Should the Timing of Inspections be Predictable?*

Ian Ball[†] Jan Knoepfle[‡]

18 April 2023

Abstract

A principal hires an agent to work on a long-term project that culminates in a breakthrough or a breakdown. At each time, the agent privately chooses to work or shirk. Working increases the arrival rate of breakthroughs and decreases the arrival rate of breakdowns. To motivate the agent to work, the principal conducts costly inspections. She fires the agent if shirking is detected. We characterize the principal's optimal inspection policy. Periodic inspections are optimal if work primarily speeds up breakthroughs. Random inspections are optimal if work primarily delays breakdowns. Crucially, the agent's actions determine his risk attitude over the timing of punishments.

Keywords: inspections, audits, dynamic contracting, moral hazard
JEL: D82, D86, M42

---

*We thank Yifan Dai for excellent research assistance. For comments, we thank Nageeb Ali, Dirk Bergemann, Olivier Compte, Brett Green, Deniz Kattwinkel, Nicolas Lambert, Stephen Morris, Teemu Pekkarinen, Jacopo Perego, João Ramos, Ron Siegel, Ludvig Sinander, Curtis Taylor, and Alex Wolitzky. For helpful feedback, we thank audiences at VSET, Stanford, SMYE Orleans, Aalto, ESSET Gerzensee, UT Austin, Georgetown, BFI Economic Theory Conference, Penn State, Columbia Economic Theory Conference, Emory, Boston University, City University of Hong Kong, Carnegie–Mellon/Pittsburgh, UCL, USC, Essex, Paris, LSE, Oxford, and Queen Mary.

†Department of Economics, MIT, ianball@mit.edu

‡School of Economics and Finance, Queen Mary University of London, j.knoepfle@qmul.ac.uk

1 Introduction

Inspections are frequently conducted to reveal information about agents' otherwise unobserved actions. Venture capital investors conduct financial audits to ensure that entrepreneurs do not divert funds for private gain. Research grants are extended only after researchers pass intermediate reviews. Workers operating expensive machinery or risky technology undergo checks to ensure their compliance with safety regulations. Financial institutions are stress-tested to confirm that they employ proper risk management.[1]

Some inspections occur at pre-announced times; others are surprises. In this paper, we study the optimal timing of costly inspections in a dynamic moral hazard setting. We show how the productive role of the inspected agent determines whether predictable or random inspections are optimal. If the agent's main task is achieving a breakthrough—think of an entrepreneur investing in an innovative industry—then predictable inspections are optimal.[2] If the main task is avoiding a breakdown—think of a financial institution managing its risk in order to avoid default—then random inspections are optimal.

We analyze the following continuous-time model. A principal hires an agent to work on a long-term project that culminates in a successful breakthrough or a catastrophic breakdown. The principal commits to the timing of costly inspections. At each time, the agent privately works or shirks.[3] Work increases the arrival rate of breakthroughs and decreases the arrival rate of breakdowns. Each inspection yields a binary result—pass or fail—that (partially) reveals the agent's past actions, as described in detail below. If the agent fails an inspection, the principal terminates the project. We solve for the cheapest inspection policy that induces the agent to work until the project ends.

[1] Almost all economic activity involves inspections in some form, conducted either within organizations or by external auditors. In 2017, the four accounting firms Deloitte, EY, KPMG, and PwC earned over 47 billion dollars from auditing alone (The Economist, 2018).

[2] It is common for venture capitalists to disburse funds in stages after checking, at pre-announced dates, that the entrepreneur has invested previous funds appropriately; see Gompers and Lerner (2004, p. 5).

[3] In some applications, shirking represents an unproductive (even fraudulent) activity, such as diverting funds.

First, we derive the optimal policy in the special case of *perfect inspections*. That is, each inspection perfectly reveals whether the agent has previously shirked.

The form of the optimal inspection policy is driven by the agent's endogenous risk attitude over the timing of rewards and punishments. Planned inspections are carried out only if the project has not yet ended in a breakthrough or breakdown. When the agent considers the impact of future inspections, his effective discount factor reflects the probability that the inspection time is reached before the project ends. This probability depends on the agent's planned actions. The agent's effective discount factor is more convex if he expects the project to end sooner, making him effectively more impatient.[4]

If the agent's main task is achieving a breakthrough—work speeds up breakthroughs by more than it delays breakdowns—then it is optimal to inspect periodically. That is, the time between consecutive inspections is constant (Theorem 1). Shirking prolongs the project by reducing breakthroughs. If the agent deviates and shirks, then he is effectively more patient, making his effective discount factor as a function of time *less convex* than the principal's. Among inspection policies that are equally costly for the principal on-path, for the agent planning to shirk, the expected loss from termination after a failed inspection is greatest when inspections are periodic.

Conversely, if the agent's main task is avoiding a breakdown—work delays breakdowns by more than it speeds up breakthroughs—then it is optimal to conduct inspections randomly. Under the optimal policy, inspections are conducted with a constant hazard rate (Theorem 2). Shirking shortens the project by generating breakdowns. If the agent deviates and shirks, then he is effectively more impatient, making his effective discount factor as a function of time *more convex* than the principal's. Among inspection policies that are equally costly for the principal on-path, the agent's loss from shirking is greatest if inspections are random. In both cases, the ratio between the cost of the

[4] See DeJarnette et al. (2020) for an axiomatic analysis of the connection between impatience and risk-preferences over time-lotteries. To be sure, our results are not driven by primitive time preferences, but by effective time preferences determined endogenously by the agent's actions.

best periodic policy and the best random policy can be arbitrarily large.

Next, we derive the optimal inspection policy in the general case of *imperfect inspections*. In this case, shirking by the agent does not always leave behind a paper trail. The longer the agent shirks, the more likely he is to fail an inspection. But if the agent recently passed an inspection, then another inspection soon after is unlikely to give a different result. Thus, there is a new force toward spacing out inspections.

If the agent's main task is achieving a breakthrough, then periodic inspections are still optimal if inspections are imperfect (Theorem 3). Indeed, periodic inspections are already spaced out.

Conversely, if the agent's main task is avoiding a breakdown, then a memoryless inspection policy is no longer optimal if inspections are imperfect. Such a policy conducts inspections in quick succession with positive probability, which is now wasteful. Theorem 4 characterizes the optimal timing of imperfect inspections in this case. If the inspection technology is sufficiently imprecise, then the optimal inspection policy is periodic. The force towards spacing out inspections dominates. If the inspection technology is sufficiently precise, then the optimal policy leverages the benefits of randomization while also spacing out inspections. After each inspection, there is a fixed period during which no inspections are conducted. The next inspection is conducted with positive probability exactly at the end of this period. If the agent is not inspected at this point, then the next inspection arrives with a constant hazard rate thereafter. Once the inspection is conducted, the cycle repeats, beginning with the period of no inspection.

The rest of the paper is organized as follows. Section 1.1 discusses related literature. Section 2 presents the model. Section 3 studies the agent's behavior without inspections. Section 4 formulates the principal's problem recursively. Next, we solve for the optimal inspection policy with perfect inspections in Section 5 and with imperfect inspections in Section 6. Section 7 amends the inspection technology to allow the agent to recover from past shirking. The conclusion is in Section 8. The main proofs are in Appendix A. Additional results and proofs are in Appendix B.

1.1 Related Literature

The inspection technology that we study is backward-looking—the probability that the agent passes an inspection depends on his *past* actions. Lazear (2006) studies the optimal allocation of monitoring resources, but the monitoring technology in his model reveals only the agent's *current* action.[5] In that model, the agent faces a binary action choice. The agent is punished if he is monitored while he is taking the bad action. The cheapest way for the principal to induce the agent to follow a given action path is to uniformly randomize monitoring activity over the times at which the good action is induced.

Varas et al. (2020) study the timing of regulatory inspections in a model of firm reputation. Their inspection technology takes the specific form introduced in Board and Meyer-ter-Vehn (2013).[6] These inspections are backward-looking, but they have the special property that the agent's incentives at each time are independent of his past actions. As a result, it is sufficient to consider local incentive constraints. Conducting inspections at a constant hazard rate is the cheapest way for the principal to motivate the agent to always take the desired action.[7] In Varas et al. (2020), there are no breakthroughs or breakdowns, and the project lasts forever.

We consider a richer inspection technology in which the agent's past actions affect his preferences over future actions. This can capture, for example, that an employee who has shirked has a weaker incentive to continue working if he expects that his past shirking will be discovered at the next inspection. To solve for the optimal inspection policy in our model, we explicitly consider

[5] Most subsequent work on dynamic contracts analyzes monitoring of current actions; see Antinolfi and Carli (2015); Piskorski and Westerfield (2016); Chen et al. (2020); Li and Yang (2020); Dai et al. (2022); Rodivilov (2022); Wong (2022). In Halac and Prat (2016) and Dilmé and Garrett (2019), the principal's investment has a persistent effect on her monitoring capabilities, but monitoring still reveals information about current actions only. In dynamic adverse selection problems, see Chang (1990); Monnet and Quintin (2005); Wang (2005); Popov (2016); Malenko (2019) for monitoring of a state that is distributed independently across periods, and Ravikumar and Zhang (2012); Kim (2015) for inspection of a serially correlated state.

[6] The same inspection technology is used in Wagner and Knoepfle (2021).

[7] Varas et al. (2020) also consider the case in which the principal directly values information acquisition. They show that this creates a force toward periodic inspections.

dynamic, global deviations by the agent. Without breakthroughs and breakdowns, the optimal policy in our setting is generally periodic, in contrast to previous work.

We contribute to the dynamic contracting literature by identifying which kinds of tasks (breakthrough-enhancing v. breakdown-reducing) are best incentivized through predictable or random inspections. A number of papers study the dynamic incentives of agents working toward breakthroughs: Bergemann and Hege (1998, 2005); Hörner and Samuelson (2013); Green and Taylor (2016); Halac et al. (2016).[8] A robust finding is that the principal benefits from committing to a deterministic deadline. With breakthroughs, shirking prolongs the relationship. This is less valuable when there is a deadline. Similar to inspections in our setting, the reduction in shirking payoff is greatest when the deadline is deterministic because the agent's effective discount factor when he shirks is less convex than the principal's on-path discount factor. In Green and Taylor (2016), the agent has to complete two breakthroughs, and the second deadline is deterministic as above. The deadline for the first breakthrough, however, is random to optimally incentivize the agent to immediately report a breakthrough, which is privately observed by the agent. In our model, breakthroughs and breakdowns are public.

Finally, an important force in our paper is the relationship between time- and risk-preferences. Ortoleva et al. (2022) exploit this relationship in an adverse selection problem. In their model, the agent's discount rate is private information. They show how the agent's induced risk preferences can be used to screen the agent in an allocation problem. In our moral hazard setting, the agent's hidden actions determine the agent's *effective* risk preferences.

[8]For an analysis of incentives in the presence of breakdowns, without inspections, see Keller and Rady (2015); Bonatti and Hörner (2017); Hörner et al. (2021); Wagner and Klein (2022).

2 Model

2.1 Setting

Environment Time is continuous and the horizon is infinite. There are two players: a principal (she) and an agent (he). The principal hires the agent to work on a project. During the project, the agent privately chooses at each time t in $[0, \infty)$ whether to work ($a_t = 1$) or shirk ($a_t = 0$). The principal commits to the timing of costly inspections. Each inspection reveals information about the agent's past actions, as described below.

The project culminates in a public breakthrough or a public breakdown, which arrive independently at Poisson rates

$$a_t \lambda_G \quad \text{and} \quad (1 - a_t)\lambda_B,$$

where λ_G and λ_B are nonnegative parameters. The subscripts abbreviate *good* (for breakthroughs) and *bad* (for breakdowns). The principal can terminate the project at any time prior to a breakthrough or a breakdown. The game ends once the project ends in a breakthrough, in a breakdown, or by termination.

Inspection technology There is an evolving *detectability state* $\theta_t \in \{0, 1\}$ that is hidden to both players. This state is publicly revealed when the principal conducts an inspection. Initially, $\theta_0 = 0$. While the state equals 0, transitions to state 1 occur at Poisson rate $\delta(1 - a_t)$, independently of breakthroughs and breakdowns. State 1 is absorbing. One interpretation is that the state θ_t indicates whether the agent's past shirking left behind evidence. Such evidence is uncovered only at an inspection. Since state 1 is absorbing, evidence does not disappear.[9] The detectability parameter δ measures the rate at which evidence is left behind when the agent shirks.

Each inspection has two possible results. Say that the agent *passes* (resp. *fails*) an inspection if the state is revealed to be 0 (resp. 1). If the agent follows an action path $a = (a_t)_{t \geq 0}$, then it is straightforward to compute the

[9]In Section 7 we allow for transitions from state 1 to state 0.

probability $p_t(a)$ that he passes an inspection conducted at time t:

$$p_t(a) = \exp\left\{-\delta \int_0^t (1 - a_s)\,\mathrm{d}s\right\}.$$

The passage probability is a decreasing, convex function of the duration of shirking prior to time t. If the agent fails an inspection, then he will necessarily fail all subsequent inspections since state 1 is absorbing. The agent's conditional probability of passing an inspection at time t, given that he passed an inspection at an earlier time t', is

$$\frac{p_t(a)}{p_{t'}(a)} = \exp\left\{-\delta \int_{t'}^t (1 - a_s)\,\mathrm{d}s\right\}.$$

This conditional probability depends only on the duration of shirking between times t' and t.

Payoffs The principal and the agent both discount future payoffs using the exponential discount factor e^{-rt}, where $r > 0$. While the project continues, the agent's flow utility, as a function of his action a_t in $\{0, 1\}$, is given by

$$u(a_t) = (1 - a_t)u_0 + a_t u_1,$$

where $u_0, u_1 > 0$. When the project ends, the agent gets his outside option payoff, which is normalized to 0. The agent does not receive a lump sum utility from a breakthrough or a breakdown. This assumption is without loss because lump sum payoffs from breakthroughs and breakdowns can be annuitized into the agent's flow payoffs; see Section 2.3 for details.

At each time the principal inspects the agent, she pays a lump sum cost, normalized to 1. The setting is described by six parameters: $\lambda_G, \lambda_B, \delta, r, u_0, u_1$. We study the principal's cost-minimization problem: What is the cheapest inspection policy that induces the agent to work on the project at all times until achieving a breakthrough?

2.2 Principal's problem

The principal commits to a dynamic, stochastic inspection policy. Formally, an inspection policy is a collection $\mathbf{T} = (T_n)_{n=1}^{\infty}$ of random variables satisfying $0 < T_1 < T_2 < \cdots$. The n-th inspection is conducted at (random) time T_n if the project has not already ended.

Technically, the principal also commits to the timing of project termination, but the optimal termination policy is clear. In order to induce work until a breakthrough, the principal cannot terminate the project on-path. In the off-path event that the agent fails an inspection, it is optimal for the principal to immediately terminate the project, imposing the maximal punishment on the agent. Below, we consider the inspection policy only, with the understanding that each inspection policy is combined with the optimal termination policy.

Given an inspection policy $\mathbf{T}$, the agent chooses an action process $A = (A_t)_{t \geq 0}$ with right-continuous paths that is adapted to $\mathbf{T}$.[10] That is, the agent takes action A_t at time t, provided that the project has not yet ended. The principal chooses an inspection policy $\mathbf{T}$ to minimize the expected inspection cost, subject to the constraint that it is a best response for the agent to work until a breakthrough, i.e., to choose $A_t = 1$ for all t. Denote this action process by $A = \mathbf{1}$.

To state the problem formally, define the effective discount factor

$$D_t(A) = \exp\left\{ -rt - \lambda_G \int_0^t A_s \, \mathrm{d}s - \lambda_B \int_0^t (1 - A_s) \, \mathrm{d}s \right\}. \tag{1}$$

This discount factor is the product of the standard exponential discount factor e^{-rt} and the conditional probability that the project has not yet ended in a breakthrough or a breakdown by time t, given $(A_s)_{0 \leq s \leq t}$. To simplify notation, set $T_0 = 0$. Given an inspection policy $\mathbf{T}$, the agent's expected payoff from an

[10] Formally, A is adapted to the natural filtration generated by the counting process $N_t = \#\{n : T_n \leq t\}$ associated with $\mathbf{T}$.

action process A adapted to $\mathbf{T}$ is given by

$$U(A, \mathbf{T}) = \mathbb{E}\left[\sum_{n=1}^{\infty} p_{T_{n-1}}(A) \int_{T_{n-1}}^{T_n} D_t(A) u(A_t)\, \mathrm{d}t\right].$$

The expectation is over the inspection policy $\mathbf{T}$ and the action process A. For each realization of $(A, \mathbf{T})$, the expression inside brackets sums the agent's conditional expected utility over each inter-inspection interval. This expectation is over the random inspection results and the random arrivals of breakthroughs and breakdowns. In each term of the summation, we used conditional independence to factor the expectation into the product of the passage probability and the effective discount factor.

The principal chooses an inspection policy $\mathbf{T}$ to minimize the expected discounted inspection cost

$$\mathbb{E}\left[\sum_{n=1}^{\infty} p_{T_{n-1}}(\mathbf{1}) D_{T_n}(\mathbf{1})\right] = \mathbb{E}\left[\sum_{n=1}^{\infty} e^{-(\lambda_G + r) T_n}\right],$$

subject to the constraint that $U(\mathbf{1}, \mathbf{T}) \geq U(A, \mathbf{T})$ for all actions processes A adapted to the inspection policy $\mathbf{T}$.

2.3 Microfoundations for the agent's payoffs

For simplicity, we refer to the agent's actions as *work* and *shirk*, but our model can capture various applications with different microfoundations for the flow payoffs u_0 and u_1. Here are two examples.

- *Employment contracts.* Suppose the agent is a worker who receives flow wage w while employed, pays flow effort cost c if he works ($a_t = 1$), and gets a lump sum reward R for achieving a breakthrough. This is captured by our setting (without lump sum payoffs) by defining

$$u_0 = w, \qquad u_1 = w - c + \lambda_G R.$$

- *Investment funding.* Suppose the agent is an entrepreneur who runs a startup that is funded by an investor at flow rate φ. At each time t, the entrepreneur can invest the funds ($a_t = 1$) or divert them for private benefit ($a_t = 0$). The entrepreneur receives a lump sum reward R for a breakthrough and pays a lump sum cost C for a breakdown. This is captured by our setting (without lump sum payoffs) by defining

$$u_0 = \varphi - \lambda_B C, \qquad u_1 = \lambda_G R.$$

Our model assumes that breakthroughs arrive only when the agent is working, and breakdowns arrive only when the agent is shirking. This assumption is without loss. If the breakthrough and breakdown rates were instead

$$\underline{\lambda}_G + \lambda_G a_t \quad \text{and} \quad \underline{\lambda}_B + \lambda_B(1 - a_t),$$

then we could incorporate the baseline arrival rates into the discount rate by defining $r' = r + \underline{\lambda}_G + \underline{\lambda}_B$.[11] Thus, λ_G and λ_B represent the *sensitivity* of breakthroughs and breakdowns to the agent's actions.

3 Warm-up: No inspections

If the principal does not conduct inspections, then the agent's expected payoff from an action process A is

$$\mathbb{E}\left[\int_0^\infty D_t(A)u(A_t)\,\mathrm{d}t\right]. \tag{2}$$

Working forever and shirking forever respectively yield expected payoffs

$$U_1 := \frac{u_1}{\lambda_G + r}, \qquad U_0 := \frac{u_0}{\lambda_B + r}.$$

Without inspections, the agent's problem is stationary. His best response is also stationary—either working forever or shirking forever is optimal.

[11] If breakthroughs and breakdowns yielded lump sum rewards and costs as above, then we would add $\underline{\lambda}_G R - \underline{\lambda}_B C$ to both flow utilities.

Lemma 1 (No inspections)
Without inspections, working forever is a best response for the agent if and only if $U_1 \geq U_0$.

Myopically, the agent compares the flow payoffs of working and shirking. He prefers to work if $u_1 \geq u_0$. Dynamically, the agent considers the effect of his actions on the probability that the project continues. If $\lambda_G = \lambda_B$, then the agent's actions do not affect the expected length of the project, so the dynamic and myopic conditions coincide. If $\lambda_G > \lambda_B$, then working shortens the project because earlier breakthroughs outweigh later breakdowns. In this case, the dynamic constraint is more demanding than the myopic constraint.[12] Conversely, if $\lambda_G < \lambda_B$, then working lengthens the project because earlier breakthroughs are outweighed by later breakdowns. In this case, the dynamic constraint is less demanding than the myopic constraint.

To rule out uninteresting cases, we make the following standing assumption.

Assumption 1. $U_0 > U_1$.

In view of Lemma 1, Assumption 1 ensures that inspections are necessary to induce the agent to work.

4 Recursive formulation

We analyze the principal's problem recursively. Consider the principal's continuation problem after an inspection at time t. If the agent fails the inspection, i.e., $\theta_t = 1$, then the continuation problem is trivial—terminating the project is optimal. If the agent passes the inspection, i.e., $\theta_t = 0$, then the time-t continuation problem is identical to the time-0 problem.[13]

In the recursive formulation, after each passed inspection, the principal chooses the random time T until the next inspection, and the agent chooses

[12]This comparison between the dynamic and myopic constraint is for fixed flow payoffs u_0 and u_1. If these flow payoffs are microfounded by lump sum payments following a breakthrough or a breakdown, then it is natural to adjust u_0 and u_1 as λ_G and λ_B change.

[13]Therefore, the principal would use the same dynamic policy if she could commit only to the timing of the next inspection (and the associated termination decision).

the action path that he will follow until the next inspection. The principal's minimal cost, denoted K^*, satisfies the Bellman equation

$$K^* = \inf_T \mathbb{E}\left[e^{-(\lambda_G + r)T}(1 + K^*)\right], \tag{3}$$

where the infimum is taken over all random times T for which the value $V = U_1$ satisfies the agent's Bellman equation

$$V = \sup_a \mathbb{E}\left[\int_0^T D_t(a)u(a_t)\,\mathrm{d}t + p_T(a)D_T(a)V\right], \tag{4}$$

where the supremum is over all action paths $a = (a_t)_{t \geq 0}$.[14]

Equation (3) uses the discount factor $D_T(\mathbf{1}) = e^{-(\lambda_G + r)T}$ because the inspection is conducted at time T if and only if the project has not already ended. The principal pays the inspection cost 1 and faces the continuation cost K^* if the agent passes the inspection (which occurs on path). The value $V = U_1$ solves (4) if and only if working continuously is a best response for the agent. Inside the expectation in (4), we include only the agent's continuation value V if he passes the inspection. If he fails, then the project is terminated, so his continuation value is 0.

The system of nested Bellman equations in (3)–(4) is complex. The set of all action paths is large. For a given random time T, the supremum on the right side of (4) might be achieved by many alternating periods of working and shirking. To solve the principal's problem, we set up a relaxed problem that considers only a subclass of deviations by the agent. The suitable subclass depends on the parameter values. We solve this relaxed problem by constructing Lagrangian dual variables. Then we check that our relaxed solution is feasible in the original problem by solving the agent's HJB equation.

[14]It suffices to consider deterministic action paths. Between inspections, the agent does not learn any information, unless the project ends, in which case he has no choice to make.

5 Optimal timing of perfect inspections

To highlight the main force in the model, we first solve the principal's problem in the special case of *perfect inspections*. The passage probability is given by

$$p_t(a) = \begin{cases} 1 & \text{if } \int_0^t (1 - a_s)\,\mathrm{d}s = 0, \\ 0 & \text{if } \int_0^t (1 - a_s)\,\mathrm{d}s > 0. \end{cases}$$

That is, the agent passes the time-t inspection if and only if he has not shirked before time t. This passage probability is the limit of the passage probability with imperfect inspections as the detectability parameter δ tends to ∞.

We separate the analysis into two regimes according to the relative sensitivities of breakthroughs and breakdowns to the agent's action.

5.1 Speeding up breakthroughs: $\lambda_G > \lambda_B$

If $\lambda_G > \lambda_B$, then working increases the arrival rate of breakthroughs by more than it decreases the arrival rate of breakdowns. Consequently, working shortens the project. In particular, this case obtains if there are breakthroughs but no breakdowns ($\lambda_B = 0$).

Theorem 1 (Periodic perfect inspections)
Suppose that inspections are perfect and $\lambda_G > \lambda_B$. Then it is optimal to inspect periodically: for some period τ^*, the gap $T_n - T_{n-1}$ equals τ^* for each n. If $u_0 \geq u_1$, then this policy is uniquely optimal and the period τ^* is given by

$$e^{-(\lambda_B + r)\tau^*} = \frac{U_0 - U_1}{U_0}. \tag{5}$$

If the agent's primary task is achieving a breakthrough—think of a start-up entrepreneur or a researcher working toward a new discovery—then it is optimal to conduct inspections at regular intervals.

We first consider the case $u_0 \geq u_1$. With perfect inspections, once the agent has shirked, his subsequent actions have no effect on his passage probability. Thereafter, shirking is optimal because shirking lengthens the project (since

$\lambda_G > \lambda_B$) and yields weakly higher flow payoff (since $u_0 \geq u_1$). We consider the relaxed problem that requires only that the agent prefers working, rather than shirking, until the next inspection. The principal chooses a positive random time T to solve

$$\begin{aligned} &\text{minimize} \quad && \mathbb{E}\, e^{-(\lambda_G + r)T} \\ &\text{subject to} \quad && \mathbb{E}\, e^{-(\lambda_B + r)T} U_0 \geq U_0 - U_1. \end{aligned} \tag{6}$$

Each side of the constraint measures the agent's loss relative to his payoff U_0 from shirking forever in the absence of inspections. The right side, $U_0 - U_1$, captures this loss if the agent works until the next inspection. The left side captures this loss if the agent shirks until the next inspection. In this case, the agent will fail the inspection with certainty, thus forgoing the benefit U_0 from shirking forever after. Crucially, λ_B appears in the constraint but λ_G appears in the objective. If the agent plans to shirk until time T, then his effective discount factor is $e^{-(\lambda_B + r)T}$. In contrast, the principal uses the effective discount factor $e^{-(\lambda_G + r)T}$ because, on path, the agent works.

The solution of (6) becomes clear once we change variables. Instead of choosing the random time T of the next inspection, the principal can equivalently choose the random variable $X = e^{-(\lambda_G + r)T}$, which is the cost of conducting an inspection at time T. In terms of X, (6) becomes

$$\begin{aligned} &\text{minimize} \quad && \mathbb{E}\, X \\ &\text{subject to} \quad && \mathbb{E}\, X^{(\lambda_B + r)/(\lambda_G + r)} U_0 \geq U_0 - U_1. \end{aligned}$$

Figure 1 depicts the principal's problem before (left) and after (right) the change of variables, in an example with $\lambda_G < \lambda_B$.[15] The left panel plots, as a function of the inspection time t, the principal's inspection cost (black) and the agent's loss from the inspection if he shirks (orange). The principal chooses a distribution over the horizontal axis to minimize her expected inspection cost, subject to the constraint that the agent's expected loss from the inspection if he shirks is at least $U_0 - U_1$. As a function of the inspection time, the agent's loss is *less convex* than the principal's cost because $\lambda_G > \lambda_B$.

[15]In this example, $\lambda_G + r = 2$; $\lambda_B + r = 1$; $U_1 = 5/4$; and $U_0 = 2$.

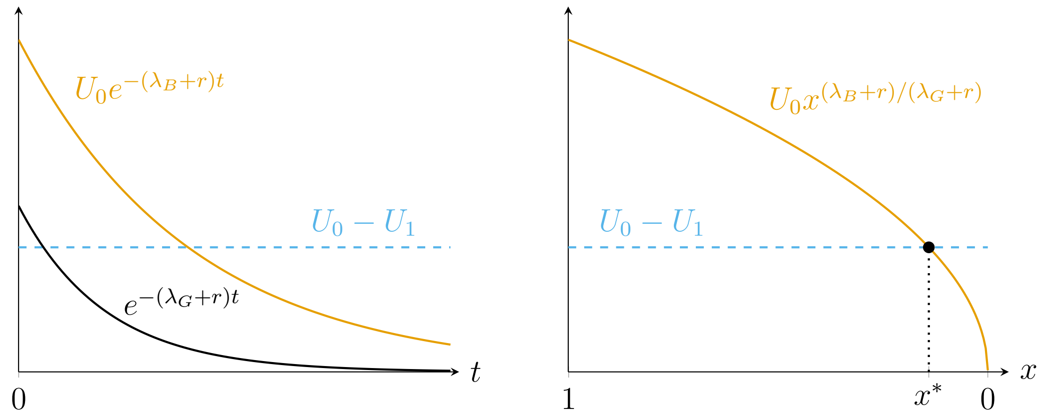

Figure 1. Shirking agent's loss from a perfect inspection with $\lambda_G > \lambda_B$

The right panel of Figure 1 puts the principal's inspection cost X on the horizontal axis $[0, 1]$, which is reversed so that time still moves from left to right. The principal chooses a distribution whose expectation is minimal (i.e., furthest right), subject to the constraint that the agent's expected loss from the inspection if he shirks is at least $U_0 - U_1$. This loss is a concave function of X. Replacing any distribution with a point mass on its expectation strictly slackens the constraint, without changing the principal's objective. Therefore, the unique solution is a unit mass on the point $x^* = e^{-(\lambda_G+r)\tau^*}$, with τ^* given by (5) so that the constraint holds with equality.

The argument above assumes that $u_0 \geq u_1$. If $u_0 < u_1$, then periodic inspections are still optimal, but the argument is more subtle. If the agent knows the time of the next inspection, then shirking all the way to the inspection is no longer the most attractive deviation. Once the agent has shirked, he is certain to fail the next inspection. As the inspection nears, the agent is increasingly myopic, eventually choosing to return to work because working yields a higher flow payoff than shirking. In the proof, we identify the binding shirk-before-work deviation.[16] The period τ^* is strictly smaller than in (5) because this new deviation binds. Shirking until the next inspection is strictly unprofitable.

[16] We describe this argument in more detail in the setting of imperfect inspections; see Theorem 3.

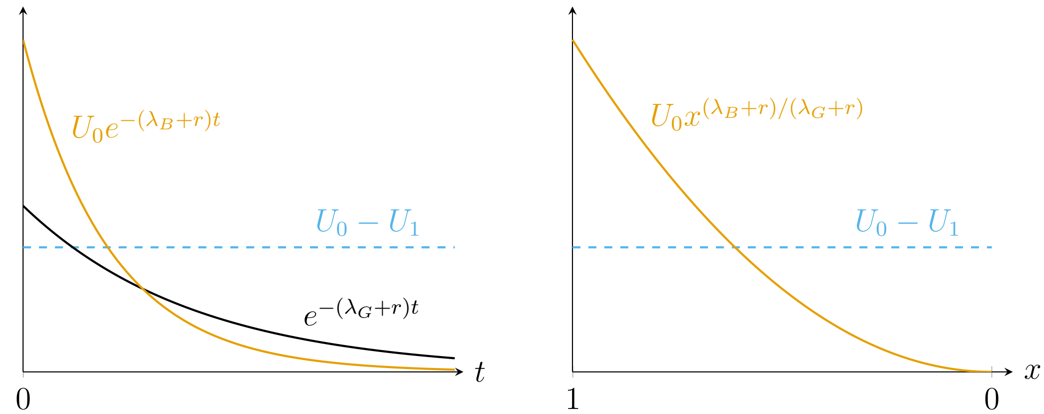

Figure 2. Shirking agent's loss from a perfect inspection with $\lambda_G < \lambda_B$

5.2 Delaying breakdowns: $\lambda_G < \lambda_B$

If $\lambda_G < \lambda_B$, then working decreases the arrival rate of breakdowns by more than it increases the arrival rate of breakthroughs. Consequently, working lengthens the project. In particular, this case obtains if there are breakdowns but no breakthroughs ($\lambda_G = 0$).

Theorem 2 (Random perfect inspections)
Suppose that inspections are perfect and $\lambda_G < \lambda_B$. Then the optimal inspection policy is unique: the gaps $(T_n - T_{n-1})_{n\geq 1}$ are independently and identically distributed according to an exponential distribution with hazard rate γ^*, where

$$\frac{\gamma^*}{\lambda_B + r} = \frac{U_0 - U_1}{U_1}. \tag{7}$$

If the agent's primary task is avoiding a breakdown—think of a worker guarding a nuclear plant or a bank managing its risk in order to avoid default—then it is optimal to conduct inspections at random times.

First, consider only the deviation in which the agent shirks until the next inspection. Figure 2 plots the same functions as Figure 1, before and after the change of variables, in an example with $\lambda_G < \lambda_B$.[17] As a function of the inspection time, the agent's loss from the inspection if he shirks is *more convex*

[17] In this example, $\lambda_G + r = 1$; $\lambda_B + r = 2$; $U_1 = 5/4$; and $U_0 = 2$.

than the principal's inspection cost. After the change of variables, the agent's loss from an inspection is a convex function of the principal's cost of conducting the inspection. Therefore, in this relaxed problem, the principal would like to spread out the distribution as much as possible by inspecting either very soon or else very far into the future.[18] But such a policy is infeasible in the original problem. If the agent is not inspected early on, then he can infer that there will be no inspection anytime soon. Instead of working continuously, the agent would strictly prefer to work briefly and then, once the risk of an inspection had passed, shirk forever after.

Therefore, we must consider a richer class of dynamic deviations by the agent. Under the optimal policy, inspections are conducted at a constant hazard rate. This policy is memoryless—the distribution of the time until the next inspection is the same, no matter the history of inspections. If the agent has worked until some time t, then he is indifferent between working continuously and shirking continuously until the next inspection. The optimal policy is the only policy for which this indifference holds for all times t.

5.3 Comparing predictable and random policies

For this subsection, suppose that $u_0 \geq u_1$.[19] Define the quantities

$$\mu = \frac{U_0 - U_1}{U_0} \quad \text{and} \quad \lambda = \frac{\lambda_G + r}{\lambda_B + r}.$$

Regardless of the relative values of λ_G and λ_B, the period τ^* given in (5) is the longest inter-inspection period that motivates the agent to work. Therefore, the inspection policy with period τ^* is optimal among all periodic policies. Similarly, the exponential policy with hazard rate γ^* given in (7) is optimal among all exponential policies. The ratio between the cost of the optimal

[18]Technically, the supremum in the relaxed problem is not achieved since $p_t(\mathbf{0})$ jumps up from 0 to 1 in the limit as $t \downarrow 0$.

[19]This is implied by Assumption 1 if $\lambda_G \leq \lambda_B$, but not if $\lambda_G > \lambda_B$.

exponential policy and the optimal periodic policy is

$$\frac{\mu(1-\mu^{\lambda})}{\lambda(1-\mu)\mu^{\lambda}}.$$

For $\lambda = 1$ (i.e., $\lambda_G = \lambda_B$), this ratio equals 1 because both the periodic and exponential policies are optimal. When the agent's action does not affect the length of the project, the perfect inspection technology is *neutral.* It does not favor periodic or random inspections. For $\lambda > 1$, this ratio is strictly greater than 1 and tends to ∞ as λ tends to ∞. Conversely for $\lambda < 1$, the ratio is strictly less than 1 and tends to 0 as λ tends to 0. Thus, the benefit of using the right timing structure—periodic versus exponential—can be arbitrarily large.

In work on dynamic contracts, it is sometimes assumed that the principal is more patient than the agent. If the principal and the agent have respective discount rates r_P and r_A, then our results still hold, but the relevant comparison for Theorems 1 and 2 is between $\lambda_G + r_P$ and $\lambda_B + r_A$, rather than between λ_G and λ_B.

6 Optimal timing of imperfect inspections

Consider the imperfect inspection technology with detectability parameter δ, described in Section 2 on p. 6. We impose the standing assumption that inspections are sufficiently precise.

Assumption 2. The detectability parameter δ satisfies the following:

A. $\delta > (\lambda_B + r)(U_0 - U_1)/U_1$;

B. $\delta > (\lambda_G + r) + (\lambda_G - \lambda_B)$.

Assumption 2.A ensures that sufficiently frequent inspections motivate the agent to work. As δ converges downward to $(\lambda_B + r)(U_0 - U_1)/U_1$, the principal's minimal inspection cost diverges. If 2.A is violated, then the principal's problem is infeasible. Assumption 2.A is more demanding if U_0/U_1 is larger, meaning that shirking is more attractive relative to working. Assumption 2.B ensures that the agent's passage probability, as a function of the duration of past shirking, is convex enough that *local* incentive constraints do not bind. If $\lambda_G = \lambda_B$, then 2.B simply requires that the detectability parameter is larger than the agent's effective discount rate.[20]

The imperfect inspection technology creates a new motive to space out inspections. The longer the agent shirks before an inspection, the less likely he is to pass. Conversely, after passing one inspection, the agent is likely to pass another inspection conducted soon after. Even if the agent shirks in between, he is unlikely to leave behind new evidence over a short time interval.

6.1 Speeding up breakthroughs: $\lambda_G \geq \lambda_B$

If $\lambda_G \geq \lambda_B$, then with perfect inspections, it is optimal to inspect periodically (Theorem 1). A periodic inspection policy maintains a gap between consecutive inspections. This structure remains optimal with imperfect inspections.

Theorem 3 (Periodic imperfect inspections)
Suppose $\lambda_G \geq \lambda_B$. Then it is optimal to inspect periodically: for some period τ^*, the gap $T_n - T_{n-1}$ equals τ^* for each n.

As in the case of perfect inspections, we work with the agent's loss relative to his payoff U_0 from shirking forever in the absence of inspections. If the agent deviates by shirking until the next inspection, and this inspection is conducted at time t, then the agent's loss relative to U_0 is

$$L_S(t) = e^{-(\lambda_B + r)t}(U_0 - U_1 e^{-\delta t}). \tag{8}$$

The agent's effective discount factor is $e^{-(\lambda_B+r)t}$. When the inspection is conducted at time t, the agent does not forgo all of U_0. With probability $e^{-\delta t}$, the agent passes the inspection and gets his continuation value U_1. Consider first the relaxed problem that requires only that the agent prefers working, rather

[20]If $\lambda_G \neq \lambda_B$, then the inequality captures an additional effect. As the agent shirks, his continuation value decreases. If shirking lengthens the project ($\lambda_G > \lambda_B$), this decrease makes shirking less attractive, tightening the constraint. If shirking shortens the project ($\lambda_G < \lambda_B$), this decrease makes shirking more attractive, loosening the constraint.

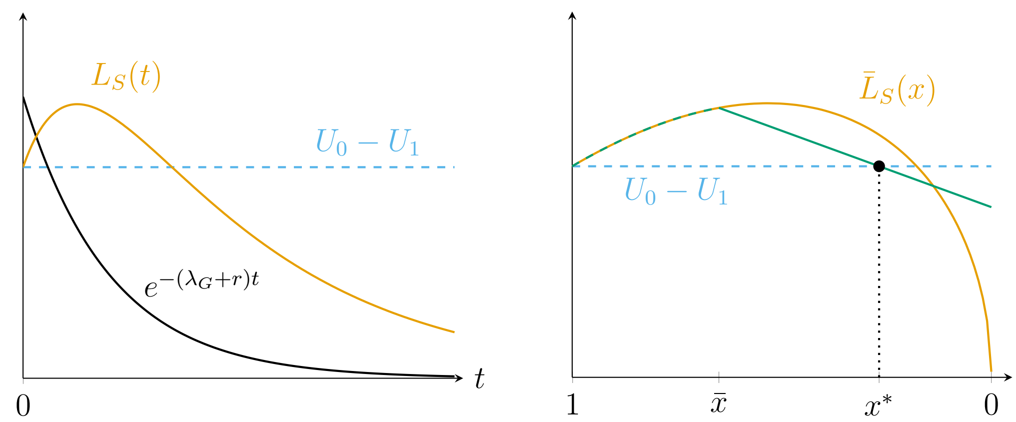

Figure 3. Shirking agent's loss from an imperfect inspection with $\lambda_G > \lambda_B$

than shirking until the next inspection. The principal chooses a positive random time T to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, e^{-(\lambda_G+r)T} \\ &\text{subject to} && \mathbb{E}\, L_S(T) \geq U_0 - U_1. \end{aligned}$$

As before, we change variables. Equivalently, the principal chooses a random variable $X = e^{-(\lambda_G+r)T}$ to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, X \\ &\text{subject to} && \mathbb{E}\, \bar{L}_S(X) \geq U_0 - U_1, \end{aligned}$$

where $\bar{L}_S(x) = x^{(\lambda_B+r)/(\lambda_G+r)}(U_0 - U_1 x^{\delta/(\lambda_G+r)})$.

In Figure 3, the left panel plots the agent's loss L_S and the principal's inspection cost as a function of the inspection time. After the change of variables, the right panel plots the agent's inspection loss $\bar{L}_S$ as a function of the principal's inspection cost. The loss function converges to $U_0 - U_1$ as t tends to 0 because very early inspections are very likely to be passed.

The binding deviation may be shirking followed by working. The right panel of Figure 3 plots in green the agent's loss, as a function of the inspection time, if prior to the next inspection the agent shirks until time $\bar{t}$ and then works thereafter. We set $\bar{x} = e^{-(\lambda_G+r)\bar{t}}$. If the inspection is conducted before time $\bar{t}$, then the agent shirks all the way until the inspection. Over this range, the agent's loss coincides with $\bar{L}_S$. For inspection times after time $\bar{t}$, the agent's loss is linear in the principal's inspection cost because, conditional upon reaching time $\bar{t}$, the project ends at the same rate under this deviation as it does on-path.

In the proof we identify the binding shirk-before-work deviation. We show that the resulting loss is concave as a function of the inspection time. In particular, concavity is preserved at the kink; otherwise, the shirk-before-work loss would lie above $\bar{L}_S$, contrary to the optimality of this deviation. The inspection period τ^* is determined by the intersection of this kinked loss curve with the threshold $U_0 - U_1$. The point $x^* = e^{-(\lambda_G + r)\tau^*}$ is shown on the plot. As in the case of perfect inspections, the periodic solution is unique if the binding deviation is to shirk all the way to the next inspection. If the binding deviation is to return to work at time $\bar{t}$, with $\bar{t} < \tau^*$, then every optimal inspection policy involves a no-inspection period of length at least $\bar{t}$. In particular, exponentially distributed inspections are strictly suboptimal.

6.2 Delaying breakdowns: $\lambda_G < \lambda_B$

If $\lambda_G < \lambda_B$, then with perfect inspections, it is optimal to conduct inspections with a constant hazard rate (Theorem 2). Under this memoryless policy, the gap between consecutive inspections can be arbitrarily small. With imperfect inspections, conducting two inspections in short succession cannot be optimal because the agent is almost certain to pass the second inspection, conditional on passing the first. The optimal policy leverages the benefits of randomization while also spacing out inspections.

To state the optimal inspection policy in this case, define the period $\hat{\tau}$ by $(\lambda_B - \lambda_G + \delta)e^{-\delta\hat{\tau}} = \lambda_B - \lambda_G$. Recall the function L_S defined in (8). Let $\hat{\delta}$ be defined by $L_S(\hat{\tau}) = U_0 - U_1$. Finally, denote by $\text{Exp}(\gamma)$ an exponentially distributed random variable with hazard rate γ.

Theorem 4 (Periodic/exponential imperfect inspections)
Suppose $\lambda_G < \lambda_B$.

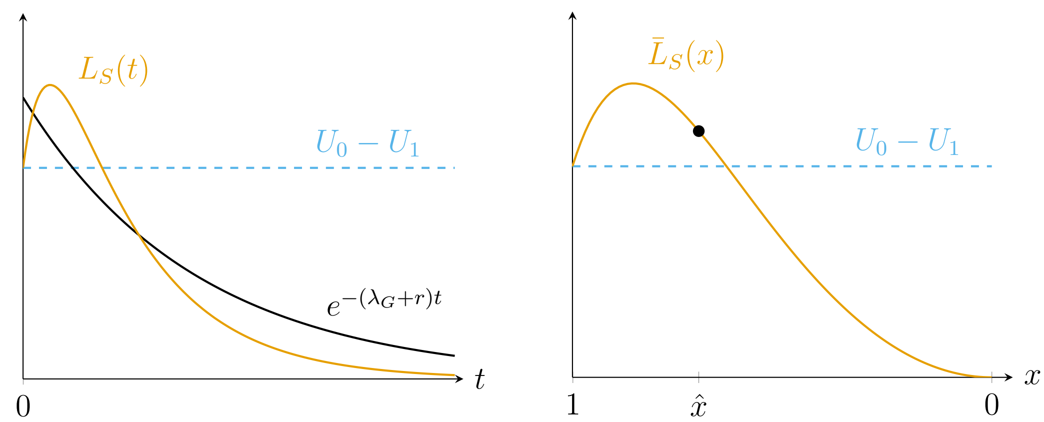

Figure 4. Shirking agent's loss from an imperfect inspection with $\lambda_G < \lambda_B$

1. *If $\delta \le \hat{\delta}$, then it is optimal to inspect periodically: for some period τ^* in $(0, \hat{\tau}]$, the gap $T_n - T_{n-1}$ equals τ^* for each n.*

2. *If $\delta > \hat{\delta}$, then the following policy is uniquely optimal. The gaps $(T_n - T_{n-1})_{n \ge 1}$ are independently and identically distributed. With some probability π^* in $(0,1)$, the gap $T_n - T_{n-1}$ equals $\hat{\tau}$. With probability $1 - \pi^*$, the gap $T_n - T_{n-1}$ has the distribution of $\hat{\tau} + \mathrm{Exp}(\gamma^*)$, where*

$$\gamma^* = \frac{(U_0 - U_1)(\lambda_B + r)(\lambda_B + r + \delta)}{U_1(\lambda_B + r + \delta) - U_0(\lambda_B + r)}. \tag{9}$$

Figure 4 plots the same loss functions as in Figure 3 in an example with $\lambda_G < \lambda_B$ and $\delta > \hat{\delta}$.[21] In general, the loss function $\bar{L}_S$ is concave and then convex. The right panel indicates the point $\hat{x} = e^{-(\lambda_G + r)\hat{\tau}}$. This point always lies before (left of) the point of inflection of $\bar{L}_S$. In this example, the solution (with the changed variables) puts positive probability on the point $\hat{x}$ and then puts density on points x right of $\hat{x}$.

If $\delta \le \hat{\delta}$, then inspecting before time $\hat{\tau}$ is necessary to motivate the agent to work. In this case, periodic inspections are optimal. The force toward spacing out inspections dominates. If $\delta > \hat{\delta}$, then inspecting with constant period $\hat{\tau}$ strictly induces the agent to work. To reduce inspection costs, the principal

[21] In this example, $\lambda_G + r = 1$; $\lambda_B + r = 2$; $U_0 = 2$; $U_1 = 5/4$; and $\delta = 5$.

could increase the period length, but this would not leverage the benefits of randomization with $\lambda_G < \lambda_B$. Instead, it is optimal for the principal to randomize the time of the next inspection, once time $\hat{\tau}$ has elapsed since the last inspection. The dynamic inspection policy proceeds as follows. The agent initially gets an inspection-free period of length $\hat{\tau}$. With positive probability, the agent is inspected exactly at time $\hat{\tau}$. If the agent is not inspected at time $\hat{\tau}$, then the next inspection arrives with a constant hazard rate thereafter. Once the inspection is conducted, the cycle repeats, beginning with a fresh period without inspections.[22]

The time $\hat{\tau}$ is decreasing in $\lambda_B - \lambda_G$ and δ, as we check in the proof. The greater the rate δ at which the agent's shirking leaves behind a paper trail, the shorter the no-inspection period. As $\lambda_B - \lambda_G$ increases, there is a stronger force towards randomization due to the relative convexity of the shirking agent's discount factor $e^{-(\lambda_B + r)T}$. As the detectability parameter δ tends to ∞, the optimal policy in Theorem 4 converges to the optimal exponential policy from Theorem 2. In particular, $\hat{\tau}$ and π^* both tend to 0.

7 Recovery

With the inspection technology in the main model, once the agent's shirking leaves behind evidence, the agent is certain to fail all subsequent inspections. We now consider how the solution changes if recovery is possible. Formally, the detectability state θ_t evolves as follows. Transitions from state 0 to state 1 occur at Poisson rate $\delta(1 - a_t)$, as before. Now transitions from 1 to 0 occur at Poisson rate ρa_t. The main model corresponds to $\rho = 0$.

The optimal policies in Theorem 3, if $u_0 > u_1$, and in Theorem 4 remain optimal if ρ is perturbed away from 0, as long as the detectability parameter δ is large enough; for a formal statement, see Appendix B.2. The reason is that the parameter ρ only makes a difference if the agent returns to work after

[22] The structure of this policy is similar to the optimal policy in Varas et al. (2020) in the case in which the principal directly values the information revealed by an inspection. In our model, however, the principal does not directly value this information. Moreover, the structure of the binding deviation is quite different in our problem.

shirking, which is not a binding deviation in the cases above.

If the recovery rate ρ is sufficiently large, however, the agent's incentives change substantially. If $\rho = 0$, then the result of an inspection depends only on the duration of shirking, not the allocation of this shirking over time. If $\rho > 0$, then for a given duration of shirking, the agent's passage probability is maximized if working is concentrated right before the inspection. With a high recovery rate ρ, window-dressing becomes effective—the agent can substantially increase his passage probability by working just before the inspection. In an education setting, this corresponds to "cramming" for an exam. To deter this deviation, it is optimal for the principal to conduct inspections with a constant hazard rate, so that the timing of past inspections provides no information about the timing of future inspections.

Theorem 5 (Exponential inspections with recovery)
Suppose that inspections have recovery rate $\rho > 0$. If $\rho + \lambda_G \geq \delta + \lambda_B$, then the following policy is uniquely optimal. The gaps $(T_n - T_{n-1})_{n \geq 1}$ are independently and identically distributed according to an exponential distribution with hazard rate γ^*, where

$$\gamma^* = \frac{(\lambda_B + r)(\lambda_G + r + \rho)(U_0 - U_1)}{U_1(\lambda_B + r + \delta) - U_0(\lambda_B + r)}. \tag{10}$$

The condition $\rho + \lambda_G \geq \delta + \lambda_B$ ensures that under the exponential policy, the agent has a stronger incentive to work if he has previously shirked. In this case the binding deviation is local—shirking briefly and then returning to work—rather than global as in our main model. If $\rho = 0$, then the condition in Theorem 5 becomes $\delta \leq \lambda_G - \lambda_B$, which is inconsistent with Assumption 2.B.[23] If $\lambda_G = \lambda_B$, then the condition in Theorem 5 reduces to the inequality $\rho \geq \delta$, i.e., the detectability state θ_t is more sensitive to the agent's action if $\theta_t = 1$ than if $\theta_t = 0$.[24] In this special case, Theorem 5 suggests that the force

[23]In fact, for $\delta \leq \lambda_G - \lambda_B + (\lambda_G + r)$, as long as the problem is feasible, it can be shown that periodic inspections are still optimal. There may be other optimal policies. In particular, if $\delta \leq \lambda_G - \lambda_B$, a slight modification of the proof of Theorem 5 shows that exponential inspections are also optimal, but the uniqueness part of the proof requires $\rho > 0$.

[24]If $\lambda_G \neq \lambda_B$, then there is an additional effect. If the agent shirks, his continuation value decreases. If shirking lengthens the project ($\lambda_G > \lambda_B$), this decrease makes working

towards the optimality of random inspections identified by Varas et al. (2020) extends to this richer inspection technology with $\rho \geq \delta$.

Finally, compare the hazard rate γ^* in Theorem 5 and Theorem 4. In Theorem 4, the binding deviation does not involve returning to work. Thus, the payoff from this deviation is independent of ρ. If $\rho + \lambda_G > \delta + \lambda_B$, then the local deviation binds. Shirking until the next inspection is strictly unprofitable, so the hazard rate in (10) must be strictly larger than the hazard rate in (9).

8 Conclusion

We study the design of inspection policies in a dynamic moral hazard setting. In contrast to previous work, dynamic deviations play a central role in our analysis. We show that predictable inspections are better suited to motivating agents who are tasked with achieving a breakthrough, such as entrepreneurs in an innovative industry. Random inspections are better suited to motivating agents who are tasked with avoiding a breakdown, such as safety personnel responsible for preventing accidents. This dichotomy is driven by the agent's effective risk attitude over time lotteries, which is determined endogenously by his actions.

more attractive, loosening the condition. If shirking shortens the project ($\lambda_G < \lambda_B$), this decrease makes working less attractive, tightening the condition.

A Main proofs

A.1 Notation

Throughout the proofs we use the notation $\lambda_0 := \lambda_B + r$ and $\lambda_1 := \lambda_G + r$. Let $U_i = u_i/\lambda_i$ for $i = 0, 1$. We express the solutions in terms of the five (strictly positive) parameters $\lambda_0, \lambda_1, \delta, U_0, U_1$.

A.2 Proof of Lemma 1

We discretize the agent's problem. Fix $\Delta > 0$. In the Δ-discretized problem, the agent can change his action only at times $k\Delta$ for $k = 0, 1, \ldots$. Let V_Δ denote the agent's supremal utility in the Δ-discretized problem. The Bellman equation reads

$$V_\Delta = \max_{i=0,1} \left\{ U_i(1 - e^{-\lambda_i \Delta}) + e^{-\lambda_i \Delta} V_\Delta \right\}.$$

The unique solution is $V_\Delta = \max\{U_0, U_1\}$. By a limiting argument,[25] it follows that the agent's value in the continuous-time problem is also $\max\{U_0, U_1\}$. Thus, working forever is optimal if and only if $U_1 \geq U_0$.

A.3 Proof outline for Theorems 1–4

The proofs of Theorems 1–4 have the following structure.

I. *Binding deviations.* Identify the candidate policy and a subclass of deviations.

II. *Relaxed problem.* Consider the relaxed problem that requires only that deviations in this subclass be unprofitable. Use Lagrangian relaxation to show that the candidate policy solves the relaxed problem.

[25] Any right-continuous function $a\colon [0, \infty) \to \{0, 1\}$ can be expressed as the pointwise limit of a sequence of step functions $a_n\colon [0, \infty) \to \{0, 1\}$ defined by $a_n(t) = a(k/2^n)$ if $(k-1)/2^n \leq t < k/2^n$, for $k = 1, 2, \ldots$. By dominated convergence, as n tends to ∞, the agent's expected utility from a_n converges to the agent's expected utility from a.

III. *Remaining deviations.* Check that the candidate policy is feasible in the original problem, i.e., that all deviations outside the subclass are unprofitable.

A.4 Proof of Theorem 1

Shirk-before-work deviation Suppose that until the next inspection, the agent plans to shirk over $[0, t)$ and work over $[t, \infty)$. Let $U_{\text{SW}}(t \mid \tau)$ denote the agent's expected payoff if the principal inspects at time τ. If $0 < t \leq \tau$,[26] then the agent is certain to fail the inspection, so

$$U_{\text{SW}}(t \mid \tau) = \int_0^t u_0 e^{-\lambda_0 s} \, \mathrm{d}s + e^{-\lambda_0 t} \int_t^\tau u_1 e^{-\lambda_1 (s-t)} \, \mathrm{d}s$$
$$= U_0(1 - e^{-\lambda_0 t}) + e^{-\lambda_0 t} U_1 (1 - e^{-\lambda_1(\tau - t)}).$$

Define the period τ^* to be the largest time τ such that

$$\sup_{t \in (0, \tau]} U_{\text{SW}}(t \mid \tau) \leq U_1. \tag{11}$$

It can be checked that τ^* is well-defined and strictly positive; moreover, $\tau = \tau^*$ satisfies (11) with equality.[27]

Let U'_{SW} denote the derivative of U_{SW} with respect to its first argument. We have

$$U'_{\text{SW}}(t \mid \tau^*) = e^{-\lambda_0 t} \left[(U_0 - U_1)\lambda_0 - U_1(\lambda_1 - \lambda_0) e^{-\lambda_1(\tau^* - t)} \right].$$

Since $\lambda_1 > \lambda_0$ (because $\lambda_G > \lambda_B$), this derivative is single-crossing in t from

[26] If $t = 0$ or $\tau = 0$, then $U_{\text{SW}}(t \mid \tau) = U_1$. If $0 < \tau < t$, then $U_{\text{SW}}(t \mid \tau) = U_{\text{SW}}(\tau \mid \tau) = U_0(1 - e^{-\lambda_0 \tau})$.

[27] For $0 < t \leq \tau$, we have

$$U_{\text{SW}}(t \mid \tau) \leq U_0(1 - e^{-\lambda_0 t}) + e^{-\lambda_0 t} U_0 (1 - e^{-\lambda_1(\tau - t)}) = U_0(1 - e^{-\lambda_1 \tau}),$$

and the right side tends to 0 as $\tau \to 0$. Hence, (11) is satisfied for τ sufficiently small. On the other hand, $U_{\text{SW}}(\tau \mid \tau) \to U_0$ as $\tau \to \infty$, so (11) is violated for τ sufficiently large. By Berge's theorem, the left side of (11) is continuous in τ.

above. Note that $U'_{\text{SW}}(\tau^* \mid \tau^*) = e^{-\lambda_0 \tau^*}(U_0\lambda_0 - U_1\lambda_1)$.[28]

1. If $U_1\lambda_1 \le U_0\lambda_0$, then $U'_{\text{SW}}(\tau^* \mid \tau^*) \ge 0$ and hence the function $U_{\text{SW}}(\cdot \mid \tau^*)$ is strictly increasing over $(0, \tau^*]$. Therefore, τ^* is given by $U_{\text{SW}}(\tau^* \mid \tau^*) = U_1$, hence
$$e^{-\lambda_0 \tau^*} = \frac{U_0 - U_1}{U_0}.$$
In this case, set $\bar{t} = \infty$ (below we consider the deviation in which the agent shirks until time $\bar{t}$).

2. If $U_1\lambda_1 > U_0\lambda_0$, then $U'_{\text{SW}}(\tau^* \mid \tau^*) < 0$. Therefore, $U_{\text{SW}}(\cdot \mid \tau^*)$ achieves its maximum over $(0, \tau^*]$ at a unique point, denoted $\bar{t}$, which lies in $(0, \tau^*)$. The times τ^* and $\bar{t}$ are given by
$$e^{-\lambda_0 \tau^*} = \frac{\lambda_1 - \lambda_0}{\lambda_1} \left(\frac{\lambda_0(U_0 - U_1)}{U_1(\lambda_1 - \lambda_0)} \right)^{\lambda_0/\lambda_1}, \qquad e^{-\lambda_0 \bar{t}} = \frac{\lambda_1 - \lambda_0}{\lambda_1}.$$

Relaxed problem With $\bar{t}$ defined above, consider the relaxed problem of choosing a positive random variable T to solve
$$\begin{aligned} &\text{minimize} && \mathbb{E}\, e^{-\lambda_1 T} \\ &\text{subject to} && \mathbb{E}\, U_{\text{SW}}(\bar{t} \mid T) \le U_1, \end{aligned}$$
where we set $U_{\text{SW}}(\infty \mid T) = U_0(1 - e^{-\lambda_0 T})$. We change variables. Define $\bar{U} \colon [0,1] \to \mathbf{R}$ by
$$\bar{U}(x) = U_{\text{SW}}(\bar{t} \mid -\lambda_1^{-1} \log x).$$
Consider the equivalent relaxed problem of choosing a random variable $X = e^{-\lambda_1 T}$ to solve
$$\begin{aligned} &\text{minimize} && \mathbb{E}\, X \\ &\text{subject to} && \mathbb{E}\, \bar{U}(X) \le U_1. \end{aligned} \tag{12}$$

[28]Technically, this is the left derivative. Throughout the proofs, derivatives evaluated at kinks should be interpreted as either left or right derivatives. The appropriate choice should be clear from context.

Set $\bar{x} = e^{-\lambda_1 \bar{t}} \in [0, 1)$. We can express $\bar{U}$ as

$$\bar{U}(x) = \begin{cases} U_0(1 - x^{\lambda_0/\lambda_1}) & \text{if } x \geq \bar{x}, \\ U_0(1 - \bar{x}^{\lambda_0/\lambda_1}) + U_1 \bar{x}^{\lambda_0/\lambda_1}(1 - x/\bar{x}) & \text{if } x < \bar{x}. \end{cases}$$

Clearly, $\bar{U}$ is strictly decreasing. Since $\lambda_1 > \lambda_0$, the function $\bar{U}$ is convex, strictly so over $[\bar{x}, 1]$.[29]

We check that the constant $x^* = e^{-\lambda_1 \tau^*}$ solves (12). By construction, $\bar{U}(x^*) = U_1$, so x^* is feasible. If a random variable X satisfies $\mathbb{E}\, X < x^*$, then X is infeasible because

$$\mathbb{E}[\bar{U}(X)] \geq \bar{U}(\mathbb{E}\, X) > \bar{U}(x^*) = U_1,$$

where the inequalities hold because $\bar{U}$ is convex and strictly decreasing.

Now we turn to uniqueness. If $U_1\lambda_1 \leq U_0\lambda_0$, then $\bar{x} = 0$. Therefore, $\bar{U}$ is strictly convex over $[0, 1]$, so x^* is the unique solution of (12). If $U_1\lambda_1 > U_0\lambda_0$, then $0 < x^* < \bar{x} < 1$. Since $\bar{U}$ is strictly convex over $[\bar{x}, 1]$, any solution X of (12) must concentrate on $[0, \bar{x}]$, hence $T = -\lambda_1^{-1} \log X$ must concentrate on $[\bar{t}, \infty]$.

Remaining deviations Suppose that the principal conducts the next inspection at time τ^*. By construction, no shirk-before-work deviation is profitable. We prove that no other deviation is profitable. By a limiting argument (see Footnote 25), it suffices to check that for each positive integer n, in the discretized problem with $\Delta = \tau^*/n$, any deviation outside the shirk-before-work class can be strictly improved upon. Suppose that for some k, with $k \leq n-2$, the agent works over $[k\Delta, (k+1)\Delta)$ and shirks over $[(k+1)\Delta, (k+2)\Delta)$. The agent can strictly improve his payoff by instead shirking over $[k\Delta, (k+1)\Delta)$

[29] Convexity is preserved even if there is a kink because $\bar{t}$ is agent-optimal. We have $\bar{x} > 0$ if and only if $U_1\lambda_1 > U_0\lambda_0$, in which case

$$D_+\bar{U}(\bar{x}) = -\frac{U_0\lambda_0}{\lambda_1}\bar{x}^{(\lambda_0-\lambda_1)/\lambda_1} > -U_1\bar{x}^{(\lambda_0-\lambda_1)/\lambda_1} = D_-\bar{U}(\bar{x}),$$

where D_+ and D_- denote right and left derivatives.

and working over $[(k+1)\Delta, (k+2)\Delta)$ since

$$U_0(1-e^{-\lambda_0\Delta}) + e^{-\lambda_0\Delta}U_1(1-e^{-\lambda_1\Delta}) > U_1(1-e^{-\lambda_1\Delta}) + e^{-\lambda_1\Delta}U_0(1-e^{-\lambda_0\Delta}).$$

Each side of this inequality is a weighted combination of U_0 and U_1 with total weight $1-e^{-\lambda_0\Delta-\lambda_1\Delta}$, but the left side puts strictly more weight on the larger utility U_0. By switching the order (but not the duration) of shirking and working, the agent's probability of passing the inspection at time τ^* does not change.

A.5 Proof of Theorem 2

Work-before-shirk deviations Suppose that until the next inspection, the agent plans to work over $[0,t)$ and shirk over $[t,\infty)$. Let $U_{\mathrm{WS}}(t \mid \tau)$ denote the agent's expected payoff if the principal inspects at time τ. If $t \geq \tau$, then $U_{\mathrm{WS}}(t \mid \tau) = U_1$. If $t < \tau$, then the agent is certain to fail the inspection, so

$$\begin{aligned} U_{\mathrm{WS}}(t \mid \tau) &= \int_0^t u_1 e^{-\lambda_1 s}\,\mathrm{d}s + e^{-\lambda_1 t}\int_t^\tau u_0 e^{-\lambda_0(s-t)}\,\mathrm{d}s \\ &= U_1(1-e^{-\lambda_1 t}) + e^{-\lambda_1 t}U_0(1-e^{-\lambda_0(\tau-t)}) \\ &= U_1 + e^{-\lambda_1 t}h(t,\tau), \end{aligned}$$

where $h(t,\tau) = U_0(1-e^{-\lambda_0(\tau-t)}) - U_1$.

Relaxed problem Consider the relaxed problem of choosing a positive random variable T to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, e^{-\lambda_1 T} \\ &\text{subject to} && \mathbb{E}\, U_{\mathrm{WS}}(t \mid T) \leq U_1, \quad t \geq 0. \end{aligned}$$

This problem is equivalent to choosing a cumulative distribution function F supported on $[0, \infty)$ to solve[30]

$$\begin{aligned} &\text{minimize} \quad \int_{[0,\infty)} e^{-\lambda_1 t} \, \mathrm{d}F(t) \\ &\text{subject to} \quad \int_{(t,\infty)} h(t,s) \, \mathrm{d}F(s) \leq 0, \quad t \geq 0. \end{aligned} \tag{13}$$

The constraints are indexed by times $t \geq 0$. Attach a nonnegative mass multiplier η_0 to the time-0 constraint and a nonnegative density multiplier $\eta(t)$ to the time-t constraint, for all $t \geq 0$. The Lagrangian becomes

$$\begin{aligned} L(F; \eta_0, \eta) = \int_{[0,\infty)} e^{-\lambda_1 t} \, \mathrm{d}F(t) + \eta_0 \int_{(0,\infty)} h(0,s) \, \mathrm{d}F(s) \\ + \int_0^\infty \left[\int_{(t,\infty)} h(t,s) \, \mathrm{d}F(s) \right] \eta(t) \, \mathrm{d}t. \end{aligned}$$

Change the order of integration in the double integral[31] and relabel the dummy variables to get

$$L(F; \eta_0, \eta) = F(0) + \int_{(0,\infty)} I(t) \, \mathrm{d}F(t),$$

where

$$I(t) = e^{-\lambda_1 t} + \eta_0 h(0,t) + \int_0^t \eta(s) h(s,t) \, \mathrm{d}s.$$

Let $\eta(t) = \bar{\eta} e^{-\lambda_1 t}$ for some $\bar{\eta} \geq 0$. Plug in this expression, integrate, and group like terms to obtain

$$\begin{aligned} I(t) = e^{-\lambda_1 t} \left(1 - \frac{\bar{\eta} U_0}{\lambda_0 - \lambda_1} - \frac{\bar{\eta}(U_0 - U_1)}{\lambda_1} \right) + e^{-\lambda_0 t} U_0 \left(-\eta_0 + \frac{\bar{\eta}}{\lambda_0 - \lambda_1} \right) \\ + (U_0 - U_1) \left(\eta_0 + \frac{\bar{\eta}}{\lambda_1} \right). \end{aligned}$$

[30] Since F is not necessarily continuous, we explicitly specify which endpoints are included in integrals against F.

[31] Here and below, we freely interchange the order of integration because h is bounded and we will only evaluate the Lagrangian for dual functions that are bounded.

To make $I(t)$ constant in t, we set

$$\bar{\eta} = \left(\frac{U_0}{\lambda_0 - \lambda_1} + \frac{U_0 - U_1}{\lambda_1} \right)^{-1} \quad \text{and} \quad \eta_0 = \frac{\bar{\eta}}{\lambda_0 - \lambda_1}.$$

Since $\lambda_1 < \lambda_0$, both $\bar{\eta}$ and η_0 are well-defined and strictly positive. With these multipliers, the Lagrangian reduces to

$$F(0) + \frac{(U_0 - U_1)\lambda_0}{(U_0 - U_1)\lambda_0 + U_1 \lambda_1} [1 - F(0)].$$

The coefficient on $1 - F(0)$ is strictly less than 1, so the Lagrangian is minimized by any cumulative distribution function F with $F(0) = 0$. By the Kuhn–Tucker conditions, a cumulative distribution function F solves (13) if and only if $F(0) = 0$ and

$$\frac{U_0}{U_0 - U_1} \int_{(t,\infty)} e^{-\lambda_0 (s-t)} \, \mathrm{d}F(s) = 1 - F(t),$$

for every time $t \geq 0$.[32] By Lemma 2 in Appendix B.3, the unique solution is the exponential distribution with hazard rate γ^* given by $\gamma^* / \lambda_0 = (U_0 - U_1)/U_1$.

Remaining deviations Suppose that the time until the next inspection is exponentially distributed with hazard rate γ^*. By construction, no work-before-shirk deviation is profitable. We prove that no other deviation is profitable. It suffices to show that once the agent has shirked, he finds it optimal to shirk until the next inspection. Once the agent has shirked, he is certain to fail the next inspection. Therefore, the project ends at Poisson rate

$$(\lambda_G + \gamma^*) a_t + (\lambda_B + \gamma^*)(1 - a_t).$$

[32] Complementary slackness requires equality for $t = 0$ and for almost every $t > 0$, but each side of this equality is right-continuous (by dominated convergence), so the equality must hold for *every* time $t \geq 0$.

The agent's problem is identical to the no-inspection problem with $\lambda_i + \gamma^*$ in place of λ_i, for $i = 0, 1$. By Lemma 1, shirking forever is optimal if and only if

$$\frac{U_1 \lambda_1}{\lambda_1 + \gamma^*} \leq \frac{U_0 \lambda_0}{\lambda_0 + \gamma^*},$$

which holds (strictly) because $U_1 < U_0$ and $\lambda_1 < \lambda_0$.

A.6 Proof of Theorems 3 and 4

Our unified proof is organized around the structure of the binding deviations.

Shirk-before-work deviations Suppose that until the next inspection, the agent plans to shirk over $[0, t)$ and work over $[t, \infty)$. Let $U_{\text{SW}}(t \mid \tau)$ denote the agent's expected payoff if the principal inspects at time τ. For $t \leq \tau$,[33] the agent's passage probability is $e^{-\delta t}$, so

$$\begin{aligned} U_{\text{SW}}(t \mid \tau) &= \int_0^t u_0 e^{-\lambda_0 s} \, \mathrm{d}s + e^{-\lambda_0 t} \int_t^\tau u_1 e^{-\lambda_1 (s-t)} \, \mathrm{d}s + e^{-\lambda_0 t - \lambda_1(\tau - t)} U_1 e^{-\delta t} \\ &= U_0(1 - e^{-\lambda_0 t}) + U_1 e^{-\lambda_0 t} \left[1 - e^{-\lambda_1(\tau - t)}(1 - e^{-\delta t})\right]. \end{aligned}$$

Define the period τ^* to be the largest time τ such that

$$\max_{t \in [0, \tau]} U_{\text{SW}}(t \mid \tau) \leq U_1. \tag{14}$$

It can be checked that τ^* is well-defined and strictly positive; moreover, $\tau = \tau^*$ satisfies (14) with equality.[34] For all t in $[0, \tau^*]$, we have $U_{\text{SW}}(t \mid \tau^*) \leq U_1 = U_{\text{SW}}(0 \mid \tau^*)$. It follows that $U'_{\text{SW}}(0 \mid \tau^*) \leq 0$, where U'_{SW} denotes the derivative with respect to the first argument.

[33] If $t > \tau$, then $U_{\text{SW}}(t \mid \tau) = U_{\text{SW}}(\tau \mid \tau) = U_0(1 - e^{-\lambda_0 \tau}) + U_1 e^{-(\lambda_0 + \delta)\tau}$.

[34] For $0 < t \leq \tau$, as (t, τ) tends to $(0, 0)$, the derivative $U'_{\text{SW}}(t \mid \tau)$ tends to $\lambda_0(U_0 - U_1) - \delta U_1$, which is strictly negative by Assumption 2.A. Thus, for τ sufficiently small, the maximum on the left side of (14) is achieved at $t = 0$ and hence (14) holds. On the other hand, $U_{\text{SW}}(\tau \mid \tau) \to U_0$ as $\tau \to \infty$, so (14) is violated for τ sufficiently large. By Berge's theorem, the left side of (14) is continuous in τ.

We claim that $U'_{\mathrm{SW}}(0 \mid \tau^*) < 0$. Suppose not. Then

$$0 = U'_{\mathrm{SW}}(0 \mid \tau^*) = (U_0 - U_1)\lambda_0 - e^{-\lambda_1 \tau^*} U_1 \delta. \tag{15}$$

Differentiating U_{SW} twice with respect to the first argument gives

$$\begin{aligned} U''_{\mathrm{SW}}(0 \mid \tau^*) &= -(U_0 - U_1)\lambda_0^2 + e^{-\lambda_1 \tau^*} U_1 \delta(\delta + 2\lambda_0 - 2\lambda_1) \\ &= (U_0 - U_1)\lambda_0(\delta + \lambda_0 - 2\lambda_1) \\ &> 0, \end{aligned}$$

where the second equality uses (15) and the inequality uses Assumption 2.B. It follows that $U_{\mathrm{SW}}(t \mid \tau^*) > U_{\mathrm{SW}}(0 \mid \tau^*) = U_1$ for t sufficiently small, contrary to the definition of τ^*.

Let $\bar{t}$ be the largest maximizer of $U_{\mathrm{SW}}(\cdot \mid \tau^*)$ over $[0, \tau^*]$. We claim that $\bar{t} > 0$. Suppose not. Then $U_{\mathrm{SW}}(t \mid \tau^*) < U_1$ for all t in $(0, \tau^*]$. Since $U'_{\mathrm{SW}}(0 \mid \tau^*) < 0$, it follows that (14) holds for some τ strictly greater than τ^*, contrary to the definition of τ^*.

We separate into two cases according to the condition

$$U_1(\lambda_0 + \delta - \lambda_1)e^{-\delta \tau^*} \geq U_0(\lambda_0 - \lambda_1). \tag{16}$$

If $\lambda_1 \geq \lambda_0$, then (16) holds (because $U_1 e^{-\delta\tau^*} < U_0$).

A. Relaxed problem: shirk-before-work Suppose that (16) holds. With $\bar{t}$ defined above, consider the relaxed problem of choosing a positive random variable T to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, e^{-\lambda_1 T} \\ &\text{subject to} && \mathbb{E}\, U_{\mathrm{SW}}(\bar{t} \mid T) \leq U_1. \end{aligned}$$

We change variables. Define functions $\bar{U}$ and $\bar{U}_S$ from $[0,1]$ to $\mathbf{R}$ by

$$\bar{U}(x) = U_{\mathrm{SW}}(\bar{t} \mid -\lambda_1^{-1} \log x), \qquad \bar{U}_S(x) = U_S(-\lambda_1^{-1} \log x \mid -\lambda_1^{-1} \log x).$$

Consider the equivalent relaxed problem of choosing a random variable $X = e^{-\lambda_1 T} \in [0,1]$ to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, X \\ &\text{subject to} && \mathbb{E}\, \bar{U}(X) \leq U_1. \end{aligned} \tag{17}$$

Set $\bar{x} = e^{-\lambda_1 \bar{t}} \in (0,1)$. We can express $\bar{U}$ as

$$\bar{U}(x) = \begin{cases} \bar{U}_S(x) & \text{if } x \geq \bar{x}, \\ \bar{U}_S(\bar{x}) + \alpha(\bar{x})(x - \bar{x}) & \text{if } x < \bar{x}, \end{cases}$$

where $\alpha(\bar{x}) = -U_1 \bar{x}^{\lambda_0/\lambda_1}(1 - \bar{x}^{\delta/\lambda_1}) < 0$. Clearly, $\bar{U}$ is strictly decreasing over $[0, \bar{x}]$. Since (16) holds, it can be shown that $\bar{U}$ is convex, strictly so over $[\bar{x}, 1]$.[35]

We check that the constant $x^* = e^{-\lambda_1 \tau^*}$ solves (17). By construction, $\bar{U}(x^*) = U_1$, so x^* is feasible. If a random variable X satisfies $\mathbb{E}\, X < x^*$, then X is infeasible because

$$\mathbb{E}[\bar{U}(X)] \geq \bar{U}(\mathbb{E}\, X) > \bar{U}(x^*) = U_1,$$

where the first inequality holds because $\bar{U}$ is convex and the second holds

[35] Differentiate $\bar{U}_S$ twice to get

$$\bar{U}_S''(x) = \lambda_1^{-2} x^{(\lambda_0 - 2\lambda_1)/\lambda_1} \left[U_1(\lambda_0 + \delta)(\lambda_0 + \delta - \lambda_1) x^{\delta/\lambda_1} - U_0 \lambda_0 (\lambda_0 - \lambda_1) \right].$$

By Assumption 2.B, $\lambda_0 + \delta - \lambda_1 > \lambda_1 > 0$. There are two cases. If $\lambda_1 \geq \lambda_0$, then $\bar{U}_S$ is globally convex. In this case, set $x_c = 0$. If $\lambda_1 < \lambda_0$, then define the point x_c in $(0,1)$ by $\bar{U}_S(x_c) - \bar{U}_S(0) = x_c \bar{U}_S'(x_c)$, or equivalently,

$$U_1(\lambda_0 + \delta - \lambda_1) x_c^{\delta/\lambda_1} = U_0(\lambda_0 - \lambda_1).$$

By (16), this point x_c exists and satisfies $x_c \leq x^* \leq \bar{x}$.

In both cases, the convexification of $\bar{U}_S$, denoted $\operatorname{cvx} \bar{U}_S$, coincides with $\bar{U}_S$ over $[x_c, 1]$. It remains to check that convexity is preserved at the kink. If $\bar{U}'$ jumped down at the kink, then for all $x \leq \bar{x}$, we would have $\bar{U}(x) \leq \operatorname{cvx} \bar{U}_S(x) \leq \bar{U}_S(x)$. In particular, taking $x = x^*$ gives

$$U_{\text{SW}}(\bar{t} \mid \tau^*) = \bar{U}(x^*) \leq \bar{U}_S(x^*) = U_{\text{SW}}(\tau^* \mid \tau^*),$$

contrary to the definition of $\bar{t}$.

because $\bar{U}$ is strictly decreasing over $[0, \bar{x}]$. Since $x^* < \bar{x}$ and $\bar{U}$ is strictly convex over $[\bar{x}, 1]$, any solution X of the relaxed problem must concentrate on $[0, \bar{x}]$, so $T = -\lambda_1^{-1} \log X$ must concentrate on $[\bar{t}, \infty]$.

B. Relaxed problem: work-before-shirk Suppose that (16) is violated. It follows that $\lambda_1 < \lambda_0$. Suppose that until the next inspection, the agent plans to work over $[0, t)$ and shirk over $[t, \infty)$. Let $U_{\mathrm{WS}}(t \mid \tau)$ denote the agent's expected payoff if the principal inspects at time τ. If $t \geq \tau$, then $U_{\mathrm{WS}}(t \mid \tau) = U_1$. If $t < \tau$, then the agent's passage probability is $e^{-\delta(\tau - t)}$, so

$$
\begin{aligned}
U_{\mathrm{WS}}(t \mid \tau) &= U_1(1 - e^{-\lambda_1 t}) + e^{-\lambda_1 t} U_0 (1 - e^{-\lambda_0(\tau - t)}) + e^{-\lambda_1 t - \lambda_0(\tau - t)} U_1 e^{-\delta(\tau - t)} \\
&= U_1 + e^{-\lambda_1 t} h(t, \tau),
\end{aligned}
$$

where $h(t, \tau) = U_0(1 - e^{-\lambda_0(\tau - t)}) - U_1(1 - e^{-(\lambda_0 + \delta)(\tau - t)})$.

Consider the relaxed problem of choosing a positive random variable T to solve

$$
\begin{aligned}
&\text{minimize} \quad && \mathbb{E}\, e^{-\lambda_1 T} \\
&\text{subject to} \quad && \mathbb{E}\, U_{\mathrm{WS}}(t \mid T) \leq U_1, \quad t \in \{0\} \cup [\tau, \infty),
\end{aligned}
$$

where the value of τ will be determined below. This problem is equivalent to choosing a cumulative distribution function F supported on $[0, \infty)$ to solve

$$
\begin{aligned}
&\text{minimize} \quad && \int_{[0,\infty)} e^{-\lambda_1 t} \, \mathrm{d}F(t) \\
&\text{subject to} \quad && \int_{(t,\infty)} h(s, t) \, \mathrm{d}F(s) \leq 0, \quad t \in \{0\} \cup [\tau, \infty).
\end{aligned}
\tag{18}
$$

The constraints are indexed by times t in $\{0\} \cup [\tau, \infty)$. Attach a nonnegative mass multiplier η_0 to the time-0 constraint and a nonnegative density multiplier $\eta(t)$ to the time-t constraint, for all $t \geq \tau$. The Lagrangian becomes

$$
\begin{aligned}
L(F; \eta_0, \eta) = \int_{[0,\infty)} e^{-\lambda_1 t} \, \mathrm{d}F(t) &+ \eta_0 \int_{(0,\infty)} h(0, s) \, \mathrm{d}F(s) \\
&+ \int_{\tau}^{\infty} \left[\int_{(t,\infty)} h(t, s) \, \mathrm{d}F(s) \right] \eta(t) \, \mathrm{d}t.
\end{aligned}
$$

Change the order of integration in the double integral and relabel the dummy variables to get

$$L(F;\eta_0,\eta) = \int_{[0,\infty)} I(t)\,\mathrm{d}F(t),$$

where

$$I(t) = e^{-\lambda_1 t} + \eta_0 h(0,t) + \int_{\tau\wedge t}^{t} \eta(s)h(s,t)\,\mathrm{d}s.$$

To get this expression for the Lagrangian, we used the fact that $I(0) = 1$, which holds because $h(0,0) = 0$.

For some $\bar\eta \geq 0$, let $\eta(t) = \bar\eta e^{-\lambda_1 t}$ for all $t \geq \tau$. Plug in this expression, integrate, and group like terms. For $t \geq \tau$, we get

$$\begin{aligned} I(t) = e^{-\lambda_1 t}&\left[1 - \frac{\bar\eta}{\lambda_1}\left(\frac{U_0\lambda_0}{\lambda_0-\lambda_1} - \frac{U_1(\lambda_0+\delta)}{\lambda_0+\delta-\lambda_1}\right)\right] \\ &+ e^{-\lambda_0 t}U_0\left[-\eta_0 + \frac{\bar\eta}{\lambda_0-\lambda_1}e^{(\lambda_0-\lambda_1)\tau}\right] \\ &+ e^{-(\lambda_0+\delta)t}U_1\left[\eta_0 - \frac{\bar\eta}{\lambda_0+\delta-\lambda_1}e^{(\lambda_0+\delta-\lambda_1)\tau}\right] \\ &+ (U_0-U_1)\left[\eta_0 + \frac{\bar\eta}{\lambda_1}e^{-\lambda_1\tau}\right]. \end{aligned} \tag{19}$$

Recall that $\hat\tau$ is defined by $e^{-\delta\hat\tau} = (\lambda_0-\lambda_1)/(\lambda_0+\delta-\lambda_1)$. Let $\tau = \tau^* \wedge \hat\tau$. Define $\bar\eta$ and η_0 by

$$\bar\eta = \lambda_1\left(\frac{U_0\lambda_0}{\lambda_0-\lambda_1} - \frac{U_1(\lambda_0+\delta)}{\lambda_0+\delta-\lambda_1}\right)^{-1}, \qquad \eta_0 = \frac{\lambda_1 e^{(\lambda_0-\lambda_1)\tau}}{U_0\lambda_0 - e^{-\delta\tau}U_1(\lambda_0+\delta)}.$$

The multiplier $\bar\eta$ is well-defined and positive because $U_1 < U_0$ and $(\lambda_0+\delta)(\lambda_0-\lambda_1) < \lambda_0(\lambda_0+\delta-\lambda_1)$. The multiplier η_0 is well-defined because

$$e^{-\delta\tau} < \frac{U_0(\lambda_0-\lambda_1)}{U_1(\lambda_0+\delta-\lambda_1)} < \frac{U_0\lambda_0 - U_1\lambda_1}{U_1(\lambda_0+\delta-\lambda_1)} < \frac{U_0\lambda_0}{U_1(\lambda_0+\delta)}, \tag{20}$$

where the first inequality holds because (16) is violated; the second inequality uses $U_0 > U_1$; and the third inequality holds by Assumption 2.

The definition of $\bar\eta$ eliminates the first line in (19). Now consider η_0. If

$\tau^* \geq \hat{\tau}$, so $\tau = \hat{\tau}$, then the second and third lines of (19) vanish as well. If $\tau^* < \hat{\tau}$, so $\tau < \hat{\tau}$, it can be checked that the derivative of (19) is zero at $t = \tau$ and is strictly positive over $[\tau, \infty)$.[36] In both cases, it can be shown that I is strictly decreasing over $[0, \tau]$.[37] We separately characterize the minimizers of the Lagrangian in the two cases.

First suppose $\tau^* < \hat{\tau}$. In this case, the integrand I is uniquely minimized at τ^*. Therefore, a point mass at τ^* is the unique minimizer of the Lagrangian. Clearly, $U_{\text{WS}}(t \mid \tau^*) = U_1$ for all $t \geq \tau = \tau^*$. To see that a point mass at τ^* solves the relaxed problem, we check that $U_{\text{WS}}(0 \mid \tau^*) = U_1$, or equivalently, $U_{\text{SW}}(\tau^* \mid \tau^*) = U_1$. We showed in the first section of the proof that $U'_{\text{SW}}(0 \mid \tau^*) < 0$. Since (16) is violated, we have

$$\begin{aligned} U'_{\text{SW}}(\tau^* \mid \tau^*) &= e^{-\lambda_0 \tau^*} \big(U_0 \lambda_0 - U_1 \lambda_1 - e^{-\delta \tau^*} U_1 (\lambda_0 + \delta - \lambda_1) \big) \\ &\geq e^{-\lambda_0 \tau^*} \lambda_1 (U_0 - U_1) \\ &> 0. \end{aligned} \tag{21}$$

The function $t \mapsto U'_{\text{SW}}(t \mid \tau^*)$ is a sum of at most three exponentials, so it has at most two zeros by Descartes' rule of signs. Therefore, $U'_{\text{SW}}(t \mid \tau^*)$ cannot cross zero from above over $(0, \tau^*)$, and hence $U_{\text{SW}}(\cdot \mid \tau^*)$ cannot have an interior maximizer over $[0, \tau^*]$. Since $\bar{t} > 0$, we must have $\bar{t} = \tau^*$. Thus, $U_{\text{SW}}(\tau^* \mid \tau^*) = U_1$.

Next suppose $\tau^* \geq \hat{\tau}$. In this case, $\operatorname{argmin}_{t \geq 0} I(t) = [\hat{\tau}, \infty)$, so the Lagrangian is minimized by any distribution supported on $[\hat{\tau}, \infty)$. By the Kuhn–Tucker conditions (see Footnote 32), a cumulative distribution function F solves the relaxed problem (18) if and only if F is supported on $[\hat{\tau}, \infty)$, and F

[36] If $\tau < \hat{\tau}$, then in (19) the coefficient on $e^{-\lambda_0 t}$ is negative and the coefficient on $e^{-(\lambda_0 + \delta)t}$ is positive. After differentiating, these signs flip, so the derivative becomes positive for $t > \tau$.

[37] Since $h(t, t) = 0$ for all t, the integrand I is differentiable at $t = \tau$, and we have $I'(\tau) = 0$. To prove that $I'(t) < 0$ for $t < \tau$, we show that that $e^{\lambda_1 t} I'(t)$ is strictly increasing over $[0, \tau]$. For $t < \tau \leq \hat{\tau}$ we have

$$\big(e^{\lambda_1 t} I'(t)\big)' = e^{-(\lambda_0 - \lambda_1)t} \eta_0 \left[U_1 (\lambda_0 + \delta) e^{-\delta t} (\lambda_0 + \delta - \lambda_1) - U_0 \lambda_0 (\lambda_0 - \lambda_1) \right] > 0,$$

where the inequality holds because $U_1(\lambda_0 + \delta) > U_0 \lambda_0$ (by Assumption 2.A) and $e^{-\delta t}(\lambda_0 + \delta - \lambda_1) > \lambda_0 - \lambda_1$ (since $t < \hat{\tau}$).

satisfies

$$\int_{(0,\infty)} \left[\frac{U_0}{U_0 - U_1} e^{-\lambda_0 s} - \frac{U_1}{U_0 - U_1} e^{-(\lambda_0+\delta)s} \right] \mathrm{d}F(s) = 1, \tag{22}$$

and for every $t \geq \hat{\tau}$,

$$\int_{(t,\infty)} \left[\frac{U_0}{U_0 - U_1} e^{-\lambda_0 (s-t)} - \frac{U_1}{U_0 - U_1} e^{-(\lambda_0+\delta)(s-t)} \right] \mathrm{d}F(s) = 1 - F(t). \tag{23}$$

By Lemma 3 in Appendix B.3, for each π in $[0,1]$ there is a unique cumulative distribution function F supported on $[\hat{\tau}, \infty)$ with $F(\hat{\tau}) = \pi$ that satisfies (23) for every $t \geq \hat{\tau}$. Among these solutions, let $\hat{\pi}$ be the unique value of π for which (22) holds. To see that such a probability $\hat{\pi}$ exists, note that left side of (22) is at most 1 if $\pi = 1$ (since $\tau^* \geq \hat{\tau}$) and is strictly greater than 1 if $\pi = 0$ (by the definition of γ^*). Moreover, $\hat{\pi}$ is unique because the integrand is strictly decreasing in s for $s \geq \tau = \hat{\tau}$ by (20).

Discount factor threshold in Theorem 4 We now relate the solution to the detectability parameter δ. Clearly, τ^* is increasing in δ. On the other hand, $\hat{\tau}$ is decreasing in δ since we have

$$\hat{\tau} = \frac{\log(\lambda_0 - \lambda_1 + \delta) - \log(\lambda_0 - \lambda_1)}{\delta},$$

and log is a concave function. Therefore, as a function of δ, the difference $\tau^* - \hat{\tau}$ is strictly increasing in δ. Observe that $\tau^* \downarrow 0$ as $\delta \downarrow (\lambda_B + r)(U_0 - U_1)/U_1$ (the lower bound in Assumption 2.A), and $\hat{\tau} \downarrow 0$ as $\delta \uparrow \infty$. We conclude that there exists a unique threshold $\hat{\delta}$ in $((\lambda_B + r)(U_0 - U_1)/U_1, \infty)$ for which τ^* coincides with $\hat{\tau}$. If $\delta \leq \hat{\delta}$, then in both relaxed problems, the solution is periodic. If $\delta \geq \hat{\delta}$, then $\tau^* \geq \hat{\tau}$, so

$$e^{-\delta \tau^*} < \frac{\lambda_0 - \lambda_1}{\lambda_0 + \delta - \lambda_1} < \frac{U_0(\lambda_0 - \lambda_1)}{U_1(\lambda_0 + \delta - \lambda_1)}.$$

Thus, (16) is violated, and hence the relaxed solution is periodic/exponential.

Finally, we check that $\hat{\delta}$ can be defined by the formula given in the main text. If $\delta = \hat{\delta}$, then $U_{\mathrm{SW}}(\tau^* \mid \tau^*) = U_1$ and hence $U_{\mathrm{SW}}(\hat{\tau} \mid \hat{\tau}) = U_1$. For any δ, we have

$$U_{\mathrm{SW}}(\hat{\tau} \mid \hat{\tau}) = U_0 - e^{-\lambda_0 \hat{\tau}} \left(U_0 - U_1 e^{-\delta \hat{\tau}}\right) = U_0 - e^{-\lambda_0 \hat{\tau}} \left(U_0 - U_1 \frac{\lambda_0 - \lambda_1}{\lambda_0 + \delta - \lambda_1}\right).$$

Since $\hat{\tau}$ is strictly decreasing δ, it follows that $U_{\mathrm{SW}}(\hat{\tau} \mid \hat{\tau})$ is strictly decreasing in δ. Therefore, $\hat{\delta}$ can be defined by the equation $U_{\mathrm{SW}}(\hat{\tau} \mid \hat{\tau}) = U_1$, as in the main text (where the formula is stated in terms of L_S rather than U_S).

Remaining deviations The same argument from "Remaining Deviations" in the proof of Theorem 2 shows that the agent has a best response in which shirking is frontloaded over inspection-free intervals. If the relaxed solution is $T = \tau^*$, then by construction no shirk-before-work deviation is profitable, so the proof is complete.

Now suppose that the relaxed solution is periodic/exponential. First we show that the agent finds it optimal to shirk over the exponential phase, no matter his action history. Over the exponential phase, the distribution of inspection times is memoryless, so the only relevant state variable is the agent's belief, denoted q, that the state is 0. The agent's belief evolves according to the differential equation $\dot{q}_t = -q_t \delta(1 - a_t)$. Therefore the HJB equation reads

$$0 = \max_{a=0,1} \left\{(1-a)u_0 + au_1 - q\delta(1-a)V'(q) - \lambda_a V(q) + \gamma^*(qU_1 - V(q))\right\}.$$

We verify that this HJB equation is solved by the value function

$$V(q) = U_1 + (q-1)(U_0 - U_1)\lambda_0/\delta. \tag{24}$$

Plug in this value function, write $\lambda_a = \lambda_1 a + \lambda_0(1-a)$, and substitute in the expression for γ^* from Theorem 4. Simplify to get

$$0 = \max_{a=0,1} a(q-1)(U_0 - U_1)\lambda_0(\lambda_0 + \delta - \lambda_1)/\delta.$$

Since $\lambda_0 + \delta \geq \lambda_0 > \lambda_1$, this equation holds. At every belief q, the agent weakly prefers shirking to working (strictly so if $q < 1$).

We have shown that given any periodic/exponential policy $(\hat{\tau}, \pi, \gamma^*)$, with $\pi < 1$, the agent has a best response that takes the following form: shirk until some time t, with $t \leq \hat{\tau}$, work until time $\hat{\tau}$, and then shirk thereafter. Let $U_{\text{SWS}}(t \mid \hat{\tau}, \pi, \gamma^*)$ denote the agent's expected payoff from such a deviation. In terms of the value function V from (24), we have

$$U_{\text{SWS}}(t \mid \hat{\tau}, \pi, \gamma^*) = U_0(1 - e^{-\lambda_0 t}) + e^{-\lambda_0 t} U_1(1 - e^{-\lambda_1(\hat{\tau} - t)}) \\ + e^{-\lambda_0 t - \lambda_1(\hat{\tau} - t)} \left[\pi U_1 e^{-\delta t} + (1 - \pi) V(e^{-\delta t}) \right].$$

Define π^* to be the smallest probability π such that

$$\max_{t \in [0, \hat{\tau}]} U_{\text{SWS}}(t \mid \hat{\tau}, \pi, \gamma^*) = U_1. \tag{25}$$

Equation (25) holds for $\pi = 1$ (since $\tau^* \geq \hat{\tau}$) and fails for $\pi = 0$ (by the definition of γ^*). By continuity, π^* is well-defined and strictly positive.

To complete the proof, it suffices to check that $U_{\text{SWS}}(\hat{\tau} \mid \hat{\tau}, \pi^*, \gamma^*) = U_1$, which implies that π^* coincides with $\hat{\pi}$. Differentiate U_{SWS} with respect to t. We have

$$U'_{\text{SWS}}(\hat{\tau} \mid \hat{\tau}, \pi^*, \gamma^*) = (U_0 - U_1)\lambda_0 e^{-\lambda_0 \hat{\tau}} > 0,$$

and

$$U'_{\text{SWS}}(0 \mid \hat{\tau}, \pi^*, \gamma^*) = (U_0 - U_1)\lambda_0 - e^{-\lambda_1 \hat{\tau}} \left[\pi^* U_1 \delta + (1 - \pi^*)(U_0 - U_1)\lambda_0 \right].$$

By Assumption 2.B, we can apply the same argument following (15), with $\pi^* U_1 \delta + (1 - \pi^*)(U_0 - U_1)\lambda_0$ in place of $U_1 \delta$, to conclude that $U'_{\text{SWS}}(0 \mid \hat{\tau}, \pi^*, \gamma^*) < 0$. Follow the argument after (21) to conclude that $U_{\text{SWS}}(\hat{\tau} \mid \hat{\tau}, \pi^*, \gamma^*) = U_1$.

References

ANTINOLFI, G. AND F. CARLI (2015): "Costly Monitoring, Dynamic Incentives, and Default," *Journal of Economic Theory*, 159, 105–119. [4]

BERGEMANN, D. AND U. HEGE (1998): "Venture Capital Financing, Moral Hazard, and Learning," *Journal of Banking & Finance*, 22, 703–735. [5]

——— (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, 719–752. [5]

BOARD, S. AND M. MEYER-TER-VEHN (2013): "Reputation for Quality," *Econometrica*, 81, 2381–2462. [4]

BONATTI, A. AND J. HÖRNER (2017): "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169, 234–269. [5]

CHANG, C. (1990): "The Dynamic Structure of Optimal Debt Contracts," *Journal of Economic theory*, 52, 68–86. [4]

CHEN, M., P. SUN, AND Y. XIAO (2020): "Optimal Monitoring Schedule in Dynamic Contracts," *Operations Research*, 68, 1285–1314. [4]

DAI, L., Y. WANG, AND M. YANG (2022): "Dynamic Contracting with Flexible Monitoring," *Available at SSRN 3496785.* [4]

DEJARNETTE, P., D. DILLENBERGER, D. GOTTLIEB, AND P. ORTOLEVA (2020): "Time Lotteries and Stochastic Impatience," *Econometrica*, 88, 619–656. [2]

DILMÉ, F. AND D. F. GARRETT (2019): "Residual Deterrence," *Journal of the European Economic Association*, 17, 1654–1686. [4]

GOMPERS, P. A. AND J. LERNER (2004): *The Venture Capital Cycle*, MIT press. [1]

Green, B. and C. R. Taylor (2016): "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects," *American Economic Review*, 106, 3660–3699. [5]

Halac, M., N. Kartik, and Q. Liu (2016): "Optimal Contracts for Experimentation," *Review of Economic Studies*, 83, 1040–1091. [5]

Halac, M. and A. Prat (2016): "Managerial Attention and Worker Performance," *American Economic Review*, 106, 3104–3132. [4]

Hörner, J., N. Klein, and S. Rady (2021): "Overcoming Free-Riding in Bandit Games," *The Review of Economic Studies*, 89, 1948–1992. [5]

Hörner, J. and L. Samuelson (2013): "Incentives for Experimenting Agents," *RAND Journal of Economics*, 44, 632–663. [5]

Keller, G. and S. Rady (2015): "Breakdowns," *Theoretical Economics*, 10, 175–202. [5]

Kim, S.-H. (2015): "Time to Come Clean? Disclosure and Inspection Policies for Green Production," *Operations Research*, 63, 1–20. [4]

Lazear, E. P. (2006): "Speeding, Terrorism, and Teaching to the Test," *Quarterly Journal of Economics*, 121, 1029–1061. [4]

Li, A. and M. Yang (2020): "Optimal Incentive Contract with Endogenous Monitoring Technology," *Theoretical Economics*, 15, 1135–1173. [4]

Malenko, A. (2019): "Optimal Dynamic Capital Budgeting," *Review of Economic Studies*, 86, 1747–1778. [4]

Monnet, C. and E. Quintin (2005): "Optimal Contracts in a Dynamic Costly State Verification Model," *Economic Theory*, 26, 867–885. [4]

Ortoleva, P., E. Safonov, and L. Yariv (2022): "Who Cares More? Allocation with Diverse Preference Intensities," Working paper. [5]

Piskorski, T. and M. M. Westerfield (2016): “Optimal Dynamic Contracts with Moral Hazard and Costly Monitoring,” *Journal of Economic Theory*, 166, 242–281. [4]

Popov, L. (2016): “Stochastic Costly State Verification and Dynamic Contracts,” *Journal of Economic Dynamics and Control*, 64, 1–22. [4]

Ravikumar, B. and Y. Zhang (2012): “Optimal Auditing and Insurance in a Dynamic Model of Tax Compliance,” *Theoretical Economics*, 7, 241–282. [4]

Rodivilov, A. (2022): “Monitoring Innovation,” *Games and Economic Behavior*, 135, 297–326. [4]

The Economist (2018): “What is an Audit For?” *The Economist*, May 26, 2018. [1]

Varas, F., I. Marinovic, and A. Skrzypacz (2020): “Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring,” *Review of Economic Studies*, 87, 2893–2937. [4, 23, 25]

Wagner, P. and J. Knoepfle (2021): “Relational Enforcement,” *Available at SSRN 3832863*. [4]

Wagner, P. A. and N. Klein (2022): “Strategic Investment and Learning with Private Information,” *Journal of Economic Theory*, 204, 105523. [5]

Wang, C. (2005): “Dynamic Costly State Verification,” *Economic Theory*, 25, 887–916. [4]

Wong, Y. F. (2022): “Dynamic Monitoring Design,” Working paper. [4]

B Online appendix

B.1 Proof of Theorem 5

Local deviations Suppose that until the next inspection, the agent plans to shirk over $[t, t+s)$ and work otherwise. Let $U_t(s \mid \tau)$ denote the agent's expected payoff if the principal inspects at time τ. For $t + s < \tau$, the agent's probability of failing the inspection is $(1 - e^{-\delta s})e^{-\rho(\tau - t - s)}$, so

$$\begin{aligned} U_t(s \mid \tau) = {} & U_1(1 - e^{-\lambda_1 t}) + e^{-\lambda_1 t} U_0 (1 - e^{-\lambda_0 s}) + e^{-\lambda_1 t - \lambda_0 s} U_1 (1 - e^{-\lambda_1(\tau - t - s)}) \\ & + e^{-\lambda_1(\tau - s) - \lambda_0 s} U_1 \left[1 - (1 - e^{-\delta s}) e^{-\rho(\tau - t - s)}\right]. \end{aligned}$$

For $t < \tau$, differentiate with respect to s and simplify to get

$$U_t'(0 \mid \tau) = e^{-\lambda_1 t} \left[(U_0 - U_1)\lambda_0 - e^{-(\lambda_1 + \rho)(\tau - t)} U_1 \delta\right]. \tag{26}$$

For $t \geq \tau$, we have $U_t(s \mid \tau) = U_1$ for all $s \geq 0$, so $U_t'(0 \mid \tau) = 0$.

Relaxed problem—local deviations Consider the relaxed problem of choosing a positive random variable T to solve

$$\begin{aligned} &\text{minimize} && \mathbb{E}\, e^{-\lambda_1 T} \\ &\text{subject to} && \mathbb{E}\, U_t'(0 \mid T) \leq 0, \quad t \geq 0. \end{aligned}$$

To see that this constraint is necessary, recall that for all times t, we must have $\mathbb{E}\, U_t(s \mid T) \leq U_1 = \mathbb{E}\, U_t(0 \mid T)$. Now differentiate under the integral sign.

After substituting in the expression for U_t above, we see that this problem is equivalent to choosing a cumulative distribution function F supported on $[0, \infty)$ to solve

$$\begin{aligned} &\text{minimize} && \int_{[0,\infty)} e^{-\lambda_1 t} \, \mathrm{d}F(t) \\ &\text{subject to} && \int_{(t,\infty)} \left[\frac{(U_0 - U_1)\lambda_0}{U_1 \delta} - e^{-(\lambda_1 + \rho)(s - t)}\right] \mathrm{d}F(s) \leq 0, \quad t \geq 0. \end{aligned} \tag{27}$$

Denote by $h(t,s)$ the term in brackets in (27). The constraints are indexed by times $t \geq 0$. Attach a nonnegative multiplier η_0 to the time-0 constraint and a nonnegative density multiplier $\eta(t)$ to the time-t constraint, for $t \geq 0$. The Lagrangian becomes

$$L(F;\eta_0,\eta) = \int_{[0,\infty)} e^{-\lambda_1 t}\,\mathrm{d}F(t) + \eta_0 \int_{[0,\infty)} h(0,s)\,\mathrm{d}F(s) + \int_0^\infty \left[\int_{(t,\infty)} h(t,s)\,\mathrm{d}F(s)\right]\eta(t)\,\mathrm{d}t.$$

Change the order of integration in the double integral and relabel the dummy variables to get

$$L(F;\eta_0,\eta) = F(0) + \int_{(0,\infty)} I(t)\,\mathrm{d}F(t),$$

where

$$I(t) = e^{-\lambda_1 t} + \eta_0 h(0,t) + \int_0^t \eta(s)h(s,t)\,\mathrm{d}s.$$

Let $\eta(t) = \bar{\eta}e^{-\lambda_1 t}$ for some $\bar{\eta} \geq 0$. Substitute in this expression, integrate, and group like terms to get

$$I(t) = e^{-\lambda_1 t}\left[1 - \frac{\bar{\eta}}{\lambda_1}\frac{(U_0-U_1)\lambda_0}{U_1\delta} - \frac{\bar{\eta}}{\rho}\right] + e^{-(\lambda_1+\rho)t}\left[-\eta_0 + \frac{\bar{\eta}}{\rho}\right] + \left(\eta_0 + \frac{\bar{\eta}}{\lambda_1}\right)\frac{(U_0-U_1)\lambda_0}{U_1\delta}.$$

To make the bracketed terms vanish, define the multipliers $\bar{\eta}$ and η_0 by

$$\bar{\eta} = \frac{U_1\delta\lambda_1\rho}{(U_0-U_1)\lambda_0\rho + U_1\delta\lambda_1}, \qquad \eta_0 = \frac{U_1\delta\lambda_1}{(U_0-U_1)\lambda_0\rho + U_1\delta\lambda_1}.$$

The integrand $I(t)$ is constant in t, so the relaxed problem is solved by any cumulative distribution function F with $F(0) = 0$ that satisfies

$$\frac{U_1\delta}{\lambda_0(U_0-U_1)} \int_{(t,\infty)} e^{-(\lambda_1+\rho)(s-t)} = 1 - F(t),$$

for every $t \geq 0$. By Assumption 2, $U_1\delta > \lambda_0(U_0-U_1)$, so we can apply Lemma 2 from Appendix B.3 to conclude that the unique solution of the relaxed problem is the exponential distribution with hazard rate

$$\gamma^* = \frac{(U_0 - U_1)\lambda_0(\lambda_1 + \rho)}{U_1(\lambda_0 + \delta) - U_0\lambda_0}.$$

Remaining deviations It remains to check that if the principal uses the exponential policy with hazard rate γ^*, then it is optimal for the agent to work until the inspection. Since the distribution of inspections is memoryless, the only revelant state variable is the agent's belief, denoted q, that the state is 0. The agent's belief evolves according to the differential equation

$$\dot{q}_t = (1 - q_t)\rho a_t - q_t\delta(1 - a_t).$$

The HJB equation reads

$$0 = \max_{a=0,1} \Big\{(1 - a)u_0 + au_1 + [(1 - q)\rho a - q\delta(1 - a)]\, V'(q) \\ - \lambda_a V(q) + \gamma^*(qU_1 - V(q))\Big\}. \tag{28}$$

We verify that this HJB equation is solved by the function

$$V(q) = U_1 + (q - 1)(U_0 - U_1)\lambda_0/\delta.$$

Plug in this value function, write $\lambda_a = \lambda_1 a + \lambda_0(1 - a)$, and substitute in the expression for γ^* from Theorem 5. Simplify to get

$$0 = \max_{a=0,1} \ (a - 1)(q - 1)(U_0 - U_1)\lambda_0(\lambda_0 + \delta - \lambda_1 - \rho)/\delta.$$

If $\lambda_1 + \rho \geq \lambda_0 + \delta$, then this equation is satisfied. In this case, at every belief q, the agent weakly prefers working to shirking (strictly so if $q < 1$).

B.2 Robustness to small recovery rate

Here we formalize the claimed robustness to perturbing ρ.

Theorem 6 (Robustness to recovery)
Fix all the problem parameters other than δ and ρ.

1. *Suppose $\lambda_G \geq \lambda_B$ and $u_1 < u_0$. Then there exist strictly positive thresholds $\bar{\delta}$ and $\bar{\rho}$ such that for all $\delta \geq \bar{\delta}$ the optimal policy in Theorem 3 remains optimal with any recovery rate ρ in $[0, \bar{\rho}]$.*

2. *Suppose $\lambda_G < \lambda_B$. Then there exist strictly positive thresholds $\bar{\delta}$ and $\bar{\rho}$ such that for all $\delta \geq \bar{\delta}$, the periodic policy in Theorem 4 remains optimal with any recovery rate ρ in $[0, \bar{\rho}]$.*

Proof. We follow the proof of Theorems 3 and 4 (Appendix A.6), indicating the appropriate modifications to accommodate a positive recovery rate ρ.

Binding shirk-before-work deviation With a positive recovery rate ρ, the agent's shirk-before-work payoff now takes a different form. For $t \leq \tau$, the agent's probability of failing the inspection is $(1 - e^{-\delta t})e^{-\rho(\tau - t)}$, so

$$
\begin{aligned}
&U_{\text{SW}}(t \mid \tau) \\
&= U_0(1 - e^{-\lambda_0 t}) + e^{-\lambda_0 t} U_1 (1 - e^{-\lambda_1(\tau - t)}) + e^{-\lambda_0 t - \lambda_1(\tau - t)} U_1 \left[1 - (1 - e^{-\delta t}) e^{-\rho(\tau - t)}\right] \\
&= U_0 - (U_0 - U_1) e^{-\lambda_0 t} - U_1 e^{-\lambda_0 t - \lambda_1(\tau - t)} (1 - e^{-\delta t}) e^{-\rho(\tau - t)}.
\end{aligned}
$$

Define τ^* and $\bar{t}$ with this new expression for U_{SW}. As before (see footnote 34), it can be shown that τ^* is well-defined.

We first identify a condition under which we get $U'_{\text{SW}}(0 \mid \tau^*) < 0$. If

$$
0 = U'_{\text{SW}}(0 \mid \tau^*) = -e^{-(\lambda_1 + \rho)\tau^*} U_1 \delta + \lambda_0 (U_0 - U_1),
$$

then

$$
\begin{aligned}
U''_{\text{SW}}(0 \mid \tau^*) &= e^{-(\lambda_1 + \rho)\tau} U_1 \delta (\delta + 2\lambda_0 - 2\lambda_1 - 2\rho) U_1 - (U_0 - U_1)\lambda_0^2 \\
&= (U_0 - U_1)\lambda_0 (\delta + \lambda_0 - 2\lambda_1 - 2\rho) U_1 \qquad (29) \\
&> 0,
\end{aligned}
$$

provided that $\delta + \lambda_0 - 2\lambda_1 - 2\rho > 0$.

Assuming $\delta + \lambda_0 - 2\lambda_1 - 2\rho > 0$, we identify a further condition under which we get $\bar{t} = \tau^*$. Since $U'_{\mathrm{SW}}(0 \mid \tau^*) < 0$, we can argue as before, using Descartes' rule of signs, provided that

$$0 < U'_{\mathrm{SW}}(\tau^* \mid \tau^*) = e^{-\lambda_0 \tau^*}\big(U_0\lambda_0 - (\lambda_1 + \rho)U_1 - e^{-\delta\tau^*}(\lambda_0 + \delta - \lambda_1 - \rho)U_1\big).$$

Combining these two sufficient conditions gives the system

$$\delta + \lambda_0 - 2\lambda_1 - 2\rho > 0, \tag{30}$$

$$\rho U_1 + e^{-\delta\tau^*}(\lambda_0 + \delta - \lambda_1 - \rho)U_1 < U_0\lambda_0 - \lambda_1 U_1. \tag{31}$$

The agent's payoff from shirking forever is unaffected by ρ. Therefore, if this system is satisfied, then the period τ^* and the relaxed problems and solutions are the same as in the model without recovery.

Remaining deviations The argument from "Remaining Deviations" in the proof of Theorem 2 can be modified to show that it is optimal for the agent to frontload shirking over no-inspection intervals. In fact, with recovery, front-loading shirking has the additional benefit of increasing the agent's passage probability.[38]

Suppose that the solution of the relaxed problem is periodic/exponential with $\pi^* < 1$. It follows that $\lambda_0 > \lambda_1$. First we check that the agent finds it optimal to shirk over the exponential inspection interval $[\hat{\tau}, \infty)$, no matter his action history. As in (28), the relevant state variable is the belief q that the the state is 0. Now the HJB equation reads

$$0 = \max_{a=0,1} \Big\{(1-a)u_0 + au_1 + [(1-q)\rho a - q\delta(1-a)]\, V'(q) \\ - \lambda_a V(q) + \gamma^*(qU_1 - V(q))\Big\}.$$

[38] If the state is 0 with probability q, then after shirking for duration Δ and then working for duration Δ, the state is 0 with probability $q_{\mathrm{SW}} = 1 - (1 - qe^{-\delta\Delta})e^{-\rho\Delta}$. If instead the agent works for duration Δ and then shirking for duration Δ, then the state is 0 with probability $q_{\mathrm{WS}} = (1 - (1-q)e^{-\rho\Delta})e^{-\delta\Delta}$. With δ and ρ both strictly positive, it can be checked that $q_{\mathrm{SW}} > q_{\mathrm{WS}}$, no matter the value of q.

We verify that this HJB equation is solved by the same value function

$$V(q) = U_1 + (q-1)(U_0 - U_1)\lambda_0/\delta.$$

Plug in this value function, write $\lambda_a = \lambda_1 a + \lambda_0(1-a)$, and substitute in the expression for γ^* from Theorem 4. Simplify to get

$$0 = \max_{a=0,1} \; a(q-1)(U_0 - U_1)\lambda_0(\lambda_0 + \delta - \lambda_1 - \rho)/\delta.$$

This equation is satisfied if $\lambda_0 + \delta - \lambda_1 - \rho > 0$, which is implied by (30). In this case, at every belief q, the agent weakly prefers shirking to working (strictly so if $q < 1$).

Finally, we turn to the shirk-work-shirk deviations. With a positive recovery rate ρ, the agent's shirk-work-shirk payoff now takes a different form. We have

$$\begin{aligned} U_{\mathrm{SWS}}(t \mid \hat{\tau}, \pi, \gamma^*) = U_0(1 - e^{-\lambda_0 t}) + e^{-\lambda_0 t} U_1 (1 - e^{-\lambda_1(\hat{\tau} - t)}) \\ + e^{-\lambda_0 t - \lambda_1(\hat{\tau}-t)} \left[\pi q(t) U_1 + (1-\pi) V(q(t))\right], \end{aligned}$$

where $q(t) = 1 - (1 - e^{-\delta t})e^{-\rho(\hat{\tau} - t)}$. Define π^* by (25), with the new expression for U_{SWS}. As before, it can be shown that π^* is well-defined.

To complete the proof, it suffices to check that $U_{\mathrm{SWS}}(\hat{\tau} \mid \hat{\tau}, \pi^*, \gamma^*) = U_1$, for then π^* agrees with $\hat{\pi}$ (which is the same in the no-recovery case). We have

$$U'_{\mathrm{SWS}}(\hat{\tau} \mid \hat{\tau}, \pi^*, \gamma^*) = e^{-\lambda_0 \hat{\tau}} \left[(U_0 - U_1)\lambda_0 - \frac{\rho(\pi^* U_1 \delta + (\pi^* - 1)(U_0 - U_1)\lambda_0)}{\lambda_0 + \delta - \lambda_1} \right].$$

This expression is strictly positive if

$$(U_0 - U_1)\lambda_0 - \frac{\rho U_1 \delta}{\lambda_0 + \delta - \lambda_1} > 0.$$

Since $\lambda_0 > \lambda_1$, this holds if

$$\rho < \frac{\lambda_0(U_0 - U_1)}{U_1}. \tag{32}$$

If (30) holds, then we can apply the same argument following (29), with $\pi^* U_1 \delta + (1 - \pi^*)(U_0 - U_1)\lambda_0$ in place of $U_1\delta$, to conclude that $U'_{\mathrm{SWS}}(0 \mid \hat{\tau}, \pi^*, \gamma^*) < 0$. The rest of the proof is completed as before using Descartes' rule of signs.

Solving for the thresholds For part 1, assuming $U_1\lambda_1 < U_0\lambda_0$, we can select thresholds $\bar{\delta}$ and $\bar{\rho}$ such that (30)–(31) hold for all $\delta \geq \bar{\delta}$ and $\rho \leq \bar{\rho}$.[39]

For part 2, we have $U_1\lambda_1 < U_0\lambda_0$ because $\lambda_1 < \lambda_0$. In this case, choose $\bar{\delta}$ and $\bar{\rho}$ as before to satisfy (30)–(31), and then use $\min\{\bar{\rho}, \lambda_0(U_0 - U_1)/U_1\}$ as the new threshold for ρ to ensure that (32) holds as well. □

B.3 Uniqueness

The proofs of uniqueness rely on the following lemmas. To represent a distribution over $[0, \infty)$, define a cumulative distribution function supported on $[0, \infty)$ to be a function $F\colon [0, \infty) \to [0, 1]$ that is weakly increasing, right-continuous, and satisfies $\sup_{x \geq 0} F(x) = 1$.

Lemma 2 (Unique fixed point—single exponential)
Fix $A > 1$ and $\alpha > 0$. For each π in $[0, 1)$, there exists exactly one cumulative distribution function F supported on $[0, \infty)$ with $F(0) = \pi$ satisfying

$$\int_{(t,\infty)} A e^{-\alpha(s-t)} \, \mathrm{d}F(s) = 1 - F(t), \tag{33}$$

for every $t \geq 0$. Namely, $F(t) = \pi + (1 - \pi)(1 - e^{-\gamma t})$, with $\gamma = \alpha/(A - 1)$.

Proof. Let F be a cumulative distribution function on $[0, \infty)$ that satisfies this system. Put $t = 0$ in (33) to get

$$\int_{(0,\infty)} A e^{-\alpha s} \, \mathrm{d}F(s) = 1 - F(0) = 1 - \pi.$$

[39] Note that the second term on the left side of (31) is at most $e^{-\delta\tau^*}(\lambda_0 + \delta - \lambda_1)U_1$. First choose $\bar{\delta}$ and $\bar{\rho}$ to ensure that τ^* is bounded from below by some positive constant c. Then take $\tau^* = c$ in (31), and strengthen the thresholds as needed to ensure that this new form of (31) and also (30) both hold.

For all $t \geq 0$, we have

$$\int_{(t,\infty)} Ae^{-\alpha(s-t)}\,\mathrm{d}F(s) = e^{\alpha t}\left[1 - \pi - \int_{(0,t]} Ae^{-\alpha s}\,\mathrm{d}F(s)\right].$$

Use the layer-cake representation and then change variables to get

$$\begin{aligned}\int_{(0,t]} Ae^{-\alpha s}\,\mathrm{d}F(s) &= Ae^{-\alpha t}[F(t) - F(0)] + A\int_{e^{-\alpha t}}^{1} [F(-\alpha^{-1}\log x) - F(0)]\,\mathrm{d}x \\ &= Ae^{-\alpha t}F(t) + \int_0^t A\alpha e^{-\alpha s}F(s)\,\mathrm{d}s - A\pi.\end{aligned}$$

Substitute these equalities into (33) to get

$$(1-\pi)e^{\alpha t} - AF(t) - \int_0^t A\alpha e^{-\alpha(s-t)}F(s)\,\mathrm{d}s + A\pi e^{\alpha t} = 1 - F(t).$$

Solve for $F(t)$ to get

$$F(t) = \pi e^{\alpha t} + \frac{1}{A-1}\left[e^{\alpha t} - 1 - \int_0^t A\alpha e^{-\alpha(s-t)}F(s)\,\mathrm{d}s\right].$$

That is, F is a fixed point of an operator defined by the the expression on the right side. Define this operator on the space of bounded functions on an interval $[0,t_1]$ with the supremum norm. For $t_1 < A\alpha/(A-1)$, this operator is a contraction, and hence has a unique fixed point, denoted F_1. For some t_2 larger than t_1, define the operator on the space of bounded functions on $[0,t_2]$, by replacing $F(t)$ with $F_1(t)$ on the right side for $t \leq t_1$. If $t_2 < t_1 + \alpha A/(A-1)$, then this operator is a contraction and hence has a unique fixed point F_2 on $[0,t_2]$ that extends F_1. Continuing in this way, each operator is a contraction provided that $t_{i+1} - t_i < A\alpha/(A-1)$. Construct a sequence (t_i) satisfying these inequalities with $t_i \uparrow \infty$. We get a sequence of fixed points F_i over $[0,t_i]$. For each fixed t, we must have $F(t) = F_i(t)$ for all i such that $t_i \leq t$. Therefore, F is unique.

It remains to check that this F is actually a cumulative distribution function. Guess that $F(t) = \pi + (1-\pi)(1-e^{-\gamma t})$ for $t \geq 0$. We have $F(0) = \pi$, and (33) is satisfied for all t if $A\alpha/(\alpha+\gamma) = 1$, hence $\gamma = \alpha/(A-1)$. This cumulative distribution function is therefore the unique solution. □

Lemma 3 (Unique fixed point—sum of exponentials)
Fix positive numbers A, B, α, β with $A - B = 1$ and $\beta B > \alpha A$. For each π in $[0, 1)$, there exists exactly one cumulative distribution function F supported on $[0, \infty)$ with $F(0) = \pi$ satisfying

$$\int_{(t,\infty)} \left[A(1-e^{-\alpha(s-t)}) - B(1-e^{-\beta(s-t)})\right] \mathrm{d}F(s) = 0, \tag{34}$$

for every $t \geq 0$. Namely, $F(t) = \pi + (1-\pi)(1-e^{-\gamma t})$, with $\gamma = \alpha\beta/(\beta B - \alpha A)$.

Proof. Let F be a cumulative distribution function on $[0, \infty)$ that satisfies this system. The integrand is continuous in (s, t) and vanishes when $s = t$. Therefore, we can calculate the total derivative of the left side with respect to t by differentiating under the integral (by dominated convergence) and ignoring the change in the left endpoint. Thus,

$$\int_{(t,\infty)} \left[-\alpha A e^{-\alpha(s-t)} + \beta B e^{-\beta(s-t)}\right] \mathrm{d}F(s) = 0, \tag{35}$$

for all $t \geq 0$. Multiply (34) by β and subtract (35). Simplify using the equality $A - B = 1$ to conclude that

$$\int_{(t,\infty)} \frac{(\beta-\alpha)A}{\beta} e^{-\alpha(s-t)} \,\mathrm{d}F(s) = 1 - F(t),$$

for all $t \geq 0$. Since $A = B + 1$ and $\beta B > \alpha A$, it follows that $\beta A > \alpha A + \beta$. Therefore, $(\beta - \alpha)A/\beta > 1$, so we can apply Lemma 2 to complete the proof, noting that

$$\gamma = \frac{\alpha}{(\beta-\alpha)A/\beta - 1} = \frac{\alpha\beta}{\beta A - \alpha A - \beta} = \frac{\alpha\beta}{\beta B - \alpha A}.$$ □