Frequent Monitoring in Dynamic Contracts

Dmitry Orlov*

University of Wisconsin-Madison

First Draft: January 2, 2015
Current Draft: August 15, 2022

Abstract

I consider a dynamic principal-agent model in which the agent does not observe the quality of his output and the principal chooses how much to monitor the agent. Monitoring improves production outcomes but can reduce the agent's incentive to work if it uncovers bad performance and leads to punishing the agent. The optimal monitoring intensity is path dependent: an agent who performed poorly in the past is monitored less going forward to reduce the risk of early termination. Conversely, an agent who performed well is monitored more as his accumulated promised compensation provides a buffer against bad monitoring outcomes and reduces the associated penalties by deferring the pay-for-performance risk to states of the world in which termination costs are low.

Keywords: repeated moral hazard, dynamic contracts, monitoring, communication, performance evaluations, interventions.

1 Introduction

Workers do not directly observe the results of their labor in many organizational settings and must rely on the supervisor for their performance information. In such settings, should a supervisor continuously monitor the agent and take corrective action when he underperforms? On the one hand, monitoring has productive value as the principal can intervene promptly. On the other, such interventions inform the agent of his poor performance, requiring future rewards to be sufficiently high to keep the worker incentivized. The worker's limited liability leads to an endogenous, information-based, monitoring

*A previous version of this paper was titled "Optimal Design of Internal Disclosure". I am grateful to Andrzej Skrzypacz, Jeffrey Zwiebel, Peter DeMarzo, Peter Koudijs, Bob Wilson, Jonathan Levin, Pavel Zryumov, Nina Varsava, Paul Pfleiderer, Gustavo Manso, and Daniel Ferreira for insightful suggestions, as well as seminar and conference participants at Stanford GSB, UNC Kenan-Flagler, Simon School of Business, Norwegian School of Economics, London School of Economics, Imperial College London, INSEAD, University of Geneva, Fuqua School of Business, SED, and SITE. Email: dmitry.orlov@wisc.edu.

cost which, when optimally managed, leads to limited monitoring of poorly-performing workers even if transfers can be made contingent on monitoring and performance outcomes.

Consider a professional services firm specializing in law or consulting. Typically, the long-term contracts for associates in these firms have up-or-out structures that provide strong incentives. Monitoring the associate is valuable: if he is about to make a poor decision, the principal can intervene and not suffer its consequences. Such intervention may, however, inform the associate that he performed poorly in that period, leading him to question his prospects at the firm and, thus, exert less effort going forward.[1] This paper shows that such monitoring costs are present even if the principal can commit to an optimal compensation plan and are more severe the worse the agent has performed in the past.

I consider a dynamic moral hazard model in which the principal (she) hires the agent (he) to find good investment projects every period. Both players are risk-neutral, and the agent is protected by limited liability. The agent's ability is common knowledge, and he exerts private effort in every period to reduce the arrival intensity of bad projects. The agent observes the arrival of some, but not all, bad projects and can choose to report them to the principal so that she can divest from them. The principal can also monitor incoming projects to uncover and publicly divest from bad projects left unreported by the agent. Monitoring outcomes are observed by the agent in real-time, informing him about that period's performance.[2] The agent does not directly observe the performance of undertaken projects, and the principal optimally chooses to communicate this information with a delay. A dynamic contract is a combination of compensation, monitoring, and performance feedback policies.[3]

Monitoring outcomes change the agent's expectation over future compensation. If the principal detects a bad project, the incentive contract penalizes the agent and brings him closer to inefficient termination. To mitigate such inefficiency, the principal must award the agent higher rents ex-ante, which is costly for her. Alternatively, by reducing the monitoring intensity and communicating the performance of unmonitored projects with a delay, the principal can postpone the pay-for-performance risk in the agent's expected compensation to states of the world when the risk of inefficient termination is lower, all while maintaining incentive provision. Optimal monitoring intensity solves the following trade-off: monitoring reduces the number of bad projects undertaken by the principal but informs the agent of his performance and increases his information rents by enabling more state-contingent deviations for him to pursue. The optimal contract relies on a combination of immediate and deferred incentives. Immediate incentives reflect how sensitive the agent's continuation value is to monitoring outcomes. Deferred incentives reflect how sensitive the agent's continuation value is to realized performance, which the principal discloses with a delay. The model delivers three main results.

[1] Murphy (1999) notes such conflicts if the worker is awarded a bonus if his performance exceeds a pre-specified level.

[2] Section B.1 of the Online Appendix shows that the results are qualitatively unchanged if the agent observes a fraction of monitoring outcomes. Unlike Piskorski and Westerfield (2016), monitoring is assumed to have no exogenous costs.

[3] Commitment to dynamic compensation is sufficient to implement optimal monitoring and communication.

First, the optimal monitoring intensity is limited in order to reduce the risk of the agent's early termination. The agent is kept less informed about current output after a bad project is uncovered either by him or the principal, as it reduces his expected stake in the firm and requires greater use of delayed incentives to avoid termination. Second, the pay-for-performance sensitivity to monitored projects is non-monotone in the agent's continuation value: it is increasing for low values, along with the monitoring intensity, but is decreasing for high values, when the monitoring reaches its capacity and deferred incentives account for a greater role in incentive provision. The optimal contract features elements of leniency toward a well-performing agent – he is monitored frequently but not penalized if monitoring detects a bad project. Third, if the agent is relatively impatient, the principal probabilistically communicates past performance if the agent's continuation value is sufficiently high. Bonuses accompany good performance evaluations.

Theoretical insights. Joint characterization of the optimal compensation and monitoring policies under repeated moral hazard is a difficult problem. On the one hand, an agent who is perfectly informed about his performance makes effort decisions sequentially. The single-deviation incentive compatibility constraint is then necessary and sufficient for global incentive compatibility. On the other hand, in the absence of performance information, the agent's multi-period effort choice becomes a multi-tasking problem since he chooses all of his actions in advance of seeing any information about his total output. The single-deviation approach is no longer sufficient in this setting, and the agent may find it strictly optimal to shirk in the future if he shirks today.[4] An information structure implied by a given monitoring policy falls in between these two extreme cases of full transparency and opacity and creates a challenge in characterizing the agent's best response to a given contract as different monitoring policies induce different optimally-binding global incentive constraints. To gain traction with this double-deviation problem for an arbitrary monitoring policy, I allow the agent to partially observe the quality of the projects that arise from his effort. In this case, if the agent shirks and generates a bad project, he can report it to the principal, suffering a penalty for generating a bad project but not carrying over private information, and thus return to the recommended path of effort. This drastically simplifies the double-deviation problem while maintaining the key feature that more information available to the agent leads to a more stringent incentive compatibility condition. The model is set in continuous time, making it easier to show the properties of the optimal monitoring policy. This setting also enables the proof of the equivalence between the main model and a version in which the parties can renegotiate the contract, but the agent only observes his effort and relies entirely on the principal to learn about his performance.

Applications. The interaction of monitoring and incentives has received considerable attention in economics. On the one hand, greater monitoring should make it easier to align incentives, as suggested

[4] Bond and Gomes (2009) and Laux (2001) show that assigning the agent to operate multiple projects simultaneously, rather than sequentially, reduces the expected compensation costs.

by Holmstrom (1979). On the other hand, both theoretical and empirical evidence, surveyed by Frey (1999), suggests that too much monitoring may crowd out the agent's effort, as it lowers the agent's self-esteem and intrinsic motivation.[5] While prior mechanisms focused on intrinsic motivation, this paper shows the detrimental incentive implications of monitoring if it runs the risk of informing the worker of his poor performance, thus lowering his expected compensation to a point where incentives can no longer be provided.[6] The immediate implication and empirical prediction of this argument is that monitoring has fewer downsides when it is aimed at workers with sufficient value vested in the firm. Such implications of monitoring are relevant in organizational settings where (a) the worker relies on the principal for information about his performance, as studied in Baker (1992), and (b) the principal's monitoring outcomes may be observed by the agent. Examples of such settings are common. In information technology firms, the development team creates the product; however, it is the sales and support teams that collect customer feedback and learn about the product's quality. If there is a fault with it, the development team needs to be involved in fixing it. In law firms, junior associates write legal briefs, while senior associates and partners evaluate the resulting legal arguments before deploying them in court. Junior associates may observe when their arguments are overruled by their supervisors. In large firms where multiple divisions compete for capital, each division observes the quality of its projects but not the output of other divisions, leaving it less informed about its relative performance.[7] The profit-maximizing allocation of capital by headquarters across divisions can convey some of this relative performance information. In all of these examples, excessive monitoring may convey performance results to the agent and may demoralize him if it uncovers bad performance. The principal must then be strategic in communicating such results to the agent, even if she can commit to an optimal compensation plan.

1.1 Related Literature

Monitoring improves the expected value of the agent's output but also increases his information rent by providing him with performance feedback. This trade-off is of long-standing interest in the economics literature. Holmstrom (1979) establishes the "informativeness principle", highlighting that additional information is beneficial to the principal ex-post if it is more indicative of the agent's private effort. Harris and Raviv (1978) show that the agent receives greater expected compensation if he is more informed ex-ante about his performance when making a private effort decision.[8] These studies point to the principal benefiting from more information being available ex-post but being hurt when this

[5] Barkema (1995) provides empirical evidence of how excessive monitoring reduces the incentives of top managers.

[6] This draws a distinction of monitoring a worker's input and output, raised in Chapter 8 of Lazear (1998).

[7] Scharfstein and Stein (2000) provide empirical evidence of over-allocation of capital to poorly performing divisions, which can be interpreted as "under" monitoring in the context of this model. Hornstein and Zhao (2011) find a positive empirical relationship between the firms' transparency and the efficiency of its internal resource allocation.

[8] Edmans and Gabaix (2011) that this translates to an ex-post incentive compatibility constraint for the agent, which drastically simplifies the shape of the optimal contract. The current paper uses an economically similar argument to gain traction with the multi-deviation problem of the agent under an arbitrary monitoring policy.

information is available to the agent ex-ante. Lizzeri, Meyer, and Persico (2002) present this argument in a two-period model showing that the principal optimally postpones all performance feedback until the end of the agency relationship unless such feedback increases the production value of the agent's effort.[9] The current paper shows that, in a dynamic setting, the optimal monitoring intensity and performance feedback are positively correlated with the agent's past performance.

Baker (1992) points out that the principal may be better informed about the relevant performance signals than the agent. Even though the agent perfectly observes his effort, he may not fully observe its implications on the firm. The literature on subjective performance evaluations, including Baker, Gibbons, and Murphy (1994) and MacLeod (2003), studies optimal contracting when the principal privately observes the agent's performance measure and can misreport it. Fuchs (2007) shows that, in a dynamic setting, it is optimal to postpone performance disclosure until the final period of the agency relationship in order to pool the agent's incentives across periods. The principal's truth-telling constraint may require money burning under the optimal contract. In the current paper, the principal also privately observes performance, but, in contrast to Fuchs (2007), she can verifiably communicate it to the agent both via monitoring and performance feedback. Delayed communication lets the principal reuse the punishments across periods to incentivize the agent and leads to a convex compensation structure, in which most of the agent's compensation occurs when he does very well, similar to Hoffmann, Inderst, and Opp (2019) and consistent with the bonus structures documented in Murphy (1999). This is different from the concave compensation profile obtained in Fuchs (2007), which punishes the agent only in the event he does very poorly, due to the fact that, in the current paper, the principal cannot misreport the agent's performance and decides only when and how to communicate it.

The principal must monitor the agent in order to intervene in a timely manner. Monitoring lets the principal act before it is too late but does not provide incremental information about the agent's actions. This is in contrast with Piskorski and Westerfield (2016) and Chen, Sun, and Xiao (2017), who model monitoring in a dynamic agency setting as the ability of the principal to observe the agent's chosen effort at a cost. Related work by Georgiadis and Szentes (2019) and Hoffmann, Inderst, and Opp (2019) study more general costly information acquisition policies in a static principal-agent environment. In the current paper, the principal decides whether to publicly observe the agent's performance immediately or with a delay, rather than what information about the agent's effort to acquire, which leads to an economic tension between monitoring intensity and effort provision. Frey (1999) provides a survey of the "crowding out" effect that monitoring has on incentives, especially when it leads to "controlling" interventions by the principal, similar to how monitoring can lead the principal to divest from bad projects in the current paper. Barkema (1995) empirically documents

[9] Manso (2011) provides an example of how such an increase in production value arises in an experimentation setting. Kaya (2020) considers the benefit of informing the agent about his ability to reduce compensation costs.

the negative consequences of monitoring in a sample of top managers in Dutch firms. In particular, over-monitoring is most costly for top managers overseen by the firm's CEO, rather than the board of directors, which is consistent with the current findings if being overseen by the board of directors is related to a higher rank and more value vested in the firm. This paper also shows that in such instances where the agent has significant value vested in the firm, the incentive costs of monitoring can be mitigated if the principal can communicate performance results with a delay – something she has the flexibility to do in a broad range of organizational settings as discussed in Ederer (2010).

The agency friction within the firm is modeled as a repeated moral hazard problem in continuous time, similar to Holmström and Milgrom (1987) and DeMarzo and Sannikov (2006), in which the agent privately controls the arrival intensity of a negative Poisson process, similar to Biais, Mariotti, Rochet, and Villeneuve (2010). This paper shows that whether the agent observes his output instantly or with a delay is important in determining the optimal contract and contributes to understanding optimal transparency in a dynamic contracting setting. By delaying performance information, the principal can defer incentive provision to states of the world in which it is cheaper to provide. This economic force also arises if the agent's actions impact output in a persistent way, as studied in Sannikov (2014), Williams (2011), and Marinovic and Varas (2018). In these environments, the principal induces contemporaneous effort from the agent by relying on a series of future performance signals and corresponding pay-for-performance sensitivities. I also show that renegotiation-proof contracts can be tractably characterized despite the double-deviation problem, similar to Strulovici (2011).

Smolin (2021) and Ray (2007) show that the principal may wish to withhold performance information from the agent to retain him if there is uncertainty about his ability. Ball (2019) and Fudenberg and Rayo (2019) find that the principal can ameliorate agency conflicts by backloading communication of valuable information to the agent. Horner and Lambert (2019) and Varas, Marinovic, and Skrzypacz (2020) characterize optimal dynamic certification policies if the agent's incentives stem from career concerns. Ederer (2010), Aoyagi (2010), and Goltsman and Mukherjee (2011) study optimal information provision in multi-stage tournaments. Lizzeri and Siniscalchi (2008) characterize optimal interventions if the principal wishes to maximize the agent's learning on the job. This paper contributes to this literature by considering the joint design of monitoring, feedback, and contingent transfers. Murphy and Cleveland (1995) point to the applied relevance of such an approach.

The rest of the paper is organized as follows. Section 2 presents the main model. Section 3 characterizes the optimal contract and its properties. Section 4 derives the optimal renegotiation-proof contract and Section 5 concludes. Formal proofs and additional results are in the Published and Online Appendices.

2 Main Model

Overview. The principal (she) employs the agent (he) to manage a flow of projects for her to undertake. The agent exerts costly private effort to reduce the rate of arrival of bad projects. He observes some of the bad projects and can report them to the principal so that she can divest from them. The agent cannot identify all bad projects, however, and the principal can monitor the quality of the remaining projects herself to screen out the ones unobserved or unreported by the agent. The agent relies on monitoring outcomes and the principal's communication to learn about the quality of projects he does not independently identify as bad. The principal designs a long-term contract comprised of compensation, monitoring, and communication policies to maximize her expected profits.

Players. Time is continuous and indexed by $t \in [0, +\infty)$. The principal is risk-neutral, discounts future consumption at rate r, and has an outside option of 0 when not employing the agent. The agent is risk-neutral, has limited liability in every instant, discounts future consumption at the same rate r, and has an initial reservation value $\underline{w} \geq 0$ at $t = 0$, which disappears if he accepts a contract offered by the principal. Reservation value $\underline{w}$ determines the agent's ex-ante bargaining power with the principal, while limited liability ensures weakly positive transfers from the principal to the agent in any contract. The agent retires at an exogenous random time η, distributed exponentially with parameter γ, at which he can claim the expected continuation value of the contract.

Private effort, reporting, and monitoring. In every instant t, the agent generates a new project which can be either good or bad. The agent produces a positive cash flow at rate α, but causes losses of magnitude -1 at discrete times according to a Poisson process with arrival intensity $\mu+\lambda+\Delta\cdot(1-a_t)$, which decreases in the agent's private effort $a_t \in \{0, 1\}$. By exerting effort a_t at a private, non-pecuniary, flow cost $a_t \cdot h$, the agent reduces the arrival intensity of bad projects by $a_t \cdot \Delta$ in instant t. Denote by X_t the total number of bad projects up to and including time t.

The process $X = (X_t)_{t\geq 0}$ can be decomposed into the sum of two independent counting processes $N = (N_t)_{t\geq 0}$ and $Y = (Y_t)_{t\geq 0}$. N is a Poisson process with intensity μ and Y is a Poisson process with intensity $\lambda + \Delta(1 - a_t)$, which depends on the agent's effort $a_t \in \{0, 1\}$. The agent privately observes the process Y. This assumption captures the idea that the agent can sometimes notice mistakes that the principal cannot detect easily. It plays a key role in the tractability of the model, as discussed in Section 4.1. When the agent observes $dY_t = 1$ at time t, he can choose to report it to the principal, captured by a binary disclosure decision $d_t \in \{0, 1\}$. The total number of bad projects reported by the agent up to and including time t is $R_t \overset{def}{=} \int_0^t d_s \, dY_s$.[10]

[10] It is without loss to assume that the agent never misreports good projects as bad ones. To see this, we can solve the model as if the agent does not engage in such misreporting and then verify that it would be strictly suboptimal to misrepresent good projects under the optimal contract.

The principal observes neither the agent's effort a_t nor the arrival of bad projects X_t. She learns of bad projects either through the agent's reports, monitoring, or observing their realized payoffs as they occur. In each instant t, the principal chooses to monitor the project with intensity $f_t \in [0, \bar{f}]$, which captures the probability of identifying a bad project that was unreported by the agent at time t. Upper bound $\bar{f} \in [0, 1]$ is constant and captures the maximum monitoring intensity the principal can implement in any given instant. To focus on the endogenous agency implications of monitoring, I assume that monitoring imposes no deadweight costs on the principal. Denote by M_t to be the total number of bad projects uncovered by monitoring up to and including time t. It is comprised of bad projects that are either unobserved by the agent, i.e., process N, or are strategically not reported by him, i.e., process $Y - R$. If the agent reports bad projects arising from process Y truthfully, i.e., $d_t = 1$, then the monitoring intensity f_t is aimed at uncovering only the bad projects stemming from process N, in which case $M = (M_t)_{t \geq 0}$, is a counting process with arrival intensity μf_t. The process $U_t \stackrel{def}{=} X_t - R_t - M_t$ denotes the number of bad projects that end up being undertaken by time t due to not being reported by the agent or screened out via monitoring by the principal. I assume the agent does not observe process U. In contrast, the principal observes U_t as it occurs.[11] Under truthful reporting by the agent, $U = (U_t)_{t \geq 0} = N - M$ is a counting process with arrival intensity $\mu \cdot (1 - f_t)$. The principal can be selective in undertaking projects. She learns that the current project is bad, i.e., $dX_t = 1$, prior to suffering its adverse impact either via either the agent's report dR_t or monitoring outcome dM_t. In this case, she can divest from it at a fixed adjustment cost $l \in [0, 1)$, reducing the negative impact of a bad project by $1 - l > 0$.

Observability and communication. The agent observes the principal's monitoring outcomes M, which convey partial information about N, but not the quality of the remaining undertaken projects, captured by process U. The agent also observes the cumulative compensation C_t he received from the contract up to and including time t, which may also be informative about U. The principal has accurate information about process U as it occurs but has discretion over whether and when to communicate it to the agent. In continuous time, it is convenient to model all feasible communication by allowing the principal to commit to a (standard) filtration[12] $\mathbb{F}^c = (\mathcal{F}_t^c)_{t \geq 0}$ governing all of the information available to the agent, including that contained in his observed compensation. A filtration $\mathbb{F}^c$ is feasible if it

(i) contains all public information available at time t given by the agent's reports, monitoring outcomes, compensation, recommended effort, and the agent's retirement time

$$\underline{\mathcal{F}}_t^c \stackrel{def}{=} \sigma\Big\{(R_s, M_s, C_s, a_s)_{s \leq t},\ \eta \leq t\Big\} \stackrel{(i)}{\subseteq} \mathcal{F}_t^c;$$

[11]The results are unchanged as long as the principal observes U_t anytime before the agent's retirement time η. It would also be equivalent to assume that the principal observes the quality of each bad project that was unreported by the agent with probability $\bar{f}$ before its cash flow is realized but can commit to investing in it nevertheless. In this case, the divested bad projects correspond to process M, while the undertaken bad projects correspond to process U.

[12]$\sigma(\xi)$ denotes the σ-algebra generated by random element ξ. Modeling strategic verifiable communication as a filtration $\mathbb{F}^c$ in a continuous-time game is introduced in Orlov, Skrzypacz, and Zryumov (2020).

(*ii*) is contained by all information available at time t, including process $(U)_{s\leq t}$, and an independent, sufficiently rich, σ-algebra $\mathcal{F}^R$ capturing a randomization device available to the principal

$$\mathcal{F}_t^c \overset{(ii)}{\subseteq} \sigma\Big\{(R_s, M_s, C_s, a_s)_{s\leq t},\ \eta \leq t,\ (U_s)_{s\leq t},\ \mathcal{F}^R\Big\} \overset{def}{=} \overline{\mathcal{F}}_t^c.$$

Filtration $\mathbb{F}^c$ specifies the flow information available to the agent. If he deviates from the recommended action path, his information set also diverges from that of the principal. Denote by $\hat{\mathbb{F}}^c = (\hat{\mathcal{F}}_t^c)_{t\geq 0}$ the agent's off-path filtration

$$\hat{\mathcal{F}}_t^c \overset{def}{=} \sigma\left\{\mathcal{F}_t^c, \left(Y_s, \hat{a}_s, \hat{d}_s\right)_{s\leq t}\right\}.$$

For a given filtration $\mathbb{F}^c$, the agent's belief about the performance of undertaken projects U is given by $\mathrm{Law}\big((U_s)_{s\leq t} \,\big|\, \hat{\mathcal{F}}_t^c\big)$. The agent exerts effort $\hat{a}_t$ before knowing his performance dY_t or monitoring outcome dM_t in instant t. Similarly, he decides whether to report $dY_t = 1$, captured by the reporting choice $\hat{d}_t$, before observing the monitoring outcome dM_t. Such timing considerations require that an agent's strategy $(\hat{a}, \hat{d})$ is predictable with respect to $\hat{\mathbb{F}}^c$, i.e., $(\hat{a}_t, \hat{d}_t)$ is measurable with respect to $\hat{\mathcal{F}}_{t-}^c$.[13] Denote by $\mathcal{A}(\mathbb{F}^c)$ to be the set of strategies available to the agent at $t = 0$ given $\mathbb{F}^c$.

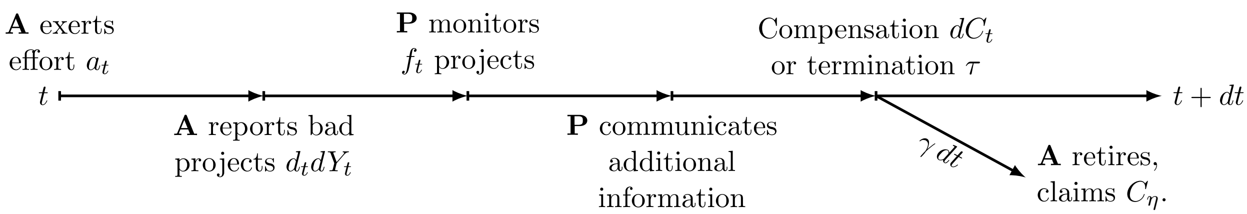

Figure 1: Timing of events within a stage game $[t, t + dt]$.

Contracts. The contract specifies a termination time τ when the production technology is liquidated and transfers are made. Formally, τ is a stopping time with respect to the contract's filtration $\mathbb{F}^c$, which occurs either due to the principal terminating the agent, or the agent retiring and claiming the continuation value of the contract. The principal also commits to a monitoring intensity process $f = (f_t)_{t\geq 0}$, recommends an effort and reporting processes $(a, d) = (a_t, d_t)_{t\geq 0}$ that are all predictable with respect to filtration $\mathbb{F}^c$, and a cumulative compensation process $C = (C_t)_{t\geq 0}$ that is adapted to $\mathbb{F}^c$ and non-decreasing due to the agent's limited liability. The principal chooses contract $\mathcal{C} = (f, a, d, C, \mathbb{F}^c)$ to maximize her expected payoff

$$\sup_{\mathcal{C}}\ \mathrm{E}_{(a,f)}\left[\int_0^\tau e^{-rt} \cdot \big(\alpha\, dt - l\, dR_t - l\, dM_t - dU_t - dC_t\big)\right] \tag{1}$$

subject to the effort and reporting strategy process $(a, d) = (a_t, d_t)_{t\geq 0}$ being incentive-compatible for the agent

[13] $\mathcal{F}_{t-}$ denotes the left limit of filtration $\mathbb{F}$ at time t given by $\mathcal{F}_{t-} \overset{def}{=} \cup_{s<t} F_s$.

$$(a,d) \in \underset{(\hat{a},\hat{d})\in\mathcal{A}(\mathbb{F}^c)}{\arg\max}\ \mathrm{E}_{(\hat{a},\hat{d})}\left[\int_0^\tau e^{-rt}\cdot(dC_t - \hat{a}_t h\, dt)\right]. \tag{2}$$

A finer, i.e., more informative, filtration $\mathbb{F}^c$ weakly increases the set of deviations $\mathcal{A}(\mathbb{F}^c)$ available to the agent, making it more difficult to satisfy the global incentive constraint (2). Since filtration $\mathbb{F}^c$ contains information in past monitoring outcomes, greater monitoring leads to more agency costs.

Model parameters. To focus on the economic trade-off between monitoring and retention, I restrict attention to parameters under which the incremental value of the agent's effort is large enough so that the principal prefers to elicit high effort $a_t = 1$ from the agent, even if it requires less monitoring, rather than closely monitor an agent who is shirking. This is captured by a restriction on model parameters:

$$\underbrace{\frac{\alpha-(\lambda+\Delta)\cdot l-\mu\cdot\bar{f}\cdot l-\mu\cdot(1-\bar{f})\cdot 1}{r+\gamma}}_{\substack{\textit{payoff under full monitoring}\\ \textit{if the agent shirks: } a_t=0}} \overset{(i)}{\leq} 0 \leq 1-l \overset{(ii)}{\leq} \underbrace{\frac{\alpha-\lambda\cdot l-\mu\cdot 1-h}{r+\gamma}-\frac{h}{\Delta}}_{\substack{\textit{payoff under no monitoring}\\ \textit{if the agent works: } a_t=1}}. \tag{3}$$

The left-hand side of (3) is the expected discounted payoff to the principal from employing the agent until his retirement time η under full monitoring, i.e., $f_t \equiv \bar{f}$, but without ever paying him, and thus, allowing him to shirk. Inequality (i) then captures the condition that the payoff from such a contract is weakly lower than the principal's outside option, normalized to 0, meaning that she would rather terminate the agent than allow him to shirk on the job.[14] Inequality (i) in (3) holds as long as arrival intensity Δ is sufficiently large. The right-hand side of (3) is equal to the expected payoff to the principal if she awards the agent with an information rent equal to h/Δ in exchange for him exerting high effort and truthfully reporting bad projects until his retirement, all while not being monitored. Inequality (ii) then states that the value of the firm under full effort and no monitoring exceeds the gain of fixing a single project. Inequality (ii) in (3) holds if the value of good projects α or the divestment cost l are sufficiently large relative to the players' discount rate r. Taken jointly, parametric condition (3) restricts attention to the region in which effort provision is valuable and the principal is better off by investing into some bad projects rather than terminating the agent.[15]

2.1 Discussion of the Model

Capturing the information implications of monitoring requires a setting in which the agent does not observe his performance perfectly otherwise. In such an environment, a single deviation may lead the agent to deviate again as he becomes relatively pessimistic about both the undertaken projects, whose quality has not been revealed, as well as his future compensation. The agent's global deviation

[14]If (i) does not hold, then the principal prefers to employ the agent while shirking until retirement rather than to terminate him. This is equivalent to adjusting the principal's reservation value to equal the l.h.s. of (3).

[15]Much of the technical analysis in the paper extends beyond the scope of condition (3).

depends in a complex way on the monitoring and compensation policies set by the contract. Allowing the agent to observe his effort-specific output process Y provides crucial tractability. The intuition is that, by the revelation principle, the optimal contract elicits truth-telling by the agent. If the agent were to shirk and generate a bad project, he would optimally report this bad project to the principal and suffer the corresponding penalty, but, crucially, not carry over this private information into the next instant. In Section 4.1, I show that the optimal renegotiation-proof contract does not depend on this simplifying assumption, i.e., on the observability of process Y,[16] serving as a robustness check for the main findings. Even though the agent reports bad projects truthfully along the path of play, and monitoring is aimed at uncovering bad projects stemming from process N, the incentive requirement for truthful reporting imposes the essential link between monitoring and incentive provision.

There are several ways to motivate the observability of monitoring outcomes by the agent. First, observing monitoring outcomes is equivalent to observing the principal's divestment decisions as long as the principal does not divest from projects not identified as bad, i.e., she does not pay the fixed divestment cost l when it is unnecessary.[17] In this case, divestment from a project conveys a perfectly informative signal of bad performance to the agent. An alternative motivation for the observability of monitoring outcomes by the agent is that it stems from information leakage within the firm. Section B.1 of the Online Appendix shows that if monitoring outcomes are revealed probabilistically, then the results are unchanged if the likelihood of such information leakage is sufficiently high.

The exogenous retirement date η lets us focus on optimal monitoring dynamics when the principal and the agent are equally patient, in which case both transfers and communication are optimally backloaded until the final date τ, similar to a finite-horizon model. Considering an exogenous retirement date η captures the basic properties of a finite-horizon model while preserving stationarity with respect to time. I show that the main results are qualitatively unchanged in a finite-horizon version of the model if the horizon is long enough or if the agent is impatient relative to the principal.[18]

3 Optimal Contract

First, I show that there exists an optimal contract in which the principal does not reveal information about process U until the agent retires. Given such communication, there exists an optimal contract admitting a simple Markovian dependence on the public history. In this contract, the agent's global incentive compatibility constraints are characterized in closed form, summarizing the tension between monitoring intensity and incentive provision. The optimal monitoring policy is the solution to the

[16] This result can also be extended to restructuring-proof contracts introduced in Cetemen, Feng, and Urgun (2019).

[17] The analysis is unchanged if the principal could invest in bad projects she uncovers via monitoring as it is equivalent to lowering the monitoring intensity in that instant.

[18] See Section 3.5 of the main text and Section B.3 of the Online Appendix for the discussion and analysis of the model in which the agent is impatient relative to the principal. See Section B.2 of the Online Appendix for the finite horizon version of the model. Both extensions support the properties of the optimal contract derived in Section 3.

single-state continuous-time dynamic program in which the principal trades-off monitoring intensity with the expected liquidation costs. If the players are sufficiently impatient, the optimal monitoring policy is given in closed form, illustrating the tension between monitoring and incentives.

3.1 Optimal Incentive Compatibility

The first simplifying observation is that the revelation principle applies, i.e., there exists an optimal contract in which the agent reports bad projects truthfully and recommends for investment only those he perceives as high quality.

Lemma 1. *There exists an optimal contract in which the agent truthfully reports bad projects arising from process Y to the principal, i.e., $d_t \equiv 1$ and $R \equiv Y$.*

Even if the agent truthfully reports bad projects stemming from process Y, the compensation, effort, monitoring, and communication policies may still depend in a complicated way on the latent performance of invested projects $(U_t)_{t\geq 0}$. Since both players are risk-neutral, equally patient, and the principal has commitment power, it is without loss to restrict attention to contracts that defer all compensation to the final date τ. This means that the compensation process in itself need not be an interim source of performance information for the agent. Lemma 2 further shows that there exists an optimal contract in which the principal does not reveal any information about $(U_t)_{t\geq 0}$ prior to the final date τ meaning that the monitoring intensity $(f_t)_{t\geq 0}$ and recommended effort $(a_t)_{t\geq 0}$ need not depend on latent performance in the optimal contract

Lemma 2. *There exists an optimal contract satisfying truthful reporting such that*

- *prior to the final date τ the principal does not communicate information about $(U_s)_{s<t\wedge\tau}$ to the agent, i.e., $\mathcal{F}^c_{t\wedge\tau-}$ is independent from $\sigma\big\{(U_s)_{s<t\wedge\tau}\big\}$;*
- *at the final date τ the principal communicates all performance information, i.e., $\mathcal{F}^c_\tau = \bar{\mathcal{F}}_\tau$.*[19]

Information expands the set of deviations available to the agent. Therefore, it is weakly optimal for the principal to postpone any communication until the agent is no longer needed, i.e., when he is let go or retires. Lemma 2 implies we can restrict attention to contracts that do not share performance information with the agent above and beyond the monitoring outcomes.[20] At the same time, the "informativeness principle" of Holmstrom (1979) states that information is valuable since the principal can pay the agent more in states which are more indicative of his effort. The principal thus finds it (weakly) optimal to disclose all performance information at time τ in order to make transfers contingent on the realization of the process $U = (U_t)_{t\leq\tau}$. Lemma 2 also implies that the recommended

[19] The optimal information set of the agent is discontinuous at time τ in the sense that $\mathcal{F}^c_{\tau-} \subset \mathcal{F}^c_\tau$.

[20] This finding is similar to Fuchs (2007) and Lizzeri et al. (2002), who show that by keeping the agent uninformed until the final period, the principal can reuse incentives across periods.

effort a_t does not depend on latent performance $(U_s)_{s<t}$ since any such dependence can be replaced with public randomization in conjunction with a modification of the agent's compensation plan.

Lemma 3. *There exists an optimal contract in which the agent is rewarded at final date τ only if no bad projects received investment, i.e, $C_\tau = 0$ if $U_\tau > 0$.*

The agent's compensation C_τ at final date τ may depend on the realization of $U = (U_t)_{t\leq\tau}$. Since the agent's effort and truthful reporting reduce the probability of investing into a bad project, the outcome $U_\tau = 0$ is most indicative that agent has neither shirked nor hidden bad projects. Since the players are risk-neutral, the principal only rewards the agent if $U_\tau = 0$, following Holmstrom (1979). The idea behind the proof of Lemma 3 is to perturb the agent's compensation by committing to the same on-path expected payoff, but delivering it only if $U_\tau = 0$ via a larger transfer in that state. This new contract sharpens incentive compatibility while keeping the expected on-path payoffs to both players the same. Lemma 3 implies that the principal imposes substantial pay-for-performance sensitivity on the agent at retirement by either rewarding him if $U_\tau = 0$ or not paying him at all. The agent is willing to exert effort because he does not know the realization of U_τ before the final date τ. Along the path of truthful reporting, the probability that $U_\tau = 0$ given the recommended effort $(a_t)_{t\geq 0}$ and monitoring policy $(f_t)_{t\geq 0}$ is given by

$$p_t \stackrel{def}{=} \mathrm{P}_a\left(\int_0^t a_s\, dU_s \;=\; 0 \;\middle|\; \mathcal{F}_t^c\right) = e^{-\int_0^t \mu a_s(1-f_s)\, ds}. \tag{4}$$

Denote by w_t the agent's expected discounted compensation net of effort costs, i.e., his continuation value, at time t along the recommended path of effort $(a_t)_{t\geq 0}$ and truthful reporting as

$$w_t \stackrel{def}{=} \mathrm{E}_a\left[e^{-r(\tau-t)} \cdot C_\tau - \int_t^\tau e^{-r(s-t)} a_s h\, ds \;\middle|\; \mathcal{F}_t^c\right]. \tag{5}$$

Process $w = (w_t)_{t\geq 0}$ is conveniently used in dynamic contracts to summarize the history of the contract, as pioneered by Spear and Srivastava (1987). When the agent retires at time $\tau = \eta$ he requests his promised utility w_t from the principal. The principal, at this point, reveals past performance $(U_t)_{t\leq\tau}$ and pays the agent only if $U_\tau = 0$ according to Lemma 3. The agent's compensation, thus, differs from the agent's continuation value at time τ, i.e., $C_\tau \neq w_\tau$, in response to this newly revealed performance information. Lemma 3 implies that the jump at τ can be characterized as[21]

$$C_\tau = \frac{w_\tau}{p_{\tau-}} \cdot \mathbb{1}\{U_\tau = 0\} \tag{6}$$

pinning down compensation C_τ as a function of the agent's continuation value w_τ in the previous instance, the "on-path" probability $p_{\tau-} = \mathrm{P}_{\tau-}(U_\tau = 0)$, and the actual realization U_τ. Compensation

[21] Because the principal shares information about U_τ at time τ, it follows that $p_\tau = \mathbb{1}\{U_\tau = 0\} \in \{0,1\}$, making $p_{\tau-}$ defined in (4) reflecting the prior belief about $U_\tau = 0$. On the other hand, continuation value w_τ is evaluated at time τ to capture the contemporaneous impact of process Y in the event that it triggers the agent's premature termination.

C_τ defined via (6) satisfies promise keeping since $\mathrm{E}\left[C_\tau \mid \mathcal{F}^c_\tau\right] = p_{\tau-} \cdot \frac{w_\tau}{p_{\tau-}} = w_\tau$. In the analysis that follows, I restrict attention to optimal contracts satisfying Lemmas 1–3.

The dynamics of the continuation value process $(w_t)_{t\geq 0}$ are determined by its sensitivities $(\psi_t)_{t\geq 0}$ to bad projects reported by the agent and sensitivities $(\phi_t)_{t\geq 0}$ to bad projects uncovered via monitoring.

Lemma 4. *There exist processes $(\phi)_{t\geq 0}$ and $(\psi)_{t\geq 0}$, both predictable with respect to $\mathbb{F}^c$, such that for $t < \tau$ the agent's continuation value satisfies the stochastic differential equation*

$$dw_t = rw_t\,dt + a_t h\,dt + \underbrace{\phi_t \cdot \Big(\mu f_t\,dt - dM_t\Big)}_{\textit{monitoring results}} + \underbrace{\psi_t \cdot \Big(\big(\lambda + \Delta \cdot (1 - a_t)\big)\,dt - dR_t\Big)}_{\textit{reporting results}}. \tag{7}$$

The agent's limited liability requires that $\phi_t \leq w_t$ and $\psi_t \leq w_t$.

If the agent reports a bad project to the principal at time t, he suffers a penalty ψ_t. If the agent does not report a bad project, but the principal uncovers a bad project via monitoring, the agent suffers a potentially different penalty ϕ_t. Lemma 4 is a standard result in the dynamic contracting literature, e.g., DeMarzo and Sannikov (2006) and Biais, Mariotti, Rochet, and Villeneuve (2010), and reduces the problem of designing a contract to a sequential choice punishments ϕ_t and ψ_t. As we can see in (7), the monitoring intensity f_t increases the rate of discovering bad projects via monitoring and, consequently, of imposing punishment ϕ_t on the agent. I can now characterize the agent's incentives under an optimal contract satisfying Lemmas 1-3 in terms of penalties ϕ_t and ψ_t defined in Lemma 4.

Proposition 1 (Incentive Compatibility)**.** *There exists an optimal contract in which the agent*

(i) reports bad projects truthfully at time t if and only if

$$\psi_t \leq f_t \cdot \underbrace{\phi_t}_{\substack{\textit{immediate}\\ \textit{penalty}}} + (1 - f_t) \cdot \underbrace{w_t}_{\substack{\textit{delayed}\\ \textit{penalty}}} \qquad \textit{(truthful-reporting IC).} \tag{8}$$

(ii) exerts high effort at time t if and only if

$$\psi_t \geq h/\Delta \qquad \textit{(effort-provision IC).} \tag{9}$$

The agent minimizes the expected punishment if he observes a bad project in instant t. If he reports it to the principal, he suffers a penalty ψ_t. If he hides it, then the principal uncovers it with probability f_t via monitoring, leading to an immediate penalty ϕ_t, defined in (7), while, with probability $1 - f_t$, the principal does not detect this bad project and ends up investing in it. While the latter seems like a good outcome for the agent, it is not as, following Lemma 3, he forfeits all of his compensation at time τ. Put differently, the agent suffers a delayed penalty equal to his entire continuation value w_t if an unreported bad project passes the principal's monitoring. The expected cost to the agent of not reporting a bad project is thus $f_t \cdot \phi_t + (1 - f_t) \cdot w_t$ and, to elicit truthful reporting, it must exceed the penalty ψ_t captured by incentive constraint (8).[22] The agent's limited liability stipulates that $\phi_t \le w_t$ implying that delayed penalties (weakly) exceed immediate ones under the optimal contract. Hence, a higher monitoring intensity f_t (weakly) tightens the agent's truthful reporting constraint by lowering the right-hand side of (8) and captures the idea that more monitoring makes it more difficult to align incentives. The principal punishes the agent at date τ for bad projects stemming from process N in order to reduce his incentive to under-report bad projects stemming from process Y.

The agent's incentive compatibility condition (9) to exert effort is straightforward once he reports bad projects truthfully. By spending a private flow cost h, the agent reduces the intensity of getting a bad project, and suffering the penalty ψ_t, by Δ. The agent, thus, exerts effort whenever the expected cost h is less than the expected penalty $\Delta \cdot \psi_t$.[23] Incentive constraints (8) and (9) imply that penalties ψ_t and ϕ_t must be sufficiently high for effort $a_t = 1$ and reporting $d_t = 1$ to be incentive-compatible.

3.2 Optimal Monitoring, Effort, and Termination

We now proceed to characterize the principal's expected value under the optimal contract and the associated monitoring policy. Denote by $v(w_0)$ the principal's expected value under an optimal contract which delivers an expected value of exactly w_0 to the agent at the beginning of the contract. Lemma 1 stipulates that there exists an optimal contract in which the agent reports bad projects truthfully and finds the recommended effort $(a_t)_{t\geq 0}$ incentive compatible. Such a contract solves

$$v(w_0) = \sup_{\mathcal{C}(w_0)} \mathrm{E}_a \left[\int_t^\tau e^{-r(s-t)} \Big(\alpha\, dt - dU_t - l\, dM_t - l\, dY_t \Big) - e^{-r(\tau - t)} \cdot C_\tau \right], \tag{10}$$

where the supremum in (10) is taken across all contracts delivering the agent an expected payoff

$$w_0 = \mathrm{E}_a \left[e^{-r(\tau-t)} \cdot C_\tau - \int_t^\tau e^{-r(s-t)} a_s h\, ds \right] = \max_{(\hat{a},\hat{d})} \mathrm{E}_{(\hat{a},\hat{d})} \left[e^{-r(\tau-t)} \cdot C_\tau - \int_t^\tau e^{-r(s-t)} \hat{a}_s h\, ds \right].$$

Proposition 1 states that there exists an optimal contract in which the necessary and sufficient global truthful-reporting and effort provision conditions are expressed as local incentive constraints (8) and (9), which depend on the principal's choice of effort a_t, monitoring intensity f_t, and penalties ϕ_t and ψ_t. The principal faces an optimal control problem in which she trades off monitoring efficiency, incentive provision, and dead-weight losses of terminating the agent before retirement.

[22] The idea of penalizing the agent by taking away all of his continuation value is similar to Piskorski and Westerfield (2016), where the principal monitors the agent's effort and fires him if he detects shirking. In the current setting, the principal penalizes the agent ex-post based on the performance of the project via Lemma 3.

[23] Incentive constraint (9) is identical to the one obtained in Biais, Mariotti, Rochet, and Villeneuve (2010) and stems from the agent observing process Y: if he deviates by shirking, he optimally reports a bad project if it occurs, and thus returns to the equilibrium path of play, rather than pursue a multi-deviation strategy in subsequent effort decisions.

Proposition 2. *Suppose λ is sufficiently small, $\Delta \geq \gamma$, and parametric condition* (3) *holds.*[24] *The principal optimally recommends the agent to exert high effort $a_t = 1$ until he is terminated or retires. For $w \geq h/\Delta$, the principal's value function satisfies the Hamilton-Jacobi-Bellman equation pinning down the optimal monitoring intensity $f(w)$ and penalties $\phi(w), \psi(w)$ given by*

$$(r+\gamma)v(w) = \max_{(f,\phi,\psi)} \Big\{ \alpha - \lambda l - \mu(1 - f + fl) - \gamma w + v'(w) \cdot \big(rw + h + \mu f \phi + \lambda \psi\big) + \mu f \cdot \big(v\left(w - \phi\right) - v(w)\big) + \lambda \cdot \big(v\left(w - \psi\right) - v(w)\big) \Big\} \tag{11}$$

where $f \in [0, \bar{f}]$ and $\phi, \psi \in [0, w]$ satisfy incentive constraints (8) *and* (9)*. For $w < h/\Delta$, the principal terminates the agent with probability $1 - \frac{w}{h/\Delta}$ and pays him nothing, while with probability $\frac{w}{h/\Delta}$ she retains him with a continuation value h/Δ resulting in*

$$v(w) = \frac{w}{h/\Delta} \cdot v\left(h/\Delta\right). \tag{12}$$

Moreover, $v(w)$ is the maximal solution to (11) *and* (12) *satisfying $\lim_{w \to \infty} v'(w) = -1$. The initial continuation value w_0^* maximizes $v(w_0^*)$ subject to it exceeding the agent's initial reservation value $\underline{w}$.*

For $w \geq h/\Delta$ the principal optimally incentivizes the agent to exert effort, while for $w < h/\Delta$ the principal prefers to probabilistically terminate the agent rather than let him shirk, which follows from the high value of the agent's effort captured by parametric condition (3). Condition $\Delta \geq \gamma$ ensures that the value function is concave at $w = h/\Delta$ as the contract leaves the termination region $[0, h/\Delta]$. The agent's continuation value evolves in response to the performance sensitivities ϕ and ψ necessary to elicit his effort and truth-telling. Monitoring is chosen to balance the immediate benefit of detecting a bad project with the risk of inefficient project liquidation stemming from the punishments necessary to motivate effort. If the agent's retirement η occurs before he is fired, then the principal pays him $\frac{w_{\eta-}}{p_{\eta-}}$ if no bad project has received investment up to the retirement date and collects her outside option.

Value function $v(w)$ is depicted in Figure 2. For $w < h/\Delta$, it is linear and given by (12). For $w > h/\Delta$, $v(w)$ is strictly concave. Figure 2 illustrates that the total value of the firm, given by $v(w) + w$, is always dominated by the value of operating the production technology at first best levels of effort and monitoring, defined as

$$\bar{V} \stackrel{def}{=} \frac{\alpha - \lambda \cdot l - \mu(1 - \bar{f}) \cdot 1 - \mu \bar{f} \cdot l - h}{r + \gamma}.$$

If the agent performs well, then his continuation value increases, and he is able to withstand more punishments in the future, leading the total value of the firm to approach first-best.

[24] The principal's Hamilton-Jacobi-Bellman equation (11) features two distinct delay terms ϕ_t and ψ_t making it difficult to analytically verify concavity of $v(w)$, which is necessary to ensure that the optimal contract does not utilize public randomization for $w_t \geq h/\Delta$. I prove concavity for $\lambda = 0$, in which case there is a single delay term ϕ_t, and then show that concavity is preserved for small enough λ. Even if the solution to (11) is not globally concave and the optimal contract requires public randomization for $w \geq h/\Delta$, the trade-off between monitoring and incentives is preserved.

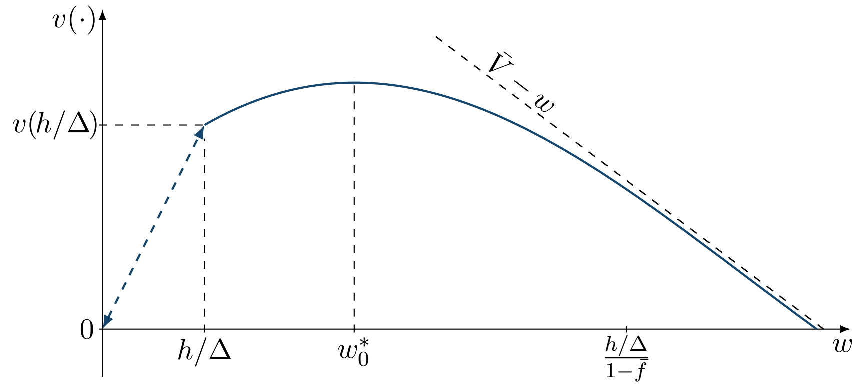

Figure 2: Principal's value function $v(\cdot)$ under the optimal contract.

Optimal termination. According to the effort provision constraint (9), the principal can incentivize the agent to exert effort only if ψ_t exceeds h/Δ, implying that the agent's continuation value w_t must also exceed h/Δ for effort provision to be feasible. If the agent's continuation value falls below h/Δ, then the principal cannot elicit effort from the agent. Instead of letting the agent shirk, the principal mixes between terminating the agent, leaving him with 0 value, and retaining him with an increased continuation value of exactly h/Δ. This way, the principal implements high effort with probability $\frac{w}{h/\Delta}$ which, under parametric condition (3), dominates the agent shirking on the job.[25]

Optimal penalties. The principal's value function $v(w)$ is strictly concave for $w > h/\Delta$ and it is optimal to set $\psi(w) \equiv h/\Delta$ for all continuation values w. A similar argument applies to $\phi(w)$ being costly to the principal since it lowers the agent's continuation value and makes it more likely that he will be inefficiently terminated before retirement. Substituting penalty $\psi(w) = h/\Delta$ into the agent's truth-telling constraint (8) we see that the optimal monitoring intensity $f(w)$ and penalty $\phi(w)$ must jointly satisfy

$$h/\Delta \leq f(w) \cdot \phi(w) + (1 - f(w)) \cdot w. \tag{13}$$

Similar to the previous argument, the concavity of $v(w)$ implies that it is optimal to minimize the absolute value of penalty $\phi(w)$ while still satisfying (13) resulting in

$$\phi(w) = \min\Big\{\phi \geq 0 : \ f(w) \cdot \phi + (1 - f(w)) \cdot w \geq h/\Delta\Big\} \quad \Rightarrow \quad \phi(w) = \max\left\{0,\, w - \frac{w - h/\Delta}{f(w)}\right\}. \tag{14}$$

Higher continuation value w is good because more incentives can be provided through delayed punishment lowering the necessary penalty $\phi(w)$. If $w \geq \frac{h/\Delta}{1-f(w)}$, the principal can set $\phi(w) = 0$ meaning that all reporting incentives are provided through the punishment deferred to the final date τ. In this case, monitoring does not introduce performance risk into the continuation value process of

[25] Zhu (2013) shows that, if the agent is impatient relative to the principal, the optimal contract may implement shirking for low continuation values w_t as the agent's higher discount rate speeds up the accumulation of the continuation value to the point where effort can be implemented.

the agent and hence does not impose an agency cost. If $w < \frac{h/\Delta}{1-f(w)}$, however, a positive penalty $\phi(w)$ is necessary since delayed punishment alone is not enough to incentivize truth-telling. The principal chooses the amount of monitoring $f(w)$ to trade-off the benefit of investment efficiency and the cost of inefficient termination associated with imposing a greater penalty $\phi(w)$ on the agent.

Optimal monitoring. The principal chooses the maximum monitoring intensity $f(w)$ accounting for the implication it has on $\phi(w)$ via (14) so that truthful reporting constraint (8) is satisfied. It follows that incentive constraint (13) optimally binds whenever $f(w) < \bar{f}$ since, otherwise, the principal can increase $f(w)$ without increasing the associated penalty $\phi(w)$. When $w < \frac{h/\Delta}{1-f}$ we can substitute the binding incentive constraint (13) into the principal's Hamilton-Jacobi-Bellman equation (11). The optimal monitoring rule $f(w)$, then, solves

$$f(w) = \underset{f \in \left[\frac{w-\delta}{w}, \bar{f}\right]}{\arg\max} \left\{ \overbrace{f(1-l)}^{\substack{\textit{benefit of} \\ \textit{monitoring}}} + \overbrace{f \cdot \underbrace{\left(w - \frac{w - h/\Delta}{f}\right)}_{=\phi} \cdot v'(w) + f \cdot \left[v \underbrace{\left(w - w + \frac{w - h/\Delta}{f} \right)}_{=w-\phi} - v(w) \right]}^{\textit{expected agency cost of monitoring}} \right\}. \quad (15)$$

The benefit of monitoring is greater investment efficiency, which is the first term in (15). The endogenous cost of monitoring is the agent's increased intensity $f(w) \cdot \mu$ and severity $\phi(w)$ of immediate punishment for bad projects, both of which increase the probability of the agent's termination before retirement and constitute the second term in (15). Taking the first-order condition of (15) with respect to f, the interior monitoring intensity f equates the benefit of divesting from a bad project with the marginal expected agency cost of monitoring

$$\underbrace{1 - l}_{\substack{\textit{marginal benefit} \\ \textit{of monitoring}}} = \underbrace{v(w) - w \cdot v'(w) - v\left(\frac{w - h/\Delta}{f}\right) + \frac{w - h/\Delta}{f} \cdot v'\left(\frac{w - h/\Delta}{f}\right)}_{\textit{marginal agency cost of monitoring}}. \quad (16)$$

As can be seen from (16), monitoring up to $\frac{w-h/\Delta}{w}$ does not exacerbate the agency problem since the principal can set $\phi(w) = 0$ and not affect the agent's continuation value even if a bad project is discovered. For $f > \frac{w-h/\Delta}{w}$, the marginal monitoring agency cost in (16) is strictly increasing in f, due to the concavity of the principal's value function, until it equates the benefit of divesting from a bad project.[26] If the agent's continuation value w is high, then the marginal agency cost of monitoring is low as the agent can sustain many consecutive punishments without being inefficiently terminated, leading to a high optimal monitoring intensity $f(w)$. Conversely, a low continuation value w exposes the agent to a high risk of inefficient termination due to the penalty for an uncovered bad project, reducing the principal's optimal monitoring. Proposition 3 establishes that optimal monitoring intensity increases in the agent's continuation value, leading to the dynamic implication that a well-performing agent is

[26]The derivative of the marginal monitoring agency cost, i.e., the right hand side of (16), with respect to monitoring intensity f is given by $-\frac{(w-h/\Delta)^2}{f^3} \cdot v''\left(\frac{w-h/\Delta}{f}\right) \geq 0$.

monitored more as there is less risk of demotivating him.

Proposition 3. *Suppose λ is sufficiently small, $\Delta \geq \gamma$, and parametric restriction* (3) *holds. Then, the optimal monitoring intensity $f(w)$ is weakly increasing in continuation value w and, moreover, is*

- *bounded from above by $\frac{w-h/\Delta}{h/\Delta}$;*
- *equal to $\bar{f}$ if $w > \bar{w}$, where $\bar{w}$ is a threshold less than $\frac{h/\Delta}{1-\bar{f}}$.*

Punishment $\phi(w)$ is increasing (decreasing) when $w < \bar{w}$ ($w > \bar{w}$). Moreover, if the players are sufficiently impatient, i.e., $r \geq \mu \cdot \bar{f} - \gamma$, then the optimal monitoring intensity is given by

$$f(w) = f^I(w) = \begin{cases} \dfrac{w - h/\Delta}{h/\Delta} & \text{if} \quad w \leq \left(1+\bar{f}\right) \cdot h/\Delta, \\ \bar{f} & \text{if} \quad w > \left(1+\bar{f}\right) \cdot h/\Delta. \end{cases} \tag{17}$$

and the corresponding optimal penalty $\phi^I(w)$ is given by (14)*.*

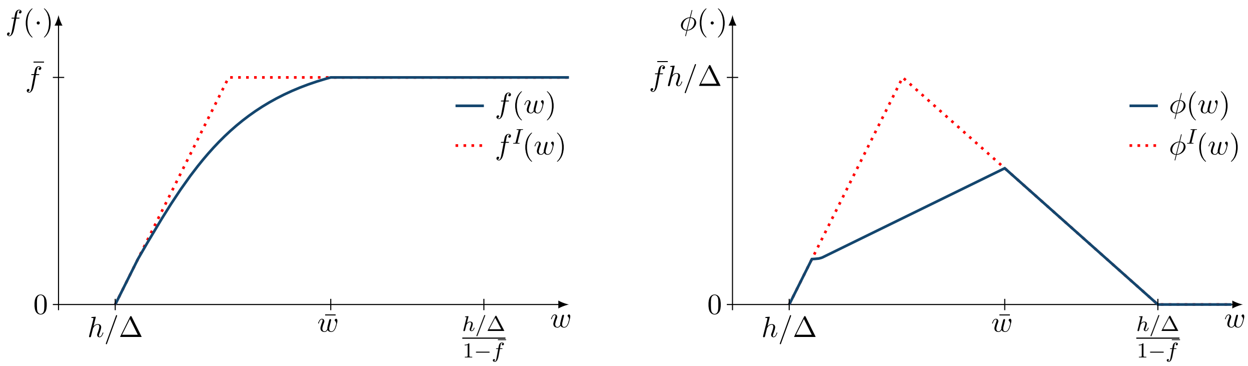

(a) Optimal monitoring intensity $f(w)$.

(b) Optimal contemporaneous penalty $\phi(w)$.

Figure 3: Optimal monitoring and pay-for-performance sensitivity.

Figure 3 depicts the key properties of the optimal monitoring intensity $f(w)$ and penalty $\phi(w)$. Figure 3b shows that the principal avoids large immediate punishments if the agent's continuation value is low since they can lead to terminating the agent prematurely. To compensate for such lack of immediate punishments, the incentive compatibility condition (13) increases the weight on delayed punishment by reducing the monitoring intensity. As a result, $f(w)$ is low for low continuation values w and is exactly equal to $\frac{w-h/\Delta}{h/\Delta}$ when w is close to h/Δ, as illustrated by Figure 3a. Even though the agent exerts high effort in this region, he is less informed about his performance, allowing the principal to backload pay-for-performance risk until the final date τ.

The principal monitors the agent more when w is higher, as can be seen in Figure 3a, because he is more likely to withstand consecutive punishments and is, thus, less likely to be terminated before retirement. As can be seen in Figure 3b, delayed incentives can be so powerful for a large w that the principal can all but eliminate the penalty $\phi(w)$ while maintaining incentive provision. The intuition is that the agent has a significant stake in the firm that is not worth risking by not reporting a bad project, similar to Piskorski and Westerfield (2016). If the continuation value w is large, then the agent is not punished at all for bad projects that are uncovered by the principal, and the principal's monitoring improves the firm's performance. At the same time, the unmonitored projects carry a lot of performance sensitivity without causing agency distortions before the agent's retirement.

Optimal monitoring intensity $f(w)$ and penalty $\phi(w)$ only depend on the agent's promised utility w_t and not the belief $p_t = \mathrm{P}_t(U_t = 0)$ given by (4) since, along the path of recommended effort, the principal and the agent share the same belief p_t, making both of them value the payment C_τ at the final date τ as exactly w_τ, as described by (6). The optimal monitoring intensity $f(w)$ is, thus, a solution to the principal's dynamic program (11) which only features the agent's continuation value w as a state variable. Probability p_t is then recovered forward using the path of past monitoring intensities $(f(w_s))_{s\le t}$ via (4). The payment C_τ made to the agent at the final date τ if $U_\tau = 0$ depends on both the agent's continuation value w_τ and belief $p_{\tau-}$ via (6). Along the path of play the arrival intensity of process U_t is given by $\mu(1 - f_t)$, leading to a higher probability p_t if the agent has been monitored a lot in the past, reducing the pay-for-performance risk that is delayed to the final date τ. If, however, the agent is monitored less, which is more likely if he performed poorly in the past, then his incentives are substantially backloaded until the final date τ.

Value of backloaded punishments. The large literature on contract design under repeated moral hazard, e.g., Spear and Srivastava (1987) and DeMarzo and Sannikov (2006), has explored the value of deferring compensation to states of the world where these monetary transfers are relatively less costly to the principal. A common manifestation of this is paying the agent when he has a sufficiently high stake in the firm, and a marginal decrease in the agent's continuation value does not substantially increase the likelihood of inefficient termination. This paper shows that the principal can also backload pay-for-performance risk, i.e., punishments, to states of the world in which the prospect of inefficient termination is less costly to the principal. Since the agency cost of the pay-for-performance risk fuels the possibility of terminating the agent prematurely, the principal postpones some of the continuation value variation to the agent's retirement date η when he is leaving the firm anyway.[27] Section 3.5 shows that, if the agent is impatient, then the same argument implies that transfers made before his retirement are accompanied by partial disclosure of latent performance U_t to maintain sharp incentives.

Delaying disclosure of performance results is the only way to backload pay-for-performance sensitivity as, if the agent observed performance, he would rationally infer its implications on his compensation in the long run. Even if the contract backloads transfers, it does not defer contemporaneous variation in the agent's continuation value. The principal must keep the agent uninformed in order to postpone

[27] Highlighting the intuition further, the pay-for-performance risk would be optimally postponed to the final date of the relationship in the finite horizon version of the model considered in Section B.2 of the Online Appendix.

the time when punishments are incorporated into the agent's expectation of his future compensation. The fact that the principal does not communicate with the agent beyond the monitoring results ensures that the contract defers both performance information and associated punishments until either the agent's termination or retirement. Moreover, Proposition 3 highlights that the principal optimally keeps the agent less informed when his continuation value is low, while she can be more transparent if the agent already has a high stake in the firm. Monitoring may inform the agent about his performance, limiting the fraction of the pay-for-performance risk the principal can backload until the final date τ.[28] The most efficient way to backload incentives is to avoid paying the agent if $U_\tau = 0$, as shown by Lemma 3. This also allows the principal to "reuse" incentives, similar to Fuchs (2007), since, if $\lambda = 0$, the principal can set $f_t = 0$ and then, by starting with a continuation value $w_0 = h/\Delta$, she can incentivize the agent to exert effort until retirement, while only giving the agent an information rent worth one unit of effort. The optimality of backloading incentives is also present in settings where the agent's actions have persistent effects on production.[29] In such environments, the optimal contract spreads the performance sensitivity over multiple periods as subsequent production outcomes are informative about past actions. This paper shows that such incentive structure can benefit the principal beyond the incremental information contained in future performance signals even if the agent's actions have no persistent implications on the production technology.

Effect of monitoring on compensation and incentives. The model carries two key predictions. First, the optimal monitoring intensity $f(w)$ weakly increases in the agent's continuation value w. Thus, as long as no bad projects are detected, the principal optimally increases the monitoring intensity. Eventually, the agent who has performed well in the past is penalized less for bad projects identified by the principal, exhibiting a form of rational favoritism – delayed incentives alone are sufficient to incentivize effort when the agent's continuation value is sufficiently large. The agent with intermediate performance, as proxied by the continuation value w close to $(1+\bar{f})\cdot h/\Delta$, faces the highest immediate pay-for-performance sensitivity. These findings complement the analysis of Piskorski and Westerfield (2016) who also find that the optimal pay-for-performance sensitivity is non-monotone in past performance in a setting where monitoring provides incremental signals about the agent's effort.

The compensation plan can also serve as a commitment device for the principal to the optimal monitoring intensity. For example, if the agent reports a bad project, the principal can publicly mix with probability $f(w)/\bar{f}$ between a compensation structure that, upon uncovering a bad project, either prescribes a penalty $\phi(w)$ to the agent or takes away all of the agent's continuation value altogether. In the former case, the principal can afford to monitor the agent with the maximum intensity $\bar{f}$, while in the latter case, uncovering a bad project leads to immediate termination, making monitoring

[28] If the principal could monitor projects without leaking their results to the agent the optimal contract would always monitor the agent, but also compensate him with a single bonus at retirement if all recommended projects were good.

[29] E.g., DeMarzo and Sannikov (2011), He (2012), Sannikov (2014), Marinovic and Varas (2018).

prohibitively costly for the principal. Such public randomization, coupled with a commitment to compensation, alleviates the need for the principal to commit to a monitoring policy explicitly.

3.3 Comparative Statics

This section illustrates the trade-off between investment efficiency and agency costs arising from monitoring by showing how optimal policies respond to the changes in the arrival intensity of bad projects μ, the severity of the agency conflict Δ, and the players' discount rate r. These comparative statics are obtained numerically due to the complexity of the first-order condition (16). All parameters are assumed to satisfy (3) and $\Delta \geq \gamma$.[30]

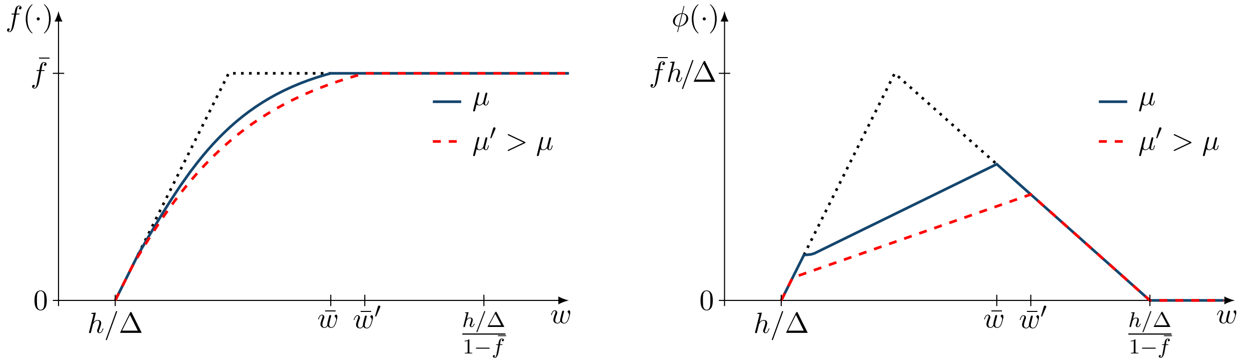

(a) Optimal monitoring intensity $f(w)$.

(b) Optimal contemporaneous penalty $\phi(w)$.

Figure 4: Optimal policies comparative static with respect to arrival intensity μ.

Arrival intensity μ captures the total flow of bad projects that the agent cannot detect himself. On the one hand, the value of monitoring increases with μ as the flow of incoming projects becomes worse. On the other hand, the agency cost of monitoring also increases the intensity of the agent suffering penalty $\phi(w)$ leading to a higher probability of inefficient project termination in the future. Figure 4a shows that the increase in the agency costs dominates, and the principal monitors the agent less if μ is higher. It highlights a potential complementarity channel in which a worse production technology also receives less input from the principal due to increased agency costs.

The degree of the agency conflict is captured by the incremental intensity Δ that the agent generates a bad project if he shirks. A lower Δ reduces the consequences of shirking for the agent, tightening the effort provision constraint (9) and making the agency conflict more severe. This reduces the optimal monitoring intensity $f(w)$ implemented by the optimal contract, as can be seen in Figure 5a, while having a non-monotone effect on the penalty $\phi(w)$, as can be seen in Figure 5b, since the total rents required by the agent increase.

If the players are less patient, the principal places greater value on immediate investment efficiency, thus increasing the optimal monitoring intensity after every history, as illustrated by Figure 6a. More-

[30]Baseline model parameters: $h = 0.05$, $\mu = 1$, $r = 0.01$, $\lambda = 0$, $\Delta = 0.5$, $\gamma = 0.01$, $\alpha = 1.3$, $l = 0.85$, $\bar{f} = 0.7$.

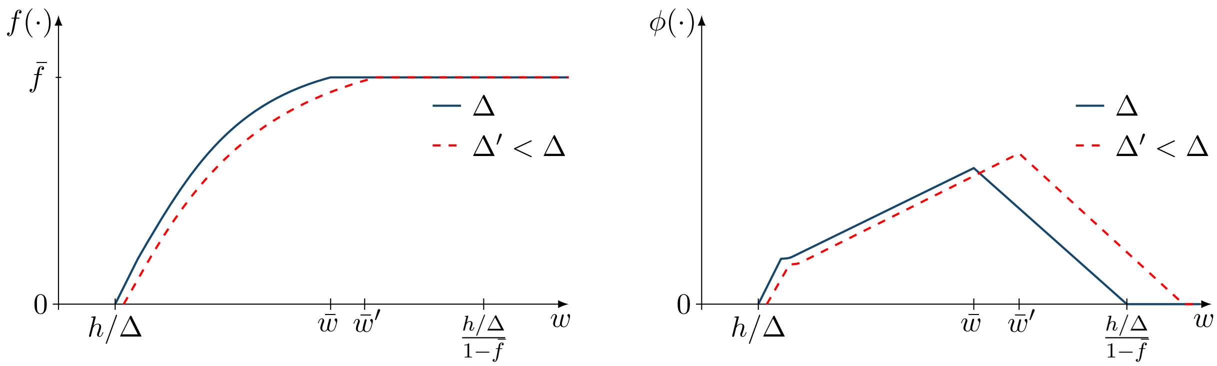

(a) Optimal monitoring intensity $f(w)$.

(b) Optimal contemporaneous penalty $\phi(w)$.

Figure 5: Optimal policies comparative static with respect to arrival intensity Δ.

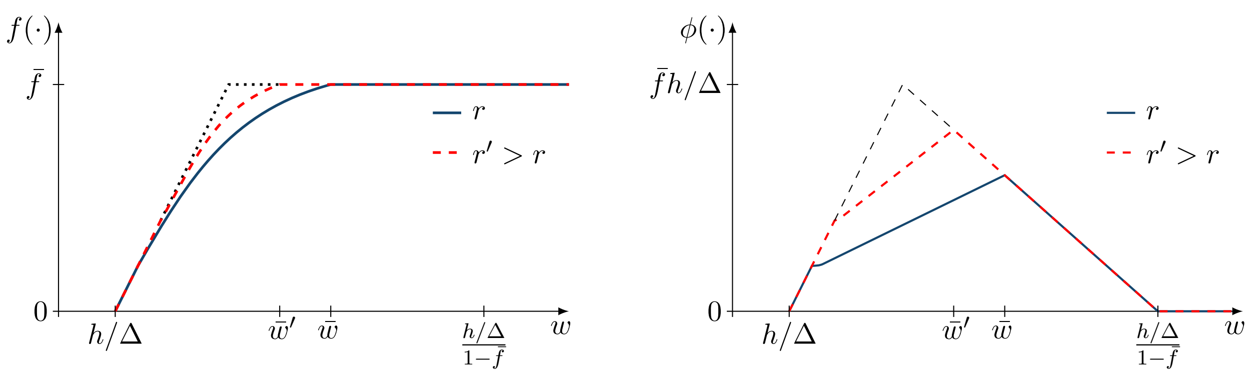

(a) Optimal monitoring intensity $f(w)$.

(b) Optimal contemporaneous penalty $\phi(w)$.

Figure 6: Optimal policies comparative static with respect to discount rate r.

over, it also shows that the optimal policy gradually converges to the "impatient" policy obtained in Proposition 3 as the discount rate r increases.[31]

3.4 Comparison of Optimal Monitoring and Full Monitoring

The endogenous information structure implied by the principal's monitoring policy is the distinguishing feature of this paper relative to prior work on dynamic contracts. The case when the agent observes output immediately corresponds to the principal fully monitoring the agent, i.e., $f(w) \equiv 1$. The agent's continuation value evolves according to

$$dw_t = rw_t\,dt + h\,dt + \hat{\phi}_t \cdot \big(dX_t - (\mu+\lambda)\,dt\big)$$

subject to the effort provision constraint $\hat{\phi}_t \geq h/\Delta$. As shown in Biais, Mariotti, Rochet, and Villeneuve (2010), this incentive constraint is binding under the optimal contract and, under full monitoring, the principal's value function $\hat{v}(w)$ satisfies the delay differential equation for $w \geq h/\Delta$

$$(r+\gamma)\cdot\hat{v}(w) = \alpha-(\lambda+\mu)\cdot l-\gamma w+\hat{v}'(w)\cdot\big(rw+h+(\mu+\lambda)\cdot h/\Delta\big)+(\mu+\lambda)\cdot\big(\hat{v}\left(w-h/\Delta\right)-\hat{v}(w)\big). \quad (18)$$

The delay differential equation (18) can be obtained directly from (11) by setting $f(w) = \bar{f} = 1$ and $\phi(w) = \psi(w) = h/\Delta$. For $w < h/\Delta$ the principal stochastically lets go of the agent, resulting in (12). At the final date τ the agent is paid out his continuation value w_τ, which has no residual dependence on U as no bad project is undertaken under full monitoring.

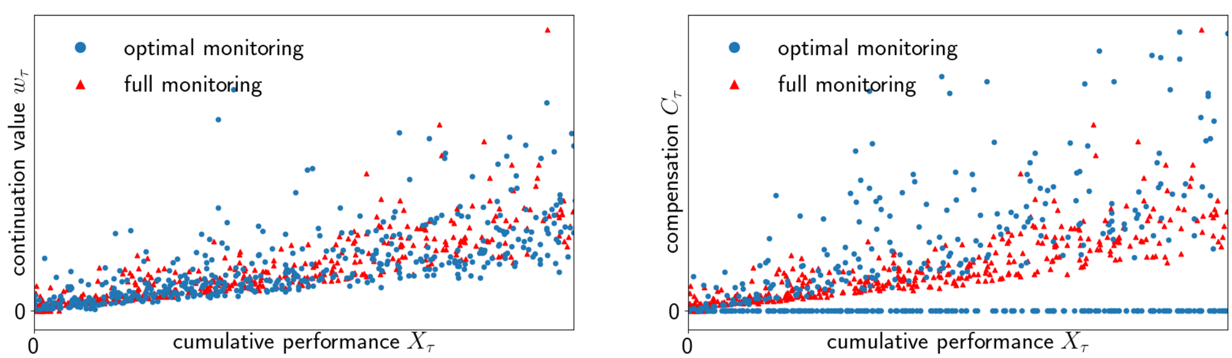

(a) Performance and continuation value.

(b) Performance and compensation.

Figure 7: Relationship between performance, continuation values, and compensation at the final date τ under full monitoring $f_t \equiv 1$ and optimal monitoring $f_t = f(w_t)$. Model parameters: $h = 0.05$, $\mu = 0.5$, $r = 0.01$, $\lambda = 0$, $\Delta = 0.5$, $\gamma = 0.01$, $\alpha = 0.8$, $l = 0.9$, $\bar{f} = 1$. [32]

Under full monitoring, the principal cannot defer incentives until the agent's retirement and imposes a penalty h/Δ for every bad project uncovered via monitoring. This leads to a more volatile contin-

[31] In the extreme case when r is so large that it violates parametric condition (3), the optimal monitoring policy switches from the impatient policy to the principal implementing maximum monitoring $\bar{f}$ for all continuation values w.

uation value process, which exposes the agent to a greater risk of early termination. To reduce the costs of such inefficient termination, the principal offers the agent a contract with a higher starting continuation value, awarding him a greater ex-ante buffer against bad projects. These greater ex-ante rents translate into greater continuation values at the final date τ, as can be seen in Figure 7a which compares the agent's continuation values at time τ between the full- and optimal monitoring contracts. Figure 7b, however, shows that compensation is significantly more volatile under the optimal monitoring contract, as the agent faces a significant pay-for-performance risk associated with latent performance U_τ accumulated as a result of deferred punishments across many periods. Such back-loading of pay-for-performance sensitivity, attained via partial monitoring, underlies the difference between continuation values and compensation in Figures 7a and 7b.

3.5 Optimal Contract for an Impatient Agent

The optimal contract in Section 3.2 is derived under the condition that the principal and agent are equally patient. In this setting, the principal can defer all compensation up until time τ with the agent receiving a positive transfer only if $U_\tau = 0$, according to Lemma 3. If, however, the agent's discount rate ρ strictly exceeds the discount rate r of the principal, then deferring all of the agent's compensation until the final date τ may be too expensive for the principal and, as illustrated by DeMarzo and Sannikov (2006) and Biais, Mariotti, Rochet, and Villeneuve (2010), the principal makes interim transfers to the agent when his continuation value is sufficiently high. Such transfers pose a challenge for the optimal contract derived in Section 3.2 since, if Lemma 3 continues to hold, then the agent can only be paid at time t if $U_t = 0$. I show that the principal easily resolves this conundrum by probabilistically disclosing good past performance $U_t = 0$ to the agent if his continuation value exceeds a threshold $\bar{w}$. Disclosure of good performance is then accompanied by a discrete bonus. This is the key distinction from DeMarzo and Sannikov (2006) and Biais, Mariotti, Rochet, and Villeneuve (2010) in that the principal utilizes both performance feedback and bonuses contingent on it when making interim transfers to the agent. See Section B.3 of the Online Appendix for more details and the formal analysis of the case when the agent is impatient relative to the principal.

4 Renegotiation and Observability

In this section, I characterize the optimal renegotiation-proof contract and show that it changes discontinuously as intensity λ converges to 0 – if $\lambda = 0$ then the principal can avoid distress with certainty which improves her ability to renegotiate the contract along the path of play. For termination to be a credible threat by the principal, it is convenient to introduce a positive value $K > 0$ that the principal

[32]Parameters differ slightly from Section 3.3 to accommodate $\bar{f} = 1$ while still satisfying parametric condition (3).

collects when she terminates the agent and liquidates the production technology.[33] I adopt the notion of weak renegotiation-proofness from Bergemann and Hege (2005) and Strulovici (2011).

Definition 1. *A contract is renegotiation-proof if there are no two public histories such that the principal-agent pair's continuation payoff following the first history Pareto dominates their continuation payoff following the second history.*

A dynamic contract may not be renegotiation proof if the principal's value function $v(w)$ is ever increasing in w since she would benefit from offering the agent a new contract with a higher continuation value. Such a lack of commitment to punishments distorts the agent's ex-ante incentives by breaking the promise-keeping property of the contract. In this model, the value function of the principal may be increasing between 0 and h/Δ; however, the contract may still be renegotiation-proof if the agent's continuation value does not enter this region.

Corollary 1. *If $\lambda = 0$, then the optimal contract is renegotiation-proof if $v'(h/\Delta) \leq 0$. If $\lambda > 0$, then the optimal contract is renegotiation-proof if $v'(0) \leq 0$.*

If $\lambda > 0$, then process Y has a positive arrival intensity even if the agent works. The penalty $\psi_t \equiv h/\Delta$ for reporting bad projects stemming from the process Y implies that the agent's continuation value may drop arbitrarily low following particularly bad realizations of the process Y. Hence, for $\lambda > 0$, the support of the agent's on-path continuation values includes 0 and, thus, the principal's payoff frontier must be downward-sloping at 0. If, however, $\lambda = 0$, then the agent does not report bad projects along the path and, by monitoring poorly-performing agents less, the principal can keep the agent's continuation value above h/Δ. In this case, the principal's payoff frontier must be decreasing for $w \geq h/\Delta$. Thus, by monitoring the agent less, the principal can significantly expand the set of possible payoffs under the optimal renegotiation-proof contract.

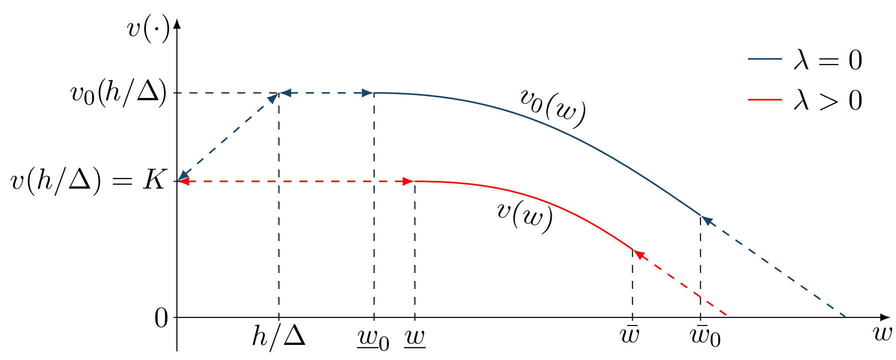

Figure 8: Principal's value function under an optimal renegotiation-proof contract.

[33]Introducing K does not change the main results as long as it is not too large relative to operating the firm.

Proposition 4. *The optimal renegotiation-proof contract is characterized by two boundaries* $\underline{w}$ *and* $\overline{w}$*. For* $w \in [\underline{w}, \overline{w}]$ *the principal's value function solves* (B.51)*. For* $w > \bar{w}$ *the value function is given by* $v(w) = v(\bar{w}) - (w - \bar{w})$*. For* $w < \bar{w}$ *two cases are possible:*

- *if* $\lambda > 0$*, then for* $w < \underline{w}$ *the principal lets go of the agent with probability* $1 - w/\underline{w}$*, and retains the agent with probability* $w/\underline{w}$ *and continuation value* $\underline{w}$*;*
- *if* $\lambda = 0$*, then for* $w < h/\Delta$ *the principal lets go of the agent with probability* $1 - w/\underline{w}$*, and retains him with probability* $w/\underline{w}$ *and continuation value* $\underline{w}$*. For* $w \in [h/\Delta, \underline{w}]$ *the principal randomizes the agent's continuation value between* h/Δ *and* $\underline{w}$*.*

The principal does not wish to terminate the agent inefficiently and, thus, has an incentive to "restart" the contract when the agent's continuation value declines. Such renegotiation, however, distorts the agent's ex-ante incentives to exert effort. For termination to be credible, the principal must be indifferent between keeping the agent employed and letting him go when the contract stipulates termination. The agent must, then, have a sufficiently high continuation value in the contract such that the principal weakly prefers termination to other outcomes. This pins down $\underline{w}$ such that for $w < \underline{w}$ it is incentive compatible for the principal to terminate the agent when prescribed by the contract. Figure 8 illustrates the distinction between $\lambda > 0$ and $\lambda = 0$ which follows from Corollary 1 where, if $\lambda = 0$, then the principal's value $v(h/\Delta)$ acts as her endogenous outside option.

4.1 Observability of Performance in a Renegotiation-proof Contract

An important assumption making the analysis of the main model in Section 3 feasible is that the agent privately observes the "effort" sensitive component of his performance $Y = (Y_t)_{t \geq 0}$. Even after shirking and observing a bad project, the agent can report it to the principal, suffer the punishment ψ_t, and return to the recommended effort path without further repercussions. The resulting incentive compatibility conditions (8) and (9) are simple and permit the characterization of the optimal monitoring policy $f(w)$ as a solution to the principal's dynamic program (11). The drawback of assuming the agent observes process Y is that it may exogenously increase his agency rent since he has access to additional performance information.

An alternative way to model the information available to the agent is to assume he only observes his effort and relies solely on monitoring outcomes and the principal's communication to infer his performance. I show in Proposition 5 below that the optimal renegotiation-proof contract can still be tractably characterized and is similar to the renegotiation-proof contract obtained in Proposition 4. The intuition for such correspondence is that, in a renegotiation-proof setting, the principal cannot commit to ex-post excessive punishments leading to the least binding incentive constraint being implemented after every history. The optimal contract, thus, keeps the agent indifferent between working and shirking until retirement, which is equivalent to effort-provision constraint (13).

Suppose the arrival intensity of process X is given by $\mu + \Delta \cdot (1 - a_t)$, i.e., $\lambda = 0$, and the agent does not observe process Y. Proposition 5 shows that the optimal renegotiation-proof contract is the same as if the agent privately observed process Y.[34]

Proposition 5. *Suppose the agent does not observe process $Y = (Y_t)_{t \geq 0}$. As long as either the agent's discount rate exceeds that of the principal's, i.e., $\rho > r$ and/or the monitoring capacity is equal to $\bar{f} = 1$, the optimal renegotiation-proof contract coincides with the one obtained in Proposition 4 with parameters μ and $\lambda = 0$. Moreover, if $\rho - r + \Delta - \gamma \geq 0$, then the optimal monitoring intensity is bounded from above by $\frac{w - h/\Delta}{h/\Delta}$.*

Incentive constraint (13) was necessary and sufficient if the agent observed process Y which means it remains sufficient if the agent doesn't as, in the former case he could have always ignored the additional information contained in Y if it led to a more profitable global deviation. What remains is then to show that incentive condition (13) is also necessary. By Lemma 3 the optimal contract pays the agent 0 at the final date τ if $U_\tau > 0$. Under an alternative effort process $\hat{a} = (\hat{a}_t)_{t \geq 0}$ the agent's off-path continuation value is equal to

$$\begin{aligned} \mathrm{E}_{\hat{a}} \left[e^{-\rho\tau} \cdot C_\tau - \int_0^\tau e^{-\rho t} \hat{a}_t h \, dt \right] &= \mathrm{E}_{\hat{a}} \left[e^{-\rho\tau} \cdot \frac{w_\tau}{p_\tau} \cdot \mathbb{1}\{U_\tau = 0\} - \int_0^\tau e^{-\rho t} \hat{a}_t h \, dt \right] \\ &= \mathrm{E}_{\hat{a}} \left[e^{-\rho\tau} \cdot w_\tau \cdot \frac{\hat{p}_\tau}{p_\tau} - \int_0^\tau e^{-\rho t} \hat{a}_t h \, dt \right] \end{aligned} \tag{19}$$

where belief $\hat{p}_t$ is given by Bayes rule as

$$\hat{p}_t \stackrel{def}{=} \mathrm{P}_{\hat{a}}(U_t = 0) = e^{-\int_0^t (\mu + (1 - \hat{a}_s) \cdot \Delta) \cdot (1 - f_s) \, ds}. \tag{20}$$

If the agent shirks at time t, he affects his private belief $\hat{p}_t$ and reduces the expected reward from the contract. This leads him to, potentially, exert less effort in subsequent periods leading to a double-deviation problem.[35] The partial derivative with respect to $\hat{p}_0$ at $\hat{p}_0 = p_0 = 1$ of (19) is given by $\mathrm{E}_{\hat{a}}\left[e^{-r\tau} w_\tau\right]$. Because the agent picks out-of-equilibrium effort $\hat{a}$ optimally, this expression is evaluated with respect to the lowest (in terms of expected cost) effort $\hat{a}$ such that

$$\mathrm{E}_{\hat{a}} \left[e^{-\rho\tau} C_\tau - \int_0^\tau e^{-\rho t} \hat{a}_t h \, dt \right] = \mathrm{E}_a \left[e^{-\rho\tau} C_\tau - \int_0^\tau e^{-\rho t} a_t h \, dt \right] = w_0. \tag{21}$$

Denote by A_t^* to be the smallest expected discounted cost of an effort profile which keeps the agent ex-ante indifferent between it and the recommended effort:

[34] Under the parametric restriction $\lambda = 0$ process $Y \equiv 0$ along the path of high effort. If $\lambda > 0$, then lack of observability of process Y corresponds to new arrival intensities $\hat{\mu}, \hat{\lambda}$ given by $\hat{\mu} = \mu + \lambda$ and $\hat{\lambda} = 0$ in the context of the renegotiation-proof contract derived in Proposition 4.

[35] Persistence in agent's deviation decisions is a complicated issue and has been studied in different settings such as He (2012), DeMarzo and Sannikov (2011), Sannikov (2014), Strulovici (2011), and Marinovic and Varas (2018) among others. In this paper the double deviation problem arises from the fact that the agent holds a more pessimistic belief $\hat{p}$ about being compensated at the final date τ if he deviates.

$$A_t^* \stackrel{def}{=} \inf_{\hat{a}\in\mathbb{F}^c} \mathrm{E}\left[\int_t^\tau e^{-\rho(s-t)}\hat{a}_s h\, ds \,\Big|\, \mathcal{F}_t^c\right], \tag{22}$$

subject to the infinum taken with respect to effort processes $\hat{a}$ satisfying (21). By the Envelope theorem, a necessary condition for effort to be incentive compatible is

$$h/\Delta \le f_t \cdot \phi_t + (1-f_t)\cdot(w_t + A_t^*). \tag{23}$$

Since the expected effort cost of the agent's effort A_t^* is weakly positive, condition (23) is weaker than (13). This is intuitive since a less informed agent is easier to motivate and it seems to suggest that the principal can improve upon the optimal contract derived earlier. In what follows, I show that under the optimal renetotiation-proof contract $A_t^* \equiv 0$ implying the equivalence of (23) and (13).

Incentive constraint (23) must bind along the path of the optimal renegotiation-proof contract. This follows from the concavity of the principal's payoff frontier in the agent's promised utility.[36] If $\bar{f} = 1$, the principal's value function is strictly concave even for large continuation values w since there always exists a chance of entering the termination region after a sequence of bad projects. If the agent is relatively impatient, i.e., $\rho > r$, then, as discussed in Section 3.5, the contract pays the agent a bonus when his continuation value enters the region $w > \bar{w}$ where the principal's value function is linear. Strict concavity of the principal's value function implies that incentive compatibility condition (23) must be binding since the principal would, otherwise, benefit from reducing penalty ϕ_t.

Now, suppose that $A_t^* > 0$ under the optimal contract and (23) binds. Consider a perturbed version of the model in which the agent's effort cost h_t scales down in proportion with the ratio $\hat{p}_t/p_t$, i.e., $h_t = \hat{p}_t/p_t$, where the on-path belief p_t and off-path belief $\hat{p}_t$ that $U_t = 0$ are given by (4). Such perturbation of the model makes deviating more profitable for the agent as it is less costly for him to work afterward, thus, increasing the payoff from shirking while keeping his payoff from working unchanged. A key advantage of such an approach is that the agent's off-path continuation value is then equal to $w_t \cdot \hat{p}_t/p_t$ since he is facing a "rescaled" contract both in terms of compensation and effort cost.[37] It is easy to show that the necessary and sufficient incentive constraint in this "rescaled" model is exactly the necessary and sufficient incentive constraint in the original model, i.e., (13). Consequently, if $A_t^* > 0$ and (23) is binding, then (13) is violated implying that the agent finds it strictly optimal to shirk even if his effort cost is scaled down in subsequent periods. Such once-and-for-all shirking deviation is available to the agent in the original model since shirking has a 0 cost of effort and is not affected by the rescaling of effort in the model. Hence, the renegotiation-proof contract is incentive-compatible only if $A_t^* = 0$, proving the equivalence of incentive constraints (13)

[36]The effect of strict concavity of the objective function of the principal is pointed out in Strulovici (2011) who shows that this leads to a unique Markov process governing the optimal renegotiation-proof contract.

[37]Rescaling of the agent's expected compensation is to him endogenously becoming more pessimistic about being paid at retirement since Lemma 3 also applies to optimal renegotiation-proof contracts.

and (23). Proposition 5 thus shows that the optimal renegotiation-proof contract does not depend on whether or not the agent observes Y, and constitutes a robustness check for the main model.

Observability in discrete models. Complete characterization of the optimal non-renegotiation-proof contract if the agent does not observe Y is beyond the scope of this paper due to the challenge in identifying the agent's optimally binding incentive compatibility constraints under various information structures implied by the principal's monitoring and communication. In Section B.4 of the Online Appendix, I provide a detailed analysis of this problem in the context of a two-period version of the model and then show numerically in a three-period extension that the optimal monitoring intensity decreases following bad performance consistent with the main findings of this paper.

5 Conclusion

This paper identifies an agency cost imposed by frequent monitoring of employees. Monitoring may inform the agent about his output and limits the principal's ability to delay pay-for-performance risk to states of the world when it is cheapest. This highlights the benefits of deferring the agent's continuation value variation that is distinct from the well-established optimality of deferring transfers. Optimal compensation stakes the agent's final date continuation value on the imperfectly observed performance and reuses incentive compensation over many periods. Furthermore, the optimal monitoring intensity depends on performance and is lower for agents who performed poorly in the past, while it may result in leniency towards agents who performed well. These results highlight the importance of the agent's informedness of his performance in compensation design and should be taken into consideration in the context of monitoring decisions, performance appraisals, and employee feedback.

6 Published Appendix: Proofs of Lemma 2, 3, and Proposition 3

This section provides the proofs for Lemmas 2, 3, and Proposition 3 of the paper. Due to space constraints of the published version of the paper, the remaining proofs and supplementary analysis are contained in Online Appendices A and B respectively.

6.1 Joint Proof of Lemmas 2 and 3

Consider a candidate optimal contract $\mathcal{C} = (a, f, C, \mathbb{F}^c)$. By Lemma 1, we can focus on contracts implementing truthful disclosure $R_t = Y_t$ by the agent. Since both players are equally patient and the principal has commitment power, it is without loss to postpone all compensation until time τ when the agent either retires or is let go by the principal.

Define $\mathcal{F}'_t \overset{def}{=} \sigma\{(a_s)_{s\le t},\ (f_s)_{s\le t},\ (M_s)_{s\le t},\ (R_s)_{s\le t},\ C_{\tau\wedge t}\}$. By construction, $\mathcal{F}'_t \subseteq \mathcal{F}^c_t$ as it includes only the information contained in the principal's effort recommendations, monitoring frequency, monitoring outcomes, and the agent's reports, but does not include any additional communication by the principal that could be part of $\mathcal{F}_t^c$. Consequently, a candidate contract $\mathcal{C}' \overset{def}{=} (a, f, C, \mathbb{F}' = (\mathcal{F}_t')_{t\geq 0})$ is incentive compatible since the agent is less informed for each history under filtration $\mathbb{F}'$ than under $\mathbb{F}^c$, implying that the set of available deviation processes $\hat{a}$ available to the agent is lower under filtration $\mathbb{F}'$ than under the original filtration $\mathbb{F}^c$. Thus, there exists an optimal contract in which $\mathbb{F}^c = \mathbb{F}'$.

A contract $\mathcal{C}$ specifies an on-path stochastic process $(a_t, f_t, R_t, M_t, U_t, C_t)_{t\geq 0}$. In what follows, I substitute the conditional dependence of processes $(a_t)_{t\geq 0}$ and $(f_t)_{t\geq 0}$ on the "hidden" process $(U_t)_{t\geq 0}$ for a process $(Z_t)_{t\geq 0}$ that has the same on-path distribution as process $(U_t)_{t\geq 0}$ but is independent of the agent's effort or processes Y and N. Given processes M and U for each $k \in \mathbb{Z}_+$ define stopping times $\tau_k^M \overset{def}{=} \inf\{t\colon\ M_t = k\}$ and $\tau_k^U \overset{def}{=} \inf\{t\colon\ U_t = k\}$. The distribution of process $(a_t, f_t, R_t, M_t, U_t)_{t\geq 0}$ is uniquely determined by the joint distribution of $(M_t, U_t)_{t\geq 0}$ and the conditional distributions

$$\text{Law}\Big[\quad (a_t, f_t, R_t, C_t)_{t\geq 0} \quad \Big| \quad \left\{\tau_k^M\right\}_{k=1}^m = \left\{t_k^M\right\}_{k=1}^m, \quad \left\{\tau_k^U\right\}_{k=1}^u = \left\{t_k^U\right\}_{k=1}^u \quad \Big].$$

Given the on-path process (a, f, R, M, U, C), I construct inductively an alternative incentive compatible contract $\mathcal{C}' = (a', f', R', C')$ under which $(a', f', R', M', C') \overset{Law}{=} (a, f, R, M, C)$, $U' \overset{Law}{=} U$, and (a', f', R', M') is independent of U'. To do so, let $\hat{N} = (\hat{N}_t)_{t\geq 0}$ be a Poisson process with arrival intensity μ that is independent from $(N, Y, a, f, M, R, C)_{t\geq 0}$. Also, let $(\xi_k, \hat{\xi}_k)_{k\in\mathbb{N}}$ be a sequence of uniformly distributed random variables independent across each other as well as from $(\hat{N}, N, Y, a, f, M, R, C)_{t\geq 0}$.

1. **Induction base.** Consider a realization of an on-path history conditional on $N_\infty = 0$:

$$(\tilde{S}_t)_{t\geq 0} \overset{def}{=} \left(\tilde{a}_t, \tilde{f}_t, \tilde{R}_t, \tilde{C}_t\right)_{t\geq 0} \overset{P}{\sim} \text{Law}\Big[\quad \left(a_t, f_t, R_t, C_t\right)_{t\geq 0} \quad \Big| \quad \tau_1^M = \infty, \quad \tau_1^U = \infty \quad \Big].$$

   Conditional on $N_t = 0$ the arrival intensity of a bad project uncovered via monitoring is $\tilde{f}_t \cdot \mu$, while the arrival intensity of bad project that gets investment is $(1 - \tilde{f}_t) \cdot \mu$. Using the realization of process $(\tilde{a}_t, \tilde{f}_t)_{t\geq 0}$ define two stopping times

$$\tilde{\tau}_1^M \overset{def}{=} \inf\left\{t\colon \quad \int_0^t \mathbb{1}\left\{\xi_{N_s} \leq f_s\right\} dN_s = 1\right\}, \qquad \tilde{\tau}_1^Z \overset{def}{=} \inf\left\{t\colon \quad \int_0^t \mathbb{1}\left\{\hat{\xi}_{\hat{N}_s} > f_s\right\} d\hat{N}_s = 1\right\}.$$

   For any $t < \min\left\{\tilde{\tau}_1^M,\ \tilde{\tau}_1^Z\right\}$ define processes

$$\tilde{M}_t \overset{def}{=} \int_0^t \mathbb{1}\left\{\xi_{N_s} \leq f_s\right\} dN_s, \quad \tilde{U}_t \overset{def}{=} \int_0^t \mathbb{1}\left\{\xi_{N_s} > f_s\right\} dN_s, \quad \tilde{Z}_t \overset{def}{=} \int_0^t \mathbb{1}\left\{\hat{\xi}_{\hat{N}_s} > f_s\right\} d\hat{N}_s. \quad (24)$$

   Values $\tilde{M}_t$ and $\tilde{U}_t$ capture the bad project that is uncovered either via monitoring or invested into. They follow the same conditional distribution as the original contract and, conditional on $(\tilde{f}_t)_{t\geq 0}$, these processes are independent. Process $\tilde{Z}_t$ follows the same distribution as process $\tilde{U}_t$, however it is independent from all other variables since it is defined via the exogenous Poisson process $\hat{N}_t$ and random variables $\hat{\xi}_k$. For $t \leq \min\left\{\tilde{\tau}_1^M,\ \tilde{\tau}_1^Z\right\}$ define the path of the alternate contract as $(S_t', U_t', Z_t') \overset{def}{=} \left(\tilde{S}_t, \tilde{U}_t, \tilde{Z}_t\right)$.

2. **Induction step:** extending the contract path to $t \geq \min\left\{\tilde{\tau}_1^M,\ \tilde{\tau}_1^Z\right\}$.

   (i) Suppose $\tilde{\tau}_1^M \leq \tilde{\tau}_1^Z$. Generate a new realization of process $\tilde{S}$ for $t \geq \tilde{\tau}_1^M$ from the conditional distribution specified by the original contract $\mathcal{C}$ conditional on (A) the history up to $\tilde{\tau}_1^M$ given by $(S'_t)_{t \leq \tilde{\tau}_1^M}$, (B) the first bad project being uncovered via monitoring at time $\tilde{\tau}_1^M$, and (C) no bad projects forever after:

$$\left(\tilde{S}_t\right)_{t\geq 0} \overset{P}{\sim} \text{Law}\left[\left(a_t, f_t, R_t, C_t\right)_{t\geq 0} \mid \left(a_t, f_t, R_t, C_t\right)_{t\leq \tilde{\tau}_1^M} = \left(S'_t\right)_{t\leq \tilde{\tau}_1^M}, \tau_1^M = \tilde{\tau}_1^M, \tau_2^M = \infty, \tau_1^U = \infty\right].$$

   Using the realization of process $(\tilde{a}_t, \tilde{f}_t)_{t\geq \tilde{\tau}_1^M}$, define $\tilde{\tau}_2^M$ and redefine $\tilde{\tau}_1^Z$ as

$$\tilde{\tau}_2^M \overset{def}{=} \inf\left\{t: \quad \int_0^t \mathbb{1}\left\{\xi_{N_s} \leq f_s\right\} dN_s = 2\right\}, \quad \tilde{\tau}_1^Z \overset{def}{=} \inf\left\{t: \quad \int_0^t \mathbb{1}\left\{\hat{\xi}_{\hat{N}_s} > f_s\right\} d\hat{N}_s = 1\right\}.$$

   For $t \in \left[\tilde{\tau}_1^M,\ \min\left\{\tilde{\tau}_2^M, \tilde{\tau}_1^Z\right\}\right]$ extend processes $\tilde{M}$, $\tilde{U}$, and $\tilde{Z}$ via (24) and extend the contract path as $(S'_t, U'_t, Z'_t) \overset{def}{=} \left(\tilde{S}_t, \tilde{U}_t, \tilde{Z}_t\right)$.

   (ii) Suppose $\tilde{\tau}_1^Z < \tilde{\tau}_1^M$. This is the step where we replace the dependency of the policy functions of the original contract on the latent performance process U with dependence on the identically distributed process Z. Generate a new realization of process $\tilde{S}$ for $t \geq \tilde{\tau}_1^Z$ from the conditional distribution specified by the original contract $\mathcal{C}$ conditional on (A) the history up to $\tilde{\tau}_1^M$ given by $(S'_t)_{t\leq \tilde{\tau}_1^Z}$, (B) the first bad project being missed by monitoring and, consequently undertaken, at time $\tilde{\tau}_1^Z$, and (C) no bad projects forever after:

$$\left(\tilde{S}_t\right)_{t\geq 0} \overset{P}{\sim} \text{Law}\left[\left(a_t, f_t, R_t, C_t\right)_{t\geq 0} \middle| \left(a_t, f_t, R_t, C_t\right)_{t\leq \tilde{\tau}_1^Z} = \left(S'_t\right)_{t\leq \tilde{\tau}_1^Z}, \tau_1^M = \infty, \tau_1^U = \tilde{\tau}_1^Z, \tau_2^U = \infty\right].$$

   Using the realization of process $(\tilde{a}_t, \tilde{f}_t)_{t\geq \tilde{\tau}_1^M}$, define $\tilde{\tau}_2^Z$ and redefine $\tilde{\tau}_1^M$ as

$$\tilde{\tau}_1^M \overset{def}{=} \inf\left\{t: \quad \int_0^t \mathbb{1}\left\{\xi_{N_s} \leq f_s\right\} dN_s = 1\right\}, \quad \tilde{\tau}_2^Z \overset{def}{=} \inf\left\{t: \quad \int_0^t \mathbb{1}\left\{\hat{\xi}_{\hat{N}_s} > f_s\right\} d\hat{N}_s = 2\right\}.$$

   For $t \in \left[\tilde{\tau}_1^M,\ \min\left\{\tilde{\tau}_1^M, \tilde{\tau}_2^Z\right\}\right]$ extend processes $\tilde{M}$, $\tilde{U}$, and $\tilde{Z}$ via (24) and extend the contract path as $(S'_t, U'_t, Z'_t) \overset{def}{=} \left(\tilde{S}_t, \tilde{U}_t, \tilde{Z}_t\right)$.

3. Having performed the above construction up to time $\min\left\{\tilde{\tau}_m^M, \tilde{\tau}_z^Z\right\}$ continue via the step of induction by generating stopping times $\tilde{\tau}_{m+1}^M$ and $\tilde{\tau}_{z+1}^Z$ and extending the contract by induction.

The sequential path-by-path construction above ensures that

$$\begin{aligned}
&\text{Law}\left[\quad (a_t, f_t, R_t, C_t, M_t, U_t)_{t\geq 0} \quad \middle| \quad \left\{\tau_k^M\right\}_{k=1}^m = \left\{t_k^M\right\}_{k=1}^m, \quad \left\{\tau_k^U\right\}_{k=1}^u = \left\{t_k^U\right\}_{k=1}^u \quad \right] \\
=&\text{Law}\left[\quad (a'_t, f'_t, R'_t, C'_t, M'_t, Z'_t)_{t\geq 0} \quad \middle| \quad \left\{\tilde{\tau}_k^M\right\}_{k=1}^m = \left\{t_k^M\right\}_{k=1}^m, \quad \left\{\tilde{\tau}_k^Z\right\}_{k=1}^u = \left\{t_k^U\right\}_{k=1}^u \quad \right],
\end{aligned}$$

with the additional property that process U' is conditionally independent from $(a'_t, f'_t, R'_t, C'_t, M'_t, Z'_t)_{t\geq 0}$. Such construction replaces the dependence of the original contract $\mathcal{C}$ on latent performance U with an exogenous, but identically distributed on-path process Z. This, however, may alter the agent's incentives as the resulting compensation C_t loses its sensitivity to the realizations of latent performance U_t. To address this issue, define compensation $C'^*_\tau \overset{def}{=} C'_\tau \cdot \mathbb{1}\{U'_\tau = 0\} / \mathrm{P}_a(U'_\tau = 0)$ and define the new contract as $\mathcal{C}' \overset{def}{=} (a', f', R', M', Z', U', C'^*)$. Contract $\mathcal{C}'$ is independent of process U', except for the last instance when the agent gets compensated only in the event that $U'_\tau = 0$. Moreover, by construction the agent's expected compensation if he exerts the recommended effort $(a'_t)_{t\geq 0}$ is equal to the expected compensation under the original contract.

Suppose the agent deviates to an alternative effort process $\hat{a}$ and reporting process $\hat{R}$. The agent's reporting strategy is given by a *predictable* process $\hat{d}_t$ such that $\hat{R}_t = \int_0^t \hat{d}_s \, dY_s$. Define the arrival intensity of process U' under the new contract at time t given the reporting strategy $(\hat{d}_t)_{t\geq 0}$ and deviated effort $(\hat{a}_t)_{t\geq 0}$ as

$$\begin{aligned}\hat{\mu}_t &\overset{def}{=} \mu \cdot (1 - f_t) + \big(\lambda + (1 - \hat{a}_t) \cdot \Delta\big) \cdot \big(1 - \hat{d}_t\big) \cdot (1 - f_t) \\ &= \big[\mu + \big(\lambda + (1 - \hat{a}_t) \cdot \Delta\big) \cdot \big(1 - \hat{d}_t\big)\big] \cdot (1 - f_t) \overset{(i)}{\geq} \mu \cdot (1 - f_t) \overset{def}{=} \mu_t.\end{aligned}$$

where μ_t is the rate of arrival of process U_t under truthful reporting and inequality (i) is strict whenever the agent deviates from truthful reporting. Then

$$\begin{aligned}\mathrm{E}_{(\hat{a},\hat{d})}\left[e^{-r\tau} \cdot C'^*_\tau\right] =& \mathrm{E}_{(\hat{a},\hat{d})}\left[\mathrm{E}\left[e^{-r\tau} \cdot C'^*_\tau \mid \hat{R}, M', Z'\right]\right] = \mathrm{E}_{(\hat{a},\hat{d})}\left[\mathrm{E}\left[e^{-r\tau} \cdot C'^*_\tau \cdot \frac{\mathbb{1}\{U'_\tau = 0\}}{\mathrm{P}_a(U'_\tau = 0)} \,\middle|\, \hat{R}, M', Z', U'\right]\right] \\ =& \mathrm{E}_{(\hat{a},\hat{d})}\left[\mathrm{E}\left[e^{-r\tau} \cdot C'^*_\tau \cdot \frac{\mathrm{P}_{(\hat{a},\hat{d})}(U'_\tau = 0)}{\mathrm{P}_a(U'_\tau = 0)} \,\middle|\, \hat{R}, M', Z', U'\right]\right] \\ =& \mathbb{E}_{(\hat{a},\hat{d})}\left[\mathbb{E}\left[e^{-r\tau} \cdot C'^*_\tau \cdot e^{-\int_0^\tau (\hat{\mu}_s - \mu_s)\, ds} \cdot \frac{\mathrm{P}_a(U'_\tau = 0)}{\mathrm{P}_a(U'_\tau = 0)} \,\middle|\, \hat{R}, M', Z'\right]\right] \\ =& \mathrm{E}_{(\hat{a},\hat{d})}\left[\mathrm{E}\left[e^{-r\tau} \cdot C'^*_\tau \cdot e^{-\int_0^\tau (\hat{\mu}_s - \mu_s)\, ds} \,\middle|\, \hat{R}, M', Z'\right]\right] \\ \overset{(i)}{\leq}& \mathrm{E}_{(\hat{a},\hat{d})}\left[\mathbb{E}\left[e^{-r\tau} \cdot C_\tau \cdot e^{-\int_0^\tau (\hat{\mu}_s - \mu_s)\, ds} \cdot \underbrace{e^{\int_0^\tau \log\left(\frac{\hat{\mu}_s}{\mu_s}\right) dI_s}}_{\geq 1} \,\middle|\, \hat{R}, M', Z'\right]\right] \overset{(ii)}{=} \mathrm{E}_{(\hat{a},\hat{d})}\left[e^{-r\tau} C_\tau\right],\end{aligned}$$

where inequality (i) holds because $\hat{\mu}_s \geq \mu_s$ for any $s \geq 0$, and (ii) is the Girsanov change of measure for Poisson processes. Moreover, by definition of C'^*_τ, inequality (i) is an equality along the recommended path of effort, implying that the new contract $\mathcal{C}'$ is incentive compatible for the joint effort and monitoring distribution (a, f) prescribed by the original contract $\mathcal{C}$. Moreover, this new contract $\mathcal{C}'$ postpones information disclosure until time τ is weakly optimal.

6.2 Proof of Proposition 3

To minimize new notation in the printed Appendix, this proof is provided for the baseline model as stated in Proposition 3. The generalized proof for when the agent is impatient or when the principal has a non-zero value to liquidating the project is in Online Appendix A.

Lemma 5. *Suppose $\lambda = 0$, $\Delta \geq \gamma$, and (3) is satisfied. Then $f(w) \leq \frac{w-\delta}{\delta}$.*

Proof. First, we show that $1 - l < v(\delta) - \delta \cdot v'(\delta)$. The principal's HJB (11) at $w = \delta$ implies

$$v'(\delta) \leq \frac{(r+\gamma)\cdot v(\delta) - \Big(\alpha - \mu - \gamma\delta\Big)}{r\delta + h}$$

since the principal can always set the monitoring intensity to 0. Then

$$v(\delta) - \delta \cdot v'(\delta) \geq v(\delta) - \delta \cdot \frac{(r+\gamma)v(\delta) - (\alpha - \mu - \gamma\delta)}{r\delta + h}.$$

Then, a sufficient condition for $1 - l \leq v(\delta) - \delta \cdot v'(\delta)$ to hold is

$$1 - l \leq \frac{v(\delta)\cdot(\Delta - \gamma)\delta}{r\delta + h} + \frac{\delta}{r\delta + h}\cdot\Big(\alpha - \mu - \gamma\delta\Big). \tag{25}$$

If $\Delta \geq \gamma$, then a sufficient condition for (25) to be satisfied is if it is satisfied for a lower bound of $v(\delta)$ obtained under no monitoring until the agent's retirement

$$1 - l \leq \frac{\alpha - \mu - h - (r+\gamma)\delta}{r+\gamma} + \frac{r\delta}{r\delta + h}\cdot\Big(\alpha - \mu - h - \gamma\delta\Big),$$

which is, in turn, always satisfied if it is individually rational for the principal to employ the agent under high effort but no monitoring in perpetuity. This shows that $1 - l \leq v(\delta) - \delta \cdot v'(\delta)$.

Now, suppose that $(w-\delta)/f < \delta$. Then, due to linearity of $v(w)$ for $w \in [0, \delta]$ it follows that

$$v\left(\frac{w-\delta}{f}\right) - \frac{w-\delta}{f}\cdot v'\left(\frac{w-\delta}{f}\right) = 0.$$

It follows from (15) that monitoring intensity $f > \frac{w-\delta}{\delta}$ is suboptimal if and only if

$$1 - l \leq v(w) - w \cdot v'(w). \tag{26}$$

Since $v(w)$ is concave, $v(w) - w\cdot v'(w)$ is increasing in w. Then, (26) is satisfied for $w \geq \delta$ if it is satisfied at $w = \delta$. In this case, it cannot be optimal for $\frac{w-\delta}{f} < \delta$, which concludes the proof. □

Lemma 6. *Suppose $\lambda = 0$. Then, the principal optimally implements $f(w) = \frac{w-\delta}{\delta}$ for $w \leq \delta + h/\mu$.*

Proof. For w close to δ, it is optimal to implement maximum monitoring capacity $f(w) = \frac{w-\delta}{\delta}$. Denote by $w^* \leq (1+\bar{f})\delta$ to be the first time when the maximum monitoring intensity is not binding, but is interior:

$$w^* \overset{def}{=} \inf\Big\{w \geq \delta : \quad 1 - l + v(\delta) - \delta \cdot v'(\delta) = v(w) - w \cdot v'(w)\Big\}. \tag{27}$$

The principal's HJB (11) at $w = w^*$ is

$$(r+\gamma) \cdot v(w^*) = \alpha - \mu + \mu(1-l) \cdot \frac{w^* - \delta}{\delta} - \gamma w^* + v'(w^*) \cdot \Big(rw^* + h + \mu \frac{(w^*-\delta)^2}{\delta}\Big) + \mu \cdot \frac{w^* - \delta}{\delta} \cdot \Big(v(\delta) - v(w^*)\Big).$$

Substitute $v(w^*)$ from (27) into the principal's HJB (11) at $w = w^*$

$$\begin{aligned} &(r+\gamma) \cdot \Big(1 - l + v(\delta) + w^* \cdot v'(w^*) - \delta \cdot v'(\delta)\Big) \\ =&\alpha - \mu + \mu(1-l) \cdot \frac{w^* - \delta}{\delta} - \gamma w^* + v'(w^*) \cdot \Big(rw^* + h + \mu \frac{(w^*-\delta)^2}{\delta}\Big) \\ +&\mu \cdot \frac{w^* - \delta}{\delta} \cdot \Big(v(\delta) - v(\delta) + \delta v'(\delta) - w^* v'(w^*) - (1-l)\Big). \end{aligned}$$

Simplify terms obtain

$$\gamma(w^* - \delta) - \gamma\Big(v(\delta) - v(w^*)\Big) + r(1-l) = \Big(v'(\delta) - v'(w^*)\Big)\Big(\mu(w^* - \delta) - h\Big).$$

From the concavity of $v(w)$ and $v'(w) \geq -1$, it follows that $v(w) \geq v(\delta) - (w - \delta)$.

(i) Case $v(\delta) - v(w^*) < 0$. Then

$$\gamma(w^* - \delta) - \gamma \cdot (v(\delta) - v(w^*)) + r(1-l) > 0.$$

(ii) Case $v(\delta) - v(w^*) > 0$. Then, by concavity of the value function, $v'(w) < 0$. Then

$$\gamma(w^* - \delta - v(\delta) + v(w^*)) + r(1-l) \geq 0.$$

It follows that $\underbrace{\big(v'(\delta) - v'(w^*)\big)}_{\geq 0} \cdot \big(\mu(w^* - \delta) - h\big) \geq 0$, implying that $w^* - \delta \geq \frac{h}{\mu}$. □

Lemma 7. *Suppose $\lambda = 0$ and $\Delta \geq \gamma$. Then, the optimal monitoring intensity $f(w)$ is weakly increasing in w, while the optimal pay-for-performance sensitivity $\phi(w)$ is increasing for $f(w) < \bar{f}$.*

Proof. The binding incentive compatibility condition is $\delta = f \cdot \phi + (1 - f) \cdot w$ can be rewritten as $f = \frac{w-\delta}{w-\phi}$ and $w - \phi = \frac{w-\delta}{f}$. The first order condition (15) with respect to f (expressed via ϕ) is

$$1 - l + v(w - \phi) - (w - \phi)v'(w - \phi) = v(w) - wv'(w) \tag{28}$$

Rewrite (28) via monitoring intensity f to obtain

$$1 - l + v\left(\frac{w-\delta}{f}\right) - \frac{w-\delta}{f}v'\left(\frac{w-\delta}{f}\right) = v(w) - wv'(w). \tag{29}$$

Differentiate (29) with respect to w to obtain

$$\frac{f(w) - f'(w)(w-\delta)}{f(w)^2} \cdot v''\left(\frac{w-\delta}{f}\right) = \frac{wv''(w)}{\frac{w-\delta}{f(w)}}. \tag{30}$$

For $f'(w) \geq 0$ it is sufficient that $w \cdot v''(w)$ to be increasing, as can be seen from

$$\frac{f(w)}{f(w)^2} \cdot v''\left(\frac{w-\delta}{f(w)}\right) \leq \frac{wv''(w)}{\frac{w-\delta}{f(w)}} \qquad \Leftrightarrow \qquad \frac{1}{f(w)} \cdot \frac{w-\delta}{f(w)} \cdot v''\left(\frac{w-\delta}{f(w)}\right) \leq wv''(w).$$

If $f(w)$ is an interior solution, then the Envelope theorem implies

$$\gamma v'(w) = -\gamma + v''(w) \cdot \Big(rw + h + \mu\big(\delta + (f-1)w\big)\Big) + \mu v'\left(\frac{w-\delta}{f}\right) - \mu v'(w).$$

Rewriting this expression using punishment ϕ obtain

$$(\gamma + \mu)v'(w) = -\gamma + v''(w)\big(rw + h + \mu f\phi\big) + \mu v'(w - \phi).$$

The same expression can also be evaluated at ϕ, given the monitoring intensity $\hat{f}$ and punishment $\hat{\phi}$ corresponding to continuation value $\hat{w} = w - \phi$:

$$(\gamma + \mu)v'(w-\phi) = -\gamma + v''(w-\phi)\big(r(w-\phi) + h + \mu\hat{f}\hat{\phi}\big) + \mu v'(w - \phi - \hat{\phi}).$$

Suppose $v''(w)w = v''(w-\phi)(w-\phi)$. Then

$$\begin{aligned} (\gamma + \mu)v'(w) &= -\gamma + \frac{v''(w-\phi)(w-\phi)}{fw}\big(rw + h + \mu f\phi\big) + \mu v'(w-\phi) \\ (\gamma + \mu)v'(w-\phi) &= -\gamma + v''(w-\phi)\big(r(w-\phi) + h + \mu\hat{f}\hat{\phi}\big) + \mu v'(w - \phi - \hat{\phi}). \end{aligned} \tag{31}$$

Differentiating (31) again with respect to w obtain

$$\begin{aligned} \gamma v''(w) = v'''(w)\Big(rw + h + \mu\big(\delta + (f(w)-1)w\big)\Big) + v''(w)\Big(r + \mu\big(f(w) - 2 + f'(w)w\big)\Big) \\ + \mu\frac{f(w) - (w-\delta)f'(w)}{f(w)^2}v''\left(\frac{w-\delta}{f(w)}\right). \end{aligned} \tag{32}$$

Substitute (30) into (32) and rearrange terms to obtain

$$
\begin{aligned}
0 = {} & v'''(w)\Big(rw + h + \mu\big(\delta + (f(w)-1)w\big)\Big) \\
& + v''(w)\left(r - \gamma + \mu\big(f(w) - 2 + f'(w)w\big) + \mu\frac{wf(w)}{w-\delta}\right).
\end{aligned}
\tag{33}
$$

Suppose there exists a $\tilde{w}$ such that

$$
\tilde{w} = \inf\big\{w > \delta : \quad v''(\tilde{w}) + \tilde{w}\cdot v'''(\tilde{w}) \le 0\big\}.
$$

It follows from (33) that $v'''(w)$ is continuous. Thus $v'''(\tilde{w}) \le -\frac{v''(\tilde{w})}{\tilde{w}}$. Then at $w = \tilde{w}$ equation (33) becomes

$$
0 \le -\frac{v''(w)}{w}\Big(rw + h + \mu\big(\delta + (f(w)-1)w\big)\Big) + v''(w)\left(r - \gamma + \mu\left(f(w) - 2 + f'(w)w + \frac{wf(w)}{w-\delta}\right)\right).
$$

Dividing both sides by $v''(w) < 0$ obtain

$$
0 \ge -\frac{h + \mu\delta}{w} - \gamma + \mu(f'(w)w - 1) + \mu\frac{wf(w)}{w-\delta}.
$$

Note that $v''(w-\delta)\cdot(w-\delta) \le v''(w)\cdot w \le 0$ for $w < \tilde{w}$ implying that $\frac{v''(w-\delta)}{v''(w)} \ge \frac{w}{w-\delta}$. Hence

$$
0 \ge -\frac{h}{\mu w} - \frac{\gamma}{\mu} - \frac{\delta}{w} + f'(w)w - 1 + \frac{wf(w)}{w-\delta}.
\tag{34}
$$

Differentiate the binding optimal incentive compatibility condition (13) with respect to w to obtain

$$
f'(w) = \frac{1 - f(w) + f(w)\phi'(w)}{\frac{w-\delta}{f(w)}} \ge \frac{1 - f(w)}{\frac{w-\delta}{f(w)}} = \frac{f(w)(1-f(w))}{w-\delta}
$$

which results in a lower bound for $f'(w)$ given by $\frac{(1-f(w))f(w)}{w-\delta}$. Substitute this lower bound for $f'(w)$ into (34) to obtain a sufficient condition for the contradiction with the existence of $\tilde{w}$:

$$
0 \le -\frac{h}{\mu}(w-\delta) - \frac{\gamma}{\mu}\cdot w(w-\delta) + f(w)\big(2 - f(w)\big)\cdot w^2 - w^2 + \delta^2
\tag{35}
$$

The right hand side of (35) is increasing in $f(w)$. This means that if it is satisfied for some monitoring intensity that is lower than $f(w)$, then it is satisfied $f(w)$. Note that $\delta = f(w)\phi(w) + (1-f(w))w$ implying that $f(w) \ge \frac{w-\delta}{w}$. Substituting $f(w) = \frac{w-\delta}{w}$ into (35) obtain

$$
0 \le \frac{-\gamma w - h}{\mu} + \frac{(w-\delta)(w+\delta)^2}{\delta^2}.
\tag{36}
$$

The derivative of the right hand side of (36) respect to w is equal to

$$\frac{d}{dw}\left[-\frac{\gamma w+h}{\mu}+\frac{(w-\delta)(w+\delta)^2}{\delta^2}\right]=-\frac{\gamma}{\mu}+\frac{(w-\delta)^2+2(w^2-\delta^2)}{\delta^2}\overset{(i)}{\geq}-\frac{\Delta}{\mu}+\frac{(w-\delta)(3w+2\delta)}{\delta^2}\overset{(ii)}{>}0,$$

where (i) holds for $\Delta \geq \gamma$ and (ii) holds for $w-\delta \geq \frac{h}{\mu}$. Condition (36), thus, holds globally if it holds at $w=\delta+\frac{h}{\mu}$:

$$0 \leq -\frac{\gamma(\delta+\frac{h}{\mu})+h}{\mu}\cdot\frac{h}{\mu}+\left(\frac{h}{\mu}\right)^2\cdot\frac{(2\delta+\frac{h}{\mu})^2}{\delta^2} \qquad \Leftrightarrow \qquad 0 \leq \Delta-\gamma+\Delta\cdot\frac{2\delta+\frac{h}{\mu}}{\delta}.$$

which is satisfied whenever $\Delta \geq \gamma$. □

Lemma 8. *Suppose $\lambda = 0$ and $r+\gamma \geq \mu \cdot \bar{f}$. Then, the optimal monitoring intensity is given by* (17).

Proof. The first order condition (16) is given by

$$1-l+v'(w)w-v(w)+v\left(\frac{w-\delta}{f}\right)-\frac{w-\delta}{f}\cdot v'\left(\frac{w-\delta}{f}\right)\geq 0.$$

We need to check that for $r \geq \mu \cdot \bar{f}-\gamma$ monitoring intensity $f^I(w)$ given by (17) satisfies the corner optimality condition

$$1-l+v'(w)w-v(w)+v\left(\frac{w-\delta}{f^I(w)}\right)-\frac{w-\delta}{f^I(w)}\cdot v'\left(\frac{w-\delta}{f^I(w)}\right)\geq 0.$$

Consider the two cases:

(i) Case $w \leq \left(1+\bar{f}\right)\cdot\delta$. Thus $f(w)=\frac{w-\delta}{\delta}$. The maximum monitoring is optimal if and only if

$$1-l+wv'(w)+v(\delta)-v(w)-\delta v'(\delta)\geq 0.$$

(ii) Case $w > \left(1+\bar{f}\right)\cdot\delta$. The maximum monitoring intensity $f(w)=\bar{f}$ is optimal if

$$1-l+v'(w)w-v(w)+v\left(\frac{w-\delta}{\bar{f}}\right)-\frac{w-\delta}{\bar{f}}\cdot v'\left(\frac{w-\delta}{\bar{f}}\right)\geq 0. \qquad (37)$$

The derivative of the above expression with respect to $\bar{f}$ is $\frac{(w-\delta)^2}{\bar{f}^3}\cdot v''\left(\frac{w-\delta}{\bar{f}}\right)\leq 0$. If (37) holds at $\tilde{f}=\frac{w-\delta}{\delta}\geq \bar{f}$, then the maximum monitoring intensity $f(w)=\bar{f}$ is sufficient. At $\tilde{f}$ inequality (37) becomes

$$1-l+wv'(w)+v(\delta)-v(w)-\delta v'(\delta)\geq 0.$$

providing the same sufficient condition to verify for $w \leq \left(1+\bar{f}\right)\cdot\delta$ and $w \geq \left(1+\bar{f}\right)\cdot\delta$.

Define the auxiliary function $g(w)\overset{def}{=}w\cdot v'(w)-v(w)$. It is sufficient to show that

$$1-l+wv'(w)-v(w)+v(\delta)-\delta\cdot v'(\delta)\geq 0 \quad \Leftrightarrow \quad 1-l+g(w)-g(\delta)\geq 0.$$

Function $g(\cdot)$ is decreasing since $g'(w) = wv''(w) \leq 0$. Thus, if it holds at $w = \bar{w} = \frac{\delta}{1-\bar{f}}$, it holds for all $w \in [\delta, \bar{w}]$. We have

$$g(\bar{w}) = \bar{w} \cdot v'(\bar{w}) - v(\bar{w}) = -\bar{w} - \frac{\alpha - \mu(1 - \bar{f} + l\bar{f}) - h}{r + \gamma} + \bar{w} = -\frac{\alpha - \mu(1 - \bar{f} + l\bar{f}) - h}{r + \gamma}.$$

To evaluate $g(\delta) = \delta \cdot v'(\delta) - v(\delta)$, note that at $w = \delta$ the principal's value function satisfies the differential equation

$$-(r + \gamma) \cdot g(\delta) = \alpha - \mu - \gamma\delta + \left(1 - \frac{\gamma}{\Delta}\right) \cdot h \cdot v'(\delta).$$

Then the sufficient condition for maximum monitoring to be optimal is

$$\begin{aligned} 1 - l + g(\bar{w}) - g(\delta) &= 1 - l - \frac{\alpha - \mu(1 - \bar{f} + l\bar{f}) - h}{r + \gamma} + \frac{\alpha - \mu - \gamma\delta + \left(1 - \frac{\gamma}{\Delta}\right) h \cdot v'(\delta)}{r + \gamma} \\ &= \frac{(r + \gamma - \mu\bar{f})(1 - l) + h \cdot \left(1 - \frac{\gamma}{\Delta}\right)(v'(\delta) + 1)}{r + \gamma} \end{aligned}$$

Since $v'(\delta) \geq -1$, it implies a sufficient condition for maximum monitoring to be given by

$$r + \gamma \geq \mu \cdot \bar{f} \qquad \Rightarrow \qquad 1 - l + g(\bar{w}) - g(\delta) \geq 0. \tag{38}$$

□

The proof of Proposition 3 is provided for $\lambda = 0$ and extends to λ sufficiently small by continuity.

References

M. Aoyagi. Information feedback in a dynamic tournament. *Games and Economic Behavior*, 70(2): 242–260, Nov. 2010.

G. Baker, R. Gibbons, and K. J. Murphy. Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, 109(4):1125–1156, 1994.

G. P. Baker. Incentive contracts and performance measurement. *Journal of Political Economy*, pages 598–614, 1992.

I. Ball. Dynamic information provision: Rewarding the past and guiding the future. 2019.

H. G. Barkema. Do top managers work harder when they are monitored? *Kyklos*, 48(1):19–42, 1995.

D. Bergemann and U. Hege. The financing of innovation: Learning and stopping. *RAND Journal of Economics*, pages 719–752, 2005.

B. Biais, T. Mariotti, J.-C. Rochet, and S. Villeneuve. Large risks, limited liability, and dynamic moral hazard. *Econometrica*, 78(1):73–118, 2010.

P. Bond and A. Gomes. Multitask principal–agent problems: Optimal contracts, fragility, and effort misallocation. *Journal of Economic Theory*, 144(1):175–211, 2009.

D. Cetemen, F. Z. Feng, and C. Urgun. Contracting with non-exponential discounting: Moral hazard and dynamic inconsistency. *Available at SSRN 3442367*, 2019.

M. Chen, P. Sun, and Y. Xiao. Optimal monitoring schedule in dynamic contracts. 2017.

P. DeMarzo and Y. Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6):2681–2724, 2006.

P. DeMarzo and Y. Sannikov. Learning, termination, and payout policy in dynamic incentive contracts. 2011.

F. Ederer. Feedback and motivation in dynamic tournaments. *Journal of Economics & Management Strategy*, 19(3):733–769, 2010.

A. Edmans and X. Gabaix. Tractability in incentive contracting. *The Review of Financial Studies*, 24 (9):2865–2894, 2011.

B. S. Frey. Institutions and morale: the crowding-out. *Economics, values, and organization*, page 437, 1999.

W. Fuchs. Contracting with repeated moral hazard and private evaluations. *American Economic Review*, 2007.

D. Fudenberg and L. Rayo. Training and effort dynamics in apprenticeship. 2019.

G. Georgiadis and B. Szentes. Optimal monitoring design. *Working Paper*, 2019.

M. Goltsman and A. Mukherjee. Interim performance feedback in multistage tournaments: The optimality of partial disclosure. *Journal of Labor Economics*, 29(2):229–265, 2011.

M. Harris and A. Raviv. Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. *American Economic Review*, 68(1):20–30, 1978.

Z. He. Dynamic compensation contracts with private savings. *Review of Financial Studies*, 25(5): 1494–1549, Apr. 2012.

F. Hoffmann, R. Inderst, and M. Opp. Only time will tell: A theory of deferred compensation and its regulation. *Working paper*, 2019.

B. Holmstrom. Moral hazard and observability. *The Bell Journal of Economics*, pages 74–91, 1979.

B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, pages 303–328, 1987.

J. Horner and N. S. Lambert. Motivational ratings. 2019.

A. S. Hornstein and M. Zhao. Corporate capital budgeting decisions and information sharing. *Journal of Economics and Management Strategy*, 20(4):1135–1170, 2011.

A. Kaya. Failure to fail fast: slow learning as agency cost. *Working Paper*, 2020.

C. Laux. Limited-liability and incentive contracting with multiple projects. *RAND Journal of Economics*, pages 514–526, 2001.

E. P. Lazear. *Personnel economics for managers*. Wiley New York, 1998.

A. Lizzeri and M. Siniscalchi. Parental guidance and supervised learning. *The Quarterly Journal of Economics*, 123(3):1161–1195, 2008.

A. Lizzeri, M. A. Meyer, and N. Persico. The incentive effects of interim performance evaluations. 2002.

W. B. MacLeod. Optimal contracting with subjective evaluation. *American Economic Review*, 93(1): 216–240, 2003.

G. Manso. Motivating innovation. *The Journal of Finance*, 66(5):1823–1860, 2011.

I. Marinovic and F. Varas. Ceo horizon, optimal pay duration, and the escalation of short-termism. *The Journal of Finance*, 2018.

K. J. Murphy. Executive compensation. *Handbook of labor economics*, 3:2485–2563, 1999.

K. R. Murphy and J. Cleveland. *Understanding performance appraisal: Social, organizational, and goal-based perspectives.* Sage, 1995.

D. Orlov, A. Skrzypacz, and P. Zryumov. Persuading the principal to wait. *The Journal of Political Economy*, 2020.

T. Piskorski and M. M. Westerfield. Optimal dynamic contracts with moral hazard and costly monitoring. *Journal of Economic Theory*, 166:242–281, 2016.

K. Ray. The retention effect of withholding performance information. *Accounting Review*, 82(2): 389–425, 2007.

Y. Sannikov. Moral hazard and long-run incentives. *Working paper, Princeton University*, 2014.

D. Scharfstein and J. Stein. The dark side of internal capital markets: Divisional rent-seeking and inefficient investment. *The Journal of Finance*, 55(6):2537–2564, Nov. 2000.

A. Smolin. Dynamic evaluation design. *American Economic Journal: Microeconomics*, 13(4):300–331, 2021.

S. E. Spear and S. Srivastava. On repeated moral hazard with discounting. *The Review of Economic Studies*, 54(4):599–617, 1987.

B. H. Strulovici. Renegotiation-proof contracts with moral hazard and persistent private information. *Available at SSRN 1755081*, 2011.

F. Varas, I. Marinovic, and A. Skrzypacz. Random inspections and periodic reviews: Optimal dynamic monitoring. *The Review of Economic Studies*, 87(6):2893–2937, 2020.

N. Williams. Persistent Private Information. *Econometrica*, 79(4):1233–1275, 2011.

J. Zhu. Optimal Contracts with Shirking. *Review of Economic Studies*, 80(2):812–839, Apr. 2013.