Fire At Will:

Civil Service Reform and Bureaucratic Accountability

Dan Alexander* Ruth Holloway†

January 16, 2026

Abstract

Recently an effort has been underway to remove the tenure protections of civil servants in the United States federal government. If successful a presidential administration could fire bureaucrats at-will – rather than requiring cause – and replace them with co-partisans. This approach to governance is not new, in fact, and is strongly reminiscent of the conditions that led, some 150 years ago, to the tenure protections civil servants have today. In this paper, we model the principal-agent relationship between administration and bureaucracy under both at-will and just-cause terms of employment. We find that the greater the proportion of cross-partisans with which an administration must work, the better welfare is under just-cause relative to at-will terms.

*University of Rochester, dan.alexander@rochester.edu

†University of Rochester, rhollow6@UR.rochester.edu

Introduction

The extent to which elected officials should exercise control over the federal bureaucracy has been contested for nearly 200 years. Civil service reforms were taken up in earnest after the conclusion of the Civil War, prior to which federal bureaucrats were instated via a spoils system wherein the jobs of political losers were reallocated to political loyalists whenever the incumbent party lost control. After the implementation of the Pendleton Act and subsequent reforms, it became accepted that at least some subset of agents of the state should not only be hired based on merit (which at one point entailed civil service exams) but that career civil servants have since been insulated for wrongful or undue partisan firing.

Recent moves by the incumbent administration to reclassify a swath of career civil servants as political appointees (referred to colloquially as Schedule F) has made it relevant once again in the eyes of the public. The Trump White House, however, is hardly the first to instate sweeping changes within the federal bureaucracy toward political ends, even within recent history – it has merely been the most overt and aggressive about using the apparatus toward such ends. Take, for example, reforms initiated under former U.S. President Bill Clinton. The ongoing changes are largely consistent with the actions taken by the Clinton administration, but are being carried out via executive order, which remains an unprecedented approach to policy-making.

Prior to the passage of the Pendleton Act in 1883, it was widely accepted that control of the bureaucracy was essential for establishing political dominance in the U.S. federal government. The spoils system enabled de facto one-party rule, first by Democrats in the antebellum period and then by Republicans for roughly two decades following the Civil War. Late 19th and early 20th century civil service reforms represented a consensus about the desirability of a politically neutral bureaucracy. Yet nearly 150 years later, the question of whether such neutrality is feasible remains unresolved. The Trump administration's aggressive approach to staffing the federal bureaucracy echoes the sentiment expressed by Reagan White House advisor Lyn Nofziger, who declared that "as far as I'm concerned, anyone who supported Reagan is competent." Part of the Heritage Foundation's "Project 2025" was the herculean task of preemptively screening thousands of individuals for loyalty to President Trump, a systematic effort to implement a hiring and firing regime governed primarily by political allegiance.

With an institutional setup that closely mirrors the spoils system taking shape, we revisit the question of whether and when political control of bureaucracy improves bureaucratic performance or undermines the stated goals of creating a careerist class in the first place – consistent policy implementation, preservation of institutional memory, and public service delivery that is effective an untainted by patronage concerns.

Standard career concern models, which forms the basis for our analysis, suggest that agents make effort choices in order to increase the likelihood of receiving future payoffs. This mechanism, however, rests on the assumption that effort is, in fact, influential. When hiring and firing is carried out on a primarily ideological basis, the relationship between retention and effort choice is broken and effort becomes fruitless for unaligned agents. Thus, when such agents are present, civil service protections are essential if principals wish to increase the likelihood of project success via effort choice.

Our core inquiry examines the circumstances under which civil service protections can improve public service outcomes, both through the mechanisms of increasing bureaucratic effort and improving agent selection. We develop a two-period career concerns model in which an administration must decide whether to retain or dismiss a bureaucrat based on project success and partisan alignment. Our comparison is between two institutional setups – at-will and just cause firing, respectively – and cross-partisan and copartisan bureaucrats.

Though the model's central result has to do with the rationality of exerting costly effort under varying levels of civil service protections, we also consider the circumstances under which political principals can improve outcomes by directly monitoring agents, which speaks directly to the police patrol vs fire alarm framework introduced by (McCubbins & Schwartz 1984). We find that when an administration and its agents are ideologically aligned, fire alarm oversight – that is, conditioning retention on project outcomes alone – incentivizes greater effort than formal monitoring within a just-cause framework.

We proceed as follows: first, by reviewing the literature on civil service reform, career concerns, and bureaucratic accountability. We then introduce our model, analyze equilibrium behavior across institutional regimes and partisan combinations, then derive welfare implications. Finally, we discuss the implications of our results for the present and future of the federal civil service.

Literature Review

Partisan firing is not just a normative concern – it makes bureaucratic effort less likely. In standard career-concerns models, agents exert effort to influence principals' beliefs about their competence in an attempt to increase the likelihood of retention (Holmström 1999). When firing decisions depend primarily on whether an agent is a co-partisan as opposed to project outcomes, there is no incentive for nonaligned bureaucrats to exert effort. In such contexts, it is not rational for even high types to exert effort – they will merely accrue a greater cost in the first period without raising their odds of proceeding to the second. Empirical work suggests that appointees exhibit shorter tenures and weaker performance on observable metrics, the former of which is problematic for long-term policy goals as well as institutional memory (Lewis 2008, Ornaghi 2019, Moreira & Perez 2025).

Our core inquiry is whether civil service protections can improve agent selection while avoiding the critique that institutions such as tenure give agents less incentive to exert effort (Osborne 1993, Johnson & Libecap 1994). Understanding when insulating civil servants can achieve these ends requires examining what motivation principals have for politicizing the bureaucracy. Throughout the literature, politicization is framed as a instrument of control intended to make bureaucrats more responsive to elected officials (McCubbins, Noll & Weingast 1987, Moe 1985). While politicization may make bureaucrats more responsive to presidential co-partisans and directives from the executive branch (Lowande 2019), it proves costly when enacting such directives becomes difficult for low capacity agents (Huber & McCarty 2004, Krause, Lewis & Douglas 2006). This tradeoff, however, is less difficult to resolve for principals who do not value social welfare highly relative to other ends, such as political patronage (Bendor & Meirowitz 2004, Lewis 2009). Our model formalizes this reality by putting complementary weights on two components of the administration's utility function, project success and partisan alignment.

The ability to exert control, however, requires that principals are able to gather information about agent performance. Lupia & McCubbins (1994) distinguish between three sources of such information: direct monitoring, agents' self-evaluations, and third-party signals. Elected officials are incentivized to minimize monitoring costs, even when other performance indicators are not reliable. Our model reflects this: when the principal can condition firing decisions on project success without supplying evidence of shirking, she never chooses to exert effort to monitor the agent.

Civil service protections can improve our outcomes of interest through multiple mechanisms: by protecting high quality agents from firing for partisan reasons, by incentivizing bureaucrats to specialize and develop competence (Bawn 1995, Aghion & Tirole 1997, Stephenson 2007, Ting 2011), or by enabling individuals with desirable traits to self-select into public service jobs (Delfgaauw & Dur 2007, Dal Bó, Finan & Rossi 2013, Gibbs 2020). The literature has largely focused on the latter two channels.

Where many principal-agent models (including ours) treat an agent's competence as fixed – their type, which nature assigns – Gailmard & Patty (2013) argue that expertise is endogenous to the institutional environment. Agents with sufficient public service motivation or investment in policy outcomes can become effective over time through learning and skill acquisition. In this framework, civil service protections provide a resolution to the problems identified by Fox & Weelden (2012) and Prendergast (1993), in which agents avoid actions that may produce observable mistakes even when those actions represent better solutions given the state of the world. When agents exercise their acquired expertise rather than simply catering to the preferences of their principal to avoid firing, their informed choices form the basis for good policy.

Distinct from the literature on incentives for learning or differentiating agents by type is altering the very pool from which agents are drawn. Gibbs (2020) argues that merit-based systems attract individuals with desirable traits – namely, public service motivation – than patronage-based systems. At the core of this reasoning is that low types will self-select out of the pool when civil service protections are in place, knowing that it will be burdensome to imitate their peers. Other work on influencing the quality of the hiring pool is focused mainly on pecuniary incentives, with mixed results about whether higher wages attract public service minded individuals (Delfgaauw & Dur 2007, Dal Bó, Finan & Rossi 2013).

The potential benefits of civil service protections, however, come at a cost. Insulating bureaucrats from dismissal, though meant to protect against wrongful firing, weakens the ability of elected officials to discipline wayward agents. Ujhelyi (2014) lays out the core tension: civil service protections may improve the quality of bureaucratic selection, but tenure exacerbates incentive misalignment between politicians and agents of the state by severing the ties between responsiveness and retention. This creates a principal-agent problem in which bureaucrats may pursue their own policy preferences as opposed to democratically chosen outcomes (McCubbins, Noll & Weingast 1987, Kiewiet & McCubbins 1991, Epstein & O'Halloran 1999). How principals balance competence and agency problems determines the optimal distribution of appointees versus career civil servants across the federal bureaucracy (Maskin & Tirole 2004, Alesina & Tabellini 2007).

The Model

Preliminaries

Consider two players: a civil-servant/bureaucrat, B, and an administration/appointee, A. The bureaucrat may put forth effort, $e \in \{0, 1\}$, towards the success of a project. The choice $e = 1$ implies the bureaucrat puts forth effort at a cost of κ_e, while $e = 0$ entails no effort and no cost. The administration may monitor, $m \in \{0, 1\}$, the bureaucrats' effort. The choice $m = 1$ implies the administration monitors at a cost of κ_m, while $m = 0$ entails no monitoring and no cost.

The sequence of play and information structure are as follows. Nature chooses B's quality (which persists across periods that B is retained), $\theta \in \{0, 1\}$. None of the players know the quality, but all are aware that B is a high type, $\theta = 1$, with probability π, and a low type, $\theta = 0$, with probability $1 - \pi$. The bureaucrat then chooses whether or not to exert effort. The administration does not observe B's choice of e. The administration then decides whether to monitor or not. Monitoring produces hard (i.e., verifiable) evidence of shirking only if $e = 0$. Specifically, if $m = 1$ and $e = 1$, I's investigation yields evidence of shirking, $h = 1$, with probability η. If $m = 1$ but $e = 0$, I's investigation will not yield evidence of shirking and instead produces the signal $h = 0$.

The project succeeds, $s = 1$, with probability given by the function $F : \{0, 1\} \times \{0, 1\} \longrightarrow [0, 1]$. Otherwise the project fails, $s = 0$. $F(e, \theta)$, often written $F_{e,\theta}$ below, is strictly increasing in both arguments, with $F(0, 0) \geq 0$ and $F(1, 1) \leq 1$.

All players observe the outcome of the project. If the project fails, A may present evidence B's shirking, but only if both $m = 1$ and the investigation yielded verifiable evidence of no effort, $h = 1$. After this, A decides whether or not to fire B. If B is fired, they are replaced. The steps above then repeat again in the second period. The game ends after the second period.

Partisanship At the start of the game, Nature also draws B's partisanship, $r \in \{0, 1\}$, where the probability of $r = 1$ is ρ, and the probability of $r = 0$ is $1 - \rho$. Partisanship is publicly observable, known to all at no cost. We assume the administration is of partisan type 1. If B is replaced, then a new civil servant steps in, with quality being distributed as above. The administration may, however, select a co-partisan with certainty when replacing B.

Utility representations B prefers to remain employed for the bureaucracy but, ceteris paribus, prefers not to exert effort. Let $\mathbb{I}_x$ denote the indicator function, equal to 1 if x obtains and 0 otherwise. We may then represent B's preferences in period $t \in \{1, 2\}$ by the following utility function, $u_{B,t}$: $\{0, 1\} \times \{0, 1\} \longrightarrow \mathbb{R}$, given by the rule $u_{B,t} = \mathbb{I}_{\text{employed}} - \mathbb{I}_{e=1} \cdot \kappa_e$.

A values successful public service provision as well as partisan alignment with B. Monitoring is costly for A. We may represent A's preferences in period $t \in \{1, 2\}$ by the following utility function, $u_{A,t} : \{0, 1\} \times \{0, 1\} \times \{0, 1\} \longrightarrow [0, 1]$, given by the rule $u_{A,t} = \mathbb{I}_{s=1} + \mathbb{I}_{\rho=1} \cdot \beta - \mathbb{I}_{m=1} \cdot \kappa_m$.

The steps below summarize the sequence of play and key notation for the first period of the game. Steps 1-4 constitute the governance stage. Step 5 is the administration's retention decision regarding the bureaucrat. If the administration dismissed the bureaucrat in the first period, a new bureaucrat is hired and the second period begins with step 1. If the administration retained the bureaucrat, the bureaucrat's type persists across periods and the second period begins with step 2. The game ends after the second-period.

1. Nature privately determines the bureaucrat's skill: $\theta \in \{0, 1\}$, $\pi := \Pr(\theta = 1)$; simultaneously nature publicly determines the bureaucrat's partisanship: $\rho \in \{0, 1\}$, $\varphi := \Pr(\rho = 1)$.
2. The bureaucrat privately makes an effort choice: $e \in \{0, 1\}$, $p := \Pr(e = 1)$; simultaneously the administration chooses whether to monitor the bureaucrat: $m \in \{0, 1\}$, $q := \Pr(m = 1)$.
3. Nature chooses whether the project succeeds: $s \in \{0, 1\}$, $F_{e,\theta} := \Pr(s = 1|e, \theta)$; simultaneously nature sends a signal about B's shirking: $h \in \{0, 1\}$, $\eta := \Pr(h = 1|m = 1 \wedge e = 0)$, $\Pr(h = 1|m = 0 \vee e = 1) = 0$.
4. Payoffs are realized
5. The administration decides whether to retain or dismiss the bureaucrat, and the previous steps repeat in period 2. The game then ends.

We consider two different assumptions regarding the grounds upon which the administration may terminate the bureaucrat's employment. Under "at-will" employment, the administration may dismiss the bureaucrat whether or not any monitoring occurred or yielded verifiable evidence of shirking. Under an assumption of "just-cause" firing, the administration may dismiss the bureaucrat only if there is verifiable evidence of the bureaucrat shirking.[1] We discuss the significance of these and other modeling assumptions further below, particularly in relation to proposed changes to the bureaucracy. Since the game is dynamic and includes private actions, we employ the solution concept of weak perfect Bayesian equilibrium (henceforth, "equilibrium").

Comments on the Model

From a technical perspective, our model is in the career concerns tradition, as the quality dimension on which A is evaluating B is hidden to both A and B. No player possesses private information

[1] This requires $m = 1$ and $h = 1$.

about types, only their actions, thereby avoiding the possibility of signaling. The presence of another dimension – partisan congruence – on which A selects B introduces complexity not present in textbook treatments.

Gathering information is costly. To obtain verifiable evidence of shirking ($h = 1$), the administration must pay κ_m. This monitoring cost is for naught if the agent chooses to exert effort or if η is not sufficiently high. The η parameter can be understood in a number of ways, be that a supervisor's managerial prowess, competence, or having a workforce sufficiently large such that agency heads and senior executives need not take on tasks typically delegated to rank-and-file bureaucrats. Taking $1-\eta$ is potentially informative about questions of agency independence, or the latitude over rulemaking and implementation delegated to a policymaker or executive agency. Conversely, η can be as simple as the likelihood of observing malfeasance during a monitoring period, whereas $1 - \eta$ can be understood as the skill of a given agent in concealing wrongdoing. Whatever the case may be, obtaining information about the agent's hidden action is not only costly, but carries the risk of type II error. We do not allow for false positives. If the administration receives evidence of low effort ($h = 1$), it knows this to be true.

In addition to output and, potentially, effort, our principal is able to condition the firing choice on the agent's ideology. With probability ρ, the bureaucrat's ideology matches that of the administration (that is, $\Pr(r = 1) = \rho$) at the outset of the game. We assume that an agent's ideology is known to all and that when replacing an incumbent bureaucrat, the administration chooses a co-partisan with certainty. This assumption is relevant across time. For instance, the spoils system all but ensured that political loyalty drove hiring decisions in the federal bureaucracy. Here, an agent's affiliation with the outgoing regime made their political persuasion common knowledge. Further, the dominance of party machines and density of political networks made simple the chore of replacing incumbent bureaucrats with loyalists.

The parameter ρ captures the mix of co-partisan and cross-partisans that A inherits when taking office. We model a single administration-bureaucrat interaction, so we may think of ρ as representing the likelihood that a given interaction is with a co-partisan bureaucrat or as reflecting multiple, independent interactions with a pool of bureaucrats. Variation in ρ could arise over time or across agencies. It is common knowledge that the Department of Homeland Security leans conservative, whereas the National Labor Relations Board skews left.Jowei et al. (2014) use campaign finance data to demonstrate that the agency ideology is an effective proxy for the political persuasion of its employees –

which lends credibility to the modeling assumption that ρ is known.

One source of generality in the model pertains to the function $F(\cdot,\cdot)$. The arguments e and θ may be complements or substitutes for one another. Whether F displays increasing or decreasing differences in e, θ determines the degree of learning possible about θ for any given level of effort.

The model abstracts from several important, well-studied features of administration-bureaucracy relationships. First, we set aside both expertise and public service motivation on the part of the bureaucrat. While this renders an inaccurate image of most members of the civil service, it allows us to focus on latent skill at public service provision and partisanship. Baseline expertise and public service motivation may also be captured by $F(0,0)$, which is the extent to which even a lower-skilled bureaucrat who is not putting forth high levels of effort will still provide basic public services. To the extent that most bureaucrats possess significant expertise, this may be reflected by a high value of π, the likelihood of a high quality draw of B. To the extent that bureaucrats may possess strong internal motivation that makes effort relatively low cost, this may be reflected in a low κ_e, the cost of effort.

Many models of the bureaucracy entail delegation decisions by the administration. Ours is not a model of delegation. The value of having a co-partisan in the civil service, however, may be read as a reduced form representation of the "ally principal."

As regards the dichotomy of police patrols and fire alarms, monitoring in our model constitutes a police patrol approach to keeping tabs on the civil service. One interpretation consistent with our model is that if we wish to replace a given bureaucrat with time left in the administration's time in office, we will need hard evidence of wrongdoing as the project's success is realized. Evidence produced from an investigation following project failure would be too late, especially for an administration that seeks to replace the existing stock of bureaucrats with co-partisans quickly.

It is worth making a distinction between patronage hires that predated the civil service reforms of the mid-19th century and politically motivated appointments common in the contemporary federal bureaucracy. Those hired under a spoils system were not ideologues per se – the emphasis remained on rewarding individuals who had proven loyal in the past (nevertheless, they were most assuredly party members). This transactional means of staffing persists among high-level appointees, the motivation at present selects on loyalty and ideological alignment. Prior service to the party is not discounted, but functions moreso as a signal of eligibility as opposed to the basis for a reward. As an example, consider the role of the Federalist Society and Heritage Foundation in reshaping how appointments are made to the judiciary and federal bureaucracy, respectively.

Equilibrium analysis

We begin by analyzing the administration's retention decision, given project success, their own decision to monitor or not, and whether or not that monitoring produced any evidence of shirking. From there, we consider the administration's decision to monitor as well as the bureaucrat's decision to put forth effort. We characterize the equilibrium cases under both terms of employment, at will and just cause, for both co- and cross-partisans. In the next section we derive welfare comparative statics. In particular, we will focus on first-period effort as well as on the probability of a high-type holding office in the second period.

When making their retention decision, the administration observes s, h, knows whether m was 0 or 1 (i.e., whether they monitored), and has a (correct-in-equilibrium) conjecture about p (the probability that $e = 1$). Note that in the second and final period, B will never put forth effort, and A will never monitor. Note further that if replacing a bureaucrat, A will select a co-partisan but must draw again in terms of quality.

Call $\mu(s, m, h|p)$ the posterior probability with which A believes that $\theta = 1$. If facing a co-partisan ($\rho = 1$), A will compare:

$$\begin{aligned}
\mathbb{E}(u_{A,2}|\rho = 1, retain) &\gtrless \mathbb{E}(u_{A,2}|\rho = 1, replace) \\
[\mu(s,h)F_{0,1} + (1 - \mu(s,h))F_{0,0}] + \beta &\gtrless [\pi F_{0,1} + (1 - \pi)F_{0,0}] + \beta \\
(\mu(s,h) - \pi)(F_{0,1} - F_{0,0}) &\gtrless 0 \\
\mu(s,h) &\gtrless \pi
\end{aligned} \tag{1}$$

If A updates positively on the skill of a co-partisan bureaucrat, it will retain. Conversely, if A updates negatively, it would like to replace the bureaucrat. A may not be able to replace B, however, if there is no monitoring or no hard evidence produced.

If facing a cross-partisan, A compares:

$$\begin{aligned}
\mathbb{E}(u_{A,2}|r = 0, retain) &\gtrless \mathbb{E}(u_{A,2}|r = 0, replace) \\
[\mu(s,h)F_{0,1} + (1 - \mu(s,h))F_{0,0}] &\gtrless [\pi F_{0,1} + (1 - \pi)F_{0,0}] + \beta \\
(\mu(s,h) - \pi)(F_{0,1} - F_{0,0}) &\gtrless \beta
\end{aligned} \tag{2}$$

We may recognize the expression on the left hand side of the inequality from the calculations for co-partisans. This expression is the incremental benefit of retaining rather than replacing the bureaucrat for obtaining project success in the second period. Condition 2 reveals the administration's higher threshold for retaining a cross-partisan than for retaining a co-partisan.

The highest that the incremental benefit could be is $(1-\pi)(F_{0,1})$ belief that a cross-partisan is high quality would need to be quite high in order to justify the opportunity cost of having a co-partisan in the bureaucracy. To embody our assumptions about the partisan motivations of administrations around the time of the Pendleton Act and today, we invoke the following assumption.

Assumption 1. $\beta > (1-\pi)(F_{0,1})$

This assumption guarantees that the administration will never want to retain a cross-partisan. As above, it may be that the administration would not be able to remove a cross-partisan, due to lack of evidence of wrongdoing under just-cause employment. In addition to being substantively aligned with the environment we seek to model, this assumption greatly simplifies the analysis.

After observing project success, the posterior beliefs, $\mu(s, m, h|p)$, about quality as a function of monitoring and signals about wrong-doing are as follows:

$$\begin{aligned}
\mu(1,0,0) =& \frac{\pi(pF_{1,1} + (1-p)F_{0,1})}{\pi(pF_{1,1} + (1-p)F_{0,1}) + (1-\pi)(pF_{1,0} + (1-p)F_{0,0})} \\
\mu(1,1,0) =& \frac{\pi(pF_{1,1} + (1-p)F_{0,1}(1-\eta))}{\pi(pF_{1,1} + (1-p)F_{0,1}(1-\eta)) + (1-\pi)(pF_{1,0} + (1-p)F_{0,0}(1-\eta))} \\
\mu(1,1,1) =& \frac{\pi F_{0,1}}{\pi F_{0,1} + (1-\pi)F_{0,0}}
\end{aligned}$$

In each of these cases, $\mu > \pi$. As such, if $s = 1$, A would wish to retain a co-partisan, regardless of monitoring or even hard evidence of shirking. A wants to dismiss B when the project fails, which perhaps better explains the incentive to monitor. The benefit of monitoring is not its role in shaping posterior beliefs, but rather it allows A to replace B should the administration wish to.

In Table 1, we summarize A's optimal retention strategy under both types of employment and in the context of both partisan types. These results are formalized in Lemma 1 in the appendix.

Given the probability of monitoring, q, and the backwards induction regarding retention/replacement, we ask when each partisan type of bureaucrat would put forth effort or not. Imposing the optimal retention strategy that A will employ, B chooses $e = 1$ if

$$\mathbb{E}U_B(e = 1|q, \rho, \pi) \geq \mathbb{E}U_B(e = 0|q, \rho, \pi).$$

Table 1: A's Optimal Approach to Retention/Dismissal

	At-Will	Just-Cause
Co-partisan	Replace if $s = 0$; retain if $s = 1$	Replace if $s = 0, h = 1$; retain if $s = 1$
Cross-partisan	Replace if $s = 0, 1$	Replace if $s = 0, 1$ and $h = 1$; retain o/w

Table 2 summarizes B's optimal effort choice for each of the four cases: {co-partisan, cross-partisan} $\times$ {at-will, just-cause}, formalized in Lemma 2 in the Appendix.

Table 2: B's Optimal Effort Choice

	B chooses $e = 1$ if $\kappa_e <$ At-Will	Just-Cause
Co-partisan	$\mathbb{E}(s\|e = 1) - \mathbb{E}(s\|e = 0)$	$q\eta(1 - \mathbb{E}(s\|e = 0))$
Cross-partisan	0	$q\eta$

Comparing the thresholds for just-cause reveals, conditional on a given q, B puts forth effort for a wider range of effort costs when B is a cross-partisan than when B is a co-partisan. It is worth noting, however, that q is an endogenous probability, and so will differ between the cases. Nonetheless, we see greater incentive for effort when it is the only action – not even project success – that will save B's employment.

Finally, we ask when monitoring is worthwhile for A, given effort probability p. The administration (A) chooses $m = 1$ if

$$\mathbb{E}U_{A_{W/C,X/O}}(m = 1|p, \pi) - \kappa_m \geq \mathbb{E}U_{A_{W/C,X/O}}(m = 0|p, \pi).$$

We can reference here equations (1) and (2), where each side of those equations are the relevant expected utilities. Lemma 3 in the Appendix formalizes our results. Table 3 presents the conditions under which A would choose to monitor.

Comparing the thresholds for just-cause reveals, conditional on a given p, A monitors for a wider range of monitoring costs for a cross-partisan than a co-partisan. Again, p is an equilibrium quantity and so will differ between the cases. It remains true, though, that monitoring is compatible with higher costs of doing so when B is a cross-partisan than when B is a co-partisan.

We may now characterize the equilibrium cases across different terms of employment and partisan

Table 3: A's Optimal Monitoring Choice

	A chooses $m = 1$ if $\kappa_m <$	
	At-Will	Just-Cause
Co-partisan	0	$\pi(1-\pi)(1-p)\eta(F_{0,1} - F_{0,0})^2$
Cross-partisan	0	$(1-p)\eta\beta$

types. Figure 1 depicts cases for each of the four combinations of terms and types. Proposition 1 in the Appendix derives these cases and states the equilibrium formally.

Figure 1: Equilibrium cases across employment terms and partisanship

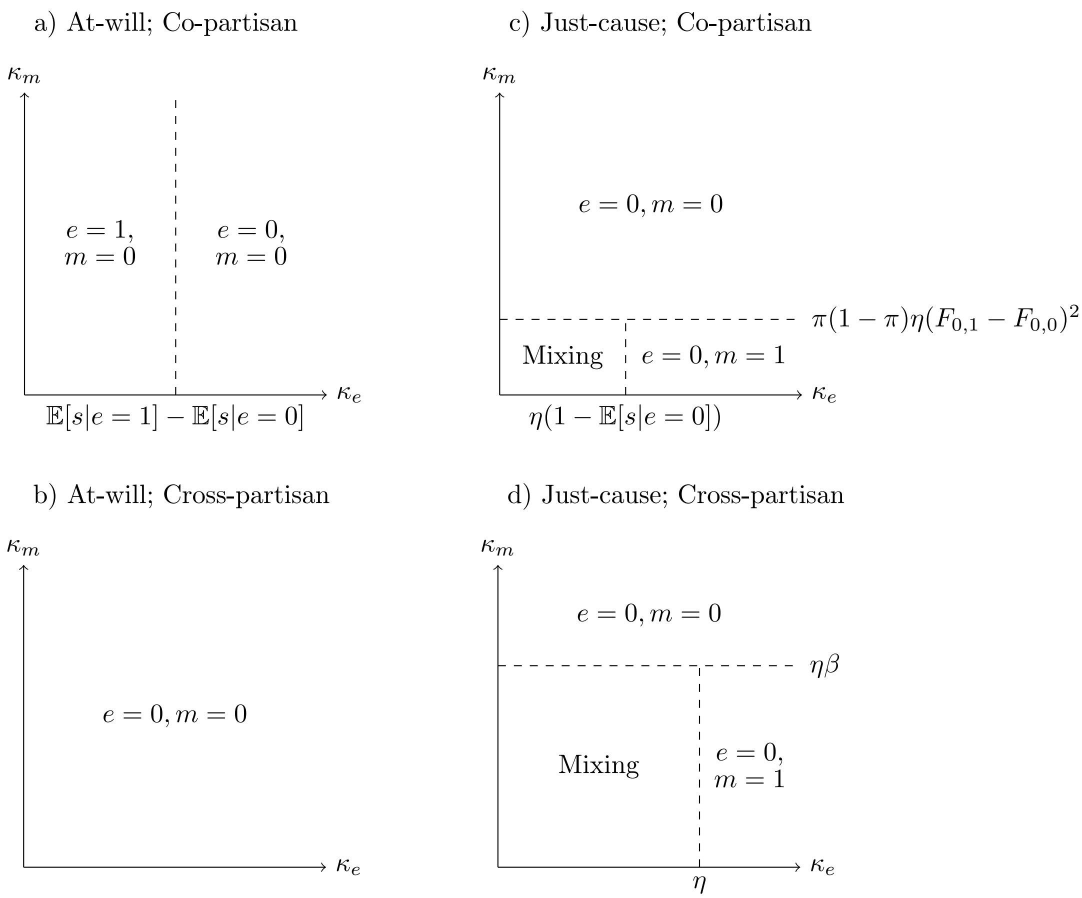

Notes: Figures constructed using $\pi = 3/4$, $\eta = 2/3$, $\beta = 1/8$, and $F_{e,\theta} = (e+\theta)/2$, $e, \theta \in \{0, 1\}$.

Welfare Calculations

We wish to compare first-period effort and the second-period benefit in terms of selection across cases. To focus on cases in which there is at least some effort under just-cause terms of employment, we employ the following assumptions on costs. In practice, we will state: "monitoring and effort costs are not too large."

Assumption 2. $\kappa_e < \beta$; $\kappa_m < \eta\beta$

Because of Assumption 1, selection of high types only occurs under co-partisans (the upper half of Figure 1). Meanwhile, effort – and especially effort partially induced by monitoring – occurs only under just-cause terms (the right half of Figure 1). The following result states that just-cause terms become more desirable relative to at-will terms as the proportion of cross-partisans, $(1 - \rho)$, increases.

Proposition 2. *When $\rho = 0$, just-cause employment terms maximize the sum of the probability of first-period public service success and second-period public service success. As ρ increases, at-will terms become more attractive relative to just-cause.*

As the administration faces a greater share of cross-partisans, just-cause protections ensure effort in the first-period. Conversely, as the administration faces a greater share of co-partisans, the selection facilitated by at-will employment terms makes such terms increasingly attractive. In particular, if the cost of effort is not too large to preclude effort by a co-partisan under at-will terms, an administration facing co-partisans does best when it is able to replace bureaucrats at will.

Discussion

Though we have characterized the retention decision as a 2x2 matrix, the principal, sans constraints, always chooses one action for the cross-partisan (fire) and one for the copartisan (keep in the even of project success). In order to achieve the former end in the just cause case, the administration must commit to monitoring. This, ultimately, is the way to understand the motivation to monitor – not that the principal derives utility from being more sure of the agent's type, but that they may want to dismiss the incumbent bureaucrat and need hard evidence to do so under certain institutional arrangements.

Our results are consistent with the police patrol/fire alarm framework presented by (McCubbins & Schwartz 1984). Using project success as an indicator of system failures (whether that arises from shirking or ill-equipped policymakers) is akin to the fire alarm style of oversight. It is not only more efficient to condition the retention decision on project outcomes, but it in fact all but ensures greater effort in the first period. When monitoring – our analogue to the "police patrol" style of oversight – becomes rational, there are no pure strategies in which the agent exerts effort. Under some circumstances, more information produces less desirable outcomes.

The scope of this finding is limited, however, to cases where the incumbent bureaucrat and administration are ideologically aligned. Under at-will employment, it is never rational for the cross-partisan to exert effort. Introducing protections against wrongful firing for agents that are ideologically opposed to the administration introduces the possibility of a positive effort choice (in mixed strategies) in the first period. Although the standard career concerns model breaks down when agents are not ideologically aligned with the administration, this case shows that information about the agent's effort choice can remedy some of the shortfalls introduced by the high premium on loyalty. In short, fire alarms are more efficient and effective when the principal and agent are of the same political persuasion, but police patrol oversight becomes rational whenever nonaligned actors are introduced.

The lack of pure strategy equilibria in which effort is a rational choice (in particular, the case of the out-partisan without civil service protections) seems at first glance to be concerning. This is a byproduct of the efficiency of our model, which does not take into account policy preference or public service motivation. In fact, we do not assume any intrinsic motivation on the part of the bureaucrat. Neither does the agent have an incentive to invest in policy-relevant expertise. Rather, their effort choice is made with the sole end of survival in mind. The nonexistent effort choice can be interpreted as a form of "quiet quitting," where individuals exert little to no effort, as marginal returns are slim – that is, investing more will not change the outcome, which, in the context of this model, is that a nonaligned agent will be fired at the conclusion of the first period, should circumstances permit.

When comparing performance within retention rules, the out partisan has a clear advantage whenever hard evidence is required to dismiss the incumbent bureaucrats as compared to their administration-aligned counterparts. Although both parties only exert effort in mixed strategies, the region in which an ideologically aligned agent will incur κ_e is less than a quarter of the size of the cross-partisan's equivalent region (see Figure 1). This does not hold when bureaucrats can be fired at will, where it the likelihood of first period effort increases drastically when the administration is paired with an aligned agent at the outset of the game.

Though it is not within the scope of this model, the ability to reclassify subsets of federal employees as either political appointees or career civil servants is one of the most powerful tools at the disposal of the executive branch. Going from at-will to just cause enables the principal to insulate appointed agents from being dismissed under the following administration, thus ensuring that ideologically motivated actors will remain in positions of power. Meanwhile, reclassifying careerists as appointees gives an administration broad latitude over hiring and firing.

The results from our model are incredibly informative when interpreted as agency-specific guidance: for instance, just cause retention rules make effort increasingly likely in agencies with missions that are perhaps orthogonal or antithetical to those of the administration. Thus, were our findings born out in the federal bureaucracy, just cause firing rules would produce superior outcomes in agencies like the EPA under the supervision of a Republican-controlled White House. Conversely, agencies where bureaucrats are more likely to aligned with the ends of the administration would see greater effort when retention is conditioned solely on the successful implementation of policy.

Conclusion

There are two possible avenues by which political principals can increase the likelihood of project success: first, by incentivizing agents to exert costly effort, and second, by retaining competent agents. When faced with a copartisan in the first period, the administration enjoys higher returns to effort under at-will dismissal rules than they would were more robust civil service protections in place. Should the incumbent bureaucrat be ideologically unaligned, the opposite is true: the only way to induce effort on the part of such an agent is to insulate them from partisan firing. The latter is somewhat counter-intuitive, as the administration's goal is to fire cross-partisan bureaucrats, whereas they only wish to dismiss aligned agents when projects are not successfully implemented. These protections are necessary, however, if the internal logic of career concerns models is to remain. Else, the link between effort and retention is broken, and cross-partisans have no incentive to contribute to project success.

Similarly, the calculus of monitoring shifts drastically across institutional arrangements. When the administration can fire agents for any reason, there is no rational choice but to condition retention on project success alone. Formal oversight only improves outcomes when the information obtained is sufficient to justify dismissing bureaucrats who were unsuccessful in implementing policy.

These results have concrete implications for the ongoing debate over the structure of the federal bureaucracy. Just-cause protections are most valuable whenever agents are known to be at ideological odds with the executive. Empirical work suggests that the political preferences of individual bureaucrats can be inferred, in part, by the overall partisan lean of an agency. Civil service protections are a necessary component in persuading agents to carry out executive directives. Meanwhile, the threat of dismissal alone may be sufficient to motivate copartisans to signal competence through exerting effort.

Future work would do well to unite within a common framework extant work on delegation and civil service protections in order to better gauge the optimal allocation of bureaucratic discretion across different institutional setups (in particular, hiring and firing rules). Further, microfounding the administration's weight on ideological alignment β may prove fruitful.

References

Aghion, Philippe & Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy* 105(1):1–29.

Alesina, Alberto & Guido Tabellini. 2007. "Bureaucrats or Politicians? Part I: A Single Policy Task." *American Economic Review* 97(1):169–179.

Bawn, Kathleen. 1995. "Political Control Versus Expertise: Congressional Choices about Administrative Procedures." *American Political Science Review* 89(1):62–73.

Bendor, Jonathan & Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98(2):293–310.

Dal Bó, Ernesto, Frederico Finan & Martín A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service." *The Quarterly Journal of Economics* 128(3):1169–1218.

Delfgaauw, Josse & Robert Dur. 2007. "Signaling and screening of workers' motivation." *Journal of Economic Behavior & Organization* 62(4):605–624.

Epstein, David & Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Political Economy of Institutions and Decisions Cambridge: Cambridge University Press.

Fox, Justin & Richard Van Weelden. 2012. "Costly transparency." *Journal of Public Economics* 96(1-2):142–150.

Gailmard, Sean & John W. Patty. 2013. *Learning While Governing: Expertise and Accountability in the Executive Branch*. University of Chicago Press.

Gibbs, Daniel. 2020. "Civil service reform, self-selection, and bureaucratic performance." *Economics & Politics* 32(2):279–304.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *The Review of Economic Studies* 66(1):169–182.

Huber, John D. & Nolan McCarty. 2004. "Bureaucratic Capacity, Delegation, and Political Reform." *American Political Science Review* 98(3):481–494.

Johnson, Ronald N. & Gary D. Libecap. 1994. *The Federal Civil Service System and the Problem of Bureaucracy: The Economics and Politics of Institutional Change.* National Bureau of Economic Research Series on Long-Term Factors in Economic Development Chicago, IL: University of Chicago Press.

Kiewiet, Roderick. & Matthew McCubbins. 1991. *The Logic of Delegation: Congressional Parties and the Appropriations Process.* Chicago, IL: University of Chicago Press.

Krause, George A., David E. Lewis & James W. Douglas. 2006. "Political Appointments, Civil Service Systems, and Bureaucratic Competence: Organizational Balancing and Executive Branch Revenue Forecasts in the American States." *American Journal of Political Science* 50(3):770–787.

Lewis, David E. 2008. *The Politics of Presidential Appointments: Political Control and Bureaucratic Performance.* Princeton University Press.

Lewis, David E. 2009. "Revisiting the Administrative Presidency: Policy, Patronage, and Agency Competence." *Presidential Studies Quarterly* 39(1):60–73.

Lowande, Kenneth. 2019. "Politicization and Responsiveness in Executive Agencies." *The Journal of Politics* 81(1):33–48.

Lupia, Arthur & Mathew McCubbins. 1994. "Who Controls? Information and the Structure of Legislative Decision Making." *Legislative Studies Quarterly* 19(3):361–384.

Maskin, Eric & Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *American Economic Review* 94(4):1034–1054.

McCubbins, Mathew, Roger G. Noll & Barry Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *The Journal of Law, Economics, and Organization* 3(2):243–277.

McCubbins & Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* 28(1):165–179.

Moe, Terry M. 1985. *The Politicized Presidency.* The New Directions in American Politics Brookings Institution.

Moreira, Diana & Santiago Perez. 2025. "Who Benfits from Meritocracy?" *NBER Working Papers* .

Ornaghi, Arianna. 2019. "Civil Service Reforms: Evidence from U.S. Police Departments." *Unpublished manuscript* .

Osborne, David. 1993. "Reinventing Government." *Public Productivity & Management Review* 16(4):349–356.

Prendergast, Canice. 1993. "A Theory of "Yes Men"." *The American Economic Review* 83(4):757–770.

Stephenson, Matthew C. 2007. "Bureaucratic Decision Costs and Endogenous Agency Expertise." *The Journal of Law, Economics, and Organization* 23(2):469–498.

Ting, Michael M. 2011. "Organizational Capacity." *The Journal of Law, Economics, and Organization* 27(2):245–271.

Ujhelyi, Gergely. 2014. "Civil Service Rules and Policy Choices: Evidence from US State Governments." *American Economic Journal: Economic Policy* 6(2):338–380.