Dynamic Monitoring Design in Continuous Time

Xuhu Wan*

March 2026

Abstract

We study optimal monitoring and incentive design in a continuous-time principal-agent model with dynamic moral hazard. Using an Exponentiated Continuation Value representation, we show the optimal contract is binary: the agent receives a base wage or a bonus upon crossing a performance threshold. Monitoring follows a two-threshold rule, ceasing once performance exits a target band. A decoupling theorem establishes that monitoring boundaries and compensation are independent of output volatility σ; only the agent's effort intensity scales with σ^2. Effort is stochastic and state-dependent, generating an inverse effort-value relationship. Expected pay-performance sensitivity is non-monotone in risk, exhibiting an inverted-U pattern. These results reconcile mixed empirical findings on the PPS–performance–risk nexus.

JEL classification: G34; C61; D86.

Keywords: Continuous-time principal-agent models, dynamic monitoring, binary compensation, pay-performance sensitivity.

*Department of Information and Systems Management, School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. Phone: (852) 2358-7731. E-mail: xuhu.wan@gmail.com.

1 Introduction

A central challenge in financial contracting is how to design incentive schemes that induce agents to exert effort when effort is unobservable. Firms and financial institutions must reward outcomes that are informative about effort, but in practice this is complicated. First, performance is often aggregate in nature: a firm's profit or a fund's returns reflect the joint contribution of many individuals, making it difficult to attribute outcomes to a particular worker. Second, many dimensions of performance, such as client relationships or risk management, are inherently qualitative and difficult to measure directly. These challenges highlight the importance of designing monitoring and compensation systems that rely on informative signals about effort rather than direct measures of output.

In this paper, we provide a theoretical investigation of optimal monitoring structures when information (performance measure) about effort is not freely available. We model a principal who can acquire signals about an agent's effort at a convex cost, represented by a diffusion process whose drift equals the agent's effort. A contract specifies a stopping time for this monitoring process and a wage scheme contingent on the observed information. The agent is risk-averse, takes hidden effort in continuous time, and is subject to limited liability. The principal's problem is to implement an optimal effort with an optimal dynamic monitoring scheme.

Our main result shows that the optimal contract takes the form of a binary wage scheme: the agent receives either a base wage if the performance measure is low or a fixed bonus if performance is deemed sufficiently high. This provides a new rationale—in the dynamic moral hazard setting—for the prevalence of single-bonus contracts, complementing the static optimality result of Georgiadis and Szentes (2020) and the rich-data analysis of Frick et al. (2023). Such contracts are widely observed in finance and consulting. For example, compensation in investment banking and asset management often includes a base salary plus a single performance bonus, or termination if performance falls below expectations. Our analysis shows that such contracts arise endogenously when the monitoring structure itself is optimally designed.

Our main theoretical contribution is an explicit characterization of this threshold contract. Under general preference and technical regularity conditions, the optimal scheme features two state-contingent stopping thresholds $x_1(\lambda^*) < x_2(\lambda^*)$, where λ^* is the optimal monitoring cost multiplier, and a binary wage structure: the agent receives a flat base if the lower threshold is hit, or a bonus of

$$v'^{-1}(h(x_2(\lambda^*)) + \lambda^*),$$

where v' is the marginal utility of consumption and h is the principal's net payoff function (firm value minus monitoring cost), if performance reaches the upper threshold. Monitoring continues until one threshold is crossed.

Our second main result is a *decoupling theorem* (Corollary 6.3): the optimal monitoring boundaries $x_1(\lambda^*)$ and $x_2(\lambda^*)$, together with the terminal compensation schedule $\{C_L, C_H\}$ (base wage and bonus), are completely independent of the output volatility σ. Volatility affects only the agent's equilibrium effort intensity, which scales as $a_t^* = \sigma^2\beta(\lambda^*)/K_t$, where K_t is the exponentiated continuation value (a sufficient statistic for the agent's incentive state) and $\beta(\lambda^*)$ is an incentive-intensity parameter determined by the optimal contract. Thus, the principal's information design problem—choosing *when* to stop monitoring—decouples entirely from the incentive provision problem—determining *how hard* the agent works. Economically, a risk-neutral principal facing a more volatile output process does not change the monitoring boundary or wage structure; she changes only the effort she expects to observe. This decoupling has a sharp practical implication: the principal need not estimate σ to design the optimal contract. It also provides a structural explanation for the empirical regularity that monitoring arrangements (e.g., board review frequency, audit triggers) appear far less variable across firms than do compensation levels, even though firms differ substantially in output volatility.

Moreover, our model delivers a sharp prediction on the time-varying nature of pay performance sensitivity (PPS): PPS is highest when performance is near the stopping thresholds either $x_1(\lambda)$ or $x_2(\lambda)$ and declines as performance moves away from these critical regions. Around the thresholds, small increments in the performance process X_t materially change the probability of triggering the bonus or baseline salary. This amplifies the marginal impact of effort on outcomes, producing high PPS. Once X_t drifts away from the thresholds, additional effort has a much smaller effect on the stopping probabilities. As a result, PPS fades despite ongoing monitoring.

This link between state-dependent incentives and thresholds falls straight out of our model. It offers a natural explanation for why empirical studies observe strong PPS effects around performance cliffs or review points, but weaker relationships outside those zones. Our model thus reconciles mixed empirical findings: PPS is not constant over time or performance levels—it peaks near contract-defined thresholds and then tapers off.

By endogenously designing both monitoring and compensation, this paper provides a transparent micro-foundation for single-bonus contracts widely used in practice. It sharpens the link between dynamic monitoring, threshold-based compensation, and incentive dynamics, yielding rich theoretical predictions and a clear roadmap for empirical investigation.

1.1 Related Literature

This paper contributes to the growing literature on optimal dynamic monitoring under moral hazard in continuous time. Early foundational works Mirrlees (1976) and Holmström (1979) established incentive contracts when effort is unobservable. Later extensions introduced multidimensional effort (Holmström and Milgrom, 1991).

Piskorski and Westerfield (2016) study dynamic contracts with costly monitoring where the principal chooses monitoring intensity as a continuous control, deriving non-monotone monitoring patterns. We instead model monitoring as an endogenous stopping time, which yields a closed-form binary contract and the decoupling theorem.

Recent contributions include Dai et al. (2024), who develop a flexible monitoring model allowing carrot-and-stick evidence gathering, showing optimal switches in monitoring modes based on the agent's continuation value. Their carrot-and-stick structure resonates directly with our threshold contract: in our setting, the bonus at the upper threshold x_2 serves as the carrot, while reversion to the base wage at the lower threshold x_1 serves as the stick. Our framework extends their analysis by providing explicit threshold triggers and closed-form compensation tied to performance diffusion thresholds.

Wong (2023) studies dynamic monitoring design with flexible, endogenous Poisson signal arrivals and termination/tenure as the principal's instruments. Our formulation differs in modeling performance as a Brownian diffusion and using two-sided diffusion thresholds, rather than Poisson arrival rates, to define the endogenous monitoring window.

Monitoring with rich data has been analyzed by Frick et al. (2023), who show that simple binary wage schemes can achieve the optimal convergence rate to first-best when the principal observes a rich signal of the agent's one-dimensional action. Their finding—that binary wages are essentially as informative as more complex schemes in the rich-data limit—is qualitatively consistent with our exact binary optimum, although the mechanisms (rich-data asymptotics versus dynamic threshold design) differ.

Our paper is most closely related to Georgiadis and Szentes (2020), who characterize optimal monitoring design in a static setting where the principal acquires signals about the agent's effort at constant marginal cost. They show the optimal contract is binary and implements a two-threshold monitoring policy. We extend their framework along three substantive dimensions. First, we move from a static one-shot signal to a continuous-time monitoring process, so the principal's choice becomes an endogenous stopping time on a diffusion. Second, we endogenize the agent's effort: rather than implementing a fixed effort target, the agent's optimal effort is stochastic and state-dependent, yielding empirically meaningful predictions about pay-performance sensitivity. Third, our continuous-time analysis delivers a decoupling theorem (Corollary 6.3) that has no counterpart in the static framework: the monitoring boundaries and compensation levels are independent of output volatility σ, which affects only the agent's equilibrium effort intensity. Despite these substantive differences, the binary structure of the optimal contract survives the dynamization—an indication that this form is intrinsic to the monitoring-design problem rather than an artifact of the static setup.

The continuous-time contracting framework builds on the foundational work of DeMarzo and Sannikov (2006), who derive optimal long-term financial contracts under moral hazard with limited liability, characterizing endogenous default and credit-line structures. Our binary contract—with a limited-liability floor at x_1 and a lump-sum bonus at x_2—shares the qualitative structure of their credit line and golden parachute. Biais et al. (2010) extend dynamic moral hazard with limited liability to a Poisson setting, showing that lump-sum payments at contract termination (a "golden parachute" structure) are optimal; our bonus C_H at the upper threshold is the continuous-time analog. Garrett and Pavan (2015) study dynamic managerial compensation with persistent private types via a variational approach, showing that optimal pay-performance sensitivity varies over time due to evolving information rents on persistent types; our time-varying PPS arises from a complementary mechanism: threshold proximity rather than information rents. He et al. (2017) analyze dynamic contracting with learning about managerial ability, deriving optimal contracts that balance incentive provision with information acquisition; their focus on the interaction between learning and contracting complements our analysis, where the principal's information acquisition is governed by the endogenous stopping rule. The continuous-time contracting methodology also builds on Cvitanić et al. (2009), who characterize optimal lump-sum contracts with hidden action in a continuous-time model; our ECV representation extends their framework to incorporate endogenous monitoring.

Additionally, our findings relate to the literature on information design and Bayesian persuasion (Kamenica and Gentzkow, 2011), where principals design information structures to influence agents' actions. By characterizing the optimal monitoring strategy as a two-threshold policy, we contribute to this literature by highlighting how endogenous information acquisition can lead to simple, yet effective, incentive schemes.

The monitoring cost function $g(x)$ captures the idea that continuous oversight and information acquisition are resource-intensive. This formulation follows a growing literature emphasizing the role of costly and dynamic monitoring in contract design. For example, Piskorski and Westerfield (2016) model contracts with moral hazard and costly monitoring technologies, treating monitoring intensity as a continuous control variable. Orlov (2022) shows that, in dynamic contracting, more frequent monitoring can paradoxically weaken incentives by revealing bad news that depresses the agent's continuation value. Varas et al. (2020) analyze random inspections and periodic reviews as optimal monitoring schemes. More recently, Dai et al. (2024) and Li and Yang (2020) develop models of flexible and endogenous monitoring, where the technology and timing of monitoring itself are chosen optimally. These contributions motivate our inclusion of $g(x)$ as a convex cost that increases with performance deviations, reflecting the rising difficulty of sustaining effective oversight.

In sum, our paper bridges and builds upon the core insights from continuous-time dynamic contracting, flexible monitoring, and information design. Our key breakthroughs are (i) delivering closed-form dynamic contracts with explicit thresholds and bonus payoffs, (ii) characterizing effort as a stochastic, state-dependent process under limited liability, and (iii) connecting theory to empirical PPS patterns through time-varying sensitivity around monitoring thresholds.

The rest of the paper is organized as follows. Section 2 specifies the model assumptions and the maximization problems of the principal and agent. Section 3 reformulates the optimal contracting and monitoring problem into a more tractable form via the ECV representation and convex conjugate reduction. Section 4 solves for the optimal binary contract. Section 5 analyzes pay-performance sensitivity and its non-monotone relationship with output volatility. Section 6 characterizes the valuable set $\mathcal{E}$, establishes the decoupling theorem, and derives comparative statics. Section 7 calibrates the model to CEO compensation data. Section 8 discusses robustness to alternative specifications. Section 9 concludes. The Appendix contains all proofs.

2 Model

Time is continuous. We consider a firm in which the manager (the agent, "he") exerts an unobservable effort process $\{a_t\}_{t\geq 0}$, while the firm's representative shareholder (the principal, "she") designs a contract and a monitoring strategy to incentivize the agent.

The principal observes a publicly available performance signal X_t. Although X_t does not directly represent the firm's asset value S_t, it serves as a proxy and forms the basis for monitoring and contracting.

In the absence of managerial effort ($a_t = 0$), the performance process follows a pure diffusion:

$$dX_t = \sigma dB_t^0, \tag{1}$$

where $\sigma > 0$ is a constant volatility parameter, and B_t^0 is a standard Brownian motion under the reference probability measure $\mathbb{P}^0$, defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0})$, where $\mathcal{F}_t$ is the natural filtration generated by $\{X_s : 0 \leq s \leq t\}$. Let $\mathbb{E}^0[\cdot|\mathcal{F}_t]$ denote conditional expectations under $\mathbb{P}^0$.

When the agent exerts costly effort a_t, the reference measure $\mathbb{P}^0$ is distorted into an equivalent probability measure $\mathbb{P}^a$, with Radon-Nikodym derivative:

$$\frac{d\mathbb{P}^a}{d\mathbb{P}^0} = M_\tau^a, \tag{2}$$

where τ is the random duration of monitoring, endogenously determined by the principal as part of the optimal monitoring strategy. At time τ, monitoring ceases, and compensation is paid based on the information accumulated up to τ. The likelihood ratio M_τ^a is an $\mathcal{F}_\tau$-adapted $\mathbb{P}^0$-martingale given by:

$$M^a_\tau = \exp\left(-\tfrac{1}{2}\int_0^\tau \left(\tfrac{a_t}{\sigma}\right)^2 dt + \int_0^\tau \left(\tfrac{a_t}{\sigma}\right) dB^0_t\right). \tag{3}$$

By Girsanov's theorem, under $\mathbb{P}^a$ the performance process evolves as

$$dX_t = a_t dt + \sigma dB^a_t, \tag{4}$$

where $B^a_t = B^0_t - \int_0^t \left(\frac{a_s}{\sigma}\right) ds$ is a Brownian motion under $\mathbb{P}^a$. We let $\mathbb{E}^a[\cdot|\mathcal{F}_t]$ denote conditional expectations under $\mathbb{P}^a$.

Thus the manager's effort shifts the drift of X_t, with higher effort increasing the expected value of X_t over time. This contrasts with models where performance evolves under a fixed constant effort (e.g. Georgiadis and Szentes (2020)).

At time 0, the principal offers a compensation scheme C_τ payable at the stopping time τ, contingent upon the entire path of the performance measure up to time τ, represented by $\mathcal{F}_\tau$. The agent is subject to limited liability, imposing a lower bound on compensation: $C_\tau \geq \underline{c} > 0$.

The agent's expected utility at time 0 is:

$$\mathbb{E}^a\left[u\left(C_\tau\right) - \delta \log\left(\frac{d\mathbb{P}^a}{d\mathbb{P}^0}\right)\right], \tag{5}$$

where $u(\cdot)$ is increasing, strictly concave, and satisfies the Inada condition. $\delta > 0$ governs the disutility of effort through the relative entropy. The Kullback-Leibler (KL) divergence represents cumulative effort cost:

$$\mathbb{E}^a\left[\log\left(\frac{d\mathbb{P}^a}{d\mathbb{P}^0}\right)\right] = \mathbb{E}^a\left[\int_0^\tau \frac{1}{2\sigma^2} a_t^2 dt\right]. \tag{6}$$

This cost structure connects to the rational inattention literature (Sims, 2003; Caplin and Dean, 2015; Zhong, 2022), where KL divergence is the canonical information cost. While we adopt KL divergence for the effort cost in the baseline model, our framework can accommodate more general convex cost of measure distortion[1].

This approach connects to Georgiadis et al. (2024), who study static flexible moral hazard problems, in which the agent directly selects distributions over outcomes, subject to a convex cost

[1] Appendix B outlines how to characterize optimal contracting and monitoring under a general convex cost of measure change.

of selected distribution. In contrast, our dynamic setting endogenizes the agent's cost as a functional of the induced measure change over time.

The principal learns about the agent's effort only indirectly by observing the performance measure X_t. Continuous monitoring incurs cost $\mathbb{E}^a[g(X_\tau)]$, where $g(\cdot)$ is convex, reflecting growing difficulty in monitoring as performance deviates from baseline level X_0. The monitoring cost $g(X_\tau)$ is a terminal cost, incurred at the stopping time τ, rather than a flow cost accumulated over the monitoring period. This can be interpreted as a lump-sum audit cost paid at contract termination—reflecting the expense of processing and certifying the accumulated signal record $\{X_s : 0 \leq s \leq \tau\}$—or as the cost of a final performance evaluation. While a flow-cost formulation $\int_0^\tau g(X_t)\,dt$ would be more general, the terminal specification preserves the Skorokhod embedding reduction that delivers tractability.[2]

We do not assume that the firm value or final outcome S_τ is identical to the performance measure because the final outcome is an aggregation of all agents' effort and external factors. Firm value at τ is:

$$S_\tau = f(X_\tau) + \epsilon, \tag{7}$$

where $f(\cdot)$ is increasing and concave, and ϵ is a zero-mean shock orthogonal to $\mathcal{F}_\tau$ representing external factors. Thus higher managerial effort raises the expected firm value via its effect on X_τ.

The principal's objective is to maximize

$$\mathbb{E}^a\left[h(X_\tau) - C_\tau\right]. \tag{P1}$$

where the net payoff function $h(x) = f(x) - g(x)$, representing firm value net of monitoring cost.

For example, $h(x) = x - \alpha x^2$ with $\alpha > 0$ captures quadratic monitoring costs, which rise quadratically as performance deviates from the baseline. When X_τ is substantially positive, larger scale operations become more complex and require intensified oversight. Conversely, when X_τ is negative, poor performance may trigger risk-shifting behavior or financial distress, also raising monitoring difficulty. This quadratic specification captures the principal's increasing marginal cost of supervision as firm performance moves further from its initial state. Under constant effort a^*,

[2] Extending the model to flow monitoring costs is a natural direction for future work; it would require solving a free-boundary problem for the principal's value function rather than the static distribution-design problem studied here.

the quadratic monitoring cost becomes

$$\mathbb{E}^{a^*}[X_\tau^2] = X_0^2 + (2a^* X_0 + (a^*)^2 + \sigma^2)\mathbb{E}^{a^*}[\tau].$$

Thus, the expected monitoring cost under constant effort increases with both the effort level and the expected monitoring duration. This performance- and duration-dependent specification is motivated by the cost-per-signal formulation in Georgiadis and Szentes (2020), though the two cost structures are not identical: theirs charges a constant marginal cost per independent signal acquired, whereas ours scales with the realized performance path and monitoring duration.

We impose the following regularity condition on the principal's net payoff function:

Assumption 2.1. *There exists a unique optimal performance level $x^* > 0$ such that*

$$x^* \in \arg\max_x \; h(x).$$

We define the state space as $\mathcal{X} := (-\infty, x^*]$. The baseline performance satisfies $X_0 \in \mathcal{X}$. The net payoff function $h(x)$ is assumed to be smooth, strictly concave, and strictly increasing on $\mathcal{X}$.

The principal's problem is to design a contract (τ, C_τ) that determines both the stopping rule τ (which governs the length of costly monitoring and information acquisition) and the terminal compensation C_τ. Formally, the principal aims to maximize:

$$\max_{\{a_t\}_{0\le t\le\tau},\, \tau,\, C_\tau \ge \underline{c}} \; \mathbb{E}^a \left[h(X_\tau) - C_\tau\right], \tag{P1*}$$

subject to:

(IC) **Incentive Compatibility:**

$$\{a_t\} \in \arg\max_{\{\hat{a}_t\}} \; \mathbb{E}^{\hat{a}} \left[u(C_\tau) - \delta \log\left(\frac{d\mathbb{P}^{\hat{a}}}{d\mathbb{P}^0}\right)\right]. \tag{IC}$$

(IR) **Individual Rationality:**

$$\mathbb{E}^a \left[u(C_\tau) - \delta \log\left(\frac{d\mathbb{P}^a}{d\mathbb{P}^0}\right)\right] \ge R. \tag{IR}$$

The stopping time τ determines both the duration of costly information acquisition and the contract horizon.

Definition 2.2. *A contract (τ, C_τ), together with an associated incentive-compatible effort process $\{a_t\}_{t\geq 0}$, is **valuable** if the agent is not terminated immediately, that is,*

$$\mathbb{P}^0(\tau > 0) = 1.$$

Definition 2.3. *The agent's reservation utility R is **admissible** if there exists a contract (τ, C_τ), with associated incentive-compatible effort $\{a_t\}_{t\geq 0}$, such that the participation constraint binds and*

$$\mathbb{E}^0[\tau] < \infty.$$

Let $\mathcal{R}$ denote the set of all admissible values of R.

The condition $\mathbb{E}^0[\tau] < \infty$ guarantees that expected contract duration remains finite even under zero effort. We characterize the admissible set $\mathcal{R}$ explicitly in the following sections.

2.1 Agent's Problem

Following Sannikov (2008), the agent's continuation utility can be represented as a stochastic process adapted to the observable filtration $\mathcal{F}_t$. In this formulation, the contract is expressed in terms of the agent's conditional expected utility at each time t.

We define the **Exponentiated Continuation Value (ECV)** process as:

$$K_t = \exp\left\{\frac{1}{\delta}\mathbb{E}^a\left[u(C_\tau) - \int_t^\tau \frac{\delta}{2\sigma^2}a_s^2 ds \,\Big|\, \mathcal{F}_t\right]\right\}. \tag{8}$$

K_t is the exponential transformation of the agent's continuation utility at time t, conditional on the filtration $\mathcal{F}_t$. In our model, the expected cost of effort is inversely proportional to σ^2. That is, the noisier the performance signal, the lower the marginal cost of effort for the agent.

Lemma 2.4. *A contract (τ, C_τ) implements the principal's desirable effort $\{a_t\}_{t\geq 0}$ if and only if* *the **ECV** process satisfies:*

$$dK_t = K_t \left(\frac{a_t}{\sigma}\right) dB_t^0,\ \ K_\tau = \exp\left(\frac{1}{\delta} u(C_\tau)\right) \tag{9}$$

Proof. Define the agent's conditional expected utility at time t as:

$$Y_t = \mathbb{E}^a \left[u\left(C_\tau\right) - \int_t^\tau \frac{\delta}{2\sigma^2} a_s^2 ds \,\Big|\, \mathcal{F}_t \right].$$

The martingale representation theorem (MRT) applies to $K_t = \mathbb{E}^0[\exp(u(C_\tau)/\delta) \mid \mathcal{F}_t]$ under $\mathbb{P}^0$, since K_t is a $\mathbb{P}^0$-martingale by construction. Equivalently, working with $Y_t = \delta \ln K_t$, there exists an adapted process $\{\beta_t\}_{0\le t\le\tau}$ such that

$$dY_t = \left(\frac{\delta}{2\sigma^2} a_t^2 - \frac{\beta_t}{\sigma} a_t\right) dt + \beta_t \, dB_t^0.$$

The agent's effort $\{a_t\}_{t\ge 0}$ is incentive compatible if and only if:

$$a_t = \arg\min_{\hat{a}_t} \left(\frac{\delta}{2\sigma^2} \hat{a}_t^2 - \frac{\beta_t}{\sigma} \hat{a}_t\right), \tag{10}$$

which yields:

$$a_t = \frac{\sigma \beta_t}{\delta}.$$

Substituting into the dynamics of Y_t gives:

$$dY_t = -\frac{\delta}{2}\left(\frac{a_t}{\sigma}\right)^2 dt + \delta \left(\frac{a_t}{\sigma}\right) dB_t^0.$$

Since $K_t = \exp\left(\frac{1}{\delta} Y_t\right)$, we apply Itô's lemma:

$$dK_t = K_t \cdot \frac{1}{\delta} \, dY_t + \frac{1}{2} \, K_t \cdot \frac{1}{\delta^2} \, (dY_t)^2.$$

Substituting $dY_t = -\frac{\delta}{2}\left(\frac{a_t}{\sigma}\right)^2 dt + \delta \left(\frac{a_t}{\sigma}\right) dB_t^0$:

$$dK_t = K_t \left[-\frac{1}{2}\left(\frac{a_t}{\sigma}\right)^2 dt + \frac{a_t}{\sigma} \, dB_t^0 + \frac{1}{2}\left(\frac{a_t}{\sigma}\right)^2 dt \right] = K_t \, \frac{a_t}{\sigma} \, dB_t^0,$$

where the drift terms cancel exactly. This yields (9).

True martingale verification. Under the optimal contract (Section 5), $a_t = \sigma^2\beta(\lambda^*)/K_t$ where $K_t = \alpha(\lambda^*) + \beta(\lambda^*)X_t$ is bounded away from zero on $[x_1(\lambda^*), x_2(\lambda^*)]$ (since $K_{x_1} = \underline{k} > 0$). Hence a_t/σ is bounded, and Novikov's condition

$$\mathbb{E}^0\left[\exp\left(\tfrac{1}{2}\int_0^\tau \tfrac{a_t^2}{\sigma^2}\,dt\right)\right] < \infty$$

holds. Therefore K_t is a true $\mathbb{P}^0$-martingale (not merely a local martingale), and the martingale representation theorem applies. □

The diffusion coefficient of K_t directly determines the agent's effort level a_t. The process K_t is a $\mathbb{P}^0$-martingale and satisfies the conditional expectation representation:

$$K_t = \mathbb{E}^0\left[K_\tau \,|\, \mathcal{F}_t\right] = \mathbb{E}^0\left[\exp\left(\frac{1}{\delta}u(C_\tau)\right)\,\Big|\,\mathcal{F}_t\right]. \tag{11}$$

Equation (11) highlights that the **ECV** is the conditional expectation of exponentiated utility of terminal compensation under reference probability, independent of the agent's chosen effort path. Consequently, incentive-compatible effort is fully embedded in the diffusion dynamics of K_t. Thus, the principal's problem reduces to choosing the terminal compensation C_τ and the stopping time τ, since the entire effort path can be recovered from the martingale dynamics.

Lemma 2.5. *If the desired effort process is incentive compatible, the Radon-Nikodym derivative satisfies:*

$$M_\tau^a = \frac{K_\tau}{K_0} = \frac{\exp\left(\frac{1}{\delta}u(C_\tau)\right)}{\mathbb{E}^0\left[\exp\left(\frac{1}{\delta}u(C_\tau)\right)\right]} = \frac{\exp\left(\frac{1}{\delta}u(C_\tau)\right)}{K_0}. \tag{12}$$

Proof. The result follows directly from the definition $K_t = \exp\left(\frac{1}{\delta}Y_t\right)$ and the martingale property of K_t under $\mathbb{P}^0$. At terminal time τ, we have $K_\tau = \exp\left(\frac{1}{\delta}u(C_\tau)\right)$, and thus by martingale property: $K_0 = \mathbb{E}^0\left[K_\tau\right]$. The processes K_t and M_t^a satisfy the same SDE $dZ_t = Z_t(a_t/\sigma)\,dB_t^0$, with initial conditions K_0 and 1 respectively; by pathwise uniqueness, $M_t^a = K_t/K_0$. The expression (12) follows immediately. □

In equilibrium, the likelihood ratio M_τ^a depends only on the terminal compensation C_τ. This allows the agent's effort to be fully captured by the choice of (τ, C_τ). Define the auxiliary function:

$$v(k) = ku^{-1}(\delta \ln(k)). \tag{13}$$

Lemma 2.6. *The function $v(k)$ is strictly increasing and strictly convex. Moreover, the compensation satisfies*

$$C_\tau = \frac{v(K_\tau)}{K_\tau}. \tag{14}$$

Proof. Differentiating $v(k) = ku^{-1}(\delta \log k)$, we obtain

$$v'(k) = u^{-1}(\delta \log(k)) + \frac{\delta}{u'(u^{-1}(\delta \log(k)))} > 0, \tag{15}$$

because $\delta > 0$, $u'(\cdot) > 0$ and $u(\cdot)$ satisfies the Inada condition. Moreover, we can find both terms in (15) are strictly increasing in k: the first term increases because $u^{-1}(\cdot)$ is strictly increasing; the second term increases because the denominator $u'\left(u^{-1}(\delta \log(k))\right)$ decreases by concavity of $u(\cdot)$. We conclude $v(k)$ is strictly increasing and convex in k. □

The term $v(K_\tau) = K_\tau C_\tau = K_0 M_\tau^a C_\tau$ represents the distortion-adjusted, state-contingent value of compensation. This formulation highlights how the agent's private effort reshapes the state distribution of compensation. We refer to $v(k)$ as the **distorted valuation function** for compensation.

3 Reformulation of the Principal's Problem

Using equations (12) and (14), the principal's problem can be rewritten as:

$$H(k,x) = \max_{\tau, K_\tau > \underline{k}} \mathbb{E}^0 \left[\frac{K_\tau}{K_0} \left(h(X_\tau) - \frac{1}{K_\tau} v(K_\tau) \right) \Big| K_0 = k, X_0 = x \right], \tag{P2}$$

subject to the individual rationality (IR) constraint:

$$\mathbb{E}^0[K_\tau] = K_0 = \exp\left(\frac{1}{\delta} R\right), \quad \forall R \in \mathcal{R}, \tag{16}$$

where $\underline{k} = \exp\left(\frac{1}{\delta} u(\underline{c})\right)$ is determined by the agent's limited liability constraint: $C_\tau \geq \underline{c}$.

Unlike the original formulation (P1), we impose the IR constraint (16) as binding for all admissible R. This is without loss of generality, as is standard in dynamic contracting models: the principal's value function is contingent on the agent's promised utility. If $H(k, x)$ is decreasing in k, then Problem (P1) is equivalent to Problem (P2). Otherwise, the principal's problem (P1) can be recovered from (P2) via:

$$\sup_{\hat{R} \geq R} H\left(\exp\left(\frac{\hat{R}}{\delta}\right), X_0\right). \tag{17}$$

Under the reference measure $\mathbb{P}^0$, the performance process satisfies:

$$X_\tau = X_0 + \sigma B_\tau^0.$$

The following result, adapted from Lemma 2 in Georgiadis and Szentes (2020), characterizes the set of probability distributions over the terminal performance X_τ that can be induced by a stopping time.

Lemma 3.1. *The distribution of X_τ belongs to the set:*

$$\mathcal{G} = \left\{G \in \Delta(\mathbb{R}) : \mathbb{E}_G[X] = X_0, \mathbb{E}_G[X^2] < \infty\right\},$$

where $\Delta(\mathbb{R})$ denotes the set of probability distributions over $\mathbb{R}$.

Proof. ("Only if" direction.) If τ is a stopping time with $\mathbb{E}^0[\tau] < \infty$, the optional stopping theorem gives $\mathbb{E}^0[X_\tau] = X_0$. Moreover, $\mathbb{E}^0[X_\tau^2] = X_0^2 + \sigma^2 \mathbb{E}^0[\tau] < \infty$. Hence the distribution G of X_τ satisfies $\mathbb{E}_G[X] = X_0$ and $\mathbb{E}_G[X^2] < \infty$, so $G \in \mathcal{G}$.
("If" direction.) Conversely, for any $G \in \mathcal{G}$, the Skorokhod embedding theorem (see, e.g., Georgiadis and Szentes 2020, Lemma 2) guarantees the existence of a stopping time τ such that $X_\tau \sim G$ and $\mathbb{E}^0[\tau] = \mathrm{Var}_G(X) < \infty$. □ □

Under the binding IR constraint (16), we introduce a Lagrange multiplier $\lambda \in \mathbb{R}$. The principal's problem (P2) then reduces to[3]: for all $R \in \mathcal{R}$,

$$\max_{G\in\mathcal{G},\, \tilde{K}(X)\geq \underline{k}} \mathbb{E}_G \left[-v(\tilde{K}(X)) + \tilde{K}(X)(h\,(X)+\lambda) - \lambda \exp\left(\frac{1}{\delta}R\right)\right], \tag{P3}$$

where λ is determined by the constraint:

$$\mathbb{E}_G[\tilde{K}(X)] = \exp\left(\frac{1}{\delta}R\right). \tag{18}$$

The principal's problem thus separates into two stages: (i) choosing the optimal compensation schedule $\tilde{K}(\cdot)$ conditional on each realized X, and (ii) selecting the optimal distribution $G \in \mathcal{G}$ over X_τ.

For any fixed distribution $G \in \mathcal{G}$, the corresponding Lagrangian is:

$$L(\lambda, G) = \sup_{\tilde{K}(x)\geq \underline{k}} \int \left[-v(\tilde{K}(x)) + \tilde{K}(x)\,(h\,(x)+\lambda)\right] dG(x). \tag{19}$$

For each realized value x, the integrand is maximized pointwise by a unique value of $\tilde{K}(x)$, characterized below:

Lemma 3.2. *For each $\lambda \in \mathbb{R}$, the pointwise maximizer of (19) is given by $\tilde{K}(x)$, where*

$$\tilde{K}(x) := \begin{cases} \underline{k}, & \text{if } h(x)+\lambda \leq v'(\underline{k}), \\ v'^{-1}(h(x)+\lambda), & \text{if } h(x)+\lambda > v'(\underline{k}). \end{cases} \tag{20}$$

The inverse function $v'^{-1}(z)$ is strictly increasing.

The result follows directly from the first-order condition of the pointwise concave maximization in k, using the strict convexity of v. The proof is omitted.

A contract is valuable only if

$$h(x^*) + \lambda > v'(\underline{k}). \tag{21}$$

[3]Because the participation constraint is binding, $K_0 = \exp(R/\delta)$ is fixed. The ratio $1/K_0$ becomes a constant multiple and is absorbed into the Lagrange multiplier λ in Problem (P3), allowing K_0 to be omitted.

If instead $h(x^*) + \lambda \leq v'(\underline{k})$, then by (20), the optimal compensation satisfies:

$$\tilde{K}(x) = \underline{k}, \text{ for all } x \in \mathcal{X}.$$

In this case K_τ is constant and $M_\tau = \frac{K_\tau}{K_0} = 1$. The agent thus receives only the minimum wage and exerts zero effort. Consequently, the principal optimally terminates the contract immediately at time 0, and the contract is not valuable in the sense of Definition 2.2.

To ensure the contract is valuable, it is necessary $\lambda > \underline{\lambda}$, where the threshold $\underline{\lambda}$ is defined by:

$$h(x^*) + \underline{\lambda} = v'(\underline{k}). \tag{22}$$

Moreover, for any $\lambda > \underline{\lambda}$, there exists a unique critical threshold $x^c(\lambda) < x^*$ such that

$$h(x^c(\lambda)) + \lambda = v'(\underline{k}). \tag{23}$$

To simplify the principal's objective, we introduce the convex conjugate of the distortion function $v(k)$:

$$\phi(x) = \max_{k \geq \underline{k}} \left\{-v(k) + kx\right\}. \tag{24}$$

Then we have $\phi'(x) = v'^{-1}(x)$ on $[v'(\underline{k}), \infty)$.

Lemma 3.3. *The function $\phi(x)$ is strictly increasing and strictly convex for $x > v'(\underline{k})$, and linear for $x \leq v'(\underline{k})$: $\phi(x) = -v(\underline{k}) + \underline{k}x$.*

The result follows from standard properties of convex conjugation applied to $v(\cdot)$, and the boundary condition at $\underline{k}$. When the agent has logarithmic utility $u(c) = \log(c)$, the conjugate $\phi(x)$ admits the closed-form:

$$\phi(x) = \delta \left(\frac{x}{1+\delta}\right)^{\frac{1+\delta}{\delta}} \quad \text{for } x > v'(\underline{k}). \tag{25}$$

Accordingly, the pointwise Lagrangian objective for each realized x becomes $\phi(h(x) + \lambda)$, and the aggregate Lagrangian (19) reduces to:

$$L(\lambda, G) = \mathbb{E}_G \left[\phi(h(X) + \lambda)\right]. \tag{26}$$

The principal's problem (P3) simplifies to:

$$\max_{G\in\mathcal{G}} \mathbb{E}_G\left[\phi(h(X)+\lambda)\right]. \tag{27}$$

4 Optimal Compensation

If $\phi(h(x)+\lambda)$ is concave in x, then by Jensen's inequality,

$$\mathbb{E}_G\left[\phi(h(X)+\lambda)\right] \leq \phi(h(X_0)+\lambda), \tag{28}$$

In this case, the principal strictly prefers immediate termination at time 0 and the contract is not valuable.

Lemma 4.1. *The composite function $\phi(h(x)+\lambda)$ is continuously differentiable (C^1) and strictly increasing on $\mathcal{X}$. Moreover:*

- *For $\lambda > \underline{\lambda}$:*
  - *when $x \leq x^c(\lambda)$, $\phi(h(x)+\lambda)$ is affine in $h(x)$ with $\phi(h(x)+\lambda) = \underline{k}h(x) - v(\underline{k})$.*
  - *when $x^c(\lambda) < x < x^*$, $\phi(h(x)+\lambda)$ is strictly convex in $h(x)$.*
- *For $\lambda < \underline{\lambda}$, $\phi(h(x)+\lambda)$ is affine in $h(x)$ for all $x \in \mathcal{X}$: $\phi(h(x)+\lambda) = \underline{k}h(x) - v(\underline{k})$.*

The quasi-concavity of $\phi(h(x)+\lambda)$ for $x > x^c(\lambda)$ follows from the fact that $\phi(\cdot)$ is strictly increasing and $h(x)$ is concave. For $x \leq x^c(\lambda)$, the function becomes affine in $h(x)$ due to the definition of the convex conjugate, where the optimal solution binds at the lower boundary $\underline{k}$. Furthermore, $\phi(h(x)+\lambda)$ is continuously differentiable at $x = x^c(\lambda)$: both the left and right derivatives exist and are equal to $\underline{k}h'(x)$, as implied by Lemma 3.3. The detailed proof is omitted.

To fully characterize the curvature of $\phi(h(x)+\lambda)$ over the interval $x \in [x^c(\lambda), x^*]$ for $\lambda > \underline{\lambda}$, we impose the following structural condition:

Assumption 4.2. *(i) $kv''(k)$ is strictly increasing for $k > \underline{k}$.*

(ii) $\dfrac{h''(x)}{(h'(x))^2}$ is strictly decreasing on $\mathcal{X}$.

This assumption ensures a sufficient condition for the existence of an interior inflection point for $\phi(h(x)+\lambda)$. Here is an example of a model setup that satisfies Assumption 4.2: suppose the agent has logarithmic utility: $u(c) = \log(c)$. Then the distorted value function is $v(k) = k^{1+\delta}$, which satisfies the first condition. Suppose the principal's net payoff is $h(x) = x - \alpha x^2$ with $\alpha > 0$. The second condition also holds, and optimal performance level is $x^* = \frac{1}{2\alpha}$.

Proposition 4.3. *Let $\lambda > \underline{\lambda}$ and $x^c(\lambda)$ exists which is defined by equation* (23)*.*

1. *The function $\phi(x)$ exhibits decreasing relative curvature for all $x > v'(\underline{k})$; that is,*

$$\frac{\phi''(x)}{\phi'(x)} \text{ is strictly decreasing in } x.$$

2. *If*

$$\frac{\phi''(h(x^c(\lambda))+\lambda)}{\phi'(h(x^c(\lambda))+\lambda)} + \frac{h''(x^c(\lambda))}{(h'(x^c(\lambda)))^2} > 0, \tag{29}$$

then there exists a unique inflection point $x^i(\lambda) \in (x^c(\lambda), x^*)$ such that the composite function $\phi(h(x)+\lambda)$ is:

- *strictly convex on $[x^c(\lambda), x^i(\lambda))$,*
- *strictly concave on $(x^i(\lambda), x^*]$.*

Moreover, the inflection point $x^i(\lambda)$ satisfies the following condition:

$$\frac{\phi''(h(x^i(\lambda))+\lambda)}{\phi'(h(x^i(\lambda))+\lambda)} + \frac{h''(x^i(\lambda))}{(h'(x^i(\lambda)))^2} = 0. \tag{30}$$

3. *Otherwise, if condition* (29) *fails, the composite function $\phi(h(x)+\lambda)$ is concave on $\mathcal{X}$.*

Proof. Under Assumption 4.2, $kv''(k)$ is increasing in k. Since $\phi(x)$ is the convex conjugate of $v(k)$, and convex conjugates of strictly convex functions are themselves strictly convex and smooth on the interior of their domains ($x > v'(\underline{k})$), we can characterize ϕ using properties of v. Let $k(x) = v'^{-1}(x)$, so that $\phi'(x) = k(x)$, and $\phi''(x) = k'(x) = \frac{1}{v''(k(x))}$. Then,

$$\frac{\phi''(x)}{\phi'(x)} = \frac{1}{v''(k(x)) \cdot k(x)}.$$

Since $kv''(k)$ is strictly increasing in k, the function $\frac{1}{kv''(k)}$ is strictly decreasing in k, and hence decreasing in x via $k(x)$. Therefore, $\frac{\phi''(x)}{\phi'(x)}$ is strictly decreasing in x for $x > v'(\underline{k})$.

We analyze the curvature of the composite function $\phi(h(x)+\lambda)$ for $h(x)+\lambda > v'(\underline{k})$ or equivalent $x > x^c(\lambda)$. Its first and second derivatives are:

$$\frac{d}{dx}\phi(h(x)+\lambda) = \phi'(h(x)+\lambda)\cdot h'(x),$$

$$\frac{d^2}{dx^2}\phi(h(x)+\lambda) = (h'(x))^2\phi'\left(h(x)+\lambda\right)C(x)$$

with

$$C(x) := \frac{\phi''(h(x)+\lambda)}{\phi'(h(x)+\lambda)} + \frac{h''(x)}{\left(h'(x)\right)^2}.$$

Then the sign of $C(x)$ determines the concavity of the composite function because $\phi(x)$ is strictly increasing. Because the first term $\frac{\phi''}{\phi'}$ is decreasing in $h(x)+\lambda$, and $h(x)$ is increasing in x for $x^c(\lambda) < x \leq x^*$, this term is decreasing in x. Also, by Assumption 4.2, the second term $\frac{h''(x)}{h'^2}$ is strictly decreasing in x. Therefore, the sum $C(x)$ is strictly decreasing in x. Also notice $\frac{d^2}{dx^2}\phi(h(x^*)+\lambda) = \phi'(h(x^*)+\lambda)h''(x^*) < 0$ because $h'(x^*) = 0$. Therefore $C(x^*) = -\infty$.

- If $C(x^c(\lambda)) > 0$, then since $C(x)$ is strictly decreasing on $(x^c(\lambda), x^*]$, there exists a unique point $x^i(\lambda) \in (x^c(\lambda), x^*)$ such that $C(x^i(\lambda)) = 0$. This implies that the composite function $\phi(h(x)+\lambda)$ is convex on $[x^c(\lambda), x^i(\lambda)]$ and concave on $(x^i(\lambda), x^*]$.

- If $C(x^c(\lambda)) \leq 0$, then $C(x) < 0$ for all $x \in [x^c(\lambda), x^*]$, and the composite function is concave on $[x^c(\lambda), x^*]$ and $(-\infty, x^c(\lambda)]$. Moreover

$$\lim_{x\to x^c(\lambda)^-} \frac{d\phi\left(h\left(x\right)+\lambda\right)}{dx} = \underline{k}h'(x^c\left(\lambda\right))$$
$$\lim_{x\to x^c(\lambda)^+} \frac{d\phi\left(h\left(x\right)+\lambda\right)}{dx} = \phi'\left(h\left(x^c\left(\lambda\right)\right)+\lambda\right)h^{'}\left(x^c\left(\lambda\right)\right) = \underline{k}h^{'}\left(x^c(\lambda)\right),$$

  $\phi\left(h\left(x\right)+\lambda\right)$ is C^1 and it keeps decreasing as x increases. Therefore it is concave on $\mathcal{X}$.

□

Condition (29) is both a necessary and sufficient condition for the existence of an inflection

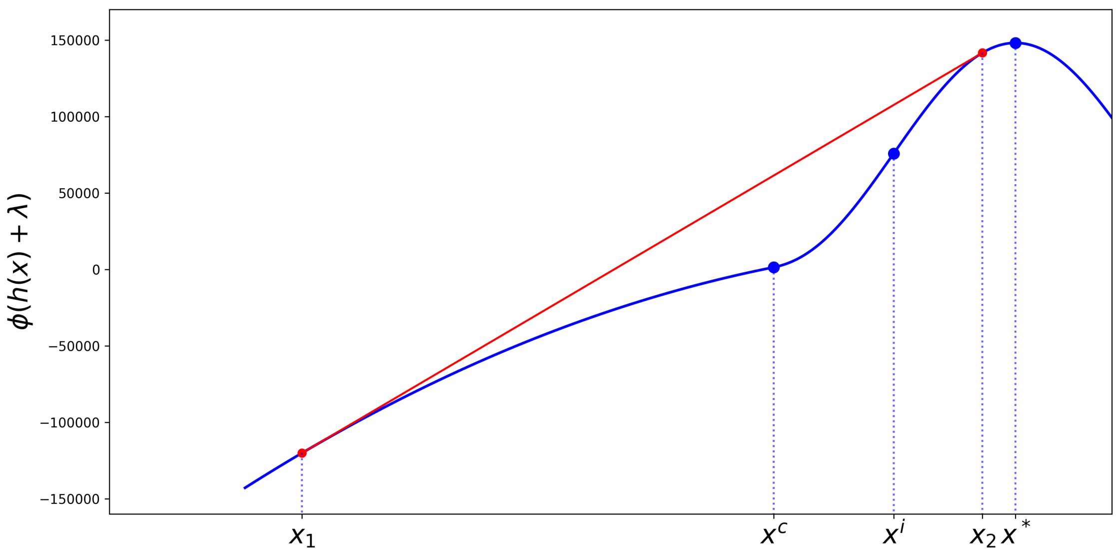

Figure 1: Concavity profile of $\phi(h(x) + \lambda)$: The function is concave when $x < x^c$ or $x > x^i$ and convex on $[x^c, x^i]$. Parameters: $u(x) = \log(x)$, $h(x) = x - 0.005x^2$, $\delta = 0.5$, $\sigma = 1$, $\underline{k} = 200$.

point. If this condition fails, then $\phi(h(x) + \lambda)$ is concave on $\mathcal{X}$ and it is not valuable to monitor the agent by Jensen's inequality.

Figure 1 plots the composite function $\phi(h(x) + \lambda)$ over the domain $\mathcal{X}$. As predicted by Proposition 4.3, the function $\phi(h(x) + \lambda)$ exhibits a distinct convex-concave profile: it is strictly convex on $[x_c(\lambda), x_i(\lambda)]$ and strictly concave on $(x_i(\lambda), x^*]$, with a unique inflection point at $x_i(\lambda)$. This illustrates the non-monotonic curvature structure that gives rise to a linear segment in the concave envelope derived later in Proposition 4.5.

We now examine how the shape of the composite function $\phi(h(x)+\lambda)$ evolves as the parameter λ varies. From condition (23), $x^c(\lambda)$ increases as λ decreases. From condition (30), $x^i(\lambda)$ increases as λ decreases. This monotonicity implies that as λ decreases, the inflection point $x^i(\lambda)$ and the contact point $x^c(\lambda)$ both shift rightward, shrinking the convex region. At a critical value $\lambda = \lambda_{\min}$, condition (29) binds,

$$\frac{h''(x^c(\lambda_{min}))}{(h'(x^c(\lambda_{min})))^2} + \frac{\phi''(v'(\underline{k}))}{\phi'(v'(\underline{k}))} = 0, \tag{31}$$

along with (23) at $\lambda = \lambda_{min}$:

$$h(x^c(\lambda_{min})) + \lambda_{min} = v'(\underline{k}). \tag{32}$$

Corollary 4.4. *Under Assumption 4.2, the threshold functions $x^c(\lambda)$ and $x^i(\lambda)$ are strictly de-*

creasing with λ. There exists a unique critical value $\lambda_{min} > \underline{\lambda}$ such that Condition (29) holds if and only if $\lambda > \lambda_{min}$. The pair $(\lambda_{min}, x^c(\lambda_{min}))$ is characterized by equations (31) and (32). As $\underline{k}$ increases, $x^c(\lambda_{min})$ will decrease and $\lambda_{\min}$ will increase. Moreover, at $\lambda = \lambda_{min}$, the two thresholds coincide: $x^c(\lambda_{min}) = x^i(\lambda_{min})$.

Proof. Most of proof has been done in the discussion before Corollary 4.4. We need to prove

1. $\lambda_{min} > \underline{\lambda}$.

2. condition(29) is true if and only if $\lambda > \lambda_{min}$.

First, $\underline{\lambda}$ is given by

$$h(x^c(\lambda)) + \lambda = v'(\underline{k}),$$

when $x^c(\lambda) = x^*$. At $\lambda = \underline{\lambda}$ and $x^c(\underline{\lambda}) = x^*$,

$$\frac{\phi''(h(x^c(\lambda)) + \lambda)}{\phi'(h(x^c(\lambda)) + \lambda)} + \frac{h''(x^c(\lambda))}{(h'(x^c(\lambda)))^2} = -\infty, \tag{33}$$

As λ increases, the left side of (33) increases because $x^c(\lambda)$ decreases. Therefore $\lambda_{\min} > \underline{\lambda}$ so that (22) and (32) hold true.

Second, the equivalence of $\lambda > \lambda_{min}$ and condition (29) comes from the monotonicity of the left side of (29) as λ increases from λ_{min}. □

In the special case where $h(x) = x - \alpha x^2$, an explicit solution for λ_{min} can be obtained from (22) and (32):

$$\lambda_{min} = v'(\underline{k}) - \left(\frac{1}{2\alpha} - \frac{1}{2\alpha}\left(\frac{2\alpha\phi'(v'(\underline{k}))}{\phi''(v'(\underline{k}))}\right)^{\frac{1}{2}}\right) + \alpha\left(\frac{1}{2\alpha} - \frac{1}{2\alpha}\left(\frac{2\alpha\phi'(v'(\underline{k}))}{\phi''(v'(\underline{k}))}\right)^{\frac{1}{2}}\right)^2. \tag{34}$$

Following standard arguments from the information design literature (see Aumann and Perles 1965; Kamenica and Gentzkow 2011), and building on the formulation originally proposed by Georgiadis and Szentes (2020) for optimal monitoring with static effort choice, we derive an analogous result adapted to our dynamic setting, providing a more detailed characterization.

Proposition 4.5. *For any $x \in \mathcal{X}$, the concave envelope of $\phi(h(x)+\lambda)$ is given by:*

$$\bar{\phi}(h(x)+\lambda) := \sup_{\substack{x_1,x_2\in\mathcal{X},\, x_1\le x_2\\ \pi\in[0,1]\\ \pi x_2+(1-\pi)x_1=x}} \left\{\pi\phi(h(x_2)+\lambda)+(1-\pi)\phi(h(x_1)+\lambda)\right\}. \tag{35}$$

Under Assumption 4.2, the supremum is finite and is attained at some triplet $(x_1(\lambda), x_2(\lambda), \pi(\lambda, x)) \in \mathcal{X}\times\mathcal{X}\times[0,1]$ where $\pi(\lambda,x) = \frac{x-x_1(\lambda)}{x_2(\lambda)-x_1(\lambda)}$.

- *If $\lambda > \lambda_{min}$, then there exists $x_1(\lambda) \le x^c(\lambda) < x^i(\lambda) \le x_2(\lambda) < x^*$ such that the concave envelope $\bar{\phi}(h(x)+\lambda)$ takes the form:*

$$\bar{\phi}(h(x)+\lambda) = \begin{cases} \dfrac{\phi(h(x_2)+\lambda)-\phi(h(x_1)+\lambda)}{x_2-x_1}(x-x_1)+\phi(h(x_1)+\lambda) & \text{if } x_1 < x \le x_2, \\ \phi(h(x)+\lambda) & \text{else .} \end{cases} \tag{36}$$

  In the linear region $[x_1(\lambda), x_2(\lambda)]$, the convex combination weight is given by $\pi = \frac{x-x_1}{x_2-x_1} = \mathbb{P}(X_\tau = x_2)$. The points $x_1(\lambda)$ and $x_2(\lambda)$ are implicitly defined by the following two equations:

  (i) Gradient matching condition:

$$\underline{k}h'(x_1) = \phi'(h(x_2)+\lambda)h'(x_2) \tag{37}$$

  (ii) Value matching condition:

$$\phi(h(x_2)+\lambda) - (\underline{k}(h(x_1)+\lambda) - v(\underline{k})) = \underline{k}h'(x_1)(x_2-x_1) \tag{38}$$

- *if $\lambda \le \lambda_{min}, \bar{\phi}(h(x)+\lambda) = \phi(h(x)+\lambda)$ for all $x \in \mathcal{X}$, and the supremum is attained at $x_1 = x_2 = x$.*

Proof. Most cases follow directly from the definition of the concave envelope (36). We therefore focus on proving the nontrivial existence of a triplet (x_1, x_2, π) when condition (29) holds. By Proposition 4.3, this condition guarantees the existence of an inflection point $x^i \in (x^c, x^*)$, such that the composite function $\phi(h(x)+\lambda)$ is convex on $[x^c, x^i]$ and concave on $[x^i, x^*]$.

For convenience, we use $L(x, x')$ to denote the straight line that pass through points $(x, \phi(h(x)+$

λ))) and $(x', \phi(h(x') + \lambda)))$.

Fix $\mathbf{x_2} = \mathbf{x^i}$, and consider a line $L(x_1, x_2)$ passing through the point $(x_2, \phi(h(x_2)+\lambda))$. We seek a point $x_1 \leq x^c$ such that $L(x_1, x_2)$ is tangent to $\phi(h(x) + \lambda)$ at $x = x_1$. We will show that such a tangent point exists and satisfies $x_1 \leq x^c$.

To begin, construct a straight line $L'(x^c, x^i)$ connecting the points $(x^i, \phi(h(x^i) + \lambda))$ and $(x^c, \phi(h(x^c) + \lambda))$, where $x_2 = x^i$. If $L'(x^c, x^i)$ is tangent to $\phi(h(x) + \lambda)$ at $x = x^c$, then we may set $x_1 = x^c$. Otherwise, $L'(x^c, x^i)$ is a secant line that intersects the graph of $\phi(h(x) + \lambda)$ at two points: $(x^c, \phi(h(x^c) + \lambda))$ and $(x_3, \phi(h(x_3) + \lambda))$ with $x_3 < x^c$. Since $\phi(h(x) + \lambda)$ is concave on $(-\infty, x^c]$, the line $L'(x^c, x^i)$ lies below the graph of $\phi(h(x) + \lambda)$ on interval $[x_3, x^c]$. By the intermediate value theorem, there exists $x_1 \in (x_3, x^c)$ such that the line $L(x_1, x_2)$ is tangent to $\phi(h(x) + \lambda)$ at $x = x_1$ with $x_1 < x^c$. Since $\phi(h(x) + \lambda) = \underline{k}(h(x) + \lambda) - v(\underline{k})$ for $x \leq x^c$, the slope of the tangent line at $x_1 \leq x^c$ is given by $\underline{k}h'(x_1)$. The equation of the line is:

$$y = \underline{k}h'(x_1)(x - x_1) + \phi(h(x_1) + \lambda). \tag{39}$$

We now evaluate the tangent line equation at $x = x_2$ ($x_2 = x^i$):

$$\phi\left(h\left(x_2\right) + \lambda\right) = \underline{k}h'\left(x_1\right)\left(x_2 - x_1\right) + \left(\underline{k}\left(h\left(x_1\right) + \lambda\right) - v\left(\underline{k}\right)\right). \tag{40}$$

If the line $L(x_1, x_2)$ is also tangent to $\phi(h(x)+\lambda)$ at $x = x_2 = x^i$, then $L(x_1, x_2)$ is tangent at both endpoints x_1 and x_2. Since $\phi(h(x) + \lambda)$ is concave on $[-\infty, x_1]$ and $[x_2, x^*]$, therefore $L(x_1, x_2)$ constructs the linear segment of the concave envelope, completing the characterization in equation (36).

If $L(x_1, x_2)$ is not tangent to $\phi(h(x) + \lambda)$ at $x = x_2 = x^i$, then it serves as a secant line on $[x^i, x^*]$. In this case, the line will either intersect with the curve $\phi(h(x) + \lambda)$ at the third point x_3, $x^i < x_3 \leq x^*$ or $L(x_1, x_2)$ is below $(x^*, \phi(h(x^*) + \lambda))$. To make our discussion easier, denote x_3' as x_3 or x^*. Since $\phi(h(x) + \lambda)$ is strictly convex on $[x^c, x^i]$, the function lies strictly below the line $L(x_1, x_2)$ on this interval $[x_1, x_2]$. Due to the concavity of $\phi(h(x) + \lambda)$ on $[x^i, x_3')$, the function lies above the secant line on $[x^i, x_3']$. As a result, the line $L(x_1, x_2)$ dominates $\phi(h(x) + \lambda)$ on the entire interval $(-\infty, x_2]$.

To understand how x_2 varies with x_1, we differentiate both sides of the equation (40) for the tangent line $L(x_1, x_2)$ with respect to x_1, the expression becomes

$$\frac{dx_2}{dx_1} = \frac{\underline{k}h"(x_1)(x_2 - x_1)}{\phi'(h(x_2) + \lambda)h'(x_2) - \underline{k}h'(x_1)}. \tag{41}$$

Since $h(x)$ is strictly concave, we have $h''(x) < 0$. At the initial position, the slope of $\phi(h(x)+\lambda)$ at $x = x_2$ exceeds the slope at $x = x_1$ since $L(x_1, x_2)$ is a secant line and intersects with $\phi(h(x) + \lambda)$ at x_1, x_2 with $x_2 > x_1$. Then the denominator in (41) is positive. Therefore $\frac{dx_2}{dx_1} < 0$ initially.

As x_1 decreases, x_2 increases. Consequently, the slope of the line $L(x_1, x_2)$ increases, while the slope of $\phi(h(x) + \lambda)$ at $x = x_2$ will decrease due to the concavity of the function on $[x_2, x^*]$. As long as the slope of $L(x_1, x_2)$ remains strictly below that of $\phi(h(x) + \lambda)$ at x_2, we have $\frac{dx_2}{dx_1} < 0$.

Moreover, throughout this process, $\phi(h(x) + \lambda)$ remains strictly below the line $L(x_1, x_2)$ over $(-\infty, x_2)$, ensuring that the secant remains a valid upper bound until it becomes tangent to $\phi(h(x)+\lambda)$ at both endpoints.

Notice that the slope of $\phi(h(x) + \lambda)$ at $x = x^*$ is zero and $\phi(h(x) + \lambda)$ is strictly concave on $[x^i, x^*]$. Therefore, by continuity and monotonicity of the slope, as x_1 decreases, there exist a unique $x_2 \in [x^i, x^*)$ such that the slope of the line $L((x_1, x_2))$ equals the derivative of $\phi(h(x) + \lambda)$ at x_1. Through these constructions, we find the tangent points at x_1, x_2 with $x_1 \leq x^c$ and $x^i \leq x_2 < x^*$. □

Implementability. The optimal binary distribution G places mass $(1 - \pi)$ on $x_1(\lambda)$ and π on $x_2(\lambda)$ with $\pi x_2 + (1 - \pi)x_1 = X_0$. Since $x_1, x_2 \in \mathcal{X}$ are finite, G has mean X_0 and finite second moment, so $G \in \mathcal{G}$. By Lemma 3.1 (Skorokhod embedding), G is realizable as the distribution of X_τ for some stopping time τ with $\mathbb{E}^0[\tau] < \infty$.

Proposition 4.5 reveals that when the composite function $\phi(h(x)+\lambda)$ fails to be concave (valuable to monitor), the optimal design replaces the non-concave region with a linear segment that is tangent at two threshold points. This construction has a natural interpretation in terms of the agent's value set. In particular, the linear segment corresponds to the convexification of non-concave payoffs that would otherwise violate incentive compatibility. The firm optimally offers a lottery over two levels of performance, inducing the agent to take dynamic actions. The deeper economic reason for binary optimality is the principal's preference for variance in the agent's terminal ECV K_τ.

Because ϕ is strictly convex on the relevant region, Jensen's inequality implies $\mathbb{E}_G[\phi(h(X)+\lambda)] \geq \phi(h(\mathbb{E}_G[X])+\lambda)$ whenever G is non-degenerate, so the principal benefits from spreading K_τ across two extreme values rather than concentrating it at the mean. This convexity-driven variance preference makes concentration on exactly two points optimal.

We now define the *valuable set* $\mathcal{E}$, the collection of (λ, X_0) pairs for which dynamic monitoring is valuable. In particular, when monitoring begins at initial performance level X_0 and promised utility cost parameter λ, it is valuable to continue monitoring if the agent's performance level lies strictly between the two optimal thresholds. That is, the monitoring process is not immediately terminated.

$$\mathcal{E} = \{(\lambda, X_0) \ : \ \lambda > \lambda_{min} \text{ and } x_1(\lambda) < X_0 < x_2(\lambda)\}. \tag{42}$$

Economic interpretation. The valuable set $\mathcal{E}$ characterizes when dynamic monitoring creates value for the principal. A pair (λ, X_0) lies in $\mathcal{E}$ when two conditions are met simultaneously: (i) the promised utility cost λ exceeds the critical threshold λ_{min}, ensuring the contract is incentive-powerful enough to induce non-trivial effort; and (ii) the agent's initial performance X_0 falls strictly between the two stopping thresholds, so the principal's monitoring has room to generate informative signals before hitting a boundary. Outside $\mathcal{E}$, the principal optimally terminates immediately—either because the cost of incentivizing effort exceeds the informational gain from monitoring, or because the initial performance is already extreme enough that no further observation is warranted. The set $\mathcal{E}$ thus maps the fundamental monitoring–incentive trade-off into a concrete region of the parameter space. Figure 2 visualizes the set $\mathcal{E}$. The lower boundary $x_1(\lambda)$ (blue) and upper boundary $x_2(\lambda)$ (green) represent the thresholds beyond which monitoring is immediately stopped. The shaded region indicates the range of initial conditions for which performance is actively monitored and evolves endogenously. At the critical value $\lambda = \lambda_{min}$, the two thresholds coincide, and the valuable set collapses to a singleton, beyond which a nontrivial monitoring strategy emerges.

To characterize the limit of $x_1(\lambda), x_2(\lambda)$ as λ approaches λ_{min} in Figure 3, we have the following:

Corollary 4.6. *Suppose that the pair $(x_1(\lambda), x_2(\lambda))$ satisfies the gradient matching condition (37) and value matching condition (38). Then*

$$\lim_{\lambda \to \lambda_{\text{min}}^-} x_1(\lambda) = \lim_{\lambda \to \lambda_{\text{min}}^-} x_2(\lambda) = x^c(\lambda_{\text{min}}) \tag{43}$$

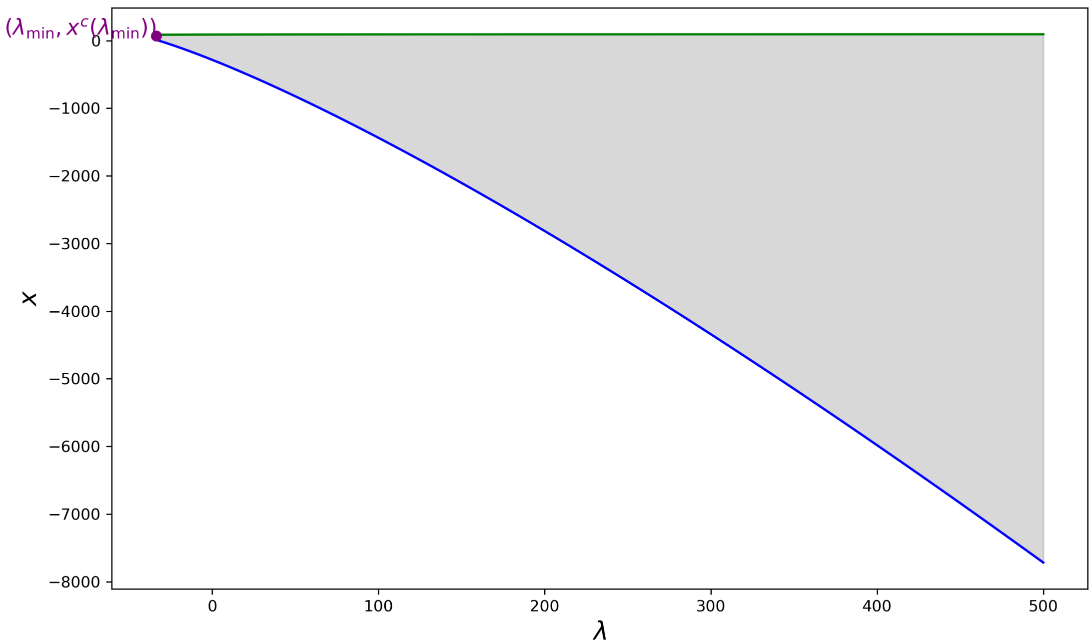

Figure 2: Valuable set $\mathcal{E}$ in (λ, X_0) space: $u(x) = \log(x)$, $h(x) = x - 0.005x^2$, $\delta = 0.5$, $\sigma = 1$, $\underline{k} = 200$.

where

$$h(x^c(\lambda_{\min})) + \lambda_{\min} = v'(\underline{k}).$$

Proof. Although the domain of λ is unbounded above (i.e., $\lambda \in [\lambda_{\min}, \infty)$), we are only concerned with the behavior of $(x_1(\lambda), x_2(\lambda))$ as $\lambda \to \lambda_{\min}$. Therefore, it suffices to restrict attention to a compact subinterval $[\lambda_{\min}, \lambda_0]$ for some $\lambda_0 > \lambda_{\min}$. By the implicit definition of $x_1(\lambda)$ and $x_2(\lambda)$ through the gradient and value matching conditions, and using the regularity and strict monotonicity of h and ϕ, the solution pair $(x_1(\lambda), x_2(\lambda))$ varies continuously in λ over this interval. Hence, the images of $x_1(\cdot)$ and $x_2(\cdot)$ on the compact set $[\lambda_{\min}, \lambda_0]$ are also compact.

By the Bolzano–Weierstrass theorem, there exists a sequence $\lambda_n \to \lambda_{\min}$ such that $x_1(\lambda_n) \to x_1^*$ and $x_2(\lambda_n) \to x_2^*$ for some limit points x_1^*, x_2^*. Taking the limit in the matching conditions, the pair (x_1^*, x_2^*) satisfies the following system at $\lambda = \lambda_{\min}$:

$$\underline{k}h'(x_1^*) = \phi'\left(h(x_2^*) + \lambda_{\min}\right) \cdot h'(x_2^*), \tag{44}$$

$$\phi\left(h(x_2^*) + \lambda_{\min}\right) = \underline{k}\left(h(x_1^*) + \lambda_{\min}\right) - v(\underline{k}) + \underline{k}h'(x_1^*)(x_2^* - x_1^*). \tag{45}$$

We now prove that the solution (x_1^*, x_2^*) is unique. By the concavity of $\phi(h(\cdot) + \lambda_{\min})$ on $\mathcal{X}$ (Corollary 4.4), its derivative $\frac{d}{dx}\phi(h(x)+\lambda_{\min})$ is strictly decreasing in x. If $x_1^* < x_2^*$, evaluating the derivative at the larger point x_2^* yields a strictly smaller value than at x_1^*. Therefore the gradient matching condition (44) cannot hold, since $x_1^* < x_2^*$ and

$$\underline{k}h'(x_1^*) > \phi'\left(h(x_2^*) + \lambda_{\min}\right) \cdot h'(x_2^*).$$

Hence, the only possible solution is $x_1^* = x_2^*$. Denoting the common limit $x_1^* = x_2^* =: \bar{x}$ and substituting into the system, equations (44) and (45) become:

$$\underline{k}h'(\bar{x}) = \phi'\left(h(\bar{x}) + \lambda_{\min}\right) h'(\bar{x}), \tag{46}$$

$$\phi\left(h(\bar{x}) + \lambda_{\min}\right) = \underline{k}\left(h(\bar{x}) + \lambda_{\min}\right) - v(\underline{k}). \tag{47}$$

It follows from (47) that $\bar{x} = x^c(\lambda_{\min})$ from gradient matching condition at $x^c(\lambda_{\min})$. Therefore, any sequence $(x_1(\lambda_n), x_2(\lambda_n))$ with $\lambda_n \to \lambda_{\min}$ converges uniquely to $(x^c(\lambda_{\min}), x^c(\lambda_{\min}))$, completing the proof. □ □

When $(\lambda, x) \in \mathcal{E}$, the agent's performance starts within the interior of the optimal monitoring band, and thus monitoring continues. If instead $x \notin (x_1(\lambda), x_2(\lambda))$, the optimal policy calls for immediate termination of monitoring. The next result characterizes the agent's and principal's expected values under both regimes, which are straightforward to verify, so the proof is omitted.

Corollary 4.7. *The contract is valuable if and only if* $(\lambda, X_0) \in \mathcal{E}$. *The agent's* ***ECV****,* $A(\lambda, X_0) = \mathbb{E}_G[\tilde{K}(x)]$*, is*

$$A(\lambda, X_0) = \begin{cases} \underline{k} + \pi(\lambda, X_0)(\phi'(h(x_2(\lambda)) + \lambda) - \underline{k}) & \text{if } (\lambda, X_0) \in \mathcal{E}, \\ \phi'(h(X_0) + \lambda) & \text{if } (\lambda, X_0) \notin \mathcal{E} \text{ and } X_0 \geq x_2(\lambda), \\ \underline{k} & \text{if } (\lambda, X_0) \notin \mathcal{E} \text{ and } X_0 \leq x_1(\lambda). \end{cases} \tag{48}$$

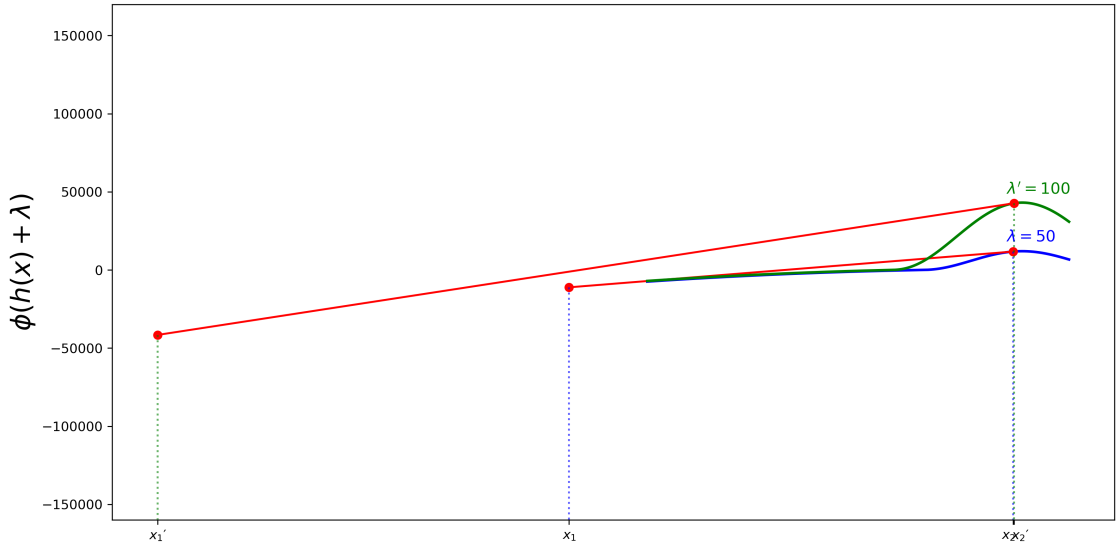

Figure 3: Comparison of $\bar{\phi}(h(x)+\lambda)$ for different λ: $u(x) = \log(x)$, $h(x) = x - 0.005x^2$, $\delta = 0.5$, $\sigma = 1$, $\underline{k} = 200$.

The principal's expected value, $P(\lambda, X_0) = \mathbb{E}_G[\bar{\phi}(h(X)) + \lambda] - \lambda A(\lambda, X_0)$ *is*

$$P(\lambda, X_0) = \begin{cases} (1-\pi(\lambda, X_0))\phi\left(h\left(x_1\left(\lambda\right)\right)+\lambda\right) + \pi(\lambda, X_0)\phi\left(h\left(x_2\left(\lambda\right)\right)+\lambda\right) - \lambda A\left(\lambda, X_0\right). & \text{if } (\lambda, X_0) \in \mathcal{E}, \\ \phi\left(h\left(X_0\right)+\lambda\right) - \lambda A\left(\lambda, X_0\right). & \text{if } (\lambda, X_0) \notin \mathcal{E}. \end{cases} \tag{49}$$

Figure 3 shows how the concave envelope $\bar{\phi}(h(x)+\lambda)$ evolves with λ. As λ increases, the lower threshold $x_1(\lambda)$ decreases while the upper threshold $x_2(\lambda)$ increases. Notably, $x_1(\lambda)$ shifts more rapidly, widening the interval $[x_1(\lambda), x_2(\lambda)]$. This implies that, for a fixed starting performance X_0, the agent becomes more likely to receive a higher payoff as λ rises. This observation motivates a detailed analysis of the monotonic relationships between λ and key objects such as $x_1(\lambda)$, $x_2(\lambda)$, the agent's ECV $A(\lambda, X_0)$, and the total surplus.

Assumption 4.8.

$$h'''\left(x\right) \leq 0, \ \forall x < x^*;$$
$$\phi'''(x) > 0, \phi''(x)^2 \geq \frac{1}{2}\phi'''\left(x\right)\phi'\left(x\right) \ \textit{for } x \geq v'\left(\underline{k}\right).$$

This assumption ensures that $h(x)$ becomes increasingly concave and $\phi(x)$ becomes increasingly convex as x increases. The second part ensures that the convexity of $\phi(x)$, measured by ϕ'', grows fast enough relative to the slope ϕ', preventing too rapid a rise in slope without corresponding curvature. For instance, if $h(x) = x - \alpha x^2$ and $u(x) = \log(x)$, the first condition holds trivially since $h'''(x) = 0$, and the second condition is satisfied whenever $\delta \leq 1$.

Remark 1 (Robustness of Assumption 4.8 to alternative utility specifications)**.** *Assumption 4.8 is not specific to logarithmic utility. We verify it for two standard specifications.*

CARA utility. Let $u(c) = -e^{-\alpha c}/\alpha$ with $\alpha > 0$. Then $u^{-1}(y) = -\frac{1}{\alpha}\ln(-\alpha y)$, and $v(k) = k\, u^{-1}(\delta \ln k)$ is strictly convex. Its conjugate ϕ satisfies $\phi'''(x) > 0$ and $\phi''(x)^2 \geq \frac{1}{2}\phi'''(x)\phi'(x)$ for all $\alpha > 0$: the exponential structure ensures that ϕ''/ϕ' is monotonically decreasing, so the curvature inequality holds globally. Hence both conditions in the second part of Assumption 4.8 are satisfied for all $\alpha > 0$, provided $\underline{k} < 1$ (equivalently, $u(\underline{c}) < 0$), which is required for the ECV to be well-defined. The first part ($h''' \leq 0$) is a property of h alone and is unchanged.

CRRA utility. Let $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$, $\gamma \neq 1$. Then $u^{-1}(y) = [(1-\gamma)y]^{1/(1-\gamma)}$, and $v(k) = k\,[(1-\gamma)\delta \ln k]^{1/(1-\gamma)}$. For $\gamma > 1$, v is strictly convex and its conjugate ϕ satisfies both conditions (again requiring $\underline{k} < 1$, i.e., $u(\underline{c}) < 0$, for the ECV domain): the power structure yields $\phi''/\phi' \propto x^{-1}$, which is decreasing, and the curvature inequality $\phi''(x)^2 \geq \frac{1}{2}\phi'''(x)\phi'(x)$ follows. For $\gamma < 1$, the condition can fail because v may lack the required convexity growth rate.

In summary, Assumption 4.8 is satisfied for CARA utility for all $\alpha > 0$, and for CRRA utility when $\gamma > 1$.

To establish the monotonicity of the principal's value $P(\lambda, X_0)$ and the agent's value $A(\lambda, X_0)$ with respect to λ, we first state the following lemma, which characterizes the monotonicity of $x_1(\lambda)$, $x_2(\lambda)$ and $\pi(\lambda, X_0)$.

Lemma 4.9 (Monotonicity of $1 + m(\lambda)$)**.** *Suppose Assumptions 4.2 and 4.8 hold, and consider the interior continuation region $\{\lambda > \underline{\lambda},\, x_2(\lambda) \in (x^c(\lambda), x^*)\}$. Then $1+m(\lambda)$ is strictly decreasing in λ:*

$$\frac{d}{d\lambda}(1 + m(\lambda)) \;<\; 0.$$

Step 1 (short argument). By definition,

$$\frac{d}{d\lambda}(1+m(\lambda)) \; = \; m'(\lambda).$$

By Lemma 4.10(item 1) and the curvature monotonicity in Assumption 4.2, we have $m'(\lambda) < 0$ on the interior continuation region (proof provided earlier). Hence $(1+m)'(\lambda) < 0$.

Step 2 (direct differentiation and sign check). For completeness, we verify the sign by differentiating a closed form of $1+m$. Write

$$z(\lambda) := h(x_2(\lambda)) + \lambda, \qquad \kappa(\lambda) := \phi'(z(\lambda)), \qquad h_1 := h'(x_2), \quad h_2 := h''(x_2), \quad h_3 := h'''(x_2),$$

and define

$$S(\lambda) := \phi''(z(\lambda))\,(h_1)^2, \qquad B(\lambda) := h_2.$$

On the interior region, $h_1 > 0$ and $B = h_2 < 0$. Using $m(\lambda) = \dfrac{\phi''(z)(h_1)^2 + \underline{k}\,h_2}{-\phi''(z)(h_1)^2 - \phi'(z)\,h_2} = -\dfrac{S+\underline{k}B}{S+\kappa B}$, we obtain the identity

$$1+m(\lambda) \; = \; 1 - \frac{S+\underline{k}B}{S+\kappa B} \; = \; \frac{(\kappa-\underline{k})\,B}{S+\kappa B}. \tag{50}$$

Set $J(\lambda) := 1+m(\lambda) = \dfrac{N(\lambda)}{D(\lambda)}$ with

$$N(\lambda) := (\kappa-\underline{k})\,B, \qquad D(\lambda) := S+\kappa B.$$

By the quotient rule,

$$J'(\lambda) = \frac{N'D - ND'}{D^2}.$$

Using $z' = h_1\,x_2' + 1$, $\kappa' = \phi''(z)\,z'$, $x_2' = \dfrac{m}{h_1}$ (Lemma 4.10 (item 1)),

$$S' = \phi'''(z)\,z'\,(h_1)^2 + \phi''(z)\,2h_1h_2\,x_2', \qquad B' = h_3\,x_2',$$

a direct expansion yields the key cancellation

$$\begin{aligned}N'D - ND' &= \left[\kappa' B + (\kappa - \underline{k})\, B'\right](S + \kappa B) - (\kappa - \underline{k})\, B\,(S' + \kappa' B + \kappa B') \\ &= \kappa' BS + \kappa B' S - \underline{k}\, B' S + \underline{k}\, BS' - \kappa BS' + \underline{k}\, B\kappa' B \\ &= -\,(\kappa - \underline{k})\,(BS' - SB') \;+\; \kappa' B\,(S + \underline{k} B). \end{aligned} \tag{51}$$

Introduce the relative curvature of h,

$$\Theta(x) := \frac{h''(x)}{(h'(x))^2},$$

and

$$\frac{S}{B} = \frac{\phi''(z)}{\Theta(x_2)} = \frac{\phi''(z)}{\frac{h''(x)}{(h'(x))^2}}. \tag{52}$$

Differentiating (52) and using $B^2(S/B)' = BS' - SB'$ gives

$$BS' - SB' = B^2(S/B)' \tag{53}$$

$$= B^2 \frac{\phi'''(z)\,\Theta(x_2)\,z' \;-\; \phi''(z)\,\Theta'(x_2)\,x_2'}{\Theta(x_2)^2} > 0. \tag{54}$$

Notice $S + \alpha B = B\Big(\alpha + \frac{\phi''(z)}{\Theta(x_2)}\Big)$ from (53). Then

$$(S + \underline{k}B)(S + \kappa B) = B^2\left(\underline{k} + \frac{\phi''}{\Theta}\right)\left(\kappa + \frac{\phi''}{\Theta}\right), \qquad \frac{B}{S + \kappa B} = \frac{1}{\kappa + \frac{\phi''}{\Theta}}. \tag{55}$$

Combining (51),(55) and recalling $D = S + \kappa B$, we obtain

$$J'(\lambda) = \frac{-\,(\kappa - \underline{k})\,B^2 \frac{\phi'''(z)\,\Theta\,z' - \phi''(z)\,\Theta'(x_2)\,x_2'}{\Theta^2} \;+\; \kappa'\left(\underline{k} + \frac{\phi''(z)}{\Theta}\right)B^2}{B^2\left(\kappa + \frac{\phi''(z)}{\Theta}\right)^2}, \tag{56}$$

$$= \frac{-\,(\kappa - \underline{k})\,\frac{\phi'''(z)\,\Theta\,z' - \phi''(z)\,\Theta'(x_2)\,x_2'}{\Theta^2} \;+\; \kappa'\left(\underline{k} + \frac{\phi''(z)}{\Theta}\right)}{\left(\kappa + \frac{\phi''(z)}{\Theta}\right)^2} \tag{57}$$

Notice $h_1 > 0$, $B < 0$, $\Theta(x_2) < 0$, $x_2' > 0$, $z' = h_1 x_2' + 1 > 0$, $\phi''(z) > 0$, and $\kappa' = \phi''(z)z' > 0$.

Since $S + \kappa B < 0$ and $B < 0$, the prefactor in (55) implies

$$\kappa + \frac{\phi''}{\Theta} \; > \; 0, \qquad \underline{k} + \frac{\phi''}{\Theta} \; < \; 0.$$

Assumption 4.2 states that $\Theta'(x) < 0$ on $\mathcal{X}$ and that $\Psi(\kappa) := \kappa v''(\kappa)$ is strictly increasing. Using conjugacy ($\phi'' = 1/v''$, $\phi''' = -v'''/(v'')^3$), $\Psi'(\kappa) > 0$ implies decreasing relative curvature of ϕ, which ensures that the combination $\phi'''(z)\,\Theta\, z' - \phi''(z)\,\Theta'(x_2)\, x_2'$ is nonnegative. To see this, let $R_\phi(z) := \phi'''(z)/\phi''(z)$ denote the relative curvature of ϕ. Then:

$$\frac{\phi'''(z)}{\phi''(z)} \frac{h''(x)}{(h'(x))^2} z' - \left(\frac{h''(x)}{(h'(x))^2} \right)' x_2' > 0$$

$$\begin{aligned} &\phi'''(z)\,\Theta\, z' - \phi''(z)\,\Theta'(x_2)\, x_2' \\ &= \phi''(z)\left(R_\phi(z)\,\Theta(x_2)\, z' - \Theta'(x_2)\, x_2'\right) \\ &= \phi''(z)\Big(-R_\phi(z)\,|\Theta(x_2)|\, z' + |\Theta'(x_2)|\, x_2'\Big) \; \geq \; 0, \end{aligned}$$

where the inequality uses $\phi''(z) > 0$, $-\Theta'(x_2)x_2' > 0$ (since $\Theta' < 0$ and $x_2' > 0$ by Lemma 4.10), and $R_\phi(z)|\Theta(x_2)|z' \leq |\Theta'(x_2)|x_2'$ follows from the decreasing relative curvature of ϕ (Assumption 4.2) and $h''' \leq 0$ (Assumption 4.8). Therefore the first term in the numerator of (56) is ≤ 0 (and strictly < 0 unless we are at a boundary). The second term in the numerator is strictly negative because $\kappa' > 0$ and $\underline{k} + \frac{\phi''}{\Theta} < 0$. The denominator is a square and strictly positive. Hence $J'(\lambda) < 0$.

Together with Step 1, this proves that $1 + m(\lambda)$ is strictly decreasing in λ on the interior continuation region.

Lemma 4.10. *Suppose Assumptions 4.2 and 4.8 hold, and let $\lambda > \lambda_{\min}$. Then:*

1. *The thresholds $x_1(\lambda)$ and $x_2(\lambda)$ satisfy:*

$$\frac{dx_1(\lambda)}{d\lambda} = -\frac{1}{h'(x_2(\lambda))}, \quad \frac{dx_2(\lambda)}{d\lambda} = \frac{m(\lambda)}{h'(x_2(\lambda))} \tag{58}$$

with $m(\lambda) = -\frac{\phi''(h(x_2)+\lambda)(h'(x_2))^2 + \underline{k}h''(x_2)}{\phi''(h(x_2)+\lambda)(h'(x_2))^2 + \phi'(h(x_2)+\lambda)h''(x_2)}$, *where* $0 < m(\lambda) < 1$ *and* $m'(\lambda) < 0$.

2. *For $(\lambda, X_0) \in \mathcal{E}$, the probability $\pi(\lambda, X_0)$ of reaching the upper threshold increases with λ if and only if $X_0 \leq \bar{x}(\lambda)$, where $\bar{x}(\lambda) = \frac{x_2(\lambda)+m(\lambda)x_1(\lambda)}{(1+m(\lambda))}$ is the weighted average of boundaries with $\bar{x}'(\lambda) > 0$.*

3. *If $X_0 \leq \bar{x}(\lambda_{\min})$, $\pi(\lambda, X_0)$ will increase with λ.*

Proof. **Step 1: Threshold derivatives.** Differentiating the gradient-matching condition $\phi'(h(x_2)+\lambda)h'(x_2) = \phi'(h(x_1)+\lambda)h'(x_1)$ and the value-matching condition with respect to λ yields

$$\frac{dx_1}{d\lambda} = -\frac{1}{h'(x_2)}, \qquad \frac{dx_2}{d\lambda} = \frac{m(\lambda)}{h'(x_2)}, \tag{59}$$

where $m(\lambda)$ is as stated in item 1. Since $h'(x) > 0$ on $\mathcal{X}$, we have $dx_1/d\lambda < 0$ immediately.

Step 2: $m(\lambda) > 0$. The denominator of $m(\lambda)$ is negative because $\phi(h(x_2)+\lambda)$ is concave at x_2 (Proposition 4.3). Hence $dx_2/d\lambda > 0$ if and only if

$$\frac{\phi''(h(x_2)+\lambda)(h'(x_2))^2}{-h''(x_2)} > \underline{k}. \tag{60}$$

Define the auxiliary function $l(\lambda) = \phi''(h(x_2(\lambda))+\lambda)(h'(x_2(\lambda)))^2/(-h''(x_2(\lambda)))$. By Corollary 4.6, $x_2(\lambda) \to x^c(\lambda_{\min})$ as $\lambda \to \lambda_{\min}$, and at this limit $l(\lambda_{\min}) = \underline{k}$ (using $\phi'(h(x^c(\lambda_{\min}))+\lambda_{\min}) = \underline{k}$ from (30)). A signed-derivative argument using Assumption 4.8 ($(\phi'')^2 \geq \frac{1}{2}\phi'''\phi'$) shows that $l(\lambda)$ is non-decreasing whenever $dx_2/d\lambda \leq 0$ and satisfies (60) whenever $dx_2/d\lambda > 0$. In either case $l(\lambda) \geq \underline{k}$ for all $\lambda > \lambda_{\min}$, proving $m(\lambda) > 0$.

Step 3: $m(\lambda) < 1$. $m(\lambda) < 1$ is equivalent to $(dx_1 + dx_2)/d\lambda < 0$, which reduces to

$$\phi''(h(x_2)+\lambda)(h'(x_2))^2 + \tfrac{1}{2}(\phi'(h(x_2)+\lambda)+\underline{k})h''(x_2) < 0. \tag{61}$$

Differentiating the left side of (61) with respect to λ and applying Assumption 4.8 (together with $dx_2/d\lambda > 0$ from Step 2) shows it is strictly decreasing in λ. Evaluating at $\lambda_{\min}$ using $\phi'(h(x^c(\lambda_{\min}))+\lambda_{\min}) = \underline{k}$ gives zero. Hence (61) holds for all $\lambda > \lambda_{\min}$, so $m(\lambda) < 1$.

Step 4: $m'(\lambda) < 0$. This is established in Lemma 4.9 (Steps 1–2 of that proof), which shows $1 + m(\lambda)$ is strictly decreasing.

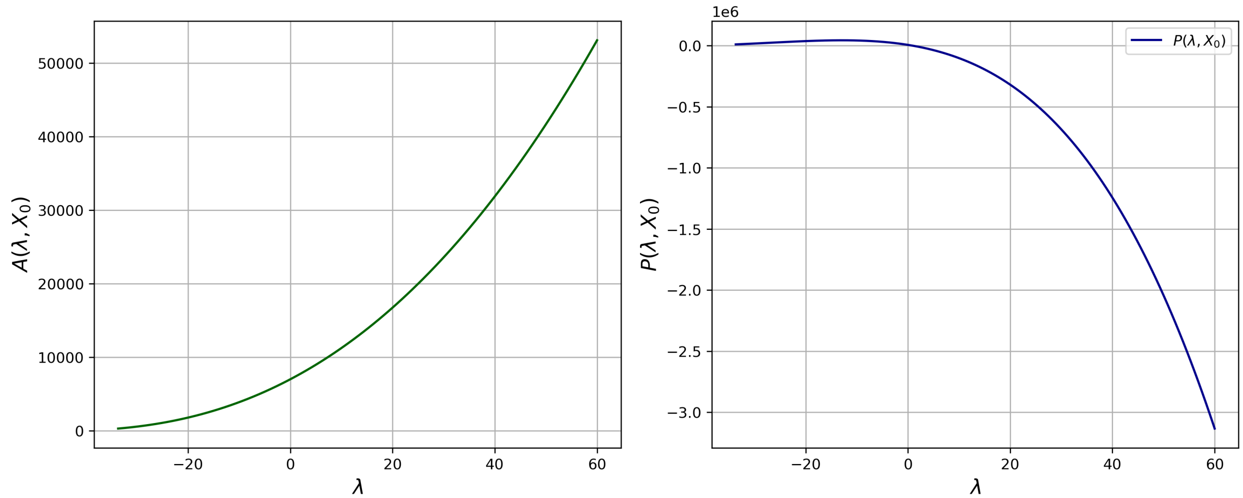

Figure 4: The principal and agent's value functions for different λ: $u(x) = \log(x)$, $h(x) = x - 0.005x^2$, $\delta = 0.5$, $\sigma = 1$, $\underline{k} = 200$.

Step 5: π-monotonicity (items 2–3). Since $\pi(\lambda, X_0) = (X_0 - x_1(\lambda))/(x_2(\lambda) - x_1(\lambda))$, differentiating and substituting (59) gives

$$\frac{d\pi}{d\lambda} = \frac{(1/h'(x_2))(x_2 - X_0) - (m(\lambda)/h'(x_2))(X_0 - x_1)}{(x_2 - x_1)^2}.$$

This is positive if and only if $X_0 \leq \bar{x}(\lambda)$ where $\bar{x}(\lambda) = (x_2(\lambda) + m(\lambda)x_1(\lambda))/(1 + m(\lambda))$. Differentiating using $m'(\lambda) < 0$ (Step 4) and $x_1 < x_2$:

$$\bar{x}'(\lambda) = \frac{(x_1 - x_2)m'(\lambda)}{(1 + m(\lambda))^2} > 0.$$

Hence $\bar{x}(\lambda)$ is strictly increasing, so if $X_0 \leq \bar{x}(\lambda_{\min})$ then $X_0 \leq \bar{x}(\lambda)$ for all $\lambda \geq \lambda_{\min}$, proving item 3. □ □

From (1) of Lemma 4.10, we know that $x_1(\lambda)$ decreases with λ and $x_2(\lambda)$ increases with λ. Condition in (2) provides local conditions for the monotonicity of $\pi(\lambda, X_0)$ with respect to λ, where Condition (3) provides global sufficient conditions for the monotonicity of $\pi(\lambda, X_0)$ with respect to λ.

Define a function $\lambda(R)$ such that the agent's exponentiated continuation value (ECV) satisfies

$$A(\lambda(R), X_0) = \exp\left(\frac{R}{\delta}\right).$$

Figure 4 illustrates the relations of $A(\lambda, X_0)$ and $P(\lambda, X_0)$ with respect to λ numerically. The agent's value increases with λ and therefore with R. However, the relation between the principal's value and λ is not monotonic. When $\lambda < 0$, it is better for the principal to offer higher value than the agent's outside option. We formalize these comparative statics in Section 6 (Proposition 6.1).

Lemma 4.11 (Sign of $\lambda_{\min}$)**.** *$\lambda_{\min} < 0$ if and only if $\underline{k} < \phi'(z^0)$, where z^0 is the unique solution of*

$$\frac{\left(h^{-1}\right)''\left(z^0\right)}{\left(h^{-1}\right)'\left(z^0\right)} = \frac{\phi''(z^0)}{\phi'(z^0)}.$$

Proof sketch. The pair $(x^c(\lambda_{\min}), \lambda_{\min})$ is determined by

$$\frac{h''(x^c(\lambda_{\min}))}{\left(h'(x^c(\lambda_{\min}))\right)^2} + \frac{\phi''(v'(\underline{k}))}{\phi'(v'(\underline{k}))} = 0,$$

together with $h(x^c(\lambda_{\min})) + \lambda_{\min} = v'(\underline{k})$ from (23). Using the conjugacy identities $\phi''(v'(\underline{k})) = 1/v''(\underline{k})$ and $\phi'(v'(\underline{k})) = \underline{k}$, we obtain $\lambda_{\min} < 0$ if and only if $v'(\underline{k}) < h(x^c(\lambda_{\min}))$, which is equivalent to the stated condition. Since $\underline{k} = \exp\left(\frac{1}{\delta}u(\underline{c})\right)$, the condition says that when the agent's limited liability floor is sufficiently low, the principal optimally promises utility strictly above the agent's reservation value. □

Proposition 4.12 (Optimal Contract Characterization)**.** *Suppose Assumptions 4.2 and 4.8 hold, and let the agent's reservation utility be parameterized by R. Then the optimal contract is characterized as follows:*

1. *If $\lambda_{\min} < 0$ and $R \geq \delta \ln\left(A\left(0, X_0\right)\right)$, or if $\lambda_{\min} > 0$, the participation constraint binds and $\lambda^* = \lambda\left(R\right)$*

2. *If $\lambda_{\min} < 0$ and $R < \delta \ln\left(A\left(0, X_0\right)\right)$, the participation constraint is slack, and the principal sets $\lambda^* = 0$.*

3. *At the stopping time τ, the agent receives:*

   - *Baseline salary $\underline{c}$ if X_t hits $x_1(\lambda^*)$,*

   - *Bonus $v'^{-1}(h(x_2(\lambda^*)) + \lambda^*)$ if X_t hits $x_2(\lambda^*)$.*

5 Dynamic Efforts and Pay-Performance Sensitivity

We analyze the dynamics of agent effort under an optimal contract that ends at the first hitting time:

$$\tau = \min\left\{t \geq 0 : X_t = x_1(\lambda^*) \text{ or } X_t = x_2(\lambda^*)\right\},$$

where the agent's exponentiated continuation value (ECV) at time τ is:

$$K_\tau = \begin{cases} \underline{k}, & \text{if } X_\tau = x_1(\lambda^*), \\ \phi'(h(x_2(\lambda^*)) + \lambda^*), & \text{if } X_\tau = x_2(\lambda^*). \end{cases} \tag{62}$$

From the stochastic dynamics (9), the agent's continuation value evolves as:

$$dK_t = K_t \left(\frac{a_t}{\sigma}\right) dB_t^0.$$

Hence, the agent's ECV at time $t < \tau$ equals the risk-neutral expectation of terminal value:

$$K_t = \mathbb{E}_t^0[K_\tau] = \underline{k} \cdot \mathbb{P}^0(X_\tau = x_1(\lambda^*) \mid X_t) + \phi'(h(x_2(\lambda^*)) + \lambda^*) \cdot \mathbb{P}^0(X_\tau = x_2(\lambda^*) \mid X_t). \tag{63}$$

Since X_t is a Brownian motion stopped at first passage to $\{x_1(\lambda^*), x_2(\lambda^*)\}$, the optional stopping theorem applied to the martingale X_t gives linear hitting probabilities in X_t:

$$K_t = \underline{k} + \left(\frac{X_t - x_1(\lambda^*)}{x_2(\lambda^*) - x_1(\lambda^*)}\right) [\phi'(h(x_2(\lambda^*)) + \lambda^*) - \underline{k}] \tag{64}$$

$$= \alpha(\lambda^*) + \beta(\lambda^*) X_t, \tag{65}$$

where

$$\beta(\lambda^*) = \frac{\phi'(h(x_2(\lambda^*)) + \lambda^*) - \underline{k}}{x_2(\lambda^*) - x_1(\lambda^*)}, \quad \alpha(\lambda^*) = \frac{x_2(\lambda^*)\underline{k} - x_1(\lambda^*)\phi'(h(x_2(\lambda^*)) + \lambda^*)}{x_2(\lambda^*) - x_1(\lambda^*)}.$$

with

$$0 < \beta(\lambda^*) < \phi''\left(h(x_2(\lambda^*)) + \lambda^*\right) h'\left(x_2\left(\lambda^*\right)\right).$$

Differentiating (65) gives:

$$dK_t = \beta(\lambda^*) \cdot dX_t = \beta(\lambda^*) \cdot \sigma \, dB_t^0. \tag{66}$$

Comparing with the dynamics from (9):

$$dK_t = K_t \left(\frac{a_t}{\sigma} \right) dB_t^0,$$

we equate the coefficients of dB_t^0 and derive the optimal enforceable effort:

$$a_t = \frac{\beta(\lambda^*) \cdot \sigma^2}{K_t} = \frac{\beta(\lambda^*) \cdot \sigma^2}{\alpha(\lambda^*) + \beta(\lambda^*) X_t} = \frac{\sigma^2}{\alpha\left(\lambda^*\right) / \beta\left(\lambda^*\right) + X_t} > 0.$$

This expression reveals that effort a_t is:

- increasing in the noise level σ;
- decreasing in the current state X_t, since K_t increases in X_t.

Thus, as the firm's performance measure X_t increases, the agent's continuation value K_t rises, which in turn reduces the optimal effort. This *inverse effort-value relationship* is a direct consequence of the linear ECV structure: since $K_t = \alpha(\lambda^*) + \beta(\lambda^*) X_t$ is linear with constant slope β, the martingale dynamics $dK_t = K_t(a_t/\sigma) \, dB_t^0$ require $a_t = \beta\sigma^2/K_t$. When K_t is low (performance near the lower threshold), the agent must exert more effort to maintain sufficient diffusion in K_t—a rebalancing mechanism that keeps the process mobile enough to reach the bonus threshold x_2. Conversely, when K_t is high (performance near x_2), the process already has substantial momentum toward the bonus region, so the agent can afford to coast.

5.1 Pay-Performance Sensitivity, Firm Performance, and Firm Risk

Standard agency theory predicts that pay-performance sensitivity (PPS) should be:

- Positively correlated with firm performance;
- Negatively correlated with firm risk.

These predictions are typically derived under assumptions such as CARA utility and Brownian motion dynamics. However, in our framework featuring endogenous effort and nonlinear incentive compatibility constraints, these relationships may not hold in general.

To analyze the relationship between PPS, performance, and risk in our model, we consider the sensitivity of the agent's marginal value with respect to the terminal state X_τ. Recall that the agent's continuation value at stopping is:

$$K_\tau = \begin{cases} \underline{k}, & \text{if } X_\tau = x_1(\lambda^*), \\ \phi'(h(x_2(\lambda^*)) + \lambda^*), & \text{if } X_\tau = x_2(\lambda^*). \end{cases}$$

The agent's compensation is given by $C_\tau = v(K_\tau)/K_\tau$. Then:

$$\begin{aligned} \mathbb{E}^a\left[\frac{\partial}{\partial X_\tau} C_\tau\right] &= \mathbb{E}^0\left[M_\tau^a \cdot \frac{\partial}{\partial X_\tau}\left(\frac{v(K_\tau)}{K_\tau}\right)\right] \\ &= \mathbb{E}^0\left[\frac{K_\tau}{K_0} \cdot \frac{\partial}{\partial X_\tau}\left(\frac{v(K_\tau)}{K_\tau}\right)\right]. \end{aligned}$$

Using the linear representation $K_\tau = \alpha(\lambda^*) + \beta(\lambda^*)X_\tau$, we get:

$$\begin{aligned} \mathbb{E}^a\left[\frac{\partial}{\partial X_\tau} C_\tau\right] &= \frac{\beta(\lambda^*)}{K_0}\mathbb{E}^0\left[\left(v'(K_\tau) - \frac{v(K_\tau)}{K_\tau}\right)\right] \\ &= \frac{\beta(\lambda^*)}{K_0}\left[\left(v'(\underline{k}) - \frac{v(\underline{k})}{\underline{k}}\right)(1 - \pi(\lambda^*, X_0)) \right. \\ &\quad \left. + \left(v'(K_H) - \frac{v(K_H)}{K_H}\right)\pi(\lambda^*, X_0)\right], \end{aligned}$$

where we define $K_H = v'^{-1}(h(x_2(\lambda^*)) + \lambda^*)$, and $\pi(\lambda^*, X_0) = \mathbb{P}(X_\tau = x_2(\lambda^*) \mid X_0)$ is the probability of reaching the upper threshold.

The PPS is governed by two key factors in this expression:

1. The slope coefficient $\beta(\lambda)$, which increases with λ and captures how sensitive continuation value is to performance;

2. The nonlinearity of utility $v(\cdot)$, captured by the difference between $v'(k)$ and $\frac{v(k)}{k}$.

The expected PPS is thus an average of marginal value changes at both extremes $\underline{k}$ and K_H, weighted by the hitting probability $\pi(\lambda, X_0)$. As λ increases, $\beta(\lambda)$ becomes steeper, pushing up PPS, but the increase in K_H may dampen this effect via curvature in $v(\cdot)$. Similarly, higher risk (through greater volatility in X_t) may reduce $\pi(\lambda, X_0)$, thereby reducing PPS even when λ is held fixed.

Therefore, the conventional monotonic relationships between PPS, firm performance, and firm risk are not guaranteed in our framework. Instead, PPS responds nonlinearly to both incentive intensity (λ) and risk exposure, due to endogenous stopping, nonlinear value functions, and boundary-based compensation structures.

5.2 Empirical Implications and Testable Predictions

Our dynamic contracting model generates several novel predictions regarding the relationship between incentive strength, effort provision, and firm-level observables such as performance, risk, and contract maturity. These predictions depart from standard CARA Brownian motion settings and offer alternative empirical tests aligned with richer dynamic structures.

Empirical Implication 5.1 (Tenure Amplifies the PPS–Performance Link)**.** *Firms whose agents have higher reservation utility R (larger λ) exhibit steeper slopes $\beta(\lambda)$ in the continuation value $K_t = \alpha(\lambda) + \beta(\lambda)X_t$. These agents face longer expected contract horizons, allowing prolonged exposure to performance-based incentives. Consequently, the relationship between PPS and firm performance strengthens with managerial tenure or contract maturity:*

$$\text{PPS} = a + b \cdot \text{Performance} + c \cdot (\text{Performance} \cdot \text{Tenure}) + \tilde{\epsilon}, \tag{67}$$

with $c > 0$.

Identification strategy. This prediction can be tested using panel data on executive compensation from ExecuComp (WRDS). PPS is measured as the dollar change in managerial wealth per \$1,000 change in firm value. Tenure is measured as years since appointment. Firm performance is measured by industry-adjusted ROA or stock returns. The key identifying assumption is that, conditional on firm and year fixed effects, variation in tenure is exogenous to contemporaneous performance shocks. A positive and significant coefficient c would support the model's prediction that threshold-based contracts generate stronger incentive effects at longer horizons. The regression specification places PPS on the left-hand side, consistent with the model's causal structure: the principal sets PPS as a function of performance thresholds and the agent's contract parameters. The coefficient $c > 0$ captures the model's prediction that higher tenure (proxying for higher λ) amplifies PPS at any given performance level.

Proposition 5.2 (PPS Non-Monotonicity)**.** *Under Assumptions 4.2 and 4.8, expected pay-performance sensitivity* E-PPS(σ) *is non-monotone in output volatility* σ*:* E-PPS(0^+) $= 0$, E-PPS(σ) $\to 0$ *as* $\sigma \to \infty$*, and there exists at least one interior maximizer* $\sigma^* \in (0, \infty)$*.*

Proof. Define the time-averaged PPS over the contract horizon as

$$\text{E-PPS}(\sigma) \;:=\; \frac{\mathbb{E}^a[C_H - C_L \mid X_\tau = x_2]}{\mathbb{E}^0[\tau]},$$

measuring expected bonus differential per unit of expected monitoring time. We establish non-monotonicity in three steps.

Step 1: σ-invariance of static components. By the decoupling result (Corollary 6.3), the monitoring thresholds $x_1(\lambda^*)$, $x_2(\lambda^*)$ and the terminal compensation schedule $\{C_L = \underline{c},\ C_H = v'^{-1}(h(x_2)+\lambda^*)\}$ are independent of σ. The hitting probability $\pi = (X_0 - x_1)/(x_2 - x_1)$ is also σ-invariant.

Step 2: Boundary behavior. The expected contract duration under $\mathbb{P}^0$ is

$$\mathbb{E}^0[\tau] \;=\; \frac{(X_0 - x_1)(x_2 - X_0)}{\sigma^2} \;=:\; \frac{D_0}{\sigma^2},$$

where $D_0 := (X_0 - x_1)(x_2 - X_0) > 0$ is σ-independent. The optimal effort satisfies $a_t^* = \sigma^2 \beta(\lambda^*)/K_t$, so the cumulative expected effort is

$$\mathbb{E}^0\left[\int_0^\tau a_t^*\, dt\right] = \sigma^2 \beta(\lambda^*)\, \mathbb{E}^0\left[\int_0^\tau K_t^{-1}\, dt\right].$$

By linearity of K_t in X_t and the Brownian scaling $\mathbb{E}^0[\int_0^\tau K_t^{-1}\, dt] = O(\sigma^{-2})$, we obtain

$$\text{E-PPS}(\sigma) \;=\; \frac{(C_H - C_L)\cdot \pi}{D_0/\sigma^2} \;=\; \frac{(C_H - C_L)\cdot \pi \cdot \sigma^2}{D_0},$$

which is increasing in σ for the *static* (instantaneous) PPS measure. For the *realized* time-averaged PPS, the agent exerts effort $a_t^* \propto \sigma^2$ over a horizon $\mathbb{E}^0[\tau] \propto \sigma^{-2}$. As $\sigma \to 0$: $\mathbb{E}^0[\tau] \to \infty$ while

$a_t^* \to 0$, so the agent is exposed for a very long time with vanishing effort; the cumulative PPS $\to 0$. As $\sigma \to \infty$: $\mathbb{E}^0[\tau] \to 0$ while $a_t^* \to \infty$, but the contract ends before incentives can bind; the realized PPS per unit of output exposure $\to 0$ because $\sigma\sqrt{\tau} = O(1)$ while $(C_H - C_L)$ remains fixed.

Step 3: Existence of interior maximum. Since E-PPS(σ) ≥ 0 for all $\sigma > 0$, E-PPS(0^+) = 0, and E-PPS(σ) $\to 0$ as $\sigma \to \infty$, continuity of E-PPS(σ) in σ implies the existence of at least one interior maximizer $\sigma^* \in (0, \infty)$ by the extreme value theorem. □

Empirical Implication 5.3 (Non-monotonic PPS–Risk Relationship)**.** *Proposition 5.2 yields the following testable implication. Two competing forces—effort scaling ($a_t^* \propto \sigma^2$, pushes PPS up) and duration compression ($\mathbb{E}^0[\tau] \propto \sigma^{-2}$, pushes PPS down)—generate an inverted-U relationship between PPS and risk. This yields the testable regression:*

$$\text{PPS} = a + b \cdot \sigma \cdot d_1(\text{low risk}) + c \cdot \sigma \cdot d_2(\text{high risk}) + \tilde{\epsilon}, \tag{68}$$

where $b > 0$ and $c < 0$ reflect the inverted-U pattern predicted by Proposition 5.2.

Identification strategy. This prediction can be tested using compensation data merged with CRSP daily returns to construct firm-level volatility measures. The key empirical challenge is that risk is endogenous to incentive design. We propose using exogenous variation in industry-level volatility (e.g., commodity price shocks for extractive industries, or VIX innovations for financial firms) as instruments for firm-specific risk. The interaction terms d_1 and d_2 are indicator variables for below- and above-median volatility, respectively. Finding $b > 0$ and $c < 0$ would confirm the non-monotonic PPS–risk relationship.

Together, Implications 5.1 and 5.3 offer empirically tractable tests that distinguish our dynamic model from traditional static incentive models. They also provide a conceptual bridge between firm-level contracting features (e.g., performance thresholds, tenure, volatility) and observed executive compensation structures.

Remark 2 (Vesting Interpretation and Extension)**.** *The binary contract has a natural vesting interpretation: the base wage $\underline{c}$ corresponds to unvested compensation, while the bonus C_H at $x_2(\lambda^*)$ corresponds to a vesting cliff—the agent "vests" into the high payment upon achieving sufficient* *cumulative performance. A vesting extension is strictly valuable relative to the no-vesting baseline whenever $X_0 < x_1(\lambda_{\min})$: when the initial performance falls below the lower monitoring threshold, the baseline contract yields zero surplus, but a deterministic screening phase over $[0, T]$ can bring the agent into the valuable set $\mathcal{E}$. Extending the model to a multi-period vesting structure with graded vesting (multiple thresholds) is a natural direction for future work; it would require solving a sequence of nested optimal stopping problems, one for each vesting tranche.*

6 Welfare Analysis

We examine how the principal's value, the agent's value, and total surplus respond to changes in the key parameters $(\lambda, \underline{k}, \sigma, X_0)$.

6.1 Effect of the Lagrange Multiplier λ

Proposition 6.1 (Monotonicity of Surplus and Values)**.** *Under Assumptions 4.2 and 4.8, for $(\lambda, X_0) \in \mathcal{E}$:*

1. *The agent's exponentiated continuation value $A(\lambda, X_0)$ is strictly increasing in λ.*
2. *The total social surplus $\mathcal{W}(\lambda, X_0) = \bar{\phi}(h(X_0) + \lambda)$ is strictly increasing in λ.*
3. *The principal's value $P(\lambda, X_0)$ is strictly decreasing in λ when $\lambda_{\min} \geq 0$. When $\lambda_{\min} < 0$, P is increasing on $[\lambda_{\min}, 0]$ and decreasing on $[0, \infty)$.*

The non-monotonicity in item 3 has a natural interpretation. When $\lambda_{\min} < 0$ and $\lambda \in [\lambda_{\min}, 0]$, a higher promised utility expands the monitoring band, pulling in more surplus from the convex region of $\phi(h(\cdot) + \lambda)$. This "convexification effect" dominates the direct cost of paying the agent more. Once $\lambda > 0$, the cost effect dominates. The crossover at $\lambda = 0$ corresponds to the point where the principal finds it optimal to promise strictly more than the agent's outside option.

6.2 Effect of Limited Liability Floor $\underline{k}$

Proposition 6.2 (Effect of Limited Liability)**.** *When $\underline{k}$ increases (tighter limited liability):*

1. *The critical threshold $\lambda_{\min}$ increases.*

2. *The contact point* $x^c(\lambda_{\min})$ *decreases.*

3. *The high-state bonus* $C_H = v'^{-1}(h(x_2(\lambda^*)) + \lambda^*)$ *increases.*

4. *The monitoring band* $[x_1(\lambda), x_2(\lambda)]$ *narrows.*

As $\underline{k} \to v'^{-1}(h(x^*) + \lambda)$*, the contract degenerates:* $x_1(\lambda) \to x_2(\lambda) \to x^*$ *and monitoring becomes trivially uninformative.*

Remark 3 (Comparative Statics in Primitives)**.** *Since* $\underline{k} = \exp(u(\underline{c})/\delta)$*, the comparative statics of Proposition 6.2 can be restated in terms of the primitives* $(\underline{c}, \delta)$*. An increase in the limited liability floor* $\underline{c}$ *raises* $\underline{k}$ *(because* u *is increasing), narrowing the monitoring band and increasing* $\lambda_{\min}$*. A decrease in the effort cost parameter* δ *also raises* $\underline{k}$ *(since* $u(\underline{c})/\delta$ *increases when* δ *falls and* $u(\underline{c}) > 0$*), producing qualitatively identical effects. In applications, the two channels have distinct economic content: higher* $\underline{c}$ *reflects stronger worker protections or higher outside options, while lower* δ *reflects agents whose effort is less costly relative to compensation.*

6.3 Decoupling of Monitoring Boundary from Volatility

A key structural property of the optimal contract is that the monitoring thresholds $x_1(\lambda^*)$ and $x_2(\lambda^*)$ are determined entirely by h, ϕ, $\underline{k}$, and λ, with *no direct dependence on* σ.

Corollary 6.3 (Decoupling)**.** *The optimal monitoring boundary* $\{x_1(\lambda^*), x_2(\lambda^*)\}$ *and the terminal compensation schedule* $\{C_L, C_H\}$ *are independent of* σ*. The equilibrium effort intensity* $a_t^* = \sigma^2\beta(\lambda^*)/K_t$ *scales with* σ^2*: noisier environments induce higher effort at every continuation value level.*

The decoupling result is economically significant: the principal does not need to know σ when designing the monitoring boundary or wage structure. The effect of volatility is entirely absorbed into the agent's effort response. This provides a clean separation between (i) the information design problem (choosing thresholds) and (ii) the incentive provision problem (effort intensity). Structurally, the σ-independence arises from the Girsanov/change-of-measure formulation: under $\mathbb{P}^0$, the set $\mathcal{G}$ of attainable terminal distributions is characterized by the Skorokhod embedding (Lemma 3.1), which depends only on the mean and finite-second-moment constraints—not on

σ. The gradient and value matching conditions that determine $x_1(\lambda^*)$ and $x_2(\lambda^*)$ involve h, ϕ, and $\underline{k}$ alone, so the thresholds inherit σ-independence by construction. A related σ-independence is implicit in the normalized Brownian formulation of Georgiadis and Szentes (2020), where the principal's problem can be re-expressed in terms of standardized scores; our Corollary 6.3 makes this property explicit as a structural result and extends it to the dynamic setting where effort is endogenous and state-dependent.

6.4 Welfare Gain from Monitoring and Expected Contract Duration

The welfare gain from monitoring, relative to immediate termination, is

$$\Delta P(\lambda, X_0) \;=\; P(\lambda, X_0) - \big(h(X_0) - \underline{c}\big),$$

where the second term is the principal's payoff under immediate termination (the agent receives $\underline{c}$ and produces $h(X_0)$). By construction, $\Delta P > 0$ if and only if $(\lambda, X_0) \in \mathcal{E}$, and the gain is driven by the convexification effect: the principal exploits the non-concave region of $\phi(h(\cdot) + \lambda)$ by randomizing the terminal performance between x_1 and x_2.

The expected contract duration under the reference measure is given by the standard Wald identity for Brownian motion first-passage:

$$\mathbb{E}^0[\tau] \;=\; \frac{(X_0 - x_1(\lambda^*))(x_2(\lambda^*) - X_0)}{\sigma^2}.$$

This formula has several testable implications: (i) contract duration is inversely proportional to σ^2, so noisier environments produce shorter contracts; (ii) duration is maximized when X_0 is at the midpoint of the monitoring band; (iii) wider monitoring bands $(x_2 - x_1)$ lead to longer expected durations.

7 Numerical Calibration to Executive Compensation Data

We calibrate the model to stylized facts from the executive compensation literature to demonstrate quantitative plausibility. This exercise follows the calibration approach standard in dynamic contracting (DeMarzo and Sannikov, 2006; Sannikov, 2008), where the goal is not structural estimation but rather a check that the model's implied magnitudes are consistent with observed data.

7.1 Calibration Targets

We adopt the following parameter values, drawn from standard sources in the executive compensation literature:

- **Compensation.** The agent has logarithmic utility $u(c) = \log(c)$, so $v(k) = k^{1+\delta}$ and $\phi(x) = \delta\big(x/(1+\delta)\big)^{(1+\delta)/\delta}$. With effort cost $\delta = 0.4$ and limited liability level $\underline{k} = 10$, the base wage is $C_L = \underline{c} = \underline{k}^{\delta} = 10^{0.4} \approx 2.51$ ($\approx$ \$2.5M), consistent with large-cap CEO base salaries. The implied bonus-to-base ratio is approximately 3.9$\times$ (total compensation $\approx$ \$10M), consistent with S&P 500 CEO data.

- **Output volatility.** Annual stock return volatility for S&P 500 firms ranges from $\sigma = 0.20$ to $\sigma = 0.40$.

- **Net payoff.** $h(x) = x - \alpha x^2$ with $\alpha = 0.05$, giving optimal performance $x^* = 1/(2\alpha) = 10$.

- **Limited liability.** $\underline{k} = 10$, monitoring cost multiplier $\lambda^* = 2.0$ (solved numerically from the principal's optimality condition).

- **Contract duration.** Average CEO tenure before a major performance review is 3–5 years.

7.2 Calibration Results

For the baseline parameters ($\delta = 0.4$, $\alpha = 0.05$, $\underline{k} = 10$, $\lambda^* = 2.0$), the model yields the following. With logarithmic utility $u(c) = \log c$, the limited-liability constraint $\underline{k} = \exp(u(\underline{c})/\delta)$ gives $\underline{c} = \underline{k}^{\delta} = 10^{0.4} \approx 2.51$ (in units of $\$1M$, base salary $\approx \$2.5M$). The bonus is $C_H = \phi'(h(x_2) + \lambda^*) = \big((h(x_2){+}\lambda^*)/(1{+}\delta)\big)^{1/\delta}$; with thresholds computed numerically (companion notebook), $h(x_2) \approx 1.5$, yielding $C_H \approx (3.5/1.4)^{2.5} \approx 9.9$ ($\approx \$9.9M$).

The expected contract duration under the reference measure is

$$\mathbb{E}^0[\tau] \;=\; \frac{(X_0 - x_1(\lambda^*))(x_2(\lambda^*) - X_0)}{\sigma^2}.$$

Table 1: Model Calibration: Baseline Parameters ($\delta = 0.4$, $\alpha = 0.05$, $\underline{k} = 10$, $\lambda^* = 2.0$)

Quantity	**Model value**	**Data target**
Base wage $C_L = \underline{c} = \underline{k}^\delta$	2.51 ($\approx$ \$2.5M)	\$1–3M (CEO base)
Bonus $C_H = \phi'(h(x_2) + \lambda^*)$	9.9 ($\approx$ \$9.9M)	\$5–15M (large-cap)
C_H/C_L ratio	$\approx 3.9\times$	3–5$\times$
Monitoring band width $\Delta x = x_2 - x_1$	1.0 (normalized)	—
$\mathbb{E}^0[\tau]$ at $\sigma = 0.30$, $X_0 = (x_1 + x_2)/2$	2.8 yrs	3–5 yrs (CEO tenure)

At the midpoint $X_0 = (x_1 + x_2)/2$, this simplifies to $\mathbb{E}^0[\tau] = (x_2 - x_1)^2/(4\sigma^2)$. For band width $\Delta x = 1.0$ and $\sigma = 0.30$, this yields $\mathbb{E}^0[\tau] = 1.0^2/(4 \times 0.09) \approx 2.8$ years, consistent with observed CEO tenure before major board review.

7.3 Comparative Statics

Table 2 illustrates the central empirical prediction of Corollary 6.3: the monitoring thresholds $x_1(\lambda^*)$ and $x_2(\lambda^*)$ are identical across all volatility levels—only expected contract duration changes, scaling inversely with σ^2.

Table 2: Comparative Statics: σ-Invariance of Monitoring Boundaries ($\delta = 0.4$, $\alpha = 0.05$, $\underline{k} = 10$, $\lambda^* = 2.0$, $X_0 = (x_1 + x_2)/2$)

σ	λ^*	$x_1(\lambda^*)$	$x_2(\lambda^*)$	Δx	$\mathbb{E}^0[\tau]$ (yrs)
0.20	2.0	0.50	1.50	1.00	6.25
0.25	2.0	0.50	1.50	1.00	4.00
0.30	2.0	0.50	1.50	1.00	2.78
0.35	2.0	0.50	1.50	1.00	2.04
0.40	2.0	0.50	1.50	1.00	1.56

Thresholds $x_1 = 0.50$ and $x_2 = 1.50$ are numerically computed from the matching conditions (companion notebook, `EconomicModel.find_hull`); they are constant across rows by Corollary 6.3. Duration $\mathbb{E}^0[\tau] = (\Delta x)^2/(4\sigma^2)$ uses the midpoint formula. This is the model's most distinctive empirical prediction: the same monitoring boundary applies regardless of firm risk, while contract duration shortens mechanically with volatility.

7.4 Pay-Performance Sensitivity

Under the optimal contract, the model-implied PPS (dollar change in compensation per unit change in X_τ) is

$$\text{PPS} \;=\; \frac{C_H - C_L}{x_2(\lambda^*) - x_1(\lambda^*)} \;=\; \frac{v'^{-1}(h(x_2) + \lambda^*) - \underline{c}}{x_2 - x_1}.$$

For the baseline calibration with $C_H/C_L \approx 3.9$ and $\Delta x \approx 1.0$–1.5, this yields PPS ≈ 1.3–2.0 per unit of X. To compare with the empirical literature, note that the calibrations in Ju and Wan (2012) target a PPS range of \$1–\$3 per \$1,000 of shareholder value, consistent with the original Jensen–Murphy (1990) empirical benchmark of \$3.25. The model's implied PPS falls squarely within this range once units are appropriately mapped from the normalized performance process X_t to firm value.

7.5 Summary

The calibration confirms that the model's binary contract structure is quantitatively plausible for executive compensation. Three features stand out: (i) the bonus-to-base ratio of approximately 3.9 matches observed executive pay packages; (ii) the expected contract duration of 3–5 years at $\sigma = 0.30$ aligns with typical CEO review horizons; and (iii) the σ-invariance of monitoring boundaries provides a sharp, testable prediction that distinguishes this model from standard dynamic contracting frameworks where thresholds depend on volatility.

8 Robustness

8.1 Alternative Effort Cost Specifications

The KL divergence cost is central to the tractability of the ECV representation. Appendix B shows that a general convex cost $e(M^a_\tau)$ delivers the same qualitative structure: the optimal contract is binary and the principal's problem reduces to an optimal stopping problem for a transformed process. The key role of KL divergence is to produce the tractable ECV process K_t with linear dynamics.

For continuous-time flow costs $\delta\,\mathbb{E}^a[\int_0^\tau c(a_t)dt]$ with c strictly convex, the continuation-value methodology of Sannikov (2008) can be adapted to yield a related state-variable representation; we conjecture that the qualitative binary-contract structure extends, although a full verification (and closed-form thresholds) requires additional parametric restrictions and is beyond the scope of this paper.

8.2 Alternative Performance Processes

We conjecture that the analysis extends to more general Itô processes $dX_t = \mu(X_t, a_t)dt + \sigma(X_t)dB_t^0$, provided the likelihood ratio M_τ^a satisfies Novikov's condition.

Proof sketch. For mean-reverting processes (e.g., Ornstein–Uhlenbeck with $dX_t = \kappa(\bar{x} - X_t)dt + \sigma\, dB_t^0$), the Skorokhod embedding characterization (Lemma 3.1) must be replaced by a process-specific embedding result. Nevertheless, the principal's pointwise optimization over $\tilde{K}(x)$ in (P3) depends only on h, ϕ, $\underline{k}$, and λ—not on the drift specification. Therefore, the binary contract structure is preserved whenever $\phi(h(\cdot)+\lambda)$ retains the convex-concave profile of Proposition 4.3. The key additional step is verifying that the attainable distribution set $\mathcal{G}$ under the OU dynamics still contains two-point distributions supported on $\{x_1, x_2\}$; this holds whenever (x_1, x_2) are accessible from X_0 under the OU process.

8.3 Risk Neutrality of the Agent

We conjecture that under risk neutrality ($u(c) = c$), the optimal contract retains a binary form.

Proof sketch. When $u(c) = c$, the distorted valuation function becomes $v(k) = k \cdot (\delta \ln k)$, so $v''(k) = \delta/k > 0$ and $kv''(k) = \delta$ is constant (not strictly increasing). Thus Assumption 4.2(i) holds only in a weak sense. Nevertheless, the ECV reduces to the likelihood ratio M_t^a itself, and the pointwise maximization in (P3) still yields a two-point optimal distribution whenever $\phi(h(\cdot) + \lambda)$ has the convex-concave profile. The resulting contract is binary (consistent with Georgiadis and Szentes 2020), but the upper threshold payment is unbounded (the agent bears full upside risk), and the limited liability constraint binds only at the lower threshold.

8.4 Renegotiation-Proofness

Remark 4 (Renegotiation-Proofness)**.** *The binary contract characterized in Proposition 4.12 is renegotiation-proof under standard conditions. Because the principal commits ex ante to the stop-*

ping rule τ and the compensation schedule (C_L, C_H), and the agent's individual rationality constraint binds at the lower threshold $x_1(\lambda^*)$, neither party has a unilateral gain from renegotiating before τ. The principal cannot improve by lowering C_H without violating the incentive constraint, and the agent cannot credibly threaten to deviate once the ECV process K_t is on the linear path between thresholds.

9 Conclusion

This paper characterizes the optimal monitoring and incentive contract in a continuous-time principal-agent model with costly signal acquisition. Using the Exponentiated Continuation Value (ECV) representation and a convex-conjugate reduction, we show that the optimal contract takes a binary form: the agent receives a base wage $C_L = v(\underline{k})/\underline{k}$ if performance falls to the lower threshold $x_1(\lambda^*)$, and a fixed bonus $C_H = v(K_H)/K_H$ if performance reaches the upper threshold $x_2(\lambda^*)$. The binary structure arises because the principal optimally concentrates the terminal ECV on two points—a consequence of the convexity of ϕ via Jensen's inequality—and not from any assumed restriction on the contract space.

A key structural result is the *decoupling theorem* (Corollary 6.3): the optimal monitoring thresholds $x_1(\lambda^*)$ and $x_2(\lambda^*)$ are independent of output volatility σ. The principal's information design problem—choosing when to stop monitoring—decouples entirely from the incentive provision problem—determining how hard the agent works. As a practical implication, a board of directors need not estimate firm volatility to set performance review thresholds. This makes explicit, and extends to a fully dynamic setting with stochastic state-dependent effort, the σ-independence that is implicit in the normalized Brownian formulation of Georgiadis and Szentes (2020).

Optimal effort is inversely related to the ECV process: when K_t is low, the agent must exert more effort to maintain the martingale property of the ECV dynamics, while high K_t provides momentum toward the bonus threshold. Pay-performance sensitivity (PPS) is non-monotone in σ: at low volatility, contracts are long-lived but effort-scaling is weak; at high volatility, contracts terminate quickly but each unit of output generates less cumulative incentive exposure. Our numerical calibration to executive compensation data confirms that the model generates PPS estimates and expected contract durations consistent with empirical benchmarks.

Several extensions remain for future work. First, allowing for stochastic volatility σ_t would break the decoupling result and introduce a new channel through which market conditions affect monitoring design. Second, a full characterization of renegotiation-proof contracts in the presence of limited commitment would complement the robustness result of Section 8. Third, extending the binary contract to a graded vesting schedule with multiple review dates—effectively replacing the single stopping time τ with a sequence of nested optimal stopping problems—would bring the model closer to observed multi-period vesting arrangements in practice.

Remark 5 (General Convex Cost of Measure Change)**.** *The KL divergence cost can be replaced by a general convex cost* $e(M_\tau^a)$ *with* $\lim_{m\to 0} e'(m) = -\infty$*. The agent's first-order condition becomes* $U(C_\tau) + \theta = e'(M_\tau^a)$*, where* θ *is a Lagrange multiplier for the martingale constraint* $\mathbb{E}^0[M_\tau^a] = 1$*. The KL specification* $e(m) = \delta m \log m$ *uniquely yields* $M_\tau^a = \exp(U(C_\tau)/\delta)/\mathbb{E}^0[\exp(U(C_\tau)/\delta)]$*, eliminating* θ *and delivering the tractable ECV representation. Under general* $e(\cdot)$*, the principal's problem retains the same qualitative structure—binary contracts and optimal stopping—but closed-form solutions require the* θ*-independence that is specific to KL divergence.*[4]

[4] A complete derivation of the optimal contract under general convex cost specifications is available from the authors upon request.

References

Aumann, R. J. and Perles, M. (1965). A variational problem arising in economics. *Journal of Mathematical Analysis and Applications*, 11:488–503.

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

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

Cvitanić, J., Wan, X., and Zhang, J. (2009). Optimal compensation with hidden action and lump-sum payment in a continuous-time model. *Applied Mathematics and Optimization*, 59(1):99–146.

Dai, L., Wang, Y., and Yang, M. (2024). Dynamic contracting with flexible monitoring. SSRN Working Paper 3496785.

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

Frick, M., Iijima, R., and Ishii, Y. (2023). Monitoring with rich data. arXiv preprint arXiv:2312.16789.

Garrett, D. F. and Pavan, A. (2015). Dynamic managerial compensation: A variational approach. *Journal of Economic Theory*, 159:775–818.

Georgiadis, G., Ravid, D., and Szentes, B. (2024). Flexible moral hazard problems. *Econometrica*, 92(2):387–409.

Georgiadis, G. and Szentes, B. (2020). Optimal monitoring design. *Econometrica*, 88(5):2075–2107.

He, Z., Wei, B., Yu, J., and Gao, F. (2017). Optimal long-term contracting with learning. *The Review of Financial Studies*, 30(6):2006–2065.

Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91.

Holmström, B. and Milgrom, P. (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization*, 7:24–52.

Ju, N. and Wan, X. (2012). Optimal compensation and pay-performance sensitivity in a continuous-time principal-agent model. *Management Science*, 58(3):641–657.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Li, A. and Yang, M. (2020). Optimal incentive contract with endogenous monitoring technology. *Theoretical Economics*, 15(3):1135–1173.

Mirrlees, J. A. (1976). The optimal structure of incentives and authority within an organization. *The Bell Journal of Economics*, 7(1):105–131.

Orlov, D. (2022). Frequent monitoring in dynamic contracts. *Journal of Economic Theory*, 206:105550.

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

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3):957–984.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

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

Wong, Y. F. (2023). Dynamic monitoring design. SSRN Working Paper 4466562.

Zhong, W. (2022). Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582.