility that some politicians are good types—intrinsically motivated to exert effort in office—is essential to our results. If all politicians are opportunistic, then although Condition FEI still characterizes when full effort can be attained, there is always an equilibrium (regardless of parameters) in which incumbents always shirk.

What if we allow for the coexistence of good types with *bad* types who always shirk in office? In addition to suitably adapting our maintained assumption of replacement costs being small enough, Condition FEI must also be strengthened for the existence of an equilibrium in which opportunistic incumbents always exert effort. The reason is that now some sequences of signals will convince voters that the incumbent is likely a bad type and must be replaced, which makes it more difficult to sustain effort from the opportunistic type along such sequences. More interestingly however, when Condition FEI fails it remains true that

²⁰ More precisely, we define the outside option as the probability of effort from a newly-installed incumbent, which can be written as $\mathbb{E}[\mathbf{a}_0] = \alpha(\emptyset) = \pi_0 + (1 - \pi_0)\sigma_P(\emptyset)$.

²¹ Recall that Condition FEI is a condition only on (κ, δ, f) , not on (c, π_0) . We show that for any (κ, δ, f) violating Condition FEI, an appropriate pair $(\bar{c}, \bar{\pi}_0)$ can be found to deliver the conclusion.

all equilibria attain eventual full effort. The mechanism still operates via selection of the good type, with an argument analogous to that of [Proposition 3](#). In particular, following the logic of [Cripps, Mailath, and Samuelson \(2004\)](#), not only will every opportunistic officeholder eventually be replaced, but so too every bad type.

Exogenous turnover. We have assumed that politicians remain eligible for office indefinitely. Plainly, this is essential for [Proposition 3](#)'s conclusion that when [Condition FEI](#) fails, some good type will eventually stay in office forever.

Suppose now that incumbents exit office exogenously with probability $1 - \rho$ in each period (e.g., due to retirement or death), for some survival probability $\rho \in (0, 1)$. Beyond implying that every incumbent leaves office at some time almost surely, politicians' discount factors are also effectively reduced to $\delta\rho$. Naturally, then, the relevant version of [Condition FEI](#) now replaces δ with $\delta\rho$ in [\(IC_{FEI}\)](#) and [\(PK_{FEI}\)](#). Our full-effort equilibrium construction remains valid when this version of [Condition FEI](#) holds. However, when this version of [Condition FEI](#) fails, eventual full effort cannot be attained in any equilibrium—the eventual exit of every incumbent creates a friction even after adjusting for the lower effective discount factor. Nevertheless, our conclusion regarding the power of selection is robust in the following sense. Holding the effective discount factor $\delta\rho$ fixed—so, in particular, not altering the aforementioned version of [Condition FEI](#)—we can ask what happens as $\rho \rightarrow 1$. Because voters are short-lived and politician incentives depend only on the effective discount factor, the set of equilibrium strategies does not change. Thus, the equilibrium distributions of the “search time” for a good type to take office and never be replaced (barring exogenous exit) is unchanging. As the survival probability $\rho \rightarrow 1$, in every equilibrium the system spends a fraction of time approaching 1 in the hands of a good type, and the long-run average effort probability converges to 1.

Recalling incumbents. We have assumed that an incumbent who is replaced can never hold office again. This “no-recall” assumption is commonplace in the literature, dating back to [Banks and Sundaram \(1993\)](#); it was also discussed by [Ferejohn \(1986\)](#) in his pure moral-hazard setting. The assumption substantially simplifies tracking the voters' outside option by making it history independent (given the focus on [personal equilibria](#)). Even if recall is allowed, [Condition FEI](#) still characterizes when there is a full-effort equilibrium, and when the condition holds, our construction of an equilibrium without eventual full effort goes through as well. While one might conjecture that allowing for recall would also not change our result that all equilibria attain eventual full effort when [Condition FEI](#) fails, we do not have a proof.

A related point is that we assume an infinite pool of politicians. With a finite pool, one would have to consider recall. Moreover, for selection to be assured in the long run, one would have to assume the finite pool contains at least one good type—so politician types could not be independent. Such an analysis is beyond the scope of the current paper.²²

Long-lived voter. As discussed in [Section 2](#), there are multiple reasons for our assumption of short-lived voters. But now consider instead only one long-lived voter who maximizes her discounted average payoff with a discount factor $\delta_V \in [0, 1)$, which can be different from the politicians' discount factor δ .

When (δ, κ, f) satisfies [Condition FEI](#) and the replacement cost c is small, there is still an equilibrium that attains full effort as well as one that fails eventual full effort. The former requires no change from our full-effort equilibrium construction. However, in our equilibrium that fails eventual full effort ([Figure 1](#)), a politician's strategy now needs to be modified to account for the voter's discount factor δ_V . Specifically, at non-initial career histories consisting only of Fail signals, the probabilities with which the opportunistic type works in order to keep the voter indifferent depend on δ_V ; but suitable probabilities exist.

When [Condition FEI](#) is violated, there is still no equilibrium with full effort. However, it is unclear whether eventual full effort attains in all, or even some, equilibria. In particular, it is unclear whether equilibria have a positive probability of retaining each good-type officeholder in the long run, because we cannot rule out that equilibrium reputation is bounded away from 1.²³ This echoes the difficulty of obtaining tight predictions on equilibrium payoffs in reputation models without replacement but with two long-lived players ([Cripps and Thomas, 1997](#); [Chan, 2000](#)).

²² In ongoing work, [Ali \(2025\)](#) studies a model similar to ours, except that he has only two politicians. He identifies conditions under which every equilibrium within a class attains eventual full effort, and shows some of those equilibria attain full effort. By contrast, in our setting [Theorem 1](#) says these two desiderata hold only under disjoint conditions. [Ali's \(2025\)](#) eventual full effort logic is not one of selection, unlike ours. One reason for the difference is [Ali's \(2025\)](#) focus on pure Markov strategies for incumbents. Our equilibrium that fails eventual full effort entails incumbents mixing; indeed, [Proposition 2](#) implies that every equilibrium of our model in which politicians use pure strategies yields eventual full effort. We see no compelling reason to exclude mixed strategies from consideration; *inter alia*, as in many models of reputation, one cannot otherwise assure equilibrium existence for all parameters—specifically, when our [Condition FEI](#) fails.

²³ In more detail: with a long-lived voter, it is no longer the case that the incumbent needs to be retained with positive probability following some signal s with $f_1(s) > f_0(s)$. As an example, consider a career history h at which the incumbent shirks: $\sigma_P(h) = 0$. Unlike when the voter is short-lived, the incumbent being retained with positive probability at h no longer implies that he will be retained with probability 1 at any career history h' with higher reputation ($\pi(h') > \pi(h)$), as a long-lived voter's discounted average payoff from keeping this incumbent can be lower at h' even when the voter's stage-payoff from doing so is higher at h' .

5. Conclusion

Our motivating question is whether the power of replacement ensures that officeholders act in voters' interests in the long run. Inspired by previous models of accountability and reputation, we have studied the question in a simple and stylized model, one that we believe is a natural benchmark. Our analysis yields a sharp, dichotomous answer: when there are some intrinsically-motivated politicians, replacement makes good long-run outcomes always possible, but only guaranteed in environments that are *not* conducive to incentive provision.

Our main result, [Theorem 1](#), provides a new perspective on the tension between sanctions and selection that has been noted by others (e.g., [Banks and Sundaram, 1993](#); [Fearon, 1999](#)). In environments conducive to incentive provision, there is an equilibrium with perpetual effort from incumbents. But politics can alternatively fall into a trap of excessive turnover, driven by voters' rational replacement of incumbents whose favorable past performance now makes them shirk, which bounds long-run outcomes away from full effort. In these environments, it is not the underlying conditions that determine long-run outcomes; instead political norms—self-fulfilling expectations—do, even absent any intertemporal voter coordination.

Conversely, in environments where strong incentives cannot be provided (i.e., there is no equilibrium with full effort), the very weakness of accountability becomes the engine of its long-run success. Selection is now inexorable: opportunistic types are weeded out, and voters are assured to eventually secure a good politician who works in their interest. To return to the view of J.S. Mill quoted in the [Introduction](#), in our model it is precisely when opportunistic officeholders cannot be fully disciplined that the worth of the State is eventually determined by the worth of the individuals composing it.

A. Proof Appendix

In our proofs below, we do *not* maintain the assumption made in the main text that $c < \min\{\pi_0, 1 - \pi_0\}$. Instead, to clarify which results require what (if any) assumption on the replacement cost c , we state the requisite assumption when needed.

A.1. Equilibrium Existence

As discussed in Section 3.1, our solution concept can be strengthened to a **personal symmetric perfect Bayesian equilibrium (PBE)** by dropping the qualification of $\sigma_V(h) < 1$ in condition (3) of our [equilibrium definition](#). The proof below shows that such a PBE exists in our game.

Proof of Proposition 1. Let \mathcal{H} be as defined for our original game. For any $\varepsilon \in [0, \frac{1}{2}]$ and $T \in \mathbb{N} \cup \{\infty\}$, define an auxiliary game $\Gamma(\varepsilon, T)$ with a voter in every period and a single politician—who will be a good type with probability π_0 and an opportunistic type with probability $1 - \pi_0$.

We begin by describing $\Gamma(0, \infty)$. The game form is almost exactly as in our original game before the first time a voter replaces the politician, except that at the initial history $h = \emptyset$, the politician chooses a continuous action $\tilde{a}_0 \in [0, 1]$, which generates a signal from $f_{\tilde{a}_0} \in \Delta S$. Say that a politician *semi-pure* strategy is $\sigma_P : \mathcal{H} \rightarrow [0, 1]$, with the interpretation that the politician is not mixing at the initial history but could be mixing (over the binary action set $\{0, 1\}$) at other histories. Voters' (mixed) strategies are given by $\sigma_V : \mathcal{H} \times S \rightarrow [0, 1]$. Furthermore, in this modified game, replacing the politician ends the game. The payoffs of the politician are exactly the same as those of the initial officeholder in our original model. If the politician is not replaced, the payoffs of voters in our auxiliary game are exactly the same as those in our original model. But if a voter replaces the politician, then her payoff is equal to $\pi_0 + (1 - \pi_0) \tilde{a}_0$, a function of the unobserved initial choice of the politician.

A key feature of the game $\Gamma(0, \infty)$ is that the politician's choice $\tilde{a}_0 \in [0, 1]$ directly enters the payoff of the voter who replaces the politician. That is why we treat the politician's choice at the initial history as a continuous action rather than a randomization over the actions 0 and 1.²⁴ With that caveat, we note that every PBE of the auxiliary game $\Gamma(0, \infty)$ in which the politician uses a semi-pure strategy—a semi-pure PBE, for short—corresponds

²⁴If the politician were to mix between actions 0 and 1 (or any other actions) at the initial history, subsequent voters' incentives would depend on the signal s_0 , because it would convey information about the realized initial action, and so the value of replacing the politician.

exactly to a personal symmetric PBE of our original model.²⁵

So it remains to see that $\Gamma(0, \infty)$ has a semi-pure PBE. The remainder of the proof now follows familiar lines (à la Fudenberg and Levine (1983) and Kreps and Wilson (1982)), but we summarize such an argument for the sake of completeness because the game $\Gamma(0, \infty)$ is not an instance of the models covered in those papers.

To show existence, define for every time horizon $T \in \mathbb{Z}_{>0}$ the game $\Gamma(\varepsilon, T)$ by further constraining players' behavior as follows. Before time T , all players (specifically, all voters and opportunistic-type politicians) choose from $[\varepsilon, 1 - \varepsilon]$ at every history rather than from $[0, 1]$; and starting at time T , all players choose from $\{\frac{1}{2}\}$.

First, consider the case of $\varepsilon > 0$ and $T < \infty$. Let us view the game $\Gamma(\varepsilon, T)$ in agent form (with a politician agent and voter agent at each public history). Observe that each agent chooses from a compact interval, only finitely many agents make a nontrivial choice, payoffs are continuous in the strategy profile, and each agent's payoff is affine in their own action. Therefore, by the Debreu-Fan-Glicksberg theorem, a Nash equilibrium $(\sigma_P^{\varepsilon, T}, \sigma_V^{\varepsilon, T})$ exists for $\Gamma(\varepsilon, T)$. Because every history is reached with strictly positive probability under this profile, a unique belief map $\pi^{\varepsilon, T}$ exists that satisfies the Bayesian property given $\sigma_P^{\varepsilon, T}$. Also because every history is reached with strictly positive probability, the Nash equilibrium property implies sequential optimality (in the restricted game).

Now, fixing $\varepsilon > 0$, consider the sequence $(\sigma_P^{\varepsilon, T}, \sigma_V^{\varepsilon, T}, \pi^{\varepsilon, T})_{T=1}^{\infty}$ from the compact metrizable space $\left[([\varepsilon, 1 - \varepsilon] \times [\varepsilon, 1 - \varepsilon]^S \times [0, 1])^{\mathcal{H}}\right]^3$. By compactness, some subsequence converges to some triple $(\sigma_P^{\varepsilon}, \sigma_V^{\varepsilon}, \pi^{\varepsilon})$.

Appealing again to compactness, $(\sigma_P^{\varepsilon}, \sigma_V^{\varepsilon}, \pi^{\varepsilon})_{\varepsilon \in (0, \frac{1}{2})}$ also has some limit point $(\sigma_P, \sigma_V, \pi)$ as $\varepsilon \rightarrow 0$. Continuity of payoffs in the strategy profile (given that the game is continuous at infinity) and continuity of the Bayesian condition (which is easy to see by clearing the denominator) imply $(\sigma_P, \sigma_V, \pi)$ is a semi-pure PBE of $\Gamma(0, \infty)$. \square

A.2. The Main Theorem

A.2.1. Proof Overview

As the main theorem's proof is somewhat involved, we first provide an overview of the arguments. To help parse the logical flow of these arguments, Figure 2 details which results

²⁵ To be precise, our notion of PBE in the auxiliary game is what Mailath (2018) refers to as an “almost perfect Bayesian equilibrium”—although Mailath (2018) restricts attention to finite games, which our auxiliary game is not, the definition can be applied essentially verbatim because we have a finite set of signals.

are inputs to which other results in the paper.

Preliminary calculations. Appendix A.2.2 begins with some preliminary calculations. Adapting standard repeated-games arguments, Lemma 1 confirms that Condition FEI exactly captures the ability to incentivize politicians to exert effort whenever in office. Using a compactness argument, Lemma 2 notes that, if Condition FEI fails, it fails uniformly. More specifically, the lemma finds a lower bound on how much a politician’s continuation value must be allowed to grow (following some signal realizations) in order to make working incentive compatible. Applying this bound, Lemma 3 finds a finite sequence of signals during which the incumbent must strictly prefer to shirk at least once.

Eventual full effort. In Appendix A.2.3, we establish that a failure of Condition FEI generates eventual full effort when replacement costs are low. The key mechanism is one of selection (Proposition 3): opportunistic types are always replaced, whereas good types have a positive probability of being kept forever, meaning the latter is sure to happen eventually. (In Appendix A.3, we observe that this selection easily implies a welfare statement—Corollary 1, which says the time- t voter’s payoff converges to first best as $t \rightarrow \infty$.)

The argument for why opportunistic types are eventually replaced has four steps. First, Lemma 4 yields a finite time horizon and a lower bound on the probability that an opportunistic incumbent is replaced over that horizon, starting from any history at which he exerts effort. Intuitively, if replacement were arbitrarily unlikely over all long horizons, permanent shirking would be profitable. Full-support monitoring allows us to convert the implied lower bound on the replacement rate under such a deviation to a lower bound on the on-path replacement rate.

Second, Lemma 5 shows that any officeholder is eventually replaced, or has his type asymptotically revealed, or asymptotically behaves indistinguishably from the good type. The logic is as in Cripps, Mailath, and Samuelson (2004). By the martingale convergence theorem, reputation converges almost surely. Since monitoring statistically identifies effort, beliefs can have an interior limit with infinitely many signals only if effort is asymptotically uninformative about type—that is, only if the opportunistic type behaves like the good type in the limit.

Third, Lemma 6 shows that that any given officeholder is in fact eventually replaced or has his type asymptotically revealed. Given Lemma 5, we need only rule out that the officeholder is never replaced and yet asymptotically works with probability approaching 1 even if opportunistic. But in the latter event, effort is eventually incentive compatible in every period, and Lemma 4 implies the officeholder is eventually replaced.

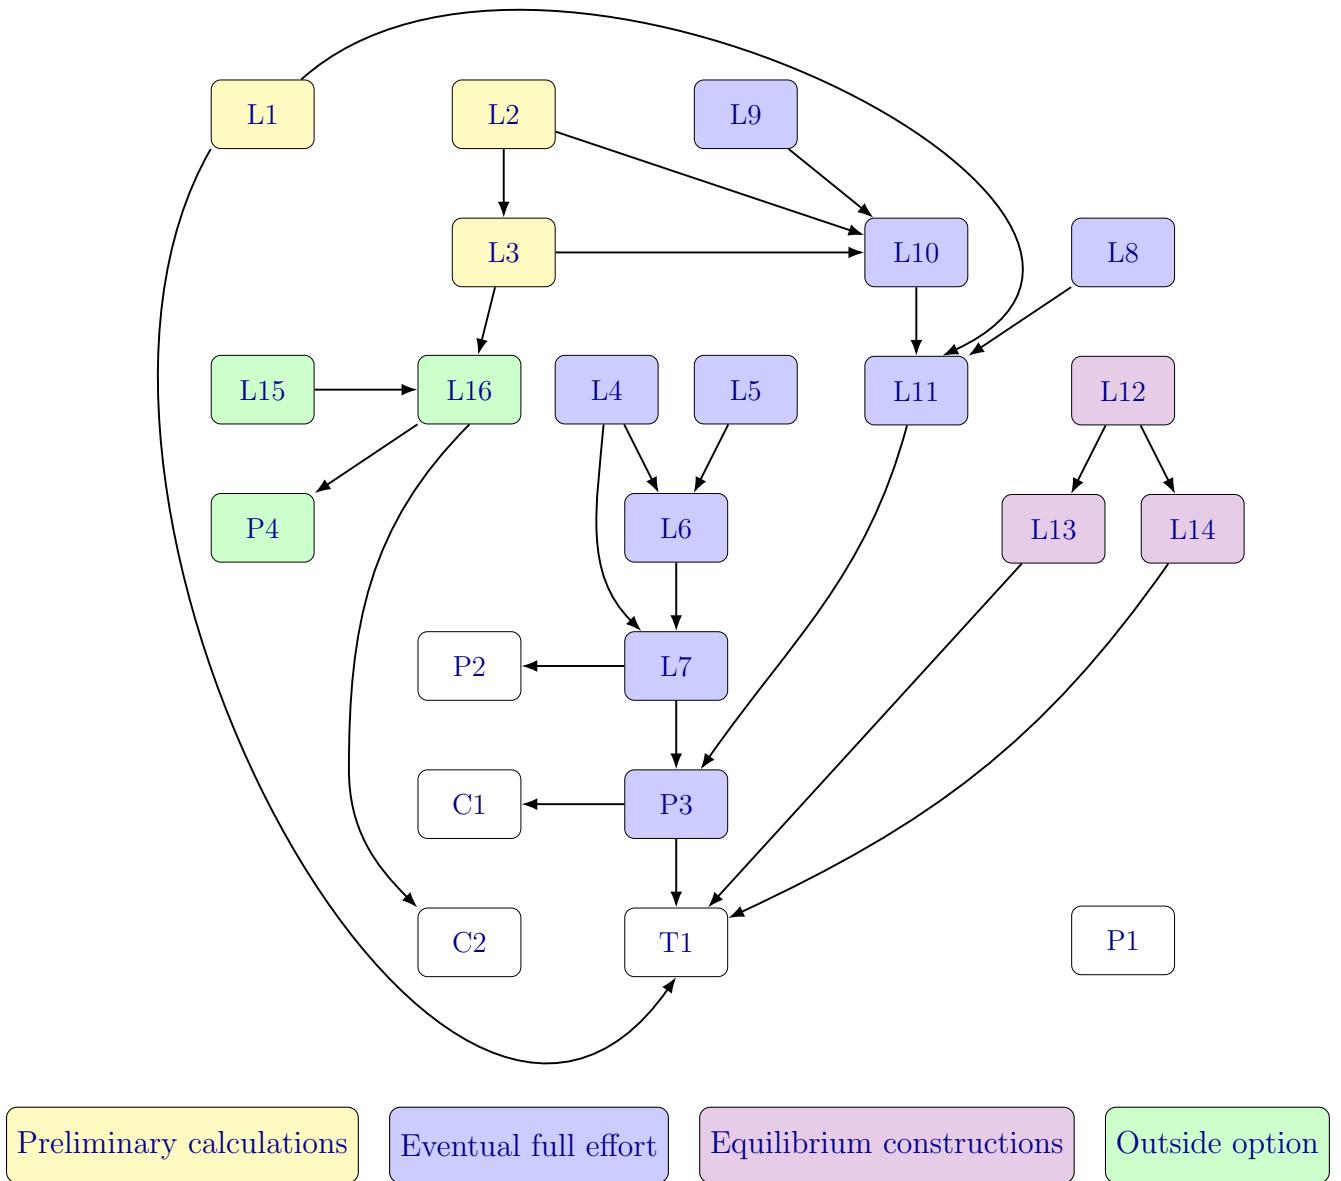


Figure 2: This flow chart depicts the logical relationships between every lemma (L), theorem (T), corollary (C), and proposition (P). An arrow pointing from one result to another indicates that the former is a direct input to the latter. Different colors correspond to different sections of the proof appendix.

Fourth, when the replacement cost is low, [Lemma 7](#) completes the argument that an opportunistic type is eventually replaced. Because of [Lemma 6](#), we need only prove that an officeholder cannot be retained forever while voters become asymptotically certain that he is opportunistic. We rule that possibility out via two cases. If the opportunistic type is eventually expected to work with positive probability, then [Lemma 4](#) guarantees eventual replacement. In the complementary case, the officeholder eventually has sufficiently low reputation and is expected to shirk if opportunistic that voters strictly prefer replacement.

We next argue through four lemmas that if [Condition FEI](#) fails, good types have a positive probability of never being replaced. [Lemma 8](#) shows that reaching a sufficiently high on-path reputation—one at which a voter would retain the incumbent even if an opportunistic type were expected to surely shirk—implies a positive probability of permanent retention.²⁶ Since beliefs are a martingale, Doob’s optional stopping theorem implies that once reputation exceeds this threshold, it has a positive probability of remaining there forever, in which event the incumbent is never replaced.

[Lemma 9](#) shows that any on-path career history is followed by some immediately subsequent on-path history with an additional “good” signal (one strictly indicative of effort). The proof considers two cases. If the incumbent surely shirks at the current history, then learning implies that after any good signal the next voter has a strict incentive to retain the incumbent. If the incumbent works with positive probability at the current history, his incentive compatibility implies that he must be rewarded with a positive probability of retention following some good signal.

[Lemma 10](#) then says an incumbent’s reputation can become arbitrarily good at some on-path histories if [Condition FEI](#) fails. The key to establishing this fact is to consider play at on-path histories with beliefs very close to a supremum that is putatively interior. As [Lemma 9](#) says some good signal results in a positive probability of retention, the hypothesis that we are near an interior supremum means opportunistic types work with high probability. But then, there is little learning no matter the signal, so next period’s belief is again close to that supremum. Iterating this argument over the (finite) horizon identified by [Lemma 3](#) generates a history at which reputation is close to the interior supremum and yet the incumbent shirks, a contradiction.

[Lemma 11](#) finally concludes that any good officeholder has a positive probability of perpetual retention when [Condition FEI](#) fails. In light of [Lemma 10](#) (and because Bayesian updating precludes only opportunists obtaining arbitrarily high reputations), the result is

²⁶ Importantly, this positive probability is unconditional on type; indeed, [Lemma 7](#) implies it is zero conditional on the opportunistic type.

immediate as long as the threshold required in [Lemma 8](#) is strictly below 1. It obviously is when the replacement cost is strictly positive. With a 0 replacement cost, we just need that the voters' outside option is strictly less than 1 in any equilibrium. But an equilibrium with outside option 1 would have incumbents working at every on-path history, which cannot happen when [Condition FEI](#) fails, by [Lemma 1](#).

Equilibrium constructions. [Appendix A.2.4](#) constructs two kinds of equilibria when [Condition FEI](#) holds. By [Lemma 1](#), [Condition FEI](#) is equivalent to the existence of a voter strategy (ignoring voter incentives) that induces an officeholder to exert effort in every period until replacement. [Lemma 12](#) shows that when such a strategy exists, there is one with a simple structure: a time- and history-independent likelihood-ratio test in which all “good” signals (those indicative of effort), and possibly some “bad” signals, pass the test.

The logic underlying the proof of [Lemma 12](#) is as follows.²⁷ Given any voter strategy that incentivizes full effort—which without loss gives the highest possible continuation value to a new officeholder—we can transform it without weakening incentives. First, rewarding all good signals with the highest continuation value relaxes both contemporaneous and past incentive constraints. Second, a replication argument implies that treating signals with the same likelihood ratio preserves incentives. Third, rearranging continuation values to be monotone in the likelihood ratio, while holding the on-path value fixed, only strengthens incentives. The resulting voter strategy takes the form of a likelihood-ratio cutoff test: replace the incumbent for signals below the cutoff, retain him with constant high probability above it, and (at most) mix between these two behaviors at the cutoff. Finally, we show that incentive constraints are relaxed by eliminating mixing and, more straightforwardly, by lowering the replacement rate after “passing” signals (which amounts to scaling up continuation values). Hence, we can further simplify to the incumbent being replaced only after a failing signal.

With [Lemma 12](#) in hand, we can explicitly construct some equilibria under [Condition FEI](#). [Lemma 13](#) constructs an equilibrium that attains full effort. On path, the officeholder always works, and the voter uses the strategy detailed by [Lemma 12](#). If any voter was supposed to have replaced a politician but did not, then this politician always shirks thereafter, her reputation is permanently zero, and all future voters replace him.²⁸ As long as the replacement cost is low, voters find it incentive compatible to replace the incumbent when they should, and every other incentive constraint holds by construction.

²⁷ The proof itself is more algebraic, expressed in terms of the function v witnessing [Condition FEI](#).

²⁸ When the replacement cost is positive ($c > 0$), the belief at the first off-path history in this construction is valid because we are using *weak* PBE. When $c = 0$, our construction can be readily modified to be a PBE.

The final input to [Theorem 1](#) is [Lemma 14](#), which constructs an equilibrium that does not attain eventual full effort. The construction builds on the likelihood ratio test of [Lemma 12](#). Call a signal a Pass if it passes the test and a Fail otherwise; all Fail signals (and possibly some Pass signals) are bad in the sense of being strictly indicative of shirking. For a given small enough replacement cost, our construction—illustrated in [Figure 1](#) for the binary-signal case—is as follows.

After the first Pass signal, continuation play proceeds as in [Lemma 13](#)’s full-effort construction: on path, the officeholder keeps exerting effort and is replaced the first time he produces a Fail signal. Prior to the first Pass signal, voters replace the incumbent with a fixed probability that makes the opportunistic type indifferent between working and shirking; such a probability exists by ([Lemma 12](#) and) the definition of Pass and Fail signals. The opportunistic type’s initial effort probability is strictly below 1, but high enough that for the relevant range of reputations, voters who expect opportunistic incumbents to surely shirk are willing to incur the replacement cost. Before the first Pass signal, effort is determined recursively by keeping voters indifferent between retaining and replacing the incumbent at every history; because Fail signals are bad, the incumbent’s reputation declines over time, which can be exactly offset by increasing the opportunistic type’s effort probability. This construction does not attain eventual full effort because every officeholder is eventually replaced and voters are therefore perpetually exposed to the strictly positive initial shirking probability.²⁹

A.2.2. Preliminary Calculations

[Table 2](#) summarizes some of the key notation used in the formal arguments that follow.

We begin by confirming that [Condition FEI](#) indeed captures incumbents’ incentives to exert full effort. Although this fact follows directly from adapting standard self-generation arguments ([Abreu, Pearce, and Stacchetti, 1990](#)), we include a proof for completeness.

Lemma 1. *The following are equivalent:*

1. *Some strategy profile attains full effort and satisfies politician incentive compatibility.*
2. *Some equilibrium of the modified game with $(\pi_0, c) = (0, 0)$ attains full effort.*
3. *Condition FEI holds.*

Proof. Let us first show condition 1 implies condition 3. To that end, suppose the strategy profile (σ_P, σ_V) attains full effort and satisfies politician incentive compatibility. Let $v : \mathcal{H} \rightarrow$

²⁹ As mentioned in [Section 3.3](#), this also implies that average voter welfare is strictly below first best, even gross of the replacement cost.

$[0, 1]$ capture the politician's continuation values under this strategy profile, and let $\mathcal{H}^* \subseteq \mathcal{H}$ denote the set of on-path histories. Consider the set

$$V := \{(v(h), ([1 - \sigma_V(h, s)] v(h, s))_{s \in S}) : h \in \mathcal{H}^*\} \subseteq [0, 1 - \kappa] \times [0, 1 - \kappa]^S,$$

and let $V_0 \subseteq [0, 1]$ denote the projection of V onto its first coordinate. These sets are nonempty because $\emptyset \in \mathcal{H}^*$. So let $\bar{v}_0 := \sup V_0 \in [0, 1]$. By definition, V contains some sequence $(v_0^n, (v^n(s))_{s \in S})_{n=1}^\infty$ such that $\lim_{n \rightarrow \infty} v_0^n = \bar{v}_0$. Dropping to a subsequence, we may assume (because $[0, 1]$ is compact) that $v^n = (v^n(s))_{s \in S}$ also converges to some $\bar{v} \in [0, 1]^S$ as $n \rightarrow \infty$. Let us show that \bar{v} witnesses [Condition FEI](#). Because it is incentive compatible for the politician to work at every history in \mathcal{H} , every n has $v_0^n = (1 - \delta)(1 - \kappa) + \delta f_1 \cdot v^n \geq (1 - \delta) + \delta f_0 \cdot v^n$, so that taking limits yields

$$\bar{v}_0 = (1 - \delta)(1 - \kappa) + \delta f_1 \cdot \bar{v} \geq (1 - \delta) + \delta f_0 \cdot \bar{v}.$$

Meanwhile, any $h \in \mathcal{H}^*$ and $s \in S$ have either $\sigma_V(h, s) = 1$ or $(h, s) \in \mathcal{H}^*$ by definition, hence $[1 - \sigma_V(h, s)] v(h, s)$ is a convex combination of 0 and something in V_0 . Thus, every $n \in \mathbb{Z}_{>0}$ and $s \in S$ have $v^n(s) \leq \bar{v}_0$, so that taking limits yields $\bar{v} \in [0, \bar{v}_0]^S$. Hence, $\bar{v}_0 = f_1 \cdot \bar{v} \geq f_0 \cdot \bar{v}$ and $\bar{v}_0 \geq \max \bar{v}(S)$, witnessing [Condition FEI](#).

To see that condition 3 implies condition 2, suppose [Condition FEI](#) holds, witnessed by $v : S \rightarrow \mathbb{R}_{\geq 0}$. Letting $v_0 := (1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s)v(s) \geq \max_{s \in S} v(s)$, we can construct a full-effort equilibrium as follows. At every history h , the politician chooses $\sigma_P(h) = 1$, the voter's belief is $\pi(h) = 0$, and following each signal $s \in S$ the voter chooses $\sigma_V(h, s) = 1 - \frac{v(s)}{v_0} \in [0, 1]$. The Bayesian property, and the incentives of the voter (who has payoff 1 for either re-election or replacement) are trivial. Toward politician incentives, note that the constant function $h \mapsto v(h) = v_0$ satisfies the recursive equation for politician continuation payoffs. Because a unique such bounded function $\mathcal{H} \rightarrow \mathbb{R}$ exists (by the Banach fixed point theorem), it must be that the politician's continuation value is v_0 at every history. But then, [\(IC_{FEI}\)](#) tells us the politician finds it optimal to work at every history, confirming equilibrium.

Finally, condition 2 obviously implies condition 1, delivering the lemma. \square

The following lemma says that if [Condition FEI](#) fails, it fails in a uniform way.

Lemma 2. *If (κ, δ, f) do not satisfy [Condition FEI](#), then some $T \in \mathbb{Z}_{>0}$ exists such that every $v \in \mathbb{R}_{\geq 0}^S$ satisfying [\(IC_{FEI}\)](#) has $(1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s)v(s) < \max_{s \in S} v(s) - \frac{1}{T}$.*

Proof. Suppose [Condition FEI](#) fails.

The set V of $v \in [0, 1]^S$ satisfying (IC_{FEI}) is compact, and the function $v \mapsto \max v(S) - [(1 - \delta)(1 - \kappa) + f_1 \cdot v]$ is continuous on it. Failure of [Condition FEI](#) then implies this function is strictly positive on V , and so bounded below by a strictly positive number.

Meanwhile, any $v \in \mathbb{R}_{\geq 0}^S$ with $\bar{v} := \max v(S) \geq 1$ has

$$(1 - \delta)(1 - \kappa) + \delta f_1 \cdot v \leq (1 - \delta)(\bar{v} - \kappa) + \delta \bar{v} = \bar{v} - (1 - \delta)\kappa,$$

so that

$$\inf_{v \in \mathbb{R}_{\geq 0}^S: v \text{ satisfies } (\text{IC}_{\text{FEI}})} \{\max v(S) - [(1 - \delta)(1 - \kappa) + \delta f_1 \cdot v]\} > 0.$$

Thus $\frac{1}{T}$ is strictly below this infimum value for T large enough. \square

The next lemma finds, when [Condition FEI](#) fails, a sequence of signals and a duration over which an incumbent cannot be incentivized to work.

Lemma 3. *Suppose [Condition FEI](#) fails, let T be as delivered by [Lemma 2](#), and fix any equilibrium and any $h_0 \in \mathcal{H}$ with $v(h_0) > 0$. Recursively define $s_0, \dots, s_{T-1} \in S$ and $h_1, \dots, h_T \in \mathcal{H}$ by letting $s_t \in \arg \max_{s \in S} \{[1 - \sigma_V(h_t, s)] v(h_t, s)\}$ and $h_{t+1} = (h_t, s_t)$ for $t \in \{0, \dots, T-1\}$. Then some $t \in \{0, \dots, T-1\}$ has $\sigma_P(h_t) = 0$.*

Proof. Assume for a contradiction that every $t \in \{0, \dots, T-1\}$ has $\sigma_P(h_t) > 0$. For each such t , the fact that the incumbent willingly exert effort tells us the vector $v^t := ([1 - \sigma_V(h_t, s)] v(h_t, s))_{s \in S} \in [0, 1]^S$ satisfies (IC_{FEI}) . Hence, [Lemma 2](#) says

$$v(h_t) = (1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s) v^t(s) < v^t(s_t) - \frac{1}{T} \leq v(h_{t+1}) - \frac{1}{T}.$$

Therefore,

$$0 \leq v(h_0) < v(h_T) - T \frac{1}{T} \leq 1 - T \frac{1}{T} = 0,$$

a contradiction. \square

For the analysis that follows, we introduce some notation to describe voters' belief updating.

Notation 1. *For any belief $\pi \in (0, 1]$, work probability $a \in [0, 1]$ for the opportunistic type, and signal $s \in S$, let*

$$\beta_a(\pi|s) := \frac{\pi f_1(s)}{[\pi + (1 - \pi)a] f_1(s) + (1 - \pi)(1 - a) f_0(s)} \in (0, 1]$$

denote the updated Bayesian belief that a voter would have after observing signal s .

For any $\pi \in (0, 1]$ and bound $\eta \in [0, 1]$, let

$$\bar{\beta}_\eta(\pi) := \sup_{a \in [\eta, 1], s \in S} \beta_a(\pi|s) \in (0, 1]$$

denote the highest level beliefs can grow to in a period when the opportunistic-type incumbent's expected effort is at least η .

For any $t \in \mathbb{Z}$ and $\eta \in [0, 1]$, let $\bar{\beta}_\eta^t : (0, 1] \rightarrow (0, 1]$ denote the t -fold composition of $\bar{\beta}_\eta$.

A.2.3. Eventual Full Effort

Primarily for the argument that a failure of Condition FEI implies all equilibria attain eventual full effort, we will also use the following notation.

Notation 2. Given any equilibrium, we define the following random variables.

For any $t \in \mathbb{Z}_{\geq 0}$, let $\mathbf{r}_t \in \{0, 1\}$ be the random variable indicating whether the first officeholder has been replaced as of the beginning of period t ; and let $\tau_r := \inf\{t \in \mathbb{Z}_{\geq 0} : \mathbf{r}_t = 1\} \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ denote the first replacement time. For any $t \in \mathbb{Z}_{\geq 0}$ with $t \leq \tau_r$, let $\mathbf{h}_t \in \mathcal{H}_t$ be the random variable denoting the first officeholder's public career history up to time t . For any $t \in \mathbb{Z}_{\geq 0}$, let $\pi_t := \pi(\mathbf{h}_{t \wedge \tau_r}) \in [0, 1]$ denote the first officeholder's reputation, and let $\tilde{\mathbf{a}}_t := (1 - \mathbf{r}_t)\sigma_P(\mathbf{h}_{t \wedge \tau_r}) \in [0, 1]$ denote the expected effort of the initial officeholder conditional on being opportunistic. Let \mathbf{r}_∞ [resp., $\pi_\infty, \tilde{\mathbf{a}}_\infty$] denote the limit as $t \rightarrow \infty$ of \mathbf{r}_t [resp., $\pi_t, \tilde{\mathbf{a}}_t$] in the event that this limit exists.

Any equilibrium induces a joint law over the stochastic process $(\mathbf{a}_t, \pi_t, \mathbf{r}_t)_{t=0}^\infty$. We take all probabilistic statements to be with respect to this law unless stated otherwise (and we suppress the underlying probability space and the specific equilibrium in our notation).

For ease of notation, we work throughout with the above random variables concerning only the first officeholder. Given our focus on personal symmetric equilibria, results we prove about the first officeholder apply without change to any officeholder.

We now show that politicians' incentives imply a lower bound on the opportunistic type's (time-averaged) hazard rate of being replaced.

Lemma 4. Some $T \in \mathbb{Z}_{>0}$ and $\eta \in (0, 1)$ exist such that, in any equilibrium and at any history $h \in \mathcal{H}$ with $\sigma_P(h) > 0$, an opportunistic-type officeholder has a probability of at least η of being replaced in the next T periods.

Proof. Take some large enough $T \in \mathbb{Z}_{>0}$ and small enough $\varepsilon \in (0, 1)$ to ensure that

$$(1 - \varepsilon)(1 - \delta^T) > (1 - \delta)(1 - \kappa) + \delta. \quad (\text{A.1})$$

Denoting $S_0 := \text{supp}(f_0)$, let $\lambda := \min_{s \in S_0} \frac{f_1(s)}{f_0(s)} \in (0, 1)$ and $\eta := \lambda^T \varepsilon$.

First, observe that at any history $h \in \mathcal{H}$ with $\sigma_P(h) > 0$, incentive compatibility requires that an opportunistic incumbent has a probability at least ε of being replaced in the next T periods, conditional on always shirking starting from h . The reason is that if this probability of replacement were strictly less than ε , then the continuation payoff from deviating to always shirking would exceed $(1 - \varepsilon)(1 - \delta^T)$, which by (A.1) strictly exceeds the maximum feasible continuation payoff from working at h .

Next, observe that for any sequence of T signals, its probability of arising (conditional on the incumbent not being replaced) under any politician strategy is at least λ^T times as high as its probability of arising if the politician were to always shirk. Therefore, the incumbent at h will be replaced in the next T periods with probability at least $\lambda^T \varepsilon = \eta$. \square

We now record a lemma that says either an officeholder's type will be revealed or learning will halt. The reasoning behind this lemma is essentially the same as that of [Cripps, Mailath, and Samuelson \(2004\)](#).

Lemma 5. *In any equilibrium, the stochastic process $(\mathbf{r}_t, \tilde{\mathbf{a}}_t, \boldsymbol{\pi}_t)_{t=0}^\infty$ satisfy the following:*

1. *The sequence $(\mathbf{r}_t, \boldsymbol{\pi}_t)_t$ converges a.s. to a random variable $(\mathbf{r}_\infty, \boldsymbol{\pi}_\infty)$.*
2. *We a.s. have $\mathbf{r}_\infty = 1$ (that is, $\tau_{\mathbf{r}} < \infty$) or $\boldsymbol{\pi}_\infty \in \{0, 1\}$ or $(\tilde{\mathbf{a}}_t)_t$ converges to 1.*

Proof. The first point follows from pathwise monotonicity of \mathbf{r} and the martingale convergence theorem for the bounded martingale $\boldsymbol{\pi}$.

For the second point, consider the event in which $\mathbf{r}_\infty \neq 1$ and $\boldsymbol{\pi}_\infty \notin \{0, 1\}$ —so $\mathbf{r}_\infty = 0$ (that is, $\tau_{\mathbf{r}} = \infty$) and $0 < \boldsymbol{\pi}_\infty < 1$. For any $\pi \in (0, 1)$, we have $\min_{s \in S, \tilde{a} \in [0, 1]} f_{\pi+(1-\pi)\tilde{a}}(s) > 0$; and some $s \in S$ has $f_1(s) \neq f_0(s)$. Therefore, when a belief lies in some small enough open neighborhood of π and the officeholder's expected effort is bounded away from 1, the probability of the belief exiting said neighborhood is bounded below. Thus, in the event that $0 < \boldsymbol{\pi}_\infty < 1$, we must have $\tilde{\mathbf{a}}_t \rightarrow 1$. \square

We now prove that in any equilibrium, an officeholder is either eventually replaced or his type is asymptotically revealed.

Lemma 6. Consider any equilibrium, and let $(\mathbf{r}_\infty, \boldsymbol{\pi}_\infty)$ be as in the statement of Lemma 5. We a.s. have $\mathbf{r}_\infty = 1$ (that is, $\tau_r < \infty$) or $\boldsymbol{\pi}_\infty \in \{0, 1\}$.

Proof. Consider the event in which $\boldsymbol{\pi}_\infty > 0$ and $\tilde{\mathbf{a}}_\infty := \lim_{t \rightarrow \infty} \tilde{\mathbf{a}}_t = 1$. In this event, we in particular have $\tilde{\mathbf{a}}_t > 0$ for sufficiently large $t \in \mathbb{Z}_{\geq 0}$. Lemma 4 delivers some $T \in \mathbb{Z}_{>0}$ and $\eta \in (0, 1)$ such that the probability of replacement in the next T periods, conditional on the opportunistic-type officeholder, is eventually at least η in this event. But then, because $\tilde{\mathbf{a}}_\infty = 1$, the good type also eventually has at least a $\frac{\eta}{2}$ probability of replacement in the next T periods. Hence, by the Borel-Cantelli Lemma, with probability 1 conditional on this event, we have $\mathbf{r}_\infty = 1$.

We have argued that, excluding a null event, if $\boldsymbol{\pi}_\infty > 0$ and $\tilde{\mathbf{a}}_\infty = 1$, then $\mathbf{r}_\infty = 1$. Combining this observation with Lemma 5 delivers the lemma. \square

The next lemma says that, when the replacement cost is small, an opportunistic officeholder is eventually replaced.

Lemma 7. If $c < \pi_0$, then with probability 1 conditional on being an opportunistic type, an officeholder is replaced.

Proof. Suppose $c < \pi_0$. Let us first show that (excepting a zero probability event) $\mathbf{r}_\infty = 1$ if $\boldsymbol{\pi}_\infty = 0$. Indeed, if $\boldsymbol{\pi}_\infty = 0$, then some $\tau \in \mathbb{Z}_{\geq 0}$ exists such that $\boldsymbol{\pi}_t < \min\{c, \frac{1}{2}\}$ for every $t \geq \tau$. Let us now show $\mathbf{r}_\infty = 1$ in two different cases. First, if $\tilde{\mathbf{a}}_t = 0$ for some $t \geq \tau$, then voter incentives require $\mathbf{r}_t = 1$ so that $\mathbf{r}_\infty = 1$ too. Second, if $\tilde{\mathbf{a}}_t > 0$ for every $t \geq \tau$, then Lemma 4 delivers some $T \in \mathbb{Z}_{>0}$ and $\eta \in (0, 1)$ such that the probability of replacement in the next T periods is at least $\frac{\eta}{2}$ from time τ onward. Then, by the Borel-Cantelli Lemma, we have $\mathbf{r}_\infty = 1$.

Lemma 6 tells us $\mathbf{r}_\infty = 1$ or $\boldsymbol{\pi}_\infty \in \{0, 1\}$ a.s., and the argument above rules out the case in which $\mathbf{r}_\infty = \boldsymbol{\pi}_\infty = 0$ with positive probability. Therefore, we a.s. have $\mathbf{r}_\infty = 1$ or $\boldsymbol{\pi}_\infty = 1$. The lemma follows because the probability of $\boldsymbol{\pi}_\infty = 1$ conditional on the officeholder being an opportunistic type is necessarily zero. \square

The next lemma shows that if an officeholder attains a high enough reputation on-path that the voter would currently retain him irrespective of opportunistic-type behavior, then he has a positive probability (unconditional on his type) of never being replaced.

Lemma 8. Fix an equilibrium, and let u_0 denote the voters' outside option. For any $t \in \mathbb{Z}_{\geq 0}$ and history $h \in \mathcal{H}$ is such that $\mathbb{P}\{\mathbf{h}_t = h \text{ and } \mathbf{r}_t = 0\} > 0$ and $\pi(h) > u_0 - c$, we have $\mathbb{P}\{\mathbf{h}_t = h, \boldsymbol{\pi}_\infty \geq \pi(h), \text{ and } \mathbf{r}_\infty = 0\} > 0$.

Proof. Define the stopping time $\tau := \inf \{\tau \in \mathbb{Z}_{\geq 0} : \tau \geq t \text{ and } \pi_\tau \leq u_0 - c\}$. Note that the random variable π_τ is well defined a.s. (even if $\tau = \infty$) by the martingale convergence theorem. Doob's optional stopping theorem implies $\mathbb{E}[\pi_\tau | \mathbf{h}_t = h] = \pi(h)$. Therefore, $\mathbb{P}\{\pi_\tau \geq \pi(h) | \mathbf{h}_t = h\} > 0$. But if $\mathbf{h}_t = h$ and $\pi_\tau \geq \pi(h) > u_0 - c$, then the voters' beliefs remain in $(u_0 - c, 1]$ from time t onward—and hence, by voter incentives, they do not replace the officeholder from time t onward. So

$$\mathbb{P}\{\mathbf{r}_\infty = 0 \text{ and } \pi_\infty \geq \pi(h) | \mathbf{h}_t = h \text{ and } \mathbf{r}_t = 0\} \geq \mathbb{P}\{\pi_\tau \geq \pi(h) | \mathbf{h}_t = h \text{ and } \mathbf{r}_t = 0\} > 0.$$

Because $\mathbb{P}\{\mathbf{h}_t = h \text{ and } \mathbf{r}_t = 0\} > 0$ by hypothesis, the result follows. \square

The following lemma says that, if a career history is consistent with an officeholder not having been replaced yet, then some good news (that is, some signal indicative of effort) exists that does not lead to immediate firing.

Lemma 9. *In any equilibrium, any on-path history h admits some $s \in S$ with $f_1(s) > f_0(s)$ and (h, s) on path.*

Proof. Let \mathcal{H}^* denote the set of on-path histories: $h \in \mathcal{H}$ if and only if some $t \in \mathbb{Z}_{\geq 0}$ has $\mathbb{P}\{\mathbf{h}_t = h \text{ and } \mathbf{r}_t = 0\} > 0$. First, consider the case in which $\sigma_P(h) = 0$. That $h \in \mathcal{H}^*$ while $\sigma_P(h) = 0$ implies $\pi(h) \geq u_0 - c$, where u_0 is the voters' outside option. Indeed, this inequality holds by voter incentives at h if $h \neq \emptyset$, and it holds by definition (since $c \geq 0$) for $h = \emptyset$. In this case, we can use any $s \in S$ with $f_1(s) > f_0(s)$ —which exists because f is informative. The Bayesian property implies (since $\pi(\cdot) > 0$ at on-path histories because f_1 has full support) that $\pi(h, s) > \pi(h)$. Hence $\pi(h, s) + [1 - \pi(h, s)] \sigma_P(h, s) > u_0 - c$, and so voters have a strict incentive to keep the incumbent at history (h, s) . Thus, $(h, s) \in \mathcal{H}^*$.

Now, turn to the case in which $\sigma_P(h) > 0$. The fact that the incumbent optimally chooses to work at h implies

$$\begin{aligned} 0 &\leq \left\{ (1 - \delta)(1 - \kappa) + \delta \sum_{s \in S} f_1(s) [1 - \sigma_V(h, s)] v(h, s) \right\} \\ &\quad - \left\{ (1 - \delta)1 + \delta \sum_{s \in S} f_0(s) [1 - \sigma_V(h, s)] v(h, s) \right\} \\ &= -(1 - \delta)\kappa + \delta \sum_{s \in S} [f_1(s) - f_0(s)] [1 - \sigma_V(h, s)] v(h, s) \\ &< \sum_{s \in S: f_1(s) > f_0(s)} \delta [f_1(s) - f_0(s)] [1 - \sigma_V(h, s)] v(h, s). \end{aligned}$$

Therefore, at least one summand is strictly positive. In particular, some $s \in S$ has $f_1(s) > f_0(s)$ and $\sigma_V(h, s) < 1$. Because $h \in \mathcal{H}^*$, the latter property tells us $(h, s) \in \mathcal{H}^*$ too. \square

The next lemma says that when **Condition FEI** fails, on-path equilibrium reputation can become arbitrarily optimistic.

Lemma 10. *Suppose **Condition FEI** fails, and fix any equilibrium. Any $\hat{\pi} \in (0, 1)$ has $\mathbb{P}\{\pi_t > \hat{\pi} \text{ and } \mathbf{r}_t = 0\} > 0$ for some $t \in \mathbb{Z}_{\geq 0}$.*

Proof. We begin with some useful definitions and notation. Let \mathcal{H}^* be the set on-path histories (as also defined at the start of the proof of [Lemma 9](#)), let $\Pi^* := \{\pi(h) : h \in \mathcal{H}^*\}$, and let $\bar{\pi} := \sup(\Pi^*) \in [\pi_0, 1]$. Define $S_1 := \{s \in S : f_1(s) > f_0(s)\}$, and note that the function $(a, \pi) \mapsto \min_{s \in S_1} \beta_a(\pi|s)$ is continuous in π on $[0, 1] \times (0, 1)$ because each of the finitely many maximands is. Hence, because $[0, 1]$ is compact and the function above takes $(1, \pi) \mapsto \pi$ for $\pi \in (0, 1)$, the function $\underline{a} : (0, 1) \rightarrow [0, 1]$ given by

$$\underline{a}(\pi) := \min \left\{ a \in [0, 1] : \min_{s \in S_1} \beta_a(\pi|s) \leq \bar{\pi} \right\}$$

is the lower envelope of a compact-valued and upper hemicontinuous correspondence—hence is well-defined and lower semicontinuous.

Our goal is to prove $\bar{\pi} = 1$; assume otherwise ($0 < \bar{\pi} < 1$) for a contradiction. Given this hypothesis, observe that any $a \in [0, 1]$ and $s \in S_1$ have $\beta_a(\bar{\pi}|s) > \bar{\pi}$; thus, $\underline{a}(\bar{\pi}) = 1$. This fact has two consequences. First, lower semicontinuity of \underline{a} tells us $\lim_{\pi \rightarrow \bar{\pi}} \underline{a}(\pi) = 1$ too. Second, consequently, some $\underline{\pi}_0 \in (0, \bar{\pi})$ has $\underline{a}(\underline{\pi}_0) > 0$ and so $\min_{s \in S_1} \beta_0(\underline{\pi}_0|s) > \bar{\pi}$.

Now, let T be as given by [Lemma 2](#). Recursively for $t \in \{1, \dots, T\}$, let us find some $\underline{\pi}_t \in (\underline{\pi}_{t-1}, \bar{\pi})$ with the property that: Every $a \in [0, 1]$ with $\inf_{\pi \in (\underline{\pi}_t, \bar{\pi}), s \in S_1} \beta_a(\pi|s) \leq \bar{\pi}$ has $\inf_{\pi \in (\underline{\pi}_t, \bar{\pi}), s \in S_1} \beta_a(\pi|s) \geq \underline{\pi}_{t-1}$. To do so, observe that $\lim_{\pi \rightarrow \bar{\pi}} \underline{a}(\pi) = 1$, and so (by continuity) every $s \in S$ has $\beta_{\underline{a}(\pi)}(\pi|s) \rightarrow \beta_1(\bar{\pi}|s) = \bar{\pi}$ as $\pi \rightarrow \bar{\pi}$. Hence, we can set $\underline{\pi}_t \in (\underline{\pi}_{t-1}, \bar{\pi})$ close enough to $\bar{\pi}$ to ensure $\beta_{\underline{a}(\underline{\pi}_t)}(\underline{\pi}_t|s) > \underline{\pi}_{t-1}$ for every $s \in S$. Let us see that such a $\underline{\pi}_t$ has the desired property. Indeed, any $a \in [0, 1]$ with $\inf_{\pi \in (\underline{\pi}_t, \bar{\pi}), s \in S_1} \beta_a(\pi|s) \leq \bar{\pi}$ will satisfy $\min_{s \in S_1} \beta_a(\underline{\pi}_t|s) \leq \bar{\pi}$ (by continuity and monotonicity of $\beta_a(\cdot|s)$), so that $a \geq \underline{a}(\underline{\pi}_t)$. Hence, any $s \in S$ satisfies $\beta_a(\underline{\pi}_t|s) \in \text{co} \{\beta_{\underline{a}(\underline{\pi}_t)}(\underline{\pi}_t|s), \beta_1(\underline{\pi}_t|s)\} = [\beta_{\underline{a}(\underline{\pi}_t)}(\underline{\pi}_t|s), \underline{\pi}_t]$. In particular, $\beta_a(\underline{\pi}_t|s) > \underline{\pi}_{t-1}$. Then, since s was arbitrary in this calculation and $\beta_a(\cdot|s)$ is increasing for every $a \in [0, 1]$ and $s \in S$, it follows that $\inf_{\pi \in (\underline{\pi}_t, \bar{\pi}), s \in S_1} \beta_a(\pi|s) > \underline{\pi}_{t-1}$.

Let us turn now to our equilibrium. Let $h_0 \in \mathcal{H}^*$ be some on-path history such that $\pi_T < \pi(h_0) < \bar{\pi}$. Define $s_0, \dots, s_{T-1} \in S$ and $h_1, \dots, h_T \in \mathcal{H}$ as defined in the statement

of [Lemma 3](#). That lemma tells us some $\tau \in \{0, \dots, T-1\}$ has $\sigma_P(h_\tau) = 0$; consider the smallest such τ . For every $t \in \{0, \dots, \tau-1\}$, the incumbent finds it optimal to exert effort at history h_t , which implies $[1 - \sigma_V(h_t, s)] v(h_t, s) > 0$ for some $s \in S$, hence for $s = s_t$. In particular, $\sigma_V(h_{t+1}) < 1$. By induction on t , it then follows that $h_1, \dots, h_\tau \in \mathcal{H}^*$.

Next, we show that $\underline{\pi}_{T-t} < \pi(h_t) < \bar{\pi}$ for every $t \in \{0, \dots, \tau\}$. We will prove this by induction; the base case of $t = 0$ holds by fiat. Toward the inductive step, suppose $t \in \{0, \dots, \tau-1\}$ has $\underline{\pi}_{T-t} < \pi(h_t) < \bar{\pi}$. Because $h_t \in \mathcal{H}^*$, [Lemma 9](#) tells us some $s'_t \in S_1$ has $(h_t, s'_t) \in \mathcal{H}^*$ too. But then, by definition of $\bar{\pi}$, it must be that $\pi(h_t, s'_t) \leq \bar{\pi}$. It then follows from the defining property of $\underline{\pi}_{T-t}$ that

$$\pi(h_{t+1}) = \pi(h_t, s_t) > \underline{\pi}_{T-t-1} = \underline{\pi}_{T-(t+1)},$$

completing the inductive step.

Finally, we now know that $h_\tau \in \mathcal{H}^*$, that $\pi(h_\tau) \in (\underline{\pi}_{T-\tau}, \bar{\pi}) \subseteq (\underline{\pi}_0, \bar{\pi})$ and that $\sigma_P(h_\tau) = 0$. [Lemma 9](#) says some $s \in S_1$ has $(h_\tau, s) \in \mathcal{H}^*$, and the definition of $\underline{\pi}_0$ means $\pi(h_\tau, s) > \bar{\pi}$, which contradicts the definition of $\bar{\pi}$. \square

Now we prove that when [Condition FEI](#) fails, any good-type officeholder has a positive probability of never being replaced.

Lemma 11. *Suppose [Condition FEI](#) fails. In any equilibrium, the probability of any given good-type officeholder never being replaced is strictly positive.*

Proof. Let us first show the result holds assuming the voters' outside option $u_0 \in [0, 1]$ satisfies $u_0 - c < 1$. Then we will verify this inequality holds.

By symmetry, we need only show the result for the first officeholder. By [Lemma 10](#) (and since $u_0 - c < 1$), there is a positive probability of reaching a history at which the first officeholder has not been replaced, and his reputation is strictly positive and strictly greater than $u_0 - c$. But then [Lemma 8](#) says the probability of never being replaced and having $\pi_\infty > 0$ is strictly positive too. Hence, the probability $\mathbb{E}[(1 - \mathbf{r}_\infty)\pi_\infty]$ of being a good type and never being replaced is strictly positive. Consequently, the probability of never being replaced *conditional on* being a good type is also strictly positive.

All that remains is to show $u_0 - c < 1$. The inequality is trivial if either $u_0 < 1$ or $c > 0$, so assume towards contradiction that $u_0 = 1$ and $c = 0$. Voter incentives then require $\sigma_P(h) = 1$ at any history $h \in \mathcal{H}$ with $\pi(h) < 1$ and $\sigma_V(h) < 1$. Because the Bayesian property then implies $\pi(h) = \pi_0 < 1$ at any history that is reached with positive probability,

the opportunistic-type officeholder works at every on-path history. But this contradicts politician incentives, by Lemma 1. \square

Now we argue that when Condition FEI fails and replacement costs are low, a good-type politician is eventually retained in office forever.

Proof of Proposition 3. The first part (which relies on $c < \pi_0$) is exactly Lemma 7, and the second (which does not require our bound on replacement cost) is Lemma 11.

Now, toward the proposition's last sentence (which again relies on $c < \pi_0$), suppose Condition FEI fails. We know every opportunistic officeholder is replaced with probability 1, whereas every good-type officeholder has a positive (and by symmetry, the same positive) probability of never being replaced. By the Borel-Cantelli Lemma, with probability 1, some officeholder is never replaced, and this officeholder is a good type. \square

A.2.4. Equilibrium Constructions

The next lemma says Condition FEI is equivalent to a more stringent condition that imposes further restrictions on the continuation values. In terms of induced behavior, if full effort can be incentivized at all, it can be done with a stationary likelihood-ratio test in which sufficiently bad news results in certain replacement.

Lemma 12. *If Condition FEI holds, then it can be witnessed by $v \in \mathbb{R}_{\geq 0}^S$ such that:*

- *Some $\bar{v} \in (0, 1 - \kappa)$ exists such that $v(S) = \{v(s) : s \in S\} = \{0, \bar{v}\}$.*
- *Every s_0 with $f_0(s_0) \leq f_1(s_0)$ has $v(s_0) = \bar{v}$.*
- *Every $s_0, s_1 \in S$ with $v(s_0) = 0$ and $v(s_1) = \bar{v}$ has $\frac{f_0(s_0)}{f_1(s_0)} > \frac{f_0(s_1)}{f_1(s_1)}$.*
- *Inequality (PK_{FEI}) holds with equality.*

Proof. Let $S_0 := \{s \in S : f_1(s) < f_0(s)\}$, and let \succsim be the complete transitive binary relation on S given by $s' \succsim s \iff \frac{f_0(s')}{f_1(s')} \leq \frac{f_0(s)}{f_1(s)}$. Suppose Condition FEI holds and is witnessed by v . We will successively show that various properties of v are without loss of generality. To that end, let $\bar{v} := \max v(S) \in \mathbb{R}_{\geq 0}$. Note that $0 < \bar{v} < 1 - \kappa$, since any v with all entries the same would violate (IC_{FEI}), any v with $\max v(S) \geq 1 - \kappa$ would violate (PK_{FEI}).

First, raising $v(s)$ to \bar{v} for each $s \in S \setminus S_0$ relaxes both constraints. So we may assume $v(s) = \bar{v}$ for every $s \in S \setminus S_0$. Second, if $\hat{S} \subseteq S_0$ is some \sim -equivalence class, then replacing

$v(s)$ with $\frac{\sum_{\hat{s} \in \hat{S}} f_1(\hat{s})v(\hat{s})}{\sum_{\hat{s} \in \hat{S}} f_1(\hat{s})}$ for every $s \in \hat{S}$ preserves both constraints. So we may assume v is constant on every \sim -equivalence class.

Next, observe that—because $v \in [0, \bar{v}]^S$ —some $s^* \in S_0$ exists such that

$$\sum_{s \succ s^*} f_1(s)\bar{v} \leq \sum_{s \in S} f_1(s)v(s) \leq \sum_{s \sim s^*} f_1(s)\bar{v};$$

and so some $\hat{v} \in [0, \bar{v}]$ exists such that

$$\left[\sum_{s \succ s^*} f_1(s) \right] \bar{v} + \left[\sum_{s \sim s^*} f_1(s) \right] \hat{v} = \sum_{s \in S} f_1(s)v(s).$$

Observe now that both inequalities in [Condition FEI](#) will still be satisfied if we replace v with the vector whose s entry is \bar{v} for $s \succ s^*$, is $\hat{v} \in [0, \bar{v}]$ for $s \sim s^*$, and is 0 for $s \prec s^*$. Indeed, this modification has no effect on [\(PK_{FEI}\)](#) but relaxes [\(IC_{FEI}\)](#). (Moreover, because we had $v|_{S \setminus S_0} = v$ before the rearrangement and $S \setminus S_0 \succ S_0$, it follows that we still have this property.) So we may further assume v takes this form.

Now, define $\hat{z} := 1 - \delta \sum_{\tilde{s} \succ s^*} f_1(\tilde{s}) > 0$ and $\bar{z} := -\delta \sum_{\tilde{s} \sim s^*} f_1(\tilde{s}) < 0$, and define $z \in \mathbb{R}^S$ by letting

$$z(s) := \begin{cases} 0 & : s \prec s^* \\ \hat{z} & : s \sim s^* \\ \bar{z} & : s \succ s^* \end{cases}$$

for each $s \in S$. Observe that the set of $\varepsilon \in \mathbb{R}$ with $\bar{v} + \varepsilon \bar{z} \geq \hat{v} + \varepsilon \hat{z} \geq 0$ is a closed interval containing zero. It is also bounded: the first inequality is violated as $\varepsilon \rightarrow \infty$, and the second is violated as $\varepsilon \rightarrow -\infty$. This interval therefore has two extreme points, and at least one of these ε is such that $v + \varepsilon z$ satisfies [\(IC_{FEI}\)](#) because the latter is affine. Finally, simple algebra shows that this $v + \varepsilon z$ satisfies [\(PK_{FEI}\)](#) too because v does. So replacing v with this modification, we may assume that $\hat{v} \in \{0, \bar{v}\}$, or equivalently that $v(S) \subseteq \{0, \bar{v}\}$. And in fact, we have $v(S) = \{0, \bar{v}\}$ because any v with all entries the same violates [\(IC_{FEI}\)](#).

Finally, because scaling up a vector v that satisfies [\(IC_{FEI}\)](#) relaxes this constraint but tightens [\(PK_{FEI}\)](#), we can scale it up to ensure [\(PK_{FEI}\)](#) holds with inequality. \square

We now prove that under [Condition FEI](#) and a small replacement cost, some equilibrium attains full effort.

Lemma 13. *If $c \leq 1 - \pi_0$ and [Condition FEI](#) holds, then some equilibrium attains full effort.*

Proof. By Lemma 12, some $\bar{v} \in (0, 1)$ and $\lambda \geq 1$ exist such that $S^* := \left\{ s \in S : \frac{f_0(s)}{f_1(s)} \leq \lambda \right\}$ satisfies

$$\bar{v} = (1 - \delta)(1 - \kappa) + \delta \sum_{s \in S^*} f_1(s)\bar{v} \geq (1 - \delta)1 + \delta \sum_{s \in S^*} f_0(s)\bar{v}.$$

Let $\mathcal{H}_{\text{Pass}} := \bigcup_{t=0}^{\infty} (S^*)^t$, the histories in which the incumbent has never generated a signal outside of S^* . We will now construct an equilibrium that attains full effort. Let the politician strategy be given by $\sigma_P(h) = \mathbf{1}_{h \in \mathcal{H}_{\text{Pass}}}$, the voters' strategies be given by $\sigma_V(h) = \mathbf{1}_{h \notin \mathcal{H}_{\text{Pass}}}$, and let the belief map be given by

$$\pi(h) = \begin{cases} \pi_0 & \text{if } h \in \mathcal{H}_{\text{Pass}} \\ \pi_0 & \text{if } h = (\tilde{h}, s) \text{ for some } \tilde{h} \in \mathcal{H}_{\text{Pass}} \text{ and } s \in S \setminus S^* \\ 0 & \text{otherwise.} \end{cases}$$

This strategy profile clearly leads to effort always being chosen until an incumbent is replaced, so we need only show this construction is in fact an equilibrium.³⁰

A incumbent's incentives at histories outside of $\mathcal{H}_{\text{Pass}}$ are satisfied because he will be replaced immediately after any signal, and so it is optimal to shirk. Meanwhile, his incentives at histories in $\mathcal{H}_{\text{Pass}}$ follow directly from the properties defining \bar{v} : his continuation value at such a history is \bar{v} , and he prefers this to the value of a one-shot deviation in which he shirks. A voter's incentives at histories inside $\mathcal{H}_{\text{Pass}}$ are satisfied because she receives effort with probability 1 whatever she does, so she prefers not to bear the nonnegative replacement cost. For her incentives at histories $h \notin \mathcal{H}_{\text{Pass}}$, her payoff from not replacing is $\pi(h) \leq \pi_0$, and her payoff from replacing is $1 - c \geq \pi_0$. Finally, the Bayesian property is straightforward: $\pi(\emptyset) = \pi_0$, and at every history h with $\sigma_V(h) < 1$, we have $\sigma_P(h) = 1$ and $\pi(h, s) = \pi(h)$ for every $s \in S$. \square

The next lemma shows how to construct an equilibrium that does not attain eventual full effort, given Condition FEI and a small replacement cost. This equilibrium has every incumbent being eventually replaced and a positive probability of shirking from new incumbents.

Lemma 14. *If $c < 1 - \pi_0$ and Condition FEI holds, then some equilibrium does not attain eventual full effort.*

³⁰ The construction is “nearly” Markovian: play depends only on beliefs and the most recent signal. In fact, because beliefs do not change on path, the equilibrium is “simple” in the sense of Banks and Sundaram (1993). Our construction in the proof of Lemma 14 is also nearly Markovian, but because beliefs do change on path it is not simple in their sense.

Proof. We begin with some useful notation and calculations. First, by Lemma 12, some $\bar{v} \in (0, 1 - \kappa)$ and $\lambda \in \left[1, \max_{s \in S} \frac{f_0(s)}{f_1(s)}\right)$ exist such that, letting

$$S^* := \left\{s \in S : \frac{f_0(s)}{f_1(s)} \leq \lambda\right\}, \text{ and } f_a^* := \sum_{s \in S^*} f_a(s) \in [0, 1) \text{ for } a \in \{0, 1\},$$

we have $\bar{v} = (1 - \delta)(1 - \kappa) + \delta f_1^* \bar{v} \geq (1 - \delta)1 + \delta f_0^* \bar{v}$. The previous inequality tells us $f_1^* > f_0^*$ so that $\tilde{v} := \bar{v} - \frac{1 - \delta}{\delta} \frac{\kappa}{f_1^* - f_0^*} < \bar{v}$. Observe also that

$$\begin{aligned} \delta(f_1^* - f_0^*)\tilde{v} &= \delta(f_1^* - f_0^*)\bar{v} - (1 - \delta)\kappa \\ &= [(1 - \delta)(1 - \kappa) + \delta f_1^* \bar{v}] - [(1 - \delta)1 + \delta f_0^* \bar{v}] \\ &\geq 0, \end{aligned}$$

so that $\tilde{v} \geq 0$. Now, defining $\hat{v} := (1 - \delta)(1 - \kappa) + \delta [f_1^* \bar{v} + (1 - f_1^*)\tilde{v}]$, note that \hat{v} is a proper weighted average of the three terms $\tilde{v} < \bar{v} < 1 - \kappa$, and so is strictly higher than \tilde{v} . Having established that $0 \leq \tilde{v} < \hat{v}$, we thus know that $x := 1 - \frac{\tilde{v}}{\hat{v}} \in (0, 1]$. Moreover,

$$\begin{aligned} &(1 - \delta) + \delta [f_0^* \bar{v} + (1 - f_0^*)\tilde{v}] - \hat{v} \\ &= (1 - \delta) [1 - (1 - \kappa)] + \delta \{[f_0^* \bar{v} + (1 - f_0^*)\tilde{v}] - [f_1^* \bar{v} + (1 - f_1^*)\tilde{v}]\} \\ &= (1 - \delta)\kappa + \delta(f_0^* - f_1^*)(\bar{v} - \tilde{v}) \\ &= 0. \end{aligned}$$

In summary, given the above calculations and substituting in $\tilde{v} = (1 - x)\hat{v}$, we have that

$$\begin{aligned} x &\in (0, 1], \\ \bar{v} &= (1 - \delta)(1 - \kappa) + \delta f_1^* \bar{v} \\ &\geq (1 - \delta)1 + \delta f_0^* \bar{v}, \text{ and} \\ \hat{v} &= (1 - \delta)(1 - \kappa) + \delta [f_1^* \bar{v} + (1 - f_1^*)(1 - x)\hat{v}] \\ &= (1 - \delta) + \delta [f_0^* \bar{v} + (1 - f_0^*)(1 - x)\hat{v}]. \end{aligned}$$

We will use these calculations to construct an equilibrium of the desired form.

To construct such an equilibrium, fix some $a_0 \in \left(\frac{c}{1 - \pi_0}, 1\right)$, which exists by hypothesis. We will next describe a strategy profile and belief map that depends on a_0 (though we will suppress this dependence in our notation). We will then argue that the profile does not

generate EFE, and that for a_0 close enough to 1 it is an equilibrium.

To describe the strategy profile and belief map, partition histories into three categories.³¹ Say a history is a *first-regime* history if no signal in S^* has ever (with this incumbent) appeared; say a history is a *second-regime* history if it is not a first-regime history, and every signal since the first S^* signal has been in S^* ; and say a history is a *third-regime* history otherwise. Consider strategies and beliefs $(\sigma_P, \sigma_V, \pi)$ as follows:

- In the first regime: Both voter and incumbent mix. Specifically, the voter always chooses $\sigma_V(h) = x$. We recursively construct politician behavior and voter beliefs at a history h in this regime as follows. Start by setting $\sigma_P(\emptyset) = a_0$ and $\pi(\emptyset) = \pi_0$. Then consider any $h = (h_-, s)$ in this regime with $h_- \in \mathcal{H}$ and $s \in S$. Given $\sigma_P(h_-) > 0$, which holds by induction, set $\pi(h) = \beta_{\sigma_P(h_-)}(\pi(h_-)|s) \in (0, \pi(h_-))$, as signal s is strictly bad news (recall $\lambda \geq 1$ and $s \notin S^*$). Set $\sigma_P(h) \in [a_0, 1)$ so that $\pi(h) + [1 - \pi(h)]\sigma_P(h) = \pi_0 + (1 - \pi_0)a_0 - c$; doing so is possible because the definition of a_0 ensures $\pi_0 + (1 - \pi_0)a_0 - c > \pi_0 \geq \pi(h)$.
- In the second regime: The voter retains the incumbent, and the incumbent works. That is, the voter always chooses $\sigma_V(h) = 0$, and the politician always chooses $\sigma_P(h) = 1$. Write $h = (h_-, s)$ for some $h_- \in \mathcal{H}$ and $s \in S$. As $\sigma_P(h_-) > 0$ by construction, set $\pi(h) = \beta_{\sigma_P(h_-)}(\pi(h_-)|s)$.
- In the third regime: The voter replaces the incumbent, and the incumbent shirks. That is, the voter always chooses $\sigma_V(h) = 1$, and the politician always chooses $\sigma_P(h) = 0$. Write $h = (h_-, s)$ for some $h_- \in \mathcal{H}$ and $s \in S$. If h_- is a second-regime history, then $\sigma_P(h_-) = 1$, so set $\pi(h) = \beta_{\sigma_P(h_-)}(\pi(h_-)|s) = \pi(h_-)$. If h_- is not a second-regime history, then $\sigma_V(h_-) = 1$ by construction, so we may set $\pi(h) = 0$.³²

Below, we will argue that the constructed $(\sigma_P, \sigma_V, \pi)$ is an equilibrium for appropriate choice of the parameter a_0 . Before doing so, we point out that the lemma would follow if it is. Specifically, EFE fails because every officeholder shirks with positive probability at his initial history, and almost surely every officeholder is eventually replaced. To confirm the latter point, note that (i) if an officeholder is not replaced at a first-regime history, then there is a probability bounded away from zero (at least $[\pi_0 + (1 - \pi_0)a_0 - c]f_1^* > \pi_0 f_1^*$) of transitioning to the second regime in the next period; (ii) from the second regime, there is a constant positive probability of transitioning to the third regime in the next period; and (iii) the officeholder is replaced at every third-regime history. So the Borel-Cantelli lemma

³¹ Recall that [Figure 1](#) depicts the strategy profile for the special case of binary signals.

³² Here is the only other place (besides that noted in [footnote 28](#)) where we use the flexibility of off-path beliefs afforded by *weak* PBE.

implies almost sure eventual replacement of any officeholder.

We now confirm that $(\sigma_P, \sigma_V, \pi)$ is an equilibrium for some a_0 . The Bayesian property is immediate from the construction. Next, consider politician incentives. These are trivial in the third regime: it is optimal for him to shirk since he will be replaced immediately no matter which signal he generates. His incentives to work in the second regime follow directly from the definition of \bar{v} and S^* . For the first regime, we can use the fact that

$$\hat{v} = (1 - \delta)(1 - \kappa) + \delta [f_1^* \bar{v} + (1 - f_1^*)(1 - x)\hat{v}] = (1 - \delta) + \delta [f_0^* \bar{v} + (1 - f_0^*)(1 - x)\hat{v}].$$

These equalities tell us that his continuation value from any first-regime history is \hat{v} , and that he is indifferent between his two effort choices there. Now, we turn to voter incentives. In the second regime, the voter gets payoff 1 from keeping the incumbent, and so does so optimally. In the first regime, the voter is indifferent because $\pi(h) + [1 - \pi(h)] \sigma_P(h) = \pi_0 + (1 - \pi_0)a_0 - c$.

All that remains is to verify voter optimality in the third regime. Because the voter's outside option is $\pi_0 + (1 - \pi_0)a_0$ and the opportunistic type always shirks in the third regime, it is optimal for the voter to replace at a third-regime history h if and only if

$$\pi(h) \leq \pi_0 + (1 - \pi_0)a_0 - c. \quad (\text{A.2})$$

Toward that inequality, let us show that some $\bar{s} \in S$ has

$$\pi(h) \leq \beta_{a_0}(\pi_0 | \bar{s}). \quad (\text{A.3})$$

First, if the previous history was not a second-regime history, then $\pi(h) = 0$ and so (A.3) holds for any $\bar{s} \in S$. So consider the complementary case, in which $h = (h_2, s_2)$ for some second-regime history h_2 and some $s_2 \in S$. By definition of the second regime, there exists some first-regime history h_1 (which comprises only signals from $S \setminus S^*$), some signal $s_1 \in S^*$, and some (possibly empty) string h_+ of signals from S^* such that $h_2 = (h_1, s_1, h_+)$. Every history \tilde{h} between³³ (h_1, s_1) and h_2 has $\sigma_P(\tilde{h}) = 1$, and so $\pi(\tilde{h}, \tilde{s}) = \pi(\tilde{h})$ for every signal \tilde{s} , implying $\pi(h) = \pi(h_2) = \pi(h_1, s_1)$. Since every $s \in S \setminus S^*$ has $f_1(s) < f_0(s)$, it follows that $0 < \beta_a(\pi | s) \leq \pi$ for every $\pi \in (0, 1]$ and $a \in [0, 1]$; hence, by induction $\pi(h_1) \leq \pi(\emptyset) = \pi_0$. Hence, by construction, the first-regime history h_1 has $\sigma_P(h_1) \geq \sigma_P(\emptyset) = a_0$. Therefore:

- If $f_1(s_1) > f_0(s_1)$, we have $\pi(h_1, s_1) = \beta_{\sigma_P(h_1)}(\pi(h_1) | s_1) \leq \beta_{a_0}(\pi(h_1) | s_1) \leq \beta_{a_0}(\pi_0 | s_1)$. So (A.3) holds for $\bar{s} = s_1$.

³³That is, for any history \tilde{h} that weakly succeeds (h_1, s_1) and weakly precedes h_2 .

- If $f_1(s_1) \leq f_0(s_1)$, we have $\pi(h_1, s_1) = \beta_{\sigma_P(h_1)}(\pi(h_1) | s_1) \leq \pi(h_1) \leq \pi_0$. So (A.3) holds for any $\bar{s} \in S$ with $f_1(\bar{s}) \geq f_0(\bar{s})$.

Given that (A.3) holds for some $\bar{s} \in S$, inequality (A.2) follows if

$$\max_{s \in S} \beta_{a_0}(\pi_0 | s) \leq \pi_0 + (1 - \pi_0)a_0 - c.$$

This inequality holds when a_0 is close enough to 1 because as $a_0 \rightarrow 1$, the right-hand side converges to $1 - c$ while the left-hand side converges to $\pi_0 < 1 - c$. \square

A.2.5. The Equivalence Theorem

Proof of Theorem 1. Recall that the theorem assumes our maintained assumption in the main text that $c < \min\{\pi_0, 1 - \pi_0\}$.

First, suppose Condition FEI fails. Proposition 3 then tells us that every equilibrium attains eventual full effort. Yet no equilibrium attains full effort because politician incentive-compatibility would be violated (Lemma 1).

Next, suppose Condition FEI holds. Lemma 13 then implies that some equilibrium attains full effort; and Lemma 14 demonstrates that not every equilibrium attains eventual full effort. \square

A.3. Other Results

A.3.1. Replacement Despite a Favorable Reputation

Proof of Proposition 2. Consider any equilibrium. We want to show it either attains eventual full effort or has $\mathbb{P}\{\boldsymbol{\pi}_{\tau_r} > \pi_0 \text{ and } \mathbf{r}_\infty = 1\} > 0$. We consider three cases.

First, if $\mathbb{P}\{\tau_r = \infty\} > 0$, then the Borel-Cantelli Lemma and symmetry of the equilibrium tell us that, with probability 1, some officeholder is never replaced. But since (by Lemma 7) that officeholder is a good type with conditional probability 1, the equilibrium attains eventual full effort.

Second, if $\mathbb{P}\{\boldsymbol{\pi}_{\tau_r} = \pi_0\} = 1$, then (because $f_1 \neq f_0$ and beliefs are a martingale) signals are uninformative of effort on path, so the equilibrium attains full effort, hence eventual full effort.

Consider now the remaining case in which $\tau_r < \infty$ a.s. and $\mathbb{P}\{\boldsymbol{\pi}_{\tau_r} = \pi_0\} < 1$. Then, Doob's optional stopping theorem implies $\mathbb{E}[\boldsymbol{\pi}_{\tau_r}] = \pi_0$. Hence, $\mathbb{P}\{\boldsymbol{\pi}_{\tau_r} > \pi_0\} > 0$, delivering the result. \square

A.3.2. Voter Welfare in the Long Run

Proof of Corollary 1. Suppose Condition FEI fails, and that $c < \pi_0$. Fix an equilibrium, and let B_t denote the event that the time- t incumbent is a commitment type who will never be replaced. Observe $B_0 \subseteq B_1 \subseteq \dots$ by definition, and Proposition 3 tells us $\mathbb{P}(\bigcup_{t=0}^{\infty} B_t) = 1$. Continuity along chains then implies $\mathbb{P}(B_t) \rightarrow 1$ as $t \rightarrow \infty$. But the time- t voter's ex-ante expected payoff is in $[\mathbb{P}(B_t), 1]$, delivering the corollary. \square

A.3.3. Voters' Outside Option

The following lemma bounds the growth on incumbent reputation, and hence the growth rate of a voter's value from keeping the incumbent, for a given lower bound on the opportunistic officeholder's expected effort.

Lemma 15. *For any $t \in \mathbb{Z}_{\geq 0}$ and $\eta \in (0, 1)$, the map $\bar{\beta}_{\eta}^t$ is increasing, and any $\pi \in (0, 1)$ has*

$$\bar{\beta}_{\eta}^t(\pi) + [1 - \bar{\beta}_{\eta}^t(\pi)] \eta \leq \frac{\pi + (1 - \pi)\eta^{t+1}}{\pi + (1 - \pi)\eta^t}.$$

Moreover, the right-hand side of the above inequality is increasing in t .

Proof. In the expression defining $\beta_a(\pi|s)$, the numerator does not depend on a , whereas the denominator is affine (hence monotone) in it, so that $\beta_a(\pi|s)$ is monotone in a too. Therefore,

$$\max_{a \in [\eta, 1]} \beta_a(\pi|s) = \max \{\beta_{\eta}(\pi|s), \beta_1(\pi|s)\} = \max \{\beta_{\eta}(\pi|s), \pi\}.$$

For each $s \in S$, observe that

$$\beta_{\eta}(\pi|s) = \frac{\pi}{\pi + (1 - \pi) \left[\eta + (1 - \eta) \frac{f_0(s)}{f_1(s)} \right]},$$

which is strictly decreasing in $\frac{f_0(s)}{f_1(s)}$. Letting $\lambda := \min_{s \in S} \frac{f_0(s)}{f_1(s)} \in [0, 1]$, it follows that

$$\bar{\beta}_{\eta}(\pi) = \max_{s \in S} \max \{\beta_{\eta}(\pi|s), \pi\} = \max \left\{ \max_{s \in S} \beta_{\eta}(\pi|s), \pi \right\} = \frac{\pi}{\pi + (1 - \pi) [\eta + (1 - \eta)\lambda]} \in (0, 1).$$

This quantity is increasing in π because $\eta + (1 - \eta)\lambda \leq 1$, and so the composition $\bar{\beta}_{\eta}^t$ is also increasing.

We now pursue the upper bound. That $\lambda \geq 0$ tells us $\bar{\beta}_{\eta}(\pi) \leq \frac{\pi}{\pi + (1 - \pi)\eta}$, implying

$$\frac{\bar{\beta}_{\eta}(\pi)}{1 - \bar{\beta}_{\eta}(\pi)} \leq \frac{\pi}{(1 - \pi)\eta}.$$

Because $\bar{\beta}_\eta$ is increasing, it follows from induction that $\frac{\bar{\beta}_\eta^t(\pi)}{1-\bar{\beta}_\eta^t(\pi)} \leq \frac{\pi}{(1-\pi)\eta^t}$ for every $t \in \mathbb{Z}_{\geq 0}$. Rearranging this inequality yields $\bar{\beta}_\eta^t(\pi) \leq \frac{\pi}{\pi+(1-\pi)\eta^t}$, and so (since $\eta \leq 1$)

$$\begin{aligned}\bar{\beta}_\eta^t(\pi) + [1 - \bar{\beta}_\eta^t(\pi)]\eta &\leq \frac{\pi}{\pi + (1-\pi)\eta^t} + \frac{(1-\pi)\eta^t}{\pi + (1-\pi)\eta^t}\eta \\ &= \frac{\pi + (1-\pi)\eta^{t+1}}{\pi + (1-\pi)\eta^t} \in (0, 1).\end{aligned}$$

All that remains now is to see that this bound increases with t . And indeed, the bound is equal to

$$\frac{\frac{\pi}{1-\pi} + \eta \cdot \eta^t}{\frac{\pi}{1-\pi} + \eta^t},$$

which (because $0 < \eta < 1$) is strictly decreasing in η^t and so strictly increasing in t . \square

The next lemma derives a parametric upper bound on the voters' outside option.

Lemma 16. *Suppose Condition FEI fails, and let T be as delivered by Lemma 2. In any equilibrium, the voters' outside option is below*

$$c + \inf_{\eta \in (0, 1)} \frac{\pi_0 + (1-\pi_0)\eta^{T+1}}{\pi_0 + (1-\pi_0)\eta^T}.$$

Proof. Fixing an equilibrium, assume for a contradiction that the voters' outside option u_0 has $u_0 - c > \frac{\pi_0 + (1-\pi_0)\eta^{T+1}}{\pi_0 + (1-\pi_0)\eta^T}$ for some $\eta \in (0, 1)$.

Let $h_0 := \emptyset$, and then let $s_0, \dots, s_{T-1} \in S$ and $h_1, \dots, h_T \in \mathcal{H}$ be as in Lemma 3.

Now, let us show by induction that every $t \in \{0, \dots, T-1\}$ has both $\sigma_P(h_t) \geq \eta$ and $\pi(h_t) \leq \bar{\beta}_\eta^t(\pi_0)$. For the base case ($t = 0$), observe that $\bar{\beta}_\eta^0(\pi_0) = \pi_0 = \pi(h_0)$, while

$$\pi_0 + (1-\pi_0)\sigma_P(h_0) = u_0 > c + \frac{\pi_0 + (1-\pi_0)\eta^T\eta}{\pi_0 + (1-\pi_0)\eta^T} \geq 0 + \frac{\pi_0 + (1-\pi_0)\eta \cdot 1}{\pi_0 + (1-\pi_0)1} = \pi_0 + (1-\pi_0)\eta \cdot 1,$$

so that $\sigma_P(h_0) \geq \eta$.

Let us turn next to the inductive step. Suppose $t \in \{0, \dots, T-2\}$ has both $\sigma_P(h_t) \geq \eta$ and $\pi(h_t) \leq \bar{\beta}_\eta^t(\pi_0)$; we aim to show $t+1$ enjoys the same properties. First, we have

$$\pi(h_{t+1}) \leq \bar{\beta}_\eta(\pi(h_t)) \leq \bar{\beta}_\eta(\bar{\beta}_\eta^t(\pi_0)) = \bar{\beta}_\eta^{t+1}(\pi_0),$$

where the first inequality holds because $\sigma_P(h_t) \geq \eta$, and the second inequality holds (given Lemma 15) because $\pi(h_t) \leq \bar{\beta}_\eta^t(\pi_0)$. Second, observe that $\sigma_P(h_t) \geq \eta > 0$ would not be

incentive compatible if $[1 - \sigma_V(h_t, s)] v(h_t, s)$ were zero for every $s \in S$. Hence, given the definition of s_t and h_{t+1} , we necessarily have $\sigma_V(h_{t+1}) < 1$. Voter incentives then imply

$$\begin{aligned} u_0 - c &\leq \pi(h_{t+1}) + [1 - \pi(h_{t+1})] \sigma_P(h_{t+1}) \\ &\leq \bar{\beta}_\eta^{t+1}(\pi_0) + [1 - \bar{\beta}_\eta^{t+1}(\pi_0)] \sigma_P(h_{t+1}). \end{aligned}$$

Meanwhile, by hypothesis and given [Lemma 15](#), we have

$$\begin{aligned} u_0 - c &> \frac{\pi_0 + (1 - \pi_0)\eta^{T+1}}{\pi_0 + (1 - \pi_0)\eta^T} \\ &\geq \bar{\beta}_\eta^{t+1}(\pi_0) + [1 - \bar{\beta}_\eta^{t+1}(\pi_0)] \eta. \end{aligned}$$

Combining the two inequality chains yields $\sigma_P(h_{t+1}) \geq \eta$, completing the inductive step.

That $\sigma_P(h_t) > 0$ for every $t \in \{0, \dots, T-1\}$ contradicts [Lemma 3](#). \square

To prove [Proposition 4](#), we show that the bound of [Lemma 16](#) can be made arbitrarily small if the prior and replacement costs are both small.³⁴

Proof of Proposition 4. Let T be as given by [Lemma 2](#). Because

$$\begin{aligned} 0 &\leq c + \inf_{\eta \in (0,1)} \frac{\pi_0 + (1 - \pi_0)\eta^{T+1}}{\pi_0 + (1 - \pi_0)\eta^T} \\ &\leq c + \frac{\pi_0 + (1 - \pi_0)\pi_0^{(T+1)/(T+2)}}{\pi_0 + (1 - \pi_0)\pi_0^{T/(T+2)}} \\ &\rightarrow 0 \text{ as } (c, \pi_0) \rightarrow (0, 0), \end{aligned}$$

it follows that $c + \inf_{\eta \in (0,1)} \frac{\pi_0 + (1 - \pi_0)\eta^{T+1}}{\pi_0 + (1 - \pi_0)\eta^T}$ converges to zero as c and π_0 do. The proposition then follows from [Lemma 16](#). \square

Finally, although not stated in the main text, we record one further result that extends our main equivalence result to the voters' outside option being bounded away from first best, so long as the replacement cost is sufficiently small.

Corollary 2. *For any $(\kappa, \delta, f, \pi_0)$, there exists $\bar{c} > 0$ such that, if $c < \bar{c}$, then the following are equivalent:*

³⁴ Although [Section 2](#) assumes $c < \min\{\pi_0, 1 - \pi_0\}$, the restriction is evidently not needed for [Proposition 4](#): if [Condition FEI](#) fails and c and π_0 are both very small, then the outside option is very small in all equilibria, irrespective of the relationship between c and π_0 .

1. All equilibria attain eventual full effort;
2. No equilibrium attains full effort;
3. *Condition FEI* fails.
4. Across all equilibria, the voters' outside option is bounded away from 1.

Proof. Let $\hat{c} := \min\{\pi_0, 1 - \pi_0\}$. If (κ, δ, f) satisfy *Condition FEI*, then let $\bar{c} := \hat{c}$. If (κ, δ, f) violate *Condition FEI*, then let $\bar{c} \in (0, \hat{c})$ be such that $\bar{c} + \inf_{\eta \in (0, 1)} \frac{\pi_0 + (1 - \pi_0)\eta^{T+1}}{\pi_0 + (1 - \pi_0)\eta^T} < 1$, where T is as given by [Lemma 2](#). Note that such \bar{c} exists because $\frac{\pi_0 + (1 - \pi_0)\eta^{T+1}}{\pi_0 + (1 - \pi_0)\eta^T} < 1$ for any given $\eta \in (0, 1)$.

[Theorem 1](#) tells us conditions 1, 2, and 3 are all equivalent; and obviously condition 4 implies 2. Lastly, [Lemma 16](#) tells us condition 3 implies 4. \square

Symbol	Description	Defined In
<i>Histories, Strategies, and Values</i>		
\mathcal{H}	Set of all (public) career histories	§2
$f_{\tilde{a}}(s)$	Probability of signal s given \tilde{a} expected effort: equals $\tilde{a}f_1(s) + (1 - \tilde{a})f_0(s)$	§2
$\sigma_P(h)$	Probability an opportunistic incumbent exerts effort at history h	§2
$\sigma_V(h)$	Probability the voter replaces the incumbent at history h	§2
$\pi(h)$	Incumbent's reputation (belief that he is the good type) at history h	§2
$\alpha(h)$	Expected effort at history h : equals $\pi(h) + [1 - \pi(h)]\sigma_P(h)$	§2
$v(h)$	Opportunistic incumbent's continuation value at history h	§2
u_0	Voters' outside option: equals $\pi_0 + (1 - \pi_0)\sigma_P(\emptyset)$	§3.4
<i>Belief Operators</i>		
$\beta_a(\pi s)$	Bayesian posterior from belief π , opportunistic type's effort a , and signal s	§A.2.2
$\bar{\beta}_\eta(\pi)$	Highest possible Bayesian posterior from belief π and opportunistic type's effort at least η	§A.2.2
$\bar{\beta}_\eta^t$	The t -fold composition of the operator $\bar{\beta}_\eta$	§A.2.2
<i>Random Variables</i>		
\mathbf{r}_t	Indicator variable for replacement: equals 1 if first ^a officeholder has been replaced by time t	§A.2.3
τ_r	Replacement time of first officeholder: equals $\inf\{t : \mathbf{r}_t = 1\} \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$	§A.2.3
π_t	First officeholder's reputation at time t (stopped at τ_r)	§A.2.3
$\tilde{\mathbf{a}}_t$	Officeholder's expected effort at time t conditional on being opportunistic	§A.2.3
$\pi_\infty, \mathbf{r}_\infty, \tilde{\mathbf{a}}_\infty$	Almost-sure limits of the stochastic processes $\pi_t, \mathbf{r}_t, \tilde{\mathbf{a}}_t$	§A.2.3

^a By equilibrium symmetry, an analog applies to other officeholders for this and other relevant random variables

Table 2: Summary of Key Notation used in the Proof Appendix

References

ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.

ACHARYA, A., E. LIPNOWSKI, AND J. RAMOS (2025): “Political Accountability under Moral Hazard,” *American Journal of Political Science*, 69, 641–652.

AKHTARI, M., D. MOREIRA, AND L. TRUCCO (2022): “Political Turnover, Bureaucratic Turnover, and the Quality of Public Services,” *American Economic Review*, 112, 442–493.

ALESINA, A. AND G. TABELLINI (2007): “Bureaucrats or Politicians? Part I: A Single Policy Task,” *American Economic Review*, 97, 169—179.

ALI, A. (2025): “Learnng Via Delegation,” Work in progress.

ANESI, V. AND P. BUISERET (2022): “Making Elections Work: Accountability with Selection and Control,” *American Economic Journal: Microeconomics*, 14, 616–644.

ASHWORTH, S. (2012): “Electoral Accountability: Recent Theoretical and Empirical Work,” *Annual Review of Political Science*, 15, 183–201.

BANKS, J. S. AND R. K. SUNDARAM (1993): “Adverse Selection and Moral Hazard in a Repeated Elections Model,” in *Political Economy: Institutions, Competition and Representation*, ed. by W. A. Barnett, M. Hinich, and N. Schofield, Cambridge University Press, chap. 12, 295–311.

BARRO, R. J. (1973): “The Control of Politicians: An Economic Model,” *Public Choice*, 14, 19–42.

BERNHARDT, D., S. DUBEY, AND E. HUGHSON (2004): “Term Limits and Pork Barrel Politics,” *Journal of Public Economics*, 88, 2383–2422.

BESLEY, T. (2005): “Political Selection,” *Journal of Economic Perspectives*, 19, 43–60.

——— (2006): *Principled Agents?: The Political Economy of Good Government*, Oxford: Oxford University Press.

BESLEY, T. AND M. GHATAK (2005): “Competition and Incentives with Motivated Agents,” *American Economic Review*, 95, 616–636.

BESLEY, T. AND M. SMART (2007): “Fiscal Restraints and Voter Welfare,” *Journal of Public Economics*, 91, 755–773.

BLACKWELL, D. (1953): “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 24, 265–272.

CHAN, J. (2000): “On the Non-Existence of Reputation Effects in Two-Person Infinitely-Repeated Games,” Unpublished.

CRIPPS, M., G. MAILATH, AND L. SAMUELSON (2004): “Imperfect Monitoring and Impermanent Reputations,” *Econometrica*, 72, 407–432.

CRIPPS, M. AND J. THOMAS (1997): “Reputation and Perfection in Repeated Common Interest Games,” *Games and Economic Behavior*, 18, 141–158.

DEB, J. AND J. FANNING (2025): “Reputation and Competitive Selection,” Working paper.

DIERMEIER, D., M. KEANE, AND A. MERLO (2005): “A Political Economy Model of Congressional Careers,” *American Economic Review*, 95, 347—373.

DUGGAN, J. (2000): “Repeated Elections with Asymmetric Information,” *Economics & Politics*, 12, 109–135.

DUGGAN, J. AND J. G. FORAND (2025): “Accountability in Markovian Elections,” *Games and Economic Behavior*, 151, 183–217.

DUGGAN, J. AND C. MARTINELLI (2017): “The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness,” *Journal of Economic Literature*, 55, 916–84.

EECKHOUT, J. (2006): “Minorities and Endogenous Segregation,” *Review of Economic Studies*, 73, 31–53.

EKMEKCI, M., O. GOSSNER, AND A. WILSON (2012): “Impermanent Types and Permanent Reputations,” *Journal of Economic Theory*, 147, 162–178.

FEARON, J. D. (1999): “Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance,” in *Democracy, Accountability, and Representation*, ed. by A. Przeworski, S. C. Stokes, and B. Manin, Cambridge University Press, Cambridge Studies in the Theory of Democracy, 55–97.

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

FUDENBERG, D. AND D. LEVINE (1983): “Subgame-perfect Equilibria of Finite-and Infinite-horizon games,” *Journal of Economic Theory*, 31, 251–268.

FUDENBERG, D. AND D. K. LEVINE (1989): “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 57, 759–778.

——— (1992): “Maintaining a Reputation when Strategies are Imperfectly Observed,” *Review of Economic Studies*, 59, 561–579.

FUDENBERG, D., D. K. LEVINE, AND J. TIROLE (1987): “Incomplete Information Bargaining with Outside Opportunities,” *Quarterly Journal of Economics*, 102, 37–50.

FUDENBERG, D. AND E. MASKIN (1986): “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533–554.

GARRETT, D. F. AND A. PAVAN (2012): “Managerial Turnover in a Changing World,” *Journal of Political Economy*, 120, 879–925.

GHOSH, P. AND D. RAY (1996): “Cooperation in Community Interaction without Information Flows,” *Review of Economic Studies*, 63, 491–519.

HÖRNER, J. (2002): “Reputation and Competition,” *American Economic Review*, 92, 644–663.

HUBER, J. D. AND M. M. TING (2021): “Civil Service and Patronage in Bureaucracies,” *Journal of Politics*, 83, 902–916.

KARTIK, N. AND R. VAN WEELDEN (2019): “Reputation Effects and Incumbency (Dis)Advantage,” *Quarterly Journal of Political Science*, 14, 131–157.

KREPS, D. AND R. WILSON (1982): “Sequential Equilibria,” *Econometrica*, 863–894.

LIU, Q. AND A. SKRZYPACZ (2014): “Limited Records and Reputation Bubbles,” *Journal of Economic Theory*, 151, 2–29.

MAILATH, G. AND L. SAMUELSON (2001): “Who Wants A Good Reputation?” *Review of Economic Studies*, 68, 415–441.

MAILATH, G. J. (2018): *Modeling Strategic Behavior: A Graduate Introduction to Game Theory and Mechanism Design*, World Scientific.

MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, Oxford University Press.

MILL, J. S. (1859): *On Liberty*, London: John W. Parker and Son.

MYERSON, R. B. (2006): “Federalism and Incentives for Success of Democracy,” *Quarterly Journal of Political Science*, 3–23.

PADRÓ-I MIQUEL, G. AND J. M. SNYDER JR. (2006): “Legislative Effectiveness and Legislative Careers,” *Legislative Studies Quarterly*, 31, 347–381.

PEI, H. (2024): “Reputation Effects under Short Memories,” *Journal of Political Economy*, 132, 3421–3460.

PHELAN, C. (2006): “Public Trust and Government Betrayal,” *Journal of Economic Theory*, 130, 27–43.

PRENDERGAST, C. (2007): “The Motivation and Bias of Bureaucrats,” *American Economic Review*, 97, 180–196.

SCHWABE, R. (2009): “Reputation and Accountability in Repeated Elections,” Unpublished.

SPILLER, P. T. AND S. URBIZTONDO (1994): “Political Appointees vs. Career Civil Servants: A Multiple Principals Theory of Political Bureaucracies,” *European Journal of Political Economy*, 10, 465–497.

TADELIS, S. (2002): “The Market for Reputations as an Incentive Mechanism,” *Journal of Political Economy*, 110, 854–882.

VAN WEELDEN, R. (2013): “Candidates, Credibility, and Re-election Incentives,” *Review of Economic Studies*, 80, 1622–1651.

WATSON, J. (1999): “Starting Small and Renegotiation,” *Journal of Economic Theory*, 85, 52–90.

WISEMAN, T. (2008): “Reputation and Impermanent Types,” *Journal of Economic Theory*, 62, 190–210.