informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0025-1909, EISSN 1526-5501

Submitted to *Management Science*

Dynamic Incentives in Service Contracting

Xu Sun
University of Miami Herbert Business School, Coral Gables, Florida 33146, USA, xxs767@miami.edu

Xiaohan Zhu
School of Management and Engineering, Nanjing University, Nanjing, China, zhuxiaohan@nju.edu.cn

Abstract. This paper addresses the design of a performance-based compensation scheme (or "contract") for a company that outsources its customer service to a third-party provider. The service process is modeled as a queue, with the service provider acting as the single server who can adjust the service rate over time. The provider operates at a base service rate when paid a fixed base wage but can increase the service rate above this base by incurring a quadratic effort cost. The company faces congestion costs due to delays in customer service. Although the company cannot directly observe the provider's actual service rate, it can encourage desired service rate adjustments by linking incentive payments to the workload, which is assumed to be contractible. The company's goal is to find a contract that minimizes the total cost, comprising both the congestion costs and the incentive payments. To derive clear insights, we model the workload process as a reflected Brownian motion with an adjustable drift rate. Using the calculus of variations, we determine the necessary conditions for an optimal contract. When congestion costs are quadratic, these conditions yield a contract that is explicit and intuitive. We also explore a scenario where the service provider is protected by limited liability. We analytically characterize the optimal contract in this context, highlighting how limited liability introduces contracting frictions.

Key words: service contracting; principal-agent framework; calculus of variation; Brownian queues

1. Introduction

Companies often choose to outsource their customer service operations to third-party providers (Zhou and Ren 2010). In 2016, the outsourced customer service market was valued at over 70 billion USD and is projected to grow at an annual rate of approximately 6% from 2017 to 2024. However, outsourcing customer service poses a significant challenge for business owners due to the difficulty in overseeing the commitment level of these providers, who might prioritize their own objectives. Unlike an in-house team, business owners have less control over the service quality rendered by outsourced providers. This creates a need for performance-based compensation schemes to motivate third-party providers to deliver high-quality service.

For example, continuous monitoring of call handling and managerial control, often implemented as an incentive-penalty system, is common in outsourced call centers (Poster 2007). Consider a situation where a company hires a freelance call center agent to handle customer requests. To align the agent's incentives with the company's objectives, a bonus-malus arrangement can be used, with payments based on agreed-upon performance metrics. If both parties can monitor or reliably estimate the workload (e.g., the time required to clear the current backlog of work or the number of jobs in the queue), incentive payments can be structured based on this contractible metric.[1] In particular, when the workload is minimal, implying that the agent is dedicating adequate effort to providing efficient customer service, the company could award a sizable bonus to the agent. Conversely, if the workload becomes excessive, the company might consider withholding the bonus or even imposing a malus payment on the agent to compensate for the service degradation. Another representative scenario involving a utility company outsourcing its IT services to a vendor was documented in Akkermans et al. (2019), where the authors consider buyer-supplier contracting using collaborative key performance indicators. In the described contractual relationship, the number of open tickets (i.e., requests yet to be resolved) serves as an observable and verifiable metric based on which a bonus-malus arrangement can be constructed. In essence, the provision of incentive payments holds the potential to incentivize the service provider to operate more efficiently, thereby granting the service "buyer" indirect influence over the customer service process.

To gain clear insights into the design of dynamic incentives in this context, we introduce a stylized model that combines the principal-agent framework with a simple queueing system. In our model, the company acts as the principal ("she"), and the service provider acts as the agent ("he"). The agent receives a base wage that is just competitive enough to incentivize him to participate and provide the desired base service rate. Under this base rate, the workload evolves as a reflected Brownian motion (RBM), representing the virtual wait time for incoming jobs. The agent can adjust the service rate, impacting the workload's drift rate, but this comes at a cost. The principal incurs congestion costs tied to the workload. While direct observation of the agent's effort is not possible, the principal can infer it from the workload. By linking payments to the workload, the principal can incentivize the agent to manage the speed of service according to her operational objectives. Specifically, she aims to motivate the agent to speed up as the workload increases. Although this may increase the agent's effort costs in the short run, it can reduce congestion levels and lead to long-term rewards for the agent under a performance-based compensation scheme.

The agent's decision-making process involves seeking a rate-control strategy to maximize his utility, which considers both received incentive payments and effort costs. This is referred to as the "incentive compatibility" (IC) constraint. If the agent finds that the highest attainable utility meets or exceeds his reservation utility, he accepts the proposed arrangement; otherwise, he declines it. Additionally, the principal needs to entice the agent to participate, leading to the "participation compatibility" (PC) constraint.[2] Given the principal's inability to directly observe the agent's actual effort, she designs a contract based on an observable and verifiable quantity, known as the contractible quantity.[3] In this study, the chosen contractible

[1] In practice, real-time workload tracking is often unfeasible, but accurate estimation is possible using indicators like queue length, as shown in works such as Ibrahim and Whitt (2009) and Ibrahim and Whitt (2011). However, our contract design does not require real-time tracking; post-computation suffices.

[2] In related literature, the PC constraint is sometimes called the "individual rationality" constraint.

[3] In this research, "contract design" specifically refers to creating performance-based incentives to motivate the agent to control the service rate in line with the principal's best interests.

quantity is the workload. However, in principle, any metric capable of reflecting the volume of pending tasks (e.g., queue length) can serve this purpose, as long as it is observable and verifiable. A contract is considered optimal if it induces a rate-control strategy that minimizes the principal's overall cost while adhering to the IC and PC constraints. For the principal, finding an optimal contract involves two steps. First, she considers each rate-control strategy the agent might employ and then deduces the least costly incentive payment plan that induces such a strategy. This is achievable because the principal knows the agent's cost structure and payoff function, allowing her to predict the agent's response to a given payment plan. Second, she identifies a rate-control strategy that, combined with the cheapest way to induce it (as determined in the first step), minimizes the principal's long-term average cost. The resulting optimal contract induces a "second-best" rate-control strategy, rather than the "first-best" policy the principal would choose if she directly controlled the service rate and incurred the effort costs herself.

Methodologically, we use a reflected Brownian motion (RBM) to describe wait-time dynamics for both practical and tractability reasons. Brownian models are widely used in the operations management/research (OM/OR) literature due to their tractability and robust mathematical foundations. An RBM effectively captures the stochastic variability in job arrivals and service completions through Brownian motion, incorporating the agent's strategic actions with an adjustable drift rate and accounting for the server's cumulative idleness through a one-sided reflecting process at the origin. The advantage of using an RBM is particularly evident in the first step of the analysis, as it allows for a clear characterization of the relationship between a rate-control strategy and the most cost-effective incentive payment plan to induce that strategy. This simplification arises because the agent's relative value function can be expressed as a solution to an ordinary differential equation, highlighting key factors influencing the agent's decision-making, such as effort cost and incentive payment.

While the first step of the analysis is greatly simplified by using the Brownian model, the second step presents significant technical difficulties. Specifically, characterizing the optimal contract involves solving an instance of the calculus of variations, an infinite-dimensional optimization problem. In classical calculus of variations, the goal is to optimize a functional—a function of functions—expressed as a definite integral over a finite interval, involving an unknown function and its derivative. However, our problem involves minimizing a functional expressed as the ratio of two integrals, each with an infinite interval. Additionally, the functional involves the antiderivative of the unknown function, its derivative, and the unknown function itself. Consequently, standard results from the calculus of variations do not directly apply.

By extending classical calculus of variations theory, we demonstrate that the second-best rate-control strategy, if it exists, must obey a Riccati equation, a first-order ordinary differential equation that is quadratic in the unknown function. For quadratic congestion costs, the solution to the Riccati equation is explicit, resulting in a neat and intuitive contract. The incentive payment plan follows a carrot-and-stick structure. The "carrot" represents the maximum bonus the agent can receive per unit of time for maintaining zero delays. The “stick” represents a performance-based penalty, interpreted as the amount of bonus withheld due to poor service levels, as measured by a large backlog of work. If the penalty exceeds the maximum bonus, the excess is interpreted as a malus payment (an amount transferred from the agent to the principal).

Driven by practical considerations, we further analyze a scenario where the agent is protected by limited liability, meaning all incentive payments from the principal to the agent are non-negative. Characterizing the optimal contract in this scenario is even more challenging, as it requires solving a variational problem with point-wise constraints. To address this challenge, we use the Lagrange multiplier approach. First, we treat the agent’s surplus (the difference between the actual utility from the contract and the agent’s reservation utility) as fixed. This allows us to integrate the Lagrange multiplier method with variational analysis to derive the necessary conditions for an optimal contract (among those that provide the desired surplus). Using these conditions, we identify a contract expected to be optimal. By varying the agent’s surplus that the principal allocates, we deduce a parametric family of contracts. Finally, we show that this parametric family contains a cost-minimizer, expected to be the optimal solution to the contracting problem with the limited liability constraint.

1.1. Contributions

We consider the paper’s contributions to be threefold:

Modeling: We formulate a model that appears to be the first dynamic principal-agent problem that cannot be effectively solved using recursive methods. To tackle this challenge, we resort to the calculus of variations, a methodology not commonly used in the OM/OR literature. We believe this approach complements the recursive method typically applied to dynamic principal-agent problems and opens the door to its potential applicability in other dynamic contracting scenarios.

Methodology: To characterize the optimal contract, we extend classical calculus of variations theory to address functionals expressed as a ratio of two definite integrals involving a function of interest and its antiderivative and derivative. We prove that the associated Euler-Lagrange equation characterizing the “stationary points” of a given functional takes the form of a Riccati equation, subject to appropriate boundary conditions (Theorem 1). We generalize these arguments further by incorporating a “Lagrange multiplier function” to characterize the optimal contract under limited liability. By fixing the agent’s surplus, we derive an equation jointly satisfied by the function of interest and the Lagrange multiplier function, subject to constraints mirroring the KKT conditions in mathematical optimization. We prove that a solution can be found via this equation and the specified constraints (Theorem 2). Additionally, there are several analytical results of independent interest in the domain of optimal control for Brownian queues. For instance, Proposition 1 extends and generalizes a result from Kim and Randhawa (2018) by establishing the well-posedness of the Bellman equation for controlling a Brownian queue under a general cost function. Furthermore, Proposition 5 reveals a phenomenon of diminishing returns on investments in flexibility, interpreted as the range of allowable rate perturbations on a Brownian queue.

Managerial Implications: Beyond the carrot-and-stick type of compensation, our study shows that risk aversion and/or limited liability can create contracting frictions. These frictions imply that achieving the first-best outcome may not be feasible with the second-best solution. Specifically, when the agent is risk-averse and not protected by limited liability, the original contracting problem can be transformed into a risk-neutral equivalent, making monetary motivation more costly. Conversely, when the agent is risk-neutral but protected by limited liability, our analysis shows that a performance-based compensation structure can still be effective, but the agent will only exert effort up to a maximum level. Although the principal can influence this maximum effort level, thereby gaining more control over system congestion, it comes at the cost of offering a greater surplus to the agent.

1.2. Organization

The remainder of this paper is organized as follows: Section 2 reviews the related literature. Section 3 formally introduces our base model, while Section 4 presents the analysis and the main results pertaining to the base model. Section 5 discusses some of the key model elements. Section 6 concerns the design of the optimal contract with limited liability. Section 7 concludes the paper. The mathematical proofs of all main results are relegated to §EC.1 of the e-companion.

2. Literature Review

Several recent papers have examined scenarios where servers act strategically, selecting their service rates to optimize a utility function, creating a game between the servers and the employing firm. Gopalakrishnan et al. (2016) consider an $M/M/N$ queue where servers balance the cost of effort with a preference for idleness. By characterizing the symmetric Nash equilibrium, they demonstrate that the servers' strategic behavior significantly influences optimal staffing and routing policies. Zhan and Ward (2019) explore how payments can motivate servers to exert effort, integrating this with staffing and routing decisions made by the system manager. This results in a joint problem of staffing, routing, and payment in the context of an $M/M/N+M$ queue, where selfish servers choose service rates to maximize their utility. Zhong et al. (2023) consider an $M/M/N/k$ queue with a finite buffer, where servers can adjust their service speed to optimize a utility function involving payment, effort cost, and idleness. They conduct an equilibrium analysis using exact and asymptotic methods to understand how job admission, staffing, and payment decisions affect servers' choices. Their findings reveal interesting non-monotonic relationships between equilibrium service rates and key model inputs such as staffing level and buffer size. Our paper differs from these studies in two key aspects. First, we treat a server's choice of effort level as "hidden actions," creating moral hazard, whereas the aforementioned papers assume the effort level is observable by the system manager. Second, our model allows servers to adjust their service speed in real time, while the other studies assume a server's effort level is fixed once chosen. In other words, the previous studies and our work focus on different time scales: one at a strategic or tactical level, and the other at an operational level.

This paper falls within the service contracting literature. Previous studies have explored various aspects of service contracts, such as capacity-setting and pricing decisions in call center outsourcing contracts (Akşin et al. 2008), contract design in service chains with observable or unobservable effort levels (Ren and Zhou 2008), and the influence of different contracts on call center capacity-setting decisions (Hasija et al. 2008). Additionally, there has been research on service contracting in situations with information asymmetry (Akan et al. 2011) and contract performance in scenarios with private information and varying contractibility of effort (Zhang et al. 2018). Jiang et al. (2012) consider queueing in their study, where a service buyer (the principal) contracts with a service provider (the agent) to minimize costs while meeting a service level target. They model the appointment dynamics as an $M/D/1$ queue and show that while a linear performance-based contract can achieve the first-best outcome without information asymmetry, it cannot coordinate the service chain with asymmetric information. However, a threshold-based contract can coordinate the chain even with information asymmetry. Our work differs from these studies in scope and methodology. We focus on incentive provisioning at the operational level, with both incentive payments and agent actions responding to the current system state. In contrast, the aforementioned studies primarily address long-term or steady-state system performance.

This paper contributes to the growing literature on dynamic principal-agent problems, which require subtle economic reasoning. A successful approach is to use a recursive formulation with the agents' promised or continuation utility as the state variable. Spear and Srivastava (1987) pioneered this method, analyzing a dynamic principal-agent model in discrete time using a recursive method. Sannikov (2008) extends this approach to continuous time, where the agent's effort influences the drift of a Brownian motion, using stochastic optimal control techniques. Biais et al. (2010) apply similar techniques to a model where the agent's effort influences the rate of a Poisson process. The promised utility framework has been successfully applied in various studies in the OM/OR literature. For example, Li et al. (2013) study a dynamic principal-agent problem for managing critical suppliers via business volume incentives, treating the suppliers' promised utility vector as the state descriptor. Wang et al. (2016) use a dynamic mechanism design framework to study voluntary disclosure of environmental hazards, resulting in explicit, easy-to-implement regulation policies. Sun and Tian (2018) examine how to incentivize an agent to adopt the desired effort level over an infinite time horizon, while Chen et al. (2020) consider optimal monitoring and payment mechanisms to incentivize effort in reducing the probability of adverse events. Most recently, Tian et al. (2021) design dynamic incentives for an agent hired to repair and maintain a machine, finding that the IC constraint is not always binding. While the promised utility framework is attractive for analyzing a wide range of dynamic contracting problems, it is challenging to apply to our problem for two reasons. First, defining the agent's promised utility is straightforward in infinite-horizon discounted settings, but in our model, the agent faces an ergodic risk-sensitive control problem, making it unclear how to define promised utility properly. Second, this framework typically requires the principal's objective to lack a state structure, which our model does not meet. Therefore, we consider an alternative approach: the method of calculus of variations, which we believe complements the promised utility approach.

In related work, Plambeck and Zenios (2000) study a dynamic principal-agent problem and show that it can be efficiently solved via dynamic programming without using the agent's promised utility as the state variable. The optimal contract they derive is history-independent and renegotiation-proof. However, their analysis relies on the assumption that the agent has exponential utility and access to a perfect financial market for borrowing and lending to smooth cash flows. Building on this framework, Plambeck and Zenios (2003) analyze a make-to-stock model where a manufacturer bears the inventory-holding and back-order costs while delegating production to a supplier who dynamically controls the production rate at some cost. Similar to their model, our work utilizes payments contingent on the observed system state. However, our paper differs fundamentally in that we do not assume the agent interacts with the financial market, which prevents us from using a recursive method to characterize the optimal contract.

Our literature search reveals limited use of the calculus of variations in the OM/OR literature. One example is Bensoussan et al. (2023), where the authors consider a decentralized supply chain involving a retailer facing random demand over time and having private inventory information. The supplier offers supply contracts to address this asymmetry. They derive the necessary conditions for optimizing long-term contracts under varied demand and belief distributions, applying these to a batch-order contract scenario. Like our work, their problem is formulated using the calculus of variations and the Gâteaux derivative. However, there are two key distinctions: their problem deals with adverse selection, while we address moral hazards, and our analysis minimizes an unconventional functional, expressed as a ratio of two definite integrals, differing from the functional optimized by Bensoussan et al. (2023).

3. Base Model

In this section, we formulate our dynamic contracting model within a principal-agent framework and present the central question we aim to address. The model assumes that the agent controls the underlying state process through unobserved actions, and both parties have complete knowledge of each other's objectives, which will be introduced shortly. The principal and agent are also aware of how the agent's actions affect the state process. Additionally, we assume an infinite time horizon to simplify the decision space and reflect the reality of ongoing contractual relationships.

3.1. System Dynamics

Customer requests are processed by a single agent who exerts a base service rate in the absence of incentives (i.e., when receiving only the base wage). With this base service rate, the system's workload, or virtual wait time, follows a driftless reflected Brownian motion (RBM). Although we do not explicitly introduce the arrival and service completion processes, we justify using an RBM to describe the wait-time dynamics of an M/G/$\cdot$ type queue in §EC.3.

The agent's adjustments to the service rate create a drift-rate control process, denoted by θ, which impacts the workload process. Specifically, with a rate-control process θ, the workload evolution over time is described by:

$$Z(t) = Z(0) - \int_0^t \theta(u)\mathrm{d}u + \sigma B(t) + L(t), \tag{1}$$

where $Z(0)$ is the initial workload, B is a standard Brownian motion introducing random shocks, and L is a non-decreasing process ensuring Z remains non-negative. To connect this to a physical queue, L can be seen as tracking the cumulative idle time of the server. The Brownian term captures the inherent stochasticity in job arrivals and service completions, with the variability level determined by the constant σ. Both the principal and agent base their decisions on the state process Z.

3.2. Economic Factors

Transitioning to the economic aspect of the model, we posit that the principal incurs congestion costs at a rate of $k(z)$ when the workload Z is at level z. These costs include various forms of dissatisfaction, such as the loss of goodwill from prolonged wait times, and are borne by the principal. We assume that k is differentiable, grows to infinity at a polynomial rate, and satisfies $k(0) = 0$. We impose the following technical assumption on the cost function k:

ASSUMPTION 1. *For any $a_1 > 0$, there exists some $a_2 > 0$ such that $k(z) > a_1 k'(z) - a_2$ for $z \geq 0$.*

Intuitively, this assumption ensures that the cost function $k(z)$ grows faster than its derivative as z increases. It is satisfied by a broad range of functions, including polynomial cost functions like $k(z) = az^b$ for $a > 0$ and $b \geq 1$. These functions are commonly used in the performance analysis and optimization of queueing systems, especially where waiting costs are nonlinear. For instance, empirical studies have shown that the marginal waiting cost for critical patients in Canadian emergency departments can be approximated with a piecewise linear function (Ding et al. 2019) and quadratic waiting cost functions influence priority policy decisions based on traffic intensity and population proportions (Ouyang et al. 2022). Thus, assuming a general congestion cost function enriches our analysis and enhances the model's practical relevance. Incentive rewards allocated to the agent correspond to costs for the principal. Both parties agree on an incentive payment plan denoted as $S(\cdot)$. If $S(z) > 0$, the principal pays the agent at the rate $S(z)$ when the workload is z. Conversely, if $S(z) < 0$, the agent pays the principal at the rate $-S(z)$. Our model assumes that incentive payments are solely based on the workload level, as the principal uses the observed workload as a signal for the agent's effort.

The agent incurs effort costs at a rate of $c(x)$ when the drift rate θ is at level x. We adopt the standard specification $c(x) := \frac{\alpha}{2}x^2$. common in the literature (e.g., Zhang et al. (2018), Huang and Gurvich (2018)). This quadratic function naturally captures the notion of increasing marginal costs of effort. Although extending the model to a general convex effort cost function is possible, it would not fundamentally alter our analysis or insights. It is important to note that this quadratic specification implies costs even for effort levels below the base value. Practically, this could mean that reducing effort below the base level causes boredom and disutility for the agent. Methodologically, inducing a slowdown is never optimal for the principal, so only the non-negative portion of the effort cost function is relevant for our analysis. Hence, we will refer to this cost as the effort cost.

3.3. The Agent's Problem

Considering these economic factors, the agent evaluates his total payoff up to time t using the following expression:

$$\mathcal{J}(t) := \int_0^t S(Z(u))\mathrm{d}u - \int_0^t c(\theta(u))\mathrm{d}u. \tag{2}$$

This represents the difference between accumulated incentive payments and effort costs up to time t. Note that the base wage is excluded as it does not impact the design of dynamic incentives, except to keep the agent content with the base service rate.

We assume the agent is risk-averse, meaning he is more sensitive to losses than gains. To capture this, we use an exponential utility function that links total payoff over time to a utility value. Mathematically, the agent's total utility up to time t is given by $(1/\gamma)\log\mathbb{E}\left[e^{\gamma\mathcal{J}(t)}\right]$, where γ the risk-sensitive parameter ($\gamma \in (-\infty, 0)$). Given an incentive payment plan S, the agent seeks a control strategy θ to maximize the long-run average utility:

$$\liminf_{t\to\infty} \frac{1}{\gamma t} \log\mathbb{E}\left[e^{\gamma\mathcal{J}(t)}\right], \tag{3}$$

where $\mathcal{J}$ is defined by equation (2). This problem is an ergodic risk-sensitive control problem. Unlike most risk-sensitive control problems that assume a positive risk-sensitive parameter for cost minimization, our model uses a negative parameter because the agent aims to maximize utility. For a comprehensive background on ergodic risk-sensitive control problems, refer to Biswas and Borkar (2023). Additionally, risk-sensitive control is closely related to robust control, which often arises in uncertain system models and is typically formulated as a two-player zero-sum game; see, e.g., Hansen et al. (2006) and Lim and Shanthikumar (2007). Thus, the agent's risk aversion can also be interpreted as uncertainty about the probabilistic law governing external events affecting the system's evolution.

Following convention, we refer to the agent's goal of maximizing (3) as the IC constraint. As γ approaches zero, expression (3) simplifies to:

$$\liminf_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\mathcal{J}(t)\right]. \tag{4}$$

This allows us to extend the range of γ to $(-\infty, 0]$ with the understanding that the agent's objective is to maximize the long-run average payoff (4) when $\gamma = 0$.

Given an infinite time horizon and state-dependent incentive payments, the agent can focus on Markov strategies, where $\theta(t)$ depends only on the current workload $Z(t)$. Thus, each rate-control strategy can be represented as a deterministic function $\theta(z)$. We define θ as an admissible control if it is non-negative, piecewise smooth, exhibits polynomial growth as z approaches infinity, and satisfies $\theta(z) \geq \chi$ for some positive constant χ and all sufficiently large z. The non-negativity requirement streamlines analysis and can be relaxed to a uniform lower bound. Piecewise smoothness is necessary for applying the calculus of variations to characterize the optimal contract, as it ensures the well-posedness of the variational problem (Olver 2012, chapter 9). We require the function value of $\theta(z)$ to be strictly bounded away from zero for all large enough z to ensure that the system is stable when the agent uses θ as his control strategy. This requirement is merely a restatement of the celebrated Foster-Lyapunov criterion for stability (Xu 2023, chapter 5). Lastly, polynomial growth of θ ensures the principal's associated relative value function exhibits polynomial growth, a condition needed to establish the verification theorem (part (ii) of Theorem 1).

In addition to controlling the service rate, the agent decides whether to accept the principal's performance-based compensation scheme. If the agent finds the plan unappealing, he can reject the contract and pursue other options, forming the PC constraint. The agent finds the arrangement acceptable if:

$$\max_{\theta} \quad \liminf_{t\to\infty} \frac{1}{\gamma t} \log \mathbb{E}\left[e^{\gamma \mathcal{J}(t)}\right] \geq 0. \tag{5}$$

The requirement that the left-hand side be greater than or equal to zero stems from the assumption that the agent's outside option can be offset by a fixed wage.

3.4. The Principal's Problem

The principal is risk-neutral and calculates the total cost up to time t by adding the incentive payments to the congestion costs:

$$\mathcal{K}(t) := \int_0^t k(Z(u))\mathrm{d}u + \int_0^t S(Z(u))\mathrm{d}u. \tag{6}$$

Similar to the agent's decision problem, this objective excludes the agent's base wage, focusing on the design of the performance-based compensation scheme.

An incentive payment plan S induces a rate control strategy θ if the agent finds θ both incentive- and participation-compatible under S. An incentive payment plan S is admissible if the induced rate-control strategy θ is admissible. The rate-control process induced by a given incentive payment plan, together with the random shocks generated by the Brownian motion, drives the evolution of the workload process, which in turn determines the congestion costs incurred over time. The principal's objective is to devise an admissible incentive payment plan that, when combined with the agent's rate-control strategy, minimizes her long-run average cost. Formally, the principal seeks to find S that minimizes

$$\limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\mathcal{K}(t)\right]. \tag{7}$$

subject to the IC constraint (3) and PC constraint (5). It is important to note that the rate-control strategy cannot be enforced; the agent will adopt it voluntarily only if it satisfies both the IC and PC constraints.

3.5. Benchmark: A Single-Party Control Problem

As a benchmark, we present a simplified version of the problem where the principal directly controls the service rate and bears the effort costs. In this case, the objective is to find an admissible θ that minimizes the long-run average cost for the principal:

$$\limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c(\theta(u))\mathrm{d}u + k(Z(u))\mathrm{d}u\right] \quad \text{subject to (1).} \tag{8}$$

Alternatively, this objective arises when the principal can directly observe the agent's actions, eliminating the need to incentivize desired effort levels. If the principal can monitor the agent's actions, she can impose severe penalties for deviations from desired actions. Consequently, compensating the agent's effort according to the function c suffices to meet the PC constraint.

4. Analysis

In this section, we implement the two-step approach outlined in the introduction to characterize the optimal contract. The first step is to devise an incentive payment plan for each rate-control strategy that induces the strategy at the least expense to the principal. This is feasible since the principal is aware of the agent's payoff structure, enabling her to predict the agent's response to a given payment plan. The second step aims to identify an incentive payment plan that minimizes the overall cost to the principal, including both congestion costs and the monetary incentives provided to the agent. To aid in the mathematical analysis and presentation of the main results, we will use the notion of relative value functions.

We begin our analysis with Problem (8), as its results are likely to inform the analysis of the optimal contracting problem. The analysis of Problem (8) will yield the first-best rate-control policy, which can be derived using dynamic programming. The following result provides a formal statement.

PROPOSITION 1. *There exists a twice differentiable function v and a constant $\rho^\star$ that solve*

$$\min_{\theta}\left\{\frac{\sigma^2}{2}v''(z) - \theta v'(z) + k(z) + c(\theta)\right\} = \rho^\star \tag{9}$$

with $v'(0)=0$, $v''(z)>0$ for all $z\geq 0$, and that $v'(z)$ grows polynomially as $z\to\infty$. Furthermore, the minimum objective value of Problem (8) *is $\rho^\star$.*

The proof of this proposition can be found in the e-companion. It reveals that equation (9) can be transformed into a Riccati equation. This type of equation has been analyzed in several studies. Çelik and Maglaras (2008) derive a Bellman equation in the form of a Riccati equation in their study of joint dynamic pricing and leadtime quotation controls in make-to-order systems. By capturing congestion-related disutilities through order rejection instead of modeling waiting costs explicitly, they obtain a closed-form solution to the Riccati equation. Ata and Barjesteh (2023) characterize a near-optimal control policy for joint pricing, outsourcing, and scheduling controls in a make-to-stock system in heavy traffic via a Bellman equation also in the form of a Riccati equation. Their solution, expressed using Bessel and Airy functions, is derived on a finite interval, with endpoints optimized. Kim and Randhawa (2018) characterize a near-optimal pricing control policy in a single-server queue using a Riccati equation defined on the entire non-negative real line. Their equation can be seen as a special case of ours, differing mainly in the assumption of a linear cost rate function. While their proof of well-posedness relies on this assumption, our proof builds on Assumption 1, accommodating a broader range of polynomial cost functions.

Equation (9) conforms to the standard form of Bellman equations for average-cost stochastic control problems. The solution can be determined explicitly or numerically, depending on the nature of the congestion cost function k. With a quadratic k, for example, the resulting v is quadratic, implying a linear rate control for Problem (8), where the service rate increases linearly with the workload. This provides a benchmark, allowing the principal to directly control the service rate.

We next return to the optimal contracting problem by solving the agent's utility maximization problem, temporarily disregarding the PC constraint. Using standard results in ergodic risk-sensitive control theory (c.f. (3.9) in Biswas and Borkar (2023)), we write the Bellman equation characterizing the agent's strategy for maximizing his utility:

$$\max_{\theta}\left\{\frac{\sigma^2}{2}U''(z)-\theta U'(z)+S(z)-c(\theta)+\frac{\gamma\sigma^2}{2}(U'(z))^2\right\}=\eta \tag{10}$$

subject to the boundary condition $U'(0)=0$. In the language of dynamic programming, $U(z)$ is the relative value function for the agent's problem, and η represents the agent's long-run average utility, provided the agent participates. The PC constraint can thus be formulated as $\eta\geq 0$. This leads to the following result, specifying the most cost-effective incentive payment plan that induces a given rate-control strategy.

PROPOSITION 2. *If a rate-control strategy $\theta(\cdot)$, induced by a specific incentive payment plan, satisfies the PC constraint, then it must hold that $\theta(0)=0$. Furthermore, there exists an incentive payment plan S that induces θ at the lowest cost to the principal. This payment plan stipulates monetary transfers according to the formula:*

$$S(z)=\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{1}{2}(\alpha+\gamma\alpha^2\sigma^2)\theta^2(z). \tag{11}$$

Proposition 2 serves as a critical building block for the subsequent analysis. It establishes the existence of an "induced cost function," which quantifies the principal's instantaneous cost of inducing a particular rate-control process θ. This means that the mission of identifying an optimal contract can be turned into an equivalent single-party decision problem. The following result contributes to this mission by establishing a connection between a specific pair of strategies employed by both parties and the long-run average cost incurred by the principal for that strategy pair.

PROPOSITION 3. *Fixing a pair of strategies* (S,θ)*, if some twice differentiable function* V *and a constant* κ *collectively solve the following equation:*

$$\frac{\sigma^2}{2}V''(z)-\theta(z)V'(z)+S(z)+k(z)=\kappa, \tag{12}$$

subject to $V'(0)=0$ *and* $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V(Z(t))]=0$*, then* κ *is the long-run average cost of the principal under the pair of strategies* (S,θ)*.*

To illustrate the value of Proposition 3, we substitute the expression for S from the right-hand side of (11) into (12), yielding

$$\frac{\sigma^2}{2}V_\theta''(z)-\theta(z)V_\theta'(z)+\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{1}{2}(\alpha+\gamma\alpha^2\sigma^2)\theta^2(z)+k(z)=\kappa(\theta), \tag{13}$$

subject to the same conditions as specified in the proposition. We have appended the subscript θ to V' and V'' and written κ as $\kappa(\theta)$ in (13) to emphasize the dependence of these quantities on the choice of θ. Suppose for now that equation (13) admits a solution pair $(V_\theta,\kappa(\theta))$ that satisfies all the stated conditions. Then Propositions 2 and 3 together imply that $\kappa(\theta)$ represents the principal's long-run average cost when the agent chooses θ induced by the cheapest incentive payment plan $S(z)$. Thus, it is instructive to view (13) and its associated conditions as a functional (i.e., a function of functions) that maps each rate-control strategy θ to the principal's long-run average cost. Consequently, the principal's problem can be reframed as selecting an appropriate θ to

$$\text{minimize}\quad \kappa(\theta). \tag{14}$$

Equation (13) departs from the standard form of Bellman equations, as it cannot be optimized point-wise due to the appearance of the first-order derivative of θ. This stands in contrast to the left-hand side of equation (9), which conforms to the standard form and allows for simpler point-wise optimization. Instead, the minimization problem defined by (14) needs to be viewed as a variational problem that optimizes an objective over a space of functions. Solving this problem yields the second-best solution, hence the optimal contract.

In preparation for presenting the main results, let us define $\phi(z)$ as the integral of $\theta(x)$ with respect to x from 0 to z, i.e., $\phi(z):=\int_0^z\theta(x)\mathrm{d}x$. For notational convenience, also define $\tilde{\alpha}_\gamma:=\alpha+\gamma\alpha^2\sigma^2$. The following theorem provides an explicit expression for $\kappa(\theta)$ and shows that the second-best rate-control strategy, if it exists, must obey a Riccati equation.

THEOREM 1. *(i) Provided* $S(z)$ *given by* (11) *is admissible, there exists a solution pair* $(V_\theta,\kappa(\theta))$ *to* (13) *subject to the specified conditions. (ii) If* $\kappa(\cdot)$ *attains its minimum at* $\theta=\theta^\star$*, then the minimizer, along with some constant* $\tilde{\kappa}$*, must obey*

$$\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\tilde{\alpha}_\gamma}{2}\theta^2(z)+k(z)=\tilde{\kappa} \tag{15}$$

subject to the condition $\theta(0) = 0$ *and the requirement that* $\theta(z)$ *is non-negative and exhibits polynomial growth as* $z \to \infty$. *The value* $\tilde{\kappa}$ *represents the principal's long-run average cost under the pair of strategies* $(S^\star, \theta^\star)$*, where* $S^\star$ *is determined by* (11) *with* $\theta = \theta^\star$.

Based on Theorem 1, if the optimal contract exists, the induced rate-control strategy is expected to satisfy, along with some constant $\tilde{\kappa}$, equation (15). Note, however, that this equation provides a necessary rather than provably sufficient condition for optimality. A closely related question arises as to whether equation (15) admits a solution pair that fulfills all the specified conditions. The answer to this question is affirmative. To see this, it suffices to observe that equations (9) and (15) are structurally the same. That is, one can replicate the entire analysis that leads to the well-posedness of equation (9) to establish the well-posedness of equation (15), without making any essential changes to the arguments used. Further, from the proof of Proposition 1, we can assert that the solution θ to equation (15) is a non-negative, differentiable, and polynomially growing function in z, satisfying $\theta(z) \geq \chi$ for some positive constant χ, and for all sufficiently large z.

To illustrate the implications of Theorem 1, consider an example where the congestion cost function is quadratic: $k(z) = k_2 z^2$ for some positive constant k_2. In this case, one can verify that equation (15), subject to the specified conditions, admits the following pair of solutions:

$$\theta(z) = \sqrt{\frac{2k_2}{\tilde{\alpha}_\gamma}}\, z \quad \text{and} \quad \tilde{\kappa} = (\sigma^2 \alpha/2)\sqrt{2k_2/\tilde{\alpha}_\gamma}.$$

The incentive payment plan that induces this rate-control strategy θ is given by:

$$S(z) = \frac{\sigma^2}{2}\alpha\sqrt{2k_2/\tilde{\alpha}_\gamma} - k_2 z^2. \tag{16}$$

This incentive payment plan has a clear carrot-and-stick structure. The "carrot" represents the maximum bonus the agent can earn per unit of time for maintaining zero delays. The "stick" is performance-based, reflecting the amount of bonus that would have been paid out but was withheld due to inadequate service levels, resulting in a backlog of work. If the performance-based penalty exceeds the maximum bonus, the excess is deducted from the agent's earnings (malus payment) and transferred to the principal. Together, these components motivate the agent to increase their speed in response to an increasing backlog.

Additionally, the system's variability, measured by σ^2, affects the principal's long-term average cost. As stochastic volatility increases, it becomes more challenging to distinguish between stochastic fluctuations and shirking behavior, complicating the principal's ability to infer the agent's hidden actions. To counteract the increased "moral hazard" effect, the principal must enhance the power of incentives (by increasing the maximum bonus) to ensure the agent takes the desired actions.

5. Discussion

Having presented a model to understand the economic drivers influencing the design of an optimal contract, we will now briefly discuss the implications of the main results obtained and some of the model's fundamentals.

5.1. Necessity vs. Sufficiency

The characterization of the optimal contract, as stated in Theorem 1, uses the calculus of variations, a mathematical technique that identifies the stationary points of a given functional by setting the functional derivative ("first variation") to zero, which ultimately leads to the Euler-Lagrange equation. Our proof of Theorem 1 relies on deriving a specific form of the Euler-Lagrange equation from the first variation. Similar to the first-order condition in standard calculus, the Euler-Lagrange equation provides a necessary but not sufficient condition for a function to be an extremum of the functional. Despite this, the Euler-Lagrange equation often offers valuable insights into the solution of the variational problem, as demonstrated in Chen et al. (2020), where the solution to this equation provides significant insights into the optimal pricing scheme. To confirm that the solution to equation (15) indeed yields the optimal contract, we must perform a second derivative test on the functional $\kappa(\theta)$ and demonstrate the positivity of its second derivative, known as the *second variation* in the calculus of variations. Unfortunately, this verification tends to be highly technical. Therefore, we have chosen not to pursue it.

Nevertheless, there are valid reasons to believe that solving equation (15) provides the second-best solution. We document two reasons. First, there is only one value of $\tilde{\kappa}$ for which equation (15), viewed as a differential equation with respect to the unknown function θ, has a solution that satisfies all the stated requirements. To elaborate, we note that $\tilde{\kappa}$ is the smallest value of κ for which the equation

$$\frac{\sigma^2}{2}\alpha\theta'(z) - \frac{\tilde{\alpha}_\gamma}{2}\theta^2(z) + k(z) = \kappa \tag{17}$$

with the boundary condition $\theta(0) = 0$ produces a non-negative solution function θ. More precisely, for all $\kappa < \tilde{\kappa}$, the solution function $\theta(z)$ to (17) approaches negative infinity as z grows large. For all $\kappa > \tilde{\kappa}$, the solution to (17), although non-negative, grows to positive infinity at least exponentially fast. Only when $\kappa = \tilde{\kappa}$ does the solution remain non-negative and grow to infinity at a polynomial rate. Second, in the special case where $\gamma = 0$, the rate-control strategy obtained by solving (15) aligns with the first-best policy derived from (9). Additionally, the constant $\tilde{\kappa}$ obtained from (15) matches the optimal objective value $\rho^\star$ as stated in Proposition 1. Therefore, the rate-control strategy obtained from (15) with $\gamma = 0$ achieves the first-best outcome. Since this strategy is also admissible, it must represent the second-best solution.

5.2. Contracting Frictions

In principal-agent problems, it is common to assume the agent is risk-averse (Dai and Chao 2013). In the principal-agent literature, it is well-established that when the agent is risk-neutral, there is no "friction" between the principal's desire to compensate the agent and the agent's willingness to accept the arrangement. Consequently, the cost to the principal of adopting the optimal contract is the same as the cost under the first-best solution. Therefore, it is unsurprising that the solution derived from Theorem 1 can achieve the first-best outcome when the agent is risk-neutral (i.e., $\gamma = 0$). However, when the agent is risk-averse, the principal must pay a risk premium to counteract the agent's aversion to risk, causing the second-best solution to differ from its first-best counterpart.

5.3. More on the Brownian Motion Setup

In queueing literature, it is well-known that an RBM naturally arises as a process-level approximation of a *critically loaded* queueing system. In this context, critical loading means that the job arrival rate matches the base service rate. While this assumption may seem a bit far-fetched, given that real-world service systems can be underloaded or overloaded, we argue that it is a valuable regime to explore. An underloaded system, where the base service capacity significantly exceeds demand, is less interesting because incentive rewards aim to alleviate congestion. However, underloaded systems rarely experience congestion and instead result in excessive idleness. Conversely, a severely overloaded system, where the base capacity falls far short of demand, suggests that the incoming load is too much for a single server to handle. This signals the principal to consider recruiting multiple agents to share the load, an interesting topic for future research. It is also plausible that the base capacity slightly differs from the demand volume. This scenario can be accurately captured by introducing a constant drift rate to the Brownian motion without major hurdles to the analysis or fundamentally altering our main results. Lastly, the assumed operational regime may result from capacity rationing, where the principal, aware of the demand volume, can screen job candidates to find one whose base service capacity matches or approximates the demand.

6. Dynamic Incentives Under Limited Liability

The implementation of malus payments may face challenges in practical scenarios due to legal restrictions. This is particularly true when the contracted agent is categorized as a *de facto* employee. In such cases, the agent may be entitled to a minimum wage as required by labor laws. For instance, suppose the base wage (intended to induce the base service rate) coincides with the minimum wage. Since an employee's pay cannot be reduced below the minimum wage, this results in a "limited liability" constraint, defined as follows:

$$S(z) \geq 0 \quad \text{for all} \quad z \geq 0. \tag{18}$$

In other words, the agent would receive, at worst, zero incentive payment.

This new constraint significantly impacts the design of incentive contracts. Consider that if the principal chooses to set η to zero in equation (10), then equation (2) would imply that θ must be a zero process. Consequently, the principal is compelled to consider only positive values of η given the newly added constraint (18). With this in mind, the rest of the section focuses on addressing the issue of contract design under the limited liability constraint.

To isolate the unique effects of limited liability on contract design, we assume both the agent and the principal are risk-neutral, setting $\gamma = 0$. The principal's problem then becomes to minimize the objective in (7), subject to the IC constraint, the PC constraint, and the limited liability constraint. As we have seen, without the limited liability constraint, the search for the optimal contract involves solving an optimization problem in an infinite-dimensional space. The addition of the limited liability constraint transforms this into a constrained infinite-dimensional optimization problem, which is more challenging than its unconstrained counterpart. To address this challenge, we adapt the Lagrange approach, a common technique for finite-dimensional constrained optimization problems, to this infinite-dimensional context.

Specifically, our solution approach draws inspiration from the celebrated Karush-Kuhn-Tucker (KKT) conditions in the theory of mathematical optimization. To apply this idea, we associate the non-negativity constraint on $S(z)$ with a Lagrange multiplier $\lambda(z)$. Because the limited liability constraint is imposed for each individual workload z, the collection of all such Lagrange multipliers forms a function of z. The introduction of the Lagrange multiplier function $\lambda(z)$ allows us to modify the objective functional by adding to it the pointwise constraints multiplied by their respective Lagrange multipliers, yielding the modified objective functional:

$$\tilde{\mathcal{K}}(t) := \int_0^t k(Z(u))\mathrm{d}u + \int_0^t S(Z(u))\mathrm{d}u - \int_0^t \lambda(Z(u))S(Z(u))\mathrm{d}u. \tag{19}$$

In addition, the KKT conditions motivate the "dual feasibility" constraint on the Lagrange multiplier function:

$$\lambda(z) \geq 0 \quad \text{for all} \quad z \geq 0, \tag{20}$$

as well as the "complementary slackness" condition:

$$\lambda(z)S(z) = 0 \quad \text{for all} \quad z \geq 0.$$

With these preparations, we can reformulate the principal's problem as one that seeks some S and a Lagrange multiplier function λ to minimize

$$\limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\tilde{\mathcal{K}}(t)\right] \tag{21}$$

subject to the IC constraint (3), the PC constraint (5), the limited liability constraint (18) and the dual feasibility constraint (20), as well as the complementary slackness condition.[4] We will refer to Problem (21) as the "Lagrange problem." The following result provides theoretical justification for considering this relaxed problem.

PROPOSITION 4. *If some $S^\star$ solves the Lagrange problem* (21) *along with a Lagrange multiplier function λ, $S^\star$ also solves the principal's original problem* (7) *with the limited liability constraint.*

To determine the long-run average cost of the principal κ under the new objective (21) and a pair of fixed strategies (S, θ), similar to equation (12), we consider the following equation:

$$\frac{\sigma^2}{2}V''(z) - \theta(z)V'(z) + S(z) + k(z) - \lambda(z)S(z) = \kappa, \tag{22}$$

[4] In the language of KKT conditions, the limited liability constraint corresponds to the "primal feasibility" condition.

subject to the boundary condition $V'(0)=0$ and the limiting condition $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V(Z(t))]=0$.

As with the base model, a key step in solving the principal's problem is to establish a connection between a rate control strategy and an incentive payment plan that induces the strategy at minimal cost. In the base model, this connection was established by solving the agent's utility-maximizing problem (10) while setting $\eta=0$. However, as explained at the beginning of the section, the introduction of the limited liability constraint would require η to be positive. For the time being, let us consider the value of η to be given and fixed. Following the analysis of Proposition 2, we can derive an incentive payment plan $S(z)$ that incentivizes a given rate control strategy θ at the lowest expense to the principal, which is expressed as follows:

$$S(z)=\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta. \tag{23}$$

Upon substituting (23) into (22), we obtain the following equation:

$$\frac{\sigma^2}{2}V_\theta''(z)-\theta(z)V_\theta'(z)+\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta+k(z)-\lambda(z)\left(\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta\right)=\kappa(\theta) \tag{24}$$

subject to the same set of conditions, namely, $V_\theta'(0)=0$ and $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V_\theta(Z(t))]=0$. As in the base model, we use the subscript θ to indicate that V' and V'' depend on the choice of θ, and we write κ as $\kappa(\theta)$ to emphasize this dependence in equation (24).

Similar to part (i) of Theorem 1, equation (24), along with the specified boundary conditions, enables us to derive an expression for $\kappa(\theta)$ as follows:

$$\kappa(\theta)=\frac{\displaystyle\int_0^\infty e^{-\frac{2}{\sigma^2}\phi(x)}\left[\frac{\sigma^2}{2}\alpha\theta'(x)-\frac{\alpha}{2}\theta^2(x)+k(x)-\lambda(x)\left(\frac{\sigma^2}{2}\alpha\theta'(x)-\frac{\alpha}{2}\theta^2(x)+\eta\right)\right]\mathrm{d}x}{\displaystyle\int_0^\infty e^{-\frac{2}{\sigma^2}\phi(x)}\mathrm{d}x}.$$

With this expression at our disposal, we can apply the calculus of variations technique akin to the proof for part (ii) of Theorem 1. In particular, we can deduce that with a fixed Lagrange multiplier function λ, the function θ that minimizes $\kappa(\theta)$ must satisfy the "stationarity" condition as encapsulated in the following equation:

$$\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+k(z)-\lambda(z)\left(\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta\right)+\frac{\sigma^2}{2}\alpha\left(\frac{\sigma^2}{2}\lambda''(z)-\theta(z)\lambda'(z)\right)=\tilde{\kappa}, \tag{25}$$

subject to the boundary condition $\theta(0)=0$. A detailed derivation for equation (25) is provided in §EC.2.

Here, $\tilde{\kappa}$ represents the long-run average cost of the principal, subtracted by η, i.e., $\tilde{\kappa}=\kappa(\theta)-\eta$. In addition, the limited liability constraint and the complementary slackness condition can be translated into

$$\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta\geq 0, \tag{26}$$

and

$$\lambda(z)\left(\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta\right)=0, \tag{27}$$

respectively. Note that we can substitute (27) into (25) to simplify the stationarity condition, yielding the following:

$$\frac{\sigma^2}{2}\alpha\theta'(z) - \frac{\alpha}{2}\theta^2(z) + \frac{\sigma^2}{2}\alpha\left(\frac{\sigma^2}{2}\lambda''(z) - \theta(z)\lambda'(z)\right) + k(z) = \tilde{\kappa}, \tag{28}$$

subject to the boundary condition $\theta(0) = 0$.

Now, our goal becomes to pinpoint a triple $(\theta, \lambda, \tilde{\kappa})$ that not only satisfies (28) but adheres to the three labeled constraints: (20), (26) and (27). This is a challenging endeavor due to the nature of (28), which does not conform to the structure of a conventional differential equation (it involves two unknown functions, θ and λ, plus one unknown constant). To be able to make further headway, our strategy is to formulate conjectures about the appropriate functional forms of θ and λ and proceed to verify that the conjectured solution meets all the stated requirements.

This solution strategy is rooted in the following insight: In the absence of the limited liability constraint, the principal retains the option to impose negative payments (penalties) as a means of exerting pressure on the agent when congestion levels are high, encouraging the agent to speed up in order to alleviate the congestion. However, in the presence of the limited liability constraint, the principal's ability to enforce penalties is compromised. This creates a situation where the principal's hands are tied, unable to levy a heavy penalty even during significant service level deterioration. What she can do best is stop providing incentive payments, and when that happens, the agent loses the impetus to increase speed further.

Based on the foregoing insight, we posit that a candidate minimizer θ has the following structural property: It is strictly increasing on the interval $[0, \bar{z})$ for some value $\bar{z}$ and satisfies the equation

$$\frac{\sigma^2}{2}\alpha\theta'(z) - \frac{\alpha}{2}\theta^2(z) + k(z) = \tilde{\kappa} \quad \text{for all} \quad z \in (0, \bar{z}), \tag{29}$$

subject to the boundary conditions $\theta(0) = 0$ and $\theta(\bar{z}) = b := \sqrt{2\eta/\alpha}$; and on the interval $[\bar{z}, \infty)$, θ is constant and equal to b. It is easy to verify that such a θ, if it exists, automatically fulfills (26). Moreover, the dual feasibility constraint (20) and the complementary slackness condition (27) imply that, for such a θ, the corresponding Lagrange multiplier function (if it exists) must satisfy the following requirement: $\lambda(z) = 0$ for $z \in [0, \bar{z})$, and

$$\frac{\sigma^2}{2}\alpha\left(\frac{\sigma^2}{2}\lambda''(z) - b\lambda'(z)\right) - \frac{\alpha}{2}b^2 + k(z) = \tilde{\kappa} \quad \text{for all} \quad z \in [\bar{z}, \infty). \tag{30}$$

Returning to equation (29), we note that the two boundary conditions $\theta(0) = 0$ and $\theta(\bar{z}) = b$, combined with equation (29), are inadequate to determine the three unknowns θ, $\tilde{\kappa}$, and $\bar{z}$. Therefore, we need to identify additional conditions. To that end, let us first consider equation (24) on $[0, \bar{z})$. By substituting (27) and (29) into (24) and letting $\tilde{\kappa} = \kappa(\theta) - \eta$, we can simplify (24) to

$$\frac{\sigma^2}{2}V_\theta''(z) - \theta(z)V_\theta'(z) = 0. \tag{31}$$

Similarly, on $[\bar{z},\infty)$, equation (24) can be simplified using (27) and the fact that θ is constant and equal to b on $[\bar{z},\infty)$, yielding

$$V_\theta''(z) - bV_\theta'(z) - \frac{\alpha}{2}b^2 + k(z) = \tilde{\kappa}. \tag{32}$$

From (31) and the boundary condition $V_\theta'(0)=0$, we can deduce that $V_\theta' \equiv 0$ on $[0,\bar{z})$. On the other hand, since V_θ must have a continuous first-order derivative, we can subject (32) to the boundary condition $V_\theta'(\bar{z})=0$, in addition to the limiting condition $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V_\theta(Z(t))]=0$.

To encapsulate, we have introduced one new unknown, V_θ satisfying (32), and added two new conditions, namely, $V_\theta'(\bar{z})=0$ and $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V_\theta(Z(t))]=0$. Together, these additions help increase both the total number of unknowns and the number of conditions to four.

THEOREM 2. *(i) There exist θ, $\tilde{\kappa}$, $\bar{z}$, and V_θ that, together, satisfy the derived conditions. (ii) For the θ, $\tilde{\kappa}$ and $\bar{z}$ found in part (i), there exists a non-negative Lagrange multiplier function λ that satisfies $\lambda(z)=0$ on the interval $z\in[0,\bar{z})$ and* (30) *on the interval $(\bar{z},\infty)$, as well as the boundary conditions $\lambda(\bar{z})=0$ and $\lambda'(\bar{z})=0$.*

An inspection of the proof of Theorem 2 unveils an interesting connection between the solution encapsulated in part (i) of the theorem and the solution to the following *single-party* stochastic control problem: Choose an admissible θ to minimize

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t c(\theta(u))\mathrm{d}u + k(Z(u))\mathrm{d}u\right] \quad \text{subject to (1) and } \theta\in[0,b]. \tag{33}$$

Notably, this problem bears resemblance to Problem (8), differing solely in the confinement of θ to the interval $[0,b]$. It turns out that the θ and $\tilde{\kappa}$ identified in part (i) of Theorem 2, coincide, respectively, with the optimal policy and the optimal objective value of Problem (33). This connection is useful in that one can compute the desired θ and $\tilde{\kappa}$ by solving the single-party stochastic control problem (8) instead. As with Problem (8), the solution to Problem (33) can be conveniently characterized and computed through a Bellman equation. This equation seeks to identify a function v that is twice differentiable and a constant $\rho^\star(b)$ that jointly satisfy the equation:

$$\min_{\theta\in[0,b]}\left\{\frac{\sigma^2}{2}v''(z) - \theta v'(z) + k(z) + c(\theta)\right\} = \rho^\star(b) \tag{34}$$

with $v'(0)=0$, $v''(z)>0$ for all $z\geq 0$, and that $v'(z)$ grows polynomially as $z\to\infty$.

Our discussions thus far have treated η as a given and fixed parameter. As a result, $\tilde{\kappa}$ given in Theorem 2 is a quantity dependent on η. To emphasize this dependence, we will henceforth write $\tilde{\kappa}$ as $\tilde{\kappa}(\eta)$. With the aim of minimizing costs in mind and given the family of contracts parameterized by η, the principal is naturally inclined to select a value for η to

$$\text{minimize} \quad \eta + \tilde{\kappa}(\eta). \tag{35}$$

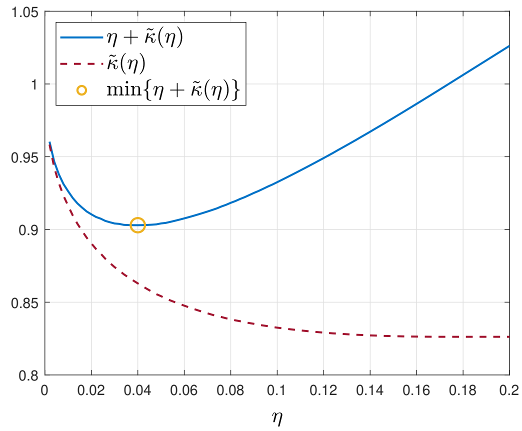

(a) Linear congestion cost

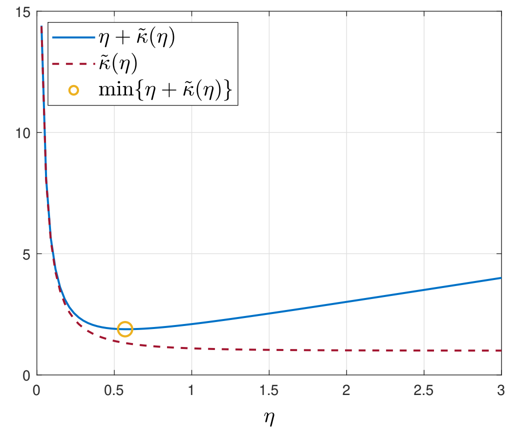

(b) Quadratic congestion cost

Figure 1 The minimum value of the long-run average cost for the principal

A related question is now in order: Does the optimization problem defined by (35) guarantee to admit a minimizer?

Before seeking an answer to this question, we numerically solved for the parameterized family of contracts, based on examples featuring linear and quadratic congestion cost functions. Let $\sigma^2 = 1$ and define the agent's effort cost as $c(\theta) = \theta^2$. Additionally, we set the linear congestion cost as $k(z) = z$ and the quadratic congestion cost as $k(z) = z^2$. The two graphs in Figure 1 illustrate the relationship between the values of $\tilde{\kappa}$ and η under linear and quadratic congestion cost functions, respectively. Both graphs demonstrate a decreasing trend in the principal's long-term average cost subtracted by η as η increases. In both cases, the function $\eta + \tilde{\kappa}(\eta)$ exhibits strict convexity in η, implying the existence of an optimal minimizing value, denoted as $\eta^\star$. This, in turn, enables us to pinpoint a potential solution for the second-best rate control strategy, $\theta^\star$, along with the corresponding incentive payment plan, $S^\star(z)$, which can be used to incentivize the adoption of $\theta^\star$.

Zooming in on the details, Figure 2 illustrates the rate-control strategy $\theta^\star$ and the corresponding incentive payment plan, $S^\star(z)$, for the specific examples. As depicted in the left panels, the incentive payment gradually decreases as congestion increases, and it suddenly drops to zero at a critical juncture, maintaining a value of zero thereafter. The right panels showcase how the agent's chosen service rate initially increases until congestion reaches the aforementioned critical point, beyond which it levels off, assuming its maximum value b.

Returning to the question about the existence of a minimizer for Problem (35), given that $b = \sqrt{2\eta/\alpha}$ (by definition) and $\tilde{\kappa}(\eta) = \rho^\star(b)$, we can see that the task of minimizing the expression in (35) is tantamount to selecting a value for b to

$$\text{minimize} \quad \frac{\alpha b^2}{2} + \rho^\star(b). \tag{36}$$

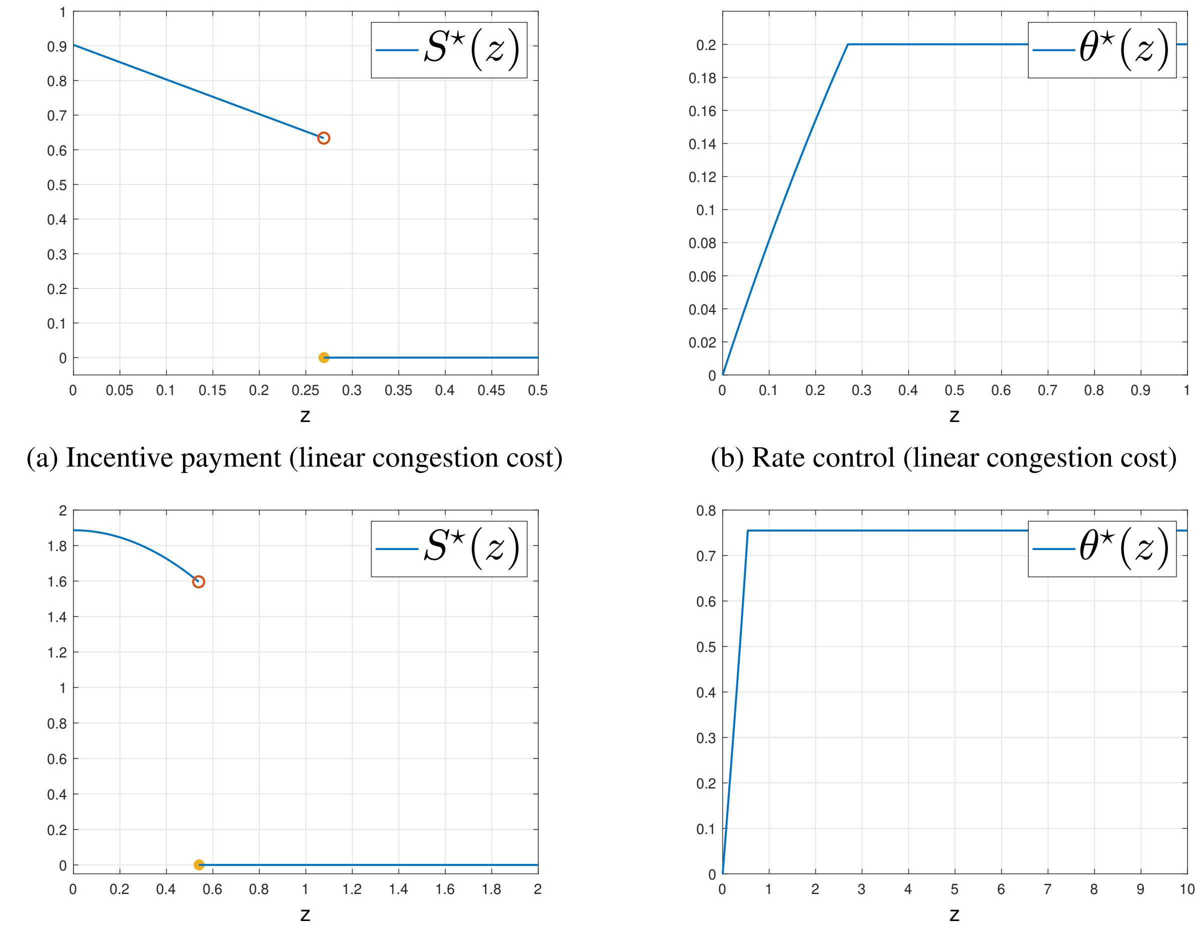

(a) Incentive payment (linear congestion cost)

(b) Rate control (linear congestion cost)

(c) Incentive payment (quadratic congestion cost)

(d) Rate control (quadratic congestion cost)

Figure 2 Illustration of the solution with linear and quadratic congestion costs when penalty payments are prohibited

The first term is quadratic, thus exhibiting convexity with respect to b. As a result, if one can demonstrate that the second term is also convex in b, we can conclude that the expression in Problem (36) is strictly convex. This, in turn, guarantees the existence of a unique minimizer for Problem (36). The following proposition serves this purpose by establishing that $\rho^\star(b)$ is indeed convex in b.

PROPOSITION 5. *The optimal objective value $\rho^\star(b)$ of Problem* (33) *is convex in b.*

Proposition 5 has implications beyond the specific contracting problem under consideration. For instance, consider a service facility with a fixed base service capacity and the option to augment it with surge capacity. Utilizing surge capacity incurs operational costs and requires an upfront investment. If the facility's workload follows a driftless RBM under the base capacity and surge capacity can be utilized partially or fully, then Problem (33) becomes a quest for the optimal utilization of surge capacity, given a maximum surge capacity limit b. Proposition 5 suggests that investing in surge capacity exhibits diminishing returns, indicating an optimal value of b that minimizes the total cost, combining operational and investment costs.

At this point, it is worthwhile to summarize our progress. Fixing the value of η, we derived the necessary conditions characterizing the optimal contract, where "optimal" means minimizing the principal's long-term average cost while ensuring the agent derives a utility of η. Using these necessary conditions, we identified a contract that fulfills them for each fixed η. By varying η, we generated a family of contracts parameterized by η. Among this family, we established the existence of a contract that minimizes the principal's objective, standing out as the "best-of-the-best." A final question remains: Is this contract the optimal solution to our contracting problem under limited liability? The answer hinges on proving the sufficiency of our "KKT conditions" for optimality. However, establishing this sufficiency is challenging even in finite-dimensional constrained optimization problems. Demonstrating that the KKT conditions suffice for optimality is intricate in most finite-dimensional cases, let alone in constrained optimization problems with infinite dimensions. Nonetheless, the contract identified as the best-of-the-best, shown in Figure 2, aligns closely with intuition about an ideal contract under limited liability. This suggests it is likely the true optimal solution. Under this contract, the principal grants high rewards for favorable states while withholding rewards for unfavorable ones. Unburdened by severe negative consequences, the agent accelerates as long as rewards continue, halting only when rewards cease.

The minimization problem (35) reveals a dilemma in designing dynamic incentives under limited liability. To raise the power of incentives and enhance the agent's responsiveness to the system state, the principal must ensure a significant disparity between compensation for desirable and undesirable states. Without the limited liability restriction, this would involve high bonuses for favorable outcomes and hefty penalties for unfavorable ones. With the limited liability constraint, the principal cannot impose penalties and can only offer zero rewards. Thus, to differentiate compensation adequately, the principal must elevate remuneration for favorable states, increasing the agent's utility from the contract, as reflected by a higher η. While increasing η promises high-powered incentives, the marginal benefit diminishes, justifying the theoretical conclusion in Proposition 5.

7. Concluding Remarks

This paper considers the optimization of contract design, where a principal hires an agent to manage customer requests in an uncertain environment. Although the agent can adjust his effort level over time, this action remains hidden and unenforceable by contract. To incentivize the agent to manage the service rate as desired, the principal links the observed system state to incentive rewards, resulting in a contracting problem. We show that finding the optimal contract involves solving a variational problem, leading to a Riccati equation with specific boundary conditions. In cases where congestion costs are quadratic, we derive explicit formulas for both the incentive payment plan and the induced service rate control. The incentive payment plan is intuitive: it represents the maximum bonus minus a performance-based penalty. Additionally, we address a scenario where negative or malus payments are restricted, leading to a contract design problem under limited liability. Our approach blends variational analysis and the Lagrange multiplier method, establishing the necessary conditions for an optimal contract. These conditions resemble the KKT conditions in conventional constrained optimization problems, allowing us to identify a parametric family of contracts that likely includes the optimal one. We demonstrate that there is indeed a best contract within this family, making it a strong candidate for the optimal contract.

There are several promising directions for future research. The first direction is to formally establish the sufficiency of the necessary conditions for contract optimality. Another avenue of exploration entails the development of dynamic incentives when dealing with multiple agents. While our current model primarily focuses on a single agent engaged with the principal, considering scenarios where multiple agents participate in competitive or collaborative roles may unveil additional managerial insights. Finally, it may be worthwhile to consider impatient jobs that can abandon the queue while waiting. By incorporating this aspect, the resulting system model is likely to offer a more faithful representation of call center operations, potentially offering new insights into contract design.

References

Akan M, Ata B, Lariviere MA (2011) Asymmetric information and economies of scale in service contracting. *Manufacturing & Service Operations Management* 13(1):58–72.

Akkermans H, Van Oppen W, Wynstra F, Voss C (2019) Contracting outsourced services with collaborative key performance indicators. *Journal of Operations Management* 65(1):22–47.

Akşin OZ, De Véricourt F, Karaesmen F (2008) Call center outsourcing contract analysis and choice. *Management Science* 54(2):354–368.

Ata B, Barjesteh N (2023) An approximate analysis of dynamic pricing, outsourcing, and scheduling policies for a multiclass make-to-stock queue in the heavy traffic regime. *Operations Research* 71(1):341–357.

Bensoussan A, Sethi S, Wang S (2023) A stationary infinite-horizon supply contract under asymmetric inventory information. *Operations Research* .

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

Biswas A, Borkar VS (2023) Ergodic risk-sensitive control—a survey. *Annual Reviews in Control* .

Browne S, Whitt W, Dshalalow J (1995) Piecewise-linear diffusion processes. *Advances in Queueing Theory, Methods, and Open Problems* 4:463–480.

Çelik S, Maglaras C (2008) Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Management Science* 54(6):1132–1146.

Chen M, Sun P, Xiao Y (2020) Optimal monitoring schedule in dynamic contracts. *Operations Research* 68(5):1285–1314.

Dai Y, Chao X (2013) Salesforce contract design and inventory planning with asymmetric risk-averse sales agents. *Operations Research Letters* 41(1):86–91.

Ding Y, Park E, Nagarajan M, Grafstein E (2019) Patient prioritization in emergency department triage systems: An empirical study of the Canadian triage and acuity scale (ctas). *Manufacturing & Service Operations Management* 21(4):723–741.

Gopalakrishnan R, Doroudi S, Ward AR, Wierman A (2016) Routing and staffing when servers are strategic. *Operations Research* 64(4):1033–1050.

Hansen LP, Sargent TJ, Turmuhambetova G, Williams N (2006) Robust control and model misspecification. *Journal of Economic Theory* 128(1):45–90.

Hasija S, Pinker EJ, Shumsky RA (2008) Call center outsourcing contracts under information asymmetry. *Management Science* 54(4):793–807.

Huang J, Gurvich I (2018) Beyond heavy-traffic regimes: Universal bounds and controls for the single-server queue. *Operations Research* 66(4):1168–1188.

Ibrahim R, Whitt W (2009) Real-time delay estimation in overloaded multiserver queues with abandonments. *Management Science* 55(10):1729–1742.

Ibrahim R, Whitt W (2011) Wait-time predictors for customer service systems with time-varying demand and capacity. *Operations Research* 59(5):1106–1118.

Jiang H, Pang Z, Savin S (2012) Performance-based contracts for outpatient medical services. *Manufacturing & Service Operations Management* 14(4):654–669.

Kim J, Randhawa RS (2018) The value of dynamic pricing in large queueing systems. *Operations Research* 66(2):409–425.

Li H, Zhang H, Fine CH (2013) Dynamic business share allocation in a supply chain with competing suppliers. *Operations Research* 61(2):280–297.

Lim AE, Shanthikumar JG (2007) Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research* 55(2):198–214.

Olver PJ (2012) The calculus of variations. *Applied Mathematics Lecture Notes. Sec* 21(4).

Ouyang H, Taník Argon N, Ziya S (2022) Assigning priorities (or not) in service systems with nonlinear waiting costs. *Management Science* 68(2):1233–1255.

Plambeck EL, Zenios SA (2000) Performance-based incentives in a dynamic principal-agent model. *Manufacturing & Service Operations Management* 2(3):240–263.

Plambeck EL, Zenios SA (2003) Incentive efficient control of a make-to-stock production system. *Operations Research* 51(3):371–386.

Poster WR (2007) Who's on the line? Indian call center agents pose as Americans for US-outsourced firms. *Industrial Relations: A Journal of Economy and Society* 46(2):271–304.

Ren ZJ, Zhou YP (2008) Call center outsourcing: Coordinating staffing level and service quality. *Management Science* 54(2):369–383.

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

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

Sun P, Tian F (2018) Optimal contract to induce continued effort. *Management Science* 64(9):4193–4217.

Tian F, Sun P, Duenyas I (2021) Optimal contract for machine repair and maintenance. *Operations Research* 69(3):916–949.

Wang S, Sun P, de Véricourt F (2016) Inducing environmental disclosures: A dynamic mechanism design approach. *Operations Research* 64(2):371–389.

Xu K (2023) Drift method: from stochastic networks to machine learning. *URL: https://web. stanford. edu/˜ kuangxu/papers/driftmethod. pdf. Last visited on* 3(09).

Zhan D, Ward AR (2019) Staffing, routing, and payment to trade off speed and quality in large service systems. *Operations Research* 67(6):1738–1751.

Zhang H, Kong G, Rajagopalan S (2018) Contract design by service providers with private effort. *Management Science* 64(6):2672–2689.

Zhong Y, Gopalakrishnan R, Ward AR (2023) Behavior-aware queueing: The finite-buffer setting with many strategic servers. *Operations Research* .

Zhou YP, Ren ZJ (2010) Service outsourcing. *Wiley Encyclopedia of Operations Research and Management Science* .

E-Companion

EC.1. Proofs of the Main Results

This part of the e-companion collects the proofs of the main results in the paper. We omit the proof of Proposition 3 because it is a fairly routine application of Itô's formula. The proofs of all the auxiliary results can be found in §EC.4.

Proof of Proposition 1. Towards establishing part (i), it is useful to define $w := v'$ and note that $\theta^\star = v'/\alpha$. Thus, we can rewrite (9) as

$$\frac{\sigma^2}{2} w'(z) - \frac{(w(z))^2}{2\alpha} + k(z) = \rho^\star.$$

For notational simplicity, in the proof of part (i), we further stipulate that $\sigma^2 = 2$ and $\alpha = 1/2$, so that we need to prove the existence of a pair $(w, \rho^\star)$ that solves

$$w'(z) - (w(z))^2 + k(z) = \rho^\star \tag{EC.1}$$

with $w(0) = 0$, $w'(z) > 0$ for all $z > 0$, and $w(z) \le C\sqrt{k(z)}$ for some $C > 0$ and all sufficiently large z. The preceding stipulations are made without loss of generality; that is, the arguments used below are still valid for general values of σ^2 and α. Now, to conduct the existence proof, we consider the following family of first-order differential equations parameterized by ρ:

$$w_\rho'(z) - (w_\rho(z))^2 + k(z) = \rho \quad \text{with} \quad w_\rho(0) = 0. \tag{EC.2}$$

Within this family, we intend to show that there exists a unique $\rho^\star$ such that the pair $(w_{\rho^\star}, \rho^\star)$ solves (EC.1) subject to the stated requirements. To that end, define

$$z_\rho := \inf\left\{ z \ge 0 : \lim_{x \uparrow z} w_\rho(x) = \pm\infty \right\}.$$

Also, let $\mathcal{C}_\rho := \{z \in [0, z_\rho) : w_\rho'(z) \le 0\}$ and define sets $\mathcal{L}$ and $\mathcal{U}$ that bisect the non-negative real line in the following way:

$$\mathcal{L} := \{\rho \ge 0 : \mathcal{C}_\rho \ne \emptyset\} \quad \text{and} \quad \mathcal{U} := \{\rho \ge 0 : \mathcal{C}_\rho = \emptyset\}.$$

To complete the proof of part (i), we need the following lemmas whose proofs are deferred to §EC.4.

LEMMA EC.1. *For any $\rho_1 < \rho_2$, we have that $w_{\rho_1}(z) < w_{\rho_2}(z)$ for $z \in [0, z_{\rho_1} \wedge z_{\rho_2})$.*

LEMMA EC.2. *For all $\rho \in \mathcal{L}$, w_ρ is quasi-concave and $\lim_{z\to\infty} w_\rho(z) = -\infty$.*

LEMMA EC.3. *Both $\mathcal{L}$ and $\mathcal{U}$ are nonempty.*

LEMMA EC.4. *Let $\rho^\star := \sup \mathcal{L}$. Then $\rho^\star \in \mathcal{U}$, and $w_{\rho^\star}'(z) > 0$ for all $z \ge 0$.*

LEMMA EC.5. *For $\rho \in (0, \rho^\star)$, we have that $w_\rho(z) \le \sqrt{2k(z) + a}$ for some positive constant a.*

Continuing our proof of Proposition 1, Suppose that $w_{\rho^\star}$ does not satisfy the stated growth condition. Then, there exists some $\tilde{z}$ such that

$$w_{\rho^\star}(\tilde{z}) \geq \sqrt{2k(\tilde{z})+a} + \epsilon.$$

On the other hand, by Lemma EC.5 we have that

$$w_\rho(\tilde{z}) \leq \sqrt{2k(\tilde{z})+a} \quad \text{for all} \quad \rho \in (0, \rho^\star).$$

This, however, leads to a contradiction because $w_\rho(z)$ is jointly continuous in ρ and z. Hence, part (i) of the proposition follows.

To establish part (ii), let θ be an admissible rate-control strategy. A routine application of Itô's formula gives the following identity

$$\mathbb{E}[v(Z(t))] - v(Z(0)) = \mathbb{E}\left[\int_0^t \Gamma v(Z(u))\mathrm{d}u + v'(0)L(t)\right],$$

where Γ is an operator and defined as $\Gamma f(z) = (\sigma^2/2)f''(z) - \theta(z)f'(z)$. Since $v'(0) = 0$ and v satisfies (9), we have that

$$\mathbb{E}[v(Z(t))] - v(Z(0)) \leq \mathbb{E}\left[\rho^\star t - \int_0^t c(\theta(u))\mathrm{d}u - \int_0^t k(Z(u))\mathrm{d}u\right].$$

Dividing both sides by t and taking the limsup, we obtain

$$\rho^\star \geq \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c(\theta(u))\mathrm{d}u + \int_0^t k(Z(u))\mathrm{d}u\right] + \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}[v(Z(t))].$$

Next, we intend to establish

$$\limsup_{t\to\infty} t^{-1}\mathbb{E}[v(Z(t))] = 0. \tag{EC.3}$$

To that end, note from part (i) that v has a polynomial growth rate. Thus, to establish (EC.3), it suffices to argue that $\limsup_{t\to\infty} t^{-1}\mathbb{E}[Z^q(t)] = 0$ for all $q > 0$. For this purpose, we need to construct a non-negative process $\tilde{Z}$ such that

$$Z \leq_{s.t.} \tilde{Z} \tag{EC.4}$$

and $\tilde{Z}$ has a steady-state distribution whose moments of all orders are finite. By the admissibility of θ, we know that $\theta(z)$ is strictly increasing in z, so there exists some $\bar{z} > 0$ and $\chi > 0$ such that $\theta(z) > \chi$ for all $z \geq \bar{z}$. This means that if we define $\tilde{Z}$ as

$$\tilde{Z}(t) = Z(0) + \int_0^t b(\tilde{Z}(u))\mathrm{d}u + \sigma B(t) + L(t),$$

where $b(z) := 0 \cdot 1_{\{0 \leq z < \bar{z}\}} - \chi \cdot 1_{\{z \geq \bar{z}\}}$ and L is a pushing process that uses a minimum effort to keep $\tilde{Z}$ non-negative, then (EC.4) holds. Moreover, $\tilde{Z}$ defined in this way is a piecewise-linear diffusion process with reflection at the origin. Using the techniques in Browne et al. (1995), we can write down its steady-state distribution and conclude that the distribution has finite moments of all orders. Taken together, we arrive at (EC.3). Hence,

$$\rho^\star \geq \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c(\theta(u))\mathrm{d}u + \int_0^t k(Z(u))\mathrm{d}u\right].$$

The desired optimality of $\theta^\star$ follows immediately by noting that all the foregoing inequalities become equalities when $\theta = \theta^\star$. This concludes the proof of part (ii). □

Proof of Proposition 2. The maximizing θ on the left-hand of (10) can be expressed as $\theta = -\frac{U'}{\alpha}$. Substituting this expression back into (10) "removes" the $\max$ operator, yielding

$$\frac{\sigma^2}{2}U''(z) + \frac{1}{2\alpha}\left(U'(z)\right)^2 + S(z) + \frac{\gamma\sigma^2}{2}(U'(z))^2 = \eta.$$

If the PC constraint is not binding, meaning that η is strictly positive, the principal can improve her objective without violating the agent's PC constraint by choosing $\tilde{S}(z) := S(s) - \eta$. This implies that the payment plan, which induces θ at a minimum cost, must possess the following property:

$$\frac{\sigma^2}{2}U''(z) + \frac{1}{2\alpha}\left(U'(z)\right)^2 + S(z) + \frac{\gamma\sigma^2}{2}(U'(z))^2 = 0. \tag{EC.5}$$

Replacing U' in (EC.5) with $-\alpha\theta$ yields (11), as desired. □

Proof of Theorem 1. To save us some ink, we will assume throughout the proof that $\gamma = 0$, resulting in $\tilde{\alpha}_\gamma = \alpha$. Starting with part (i), we observe that (13) is, in essence, a first-order linear differential equation (by viewing V' as the unknown function). Therefore, we can apply the standard formula for first-order differential equations with the condition $V'(0) = 0$ to get

$$V'(z) = \frac{2}{\sigma^2}e^{\frac{2}{\sigma^2}\phi(z)}\int_0^z e^{-\frac{2}{\sigma^2}\phi(x)}\left(\kappa - \frac{\sigma^2}{2}\alpha\theta'(x) + \frac{\alpha}{2}\theta^2(x) - k(x)\right)\mathrm{d}x, \tag{EC.6}$$

where we have defined $\phi(z) := \int_0^z \theta(x)\mathrm{d}x$. We intend to find some κ that, when plugged into the right-hand side, causes $V'(z)$ to grow polynomially as $z \to \infty$. This requirement entails that

$$\int_0^z e^{-\frac{2}{\sigma^2}\phi(x)}\left(\kappa - \frac{\sigma^2}{2}\alpha\theta'(x) + \frac{\alpha}{2}\theta^2(x) - k(x)\right)\mathrm{d}x \to 0 \quad \text{as} \quad z\to\infty.$$

In turn, this yields

$$\kappa(\theta) = \frac{\int_0^\infty e^{-\frac{2}{\sigma^2}\phi(x)}\left(\frac{\sigma^2}{2}\alpha\theta'(x) - \frac{\tilde{\alpha}_\gamma}{2}\theta^2(x) + k(x)\right)\mathrm{d}x}{\int_0^\infty e^{-\frac{2}{\sigma^2}\phi(x)}\mathrm{d}x}. \tag{EC.7}$$

By letting $\kappa = \kappa(\theta)$ in (EC.6), we obtain

$$V'(z) = \frac{2}{\sigma^2}e^{\frac{2}{\sigma^2}\phi(z)}\int_z^\infty e^{-\frac{2}{\sigma^2}\phi(x)}\left(\frac{\sigma^2}{2}\alpha\theta'(x) - \frac{\alpha}{2}\theta^2(x) + k(x)\right)\mathrm{d}x.$$

A direct application of L'Hospital's rule reveals that the right-hand side has a growth rate that is of the same order as

$$p(z) := \frac{\frac{\sigma^2}{2}\alpha\theta'(z) - \frac{\alpha}{2}\theta^2(z) + k(z)}{\theta(z)}.$$

By our assumption, $S(z)$ is admissible, so the numerator grows polynomially as $z \to \infty$, while the denominator is bounded below by some positive constant χ for all sufficiently large z. This implies that p has a polynomial growth rate, implying that V' and V exhibit polynomial growth as well. To show that $\kappa(\theta)$ is the long-run average cost of the principal, we must argue that $\limsup_{t\to\infty} t^{-1}\mathbb{E}[V(Z(t))] = 0$. The rest of the proof mimics those steps taken towards establishing part (ii) of Proposition 1. For completeness, we spell out the details below. Specifically, we need to find some non-negative process $\tilde{Z}$ such that $Z \leq_{s.t.} \tilde{Z}$ and $\tilde{Z}$ has a steady-state distribution whose moments of all orders are finite. By the hypothesis, θ is induced by an admissible payment plan, so there exist some $\bar{z} > 0$ and $\chi > 0$ such that $\theta(z) > \chi$ for all $z \geq \bar{z}$. This means we can choose $\tilde{Z}$ to be one that satisfies the following dynamics:

$$\tilde{Z}(t) = Z(0) + \int_0^t b(\tilde{Z}(u))\mathrm{d}u + \sigma B(t) + L(t),$$

where $b(z) := 0 \cdot 1_{\{0\leq z<\bar{z}\}} - \chi \cdot 1_{\{z\geq\bar{z}\}}$ and the rest of the quantities are defined as before. Note that $\tilde{Z}$ defined in this way is a piecewise-linear diffusion with reflection at the origin. Using the techniques in Browne et al. (1995), we can write down its steady-state distribution and conclude that the distribution has finite moments of all orders. This concludes the proof of part (i).

To establish part (ii), note that not only θ and ϕ determine each other, but $\phi' = \theta$ and $\phi'' = \theta'$. Thus, from (EC.7) we can see that minimizing the right-hand side of (EC.7) over all possible θ with the condition $\theta(0) = 0$ is equivalent to seeking some ϕ that minimizes the ratio

$$\frac{\int_0^\infty G(x, \phi(x), \phi'(x), \phi''(x))\mathrm{d}x}{\int_0^\infty H(\phi(x))\mathrm{d}x}, \tag{EC.8}$$

subject to the conditions $\phi(0) = 0$ and $\phi'(0) = 0$, where we defined G and H as

$$G(x,p,q,r) := e^{-\frac{2}{\sigma^2}p}\left(\frac{\sigma^2}{2}\alpha r - \frac{\alpha}{2}q^2 + k(x)\right) \quad \text{and} \quad H(p) := e^{-\frac{2}{\sigma^2}p},$$

respectively. Let ϑ be an arbitrary, twice-differentiable test function that, along with its derivative function, vanishes at the boundaries. Now, let

$$F(\epsilon) := \frac{\int_0^\infty G(x, \phi(x)+\epsilon\vartheta(x), \phi'(x)+\epsilon\vartheta'(x), \phi''(x)+\epsilon\vartheta''(x))\mathrm{d}x}{\int_0^\infty H(\phi(x)+\epsilon\vartheta(x))\mathrm{d}x}.$$

Using the first-order condition, we deduce that if ϕ is a critical point of (EC.8), then $\frac{\mathrm{d}F}{\mathrm{d}\epsilon}|_{\epsilon=0} = 0$. Define

$$g := \int_0^\infty G(x, \phi(x), \phi'(x), \phi''(x))\mathrm{d}x \quad \text{and} \quad h := \int_0^\infty H(\phi(x))\mathrm{d}x.$$

Using $\frac{\mathrm{d}F}{\mathrm{d}\epsilon}|_{\epsilon=0} = 0$ and with some rearrangement, we can deduce that

$$\frac{\int_0^\infty \left[h\left(\frac{\partial G}{\partial \phi} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial G}{\partial \phi'} + \frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial G}{\partial \phi''}\right) - g\left(\frac{\partial H}{\partial \phi}\right)\right]\vartheta(x)\mathrm{d}x}{h^2} = 0,$$

where we have omitted the arguments of the derivative functions. By the fundamental lemma of calculus of variations, the part of the integrand in the angle brackets is zero, yielding the corresponding Euler-Lagrange equation:

$$h\left(\frac{\partial G}{\partial \phi} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial G}{\partial \phi'} + \frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial G}{\partial \phi''}\right) - g\left(\frac{\partial H}{\partial \phi}\right) = 0. \tag{EC.9}$$

Using the definitions of G and H, we can deduce that

$$\begin{aligned} &\frac{\partial G}{\partial \phi} = -\frac{2}{\sigma^2}G, \quad \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial G}{\partial \phi'} = \frac{2\alpha}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}(\phi')^2 - \alpha e^{-\frac{2}{\sigma^2}\phi}\phi'', \\ &\frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial G}{\partial \phi''} = \frac{2\alpha}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}(\phi')^2 - \alpha e^{-\frac{2}{\sigma^2}\phi}\phi'' \quad \text{and} \quad \frac{\partial H}{\partial \phi} = -\frac{2}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}. \end{aligned} \tag{EC.10}$$

Plugging (EC.10) into (EC.9) and with some arrangement, we have that

$$h\left(\frac{\sigma^2}{2}\alpha\phi''(x) - \frac{\alpha}{2}(\phi'(x))^2 + k(x)\right) - g = 0.$$

Letting $\tilde{\kappa} := g/h$ in the preceding and noting that $\theta = \phi'$, we reach (15), as desired. □

REMARK EC.1. In minimizing the right-hand side of (EC.7), we use the calculus of variations. However, the functional (EC.8) is not in the canonical form of the calculus of variations because it is a ratio of two integrals. Nevertheless, we claim that using the calculus of variations directly on the functional (EC.8) is equivalent to using it on the canonical form. The idea is to convert the ratio of two integrals into a constrained optimization problem. We can define $\mathcal{L}(x, p, q, r) := G(x, p, q, r) - \tilde{\kappa}H(p)$, where $\tilde{\kappa}$ can be viewed as a Lagrange multiplier. By a slight abuse of notation, we also define

$$F(\epsilon) := \int_0^\infty \mathcal{L}(x, \phi(x) + \epsilon\vartheta(x), \phi'(x) + \epsilon\vartheta'(x), \phi''(x) + \epsilon\vartheta''(x))\mathrm{d}x.$$

Therefore, the associated Euler-Lagrange equation can be derived as

$$\frac{\partial \mathcal{L}}{\partial \phi} - \frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial \mathcal{L}}{\partial \phi'} + \frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial \mathcal{L}}{\partial \phi''} = 0. \tag{EC.11}$$

Using the definition of $\mathcal{L}$, we can deduce that

$$\begin{aligned} &\frac{\partial \mathcal{L}}{\partial \phi} = -\frac{2}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}\left(\frac{\sigma^2}{2}\alpha\phi'' - \frac{\alpha}{2}(\phi')^2 + k(x) - \tilde{\kappa}\right), \\ &\frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial \mathcal{L}}{\partial \phi'} = -\alpha\phi'' e^{-\frac{2}{\sigma^2}\phi} + \frac{2\alpha}{\sigma^2}(\phi')^2 e^{-\frac{2}{\sigma^2}\phi}, \quad \text{and} \\ &\frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial \mathcal{L}}{\partial \phi''} = -\alpha\phi'' e^{-\frac{2}{\sigma^2}\phi} + \frac{2\alpha}{\sigma^2}(\phi')^2 e^{-\frac{2}{\sigma^2}\phi}. \end{aligned} \tag{EC.12}$$

Plugging (EC.12) into (EC.11) yields that

$$-\frac{2}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}\left(\frac{\sigma^2}{2}\alpha\phi''(x)-\frac{\alpha}{2}(\phi'(x))^2+k(x)-\tilde{\kappa}\right)=0,$$

again establishing Equation (15).

Proof of Proposition 4. Let $Z^\star$ denote the state process when the principal employs $S^\star$ and the agent adopts the rate control induced by $S^\star$. By the statement of the proposition,

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S(Z(u))\mathrm{d}u+\int_0^t k(Z(u))\mathrm{d}u-\int_0^t \lambda(Z(u))S(Z(u))\mathrm{d}u\right]$$

is greater than or equal to

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S^\star(Z^\star(u))\mathrm{d}u+\int_0^t k(Z^\star(u))\mathrm{d}u-\int_0^t \lambda(Z^\star(u))S^\star(Z^\star(u))\mathrm{d}u\right]$$

for all admissible S. In turn, this implies that

$$\begin{aligned}\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S^\star(Z^\star(u))\mathrm{d}u+\int_0^t k(Z^\star(u))\mathrm{d}u\right]&\le\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S(Z(u))\mathrm{d}u+\int_0^t k(Z(u))\mathrm{d}u\right]\\&+\underbrace{\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t \lambda(Z^\star(u))S^\star(Z^\star(u))\mathrm{d}u\right]}_{\Pi_1}-\underbrace{\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t \lambda(Z(u))S(Z(u))\mathrm{d}u\right]}_{\Pi_2}.\end{aligned}$$

By the complementary slackness condition, we have $\Pi_1=0$. By the primal and dual feasibility constraints, we have $\Pi_2\ge 0$. It follows that for all admissible S,

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S^\star(Z^\star(u))\mathrm{d}u+\int_0^t k(Z^\star(u))\mathrm{d}u\right]\le\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t S(Z(u))\mathrm{d}u+\int_0^t k(Z(u))\mathrm{d}u\right].$$

The proof is therefore complete. □

Proof of Theorem 2. To prove part (i), we introduce the following proposition:

PROPOSITION EC.1. *There exists a continuously differentiable function f and a constant ρ that satisfy*

$$\min_{\theta\in[0,b]}\left\{\frac{\sigma^2}{2}f'(z)-\theta f(z)+k(z)+c(\theta)\right\}=\rho$$

with $f(0)=0$, $f'(z)>0$ for all $z\ge 0$, and $f(z)$ grows polynomially as $z\to\infty$.

This proposition is similar to Proposition 1, with the only difference being that in Proposition EC.1, θ is restricted to the range $[0,b]$, where we recall $b=\sqrt{2\eta/\alpha}$. Proposition EC.1 can be proven in almost the same fashion as in Proposition EC.1, so we omit the details of the proof.

Proposition EC.1 implies the existence of some ζ such that

$$\frac{\sigma^2}{2}f'(z)-\frac{1}{2\alpha}f^2(z)+k(z)=\rho\quad\text{for}\quad z\in[0,\zeta)$$

and

$$\frac{\sigma^2}{2}f'(z)-bf(z)+\frac{\alpha}{2}b^2+k(z)=\rho \quad \text{for} \quad z\in[\zeta,\infty).$$

In particular, ζ is such that $f(\zeta)=\alpha b$. Now, set $\tilde{\kappa}=\rho$, $\bar{z}=\zeta$, and let $\theta=\alpha f$ on $[0,\bar{z})$ and $\theta\equiv b$ on $[\bar{z},\infty)$, and $V_\theta'=f-\alpha b$ on $[\bar{z},\infty)$. It is directly verifiable that the resulting quadruple $(\theta,\tilde{\kappa},\bar{z},V_\theta)$ meets all requirements arising from the stationarity condition, the primal and dual feasibility constraints, and the complementary slackness condition. The proof of part (i) is therefore complete.

To prove part (ii), it suffices to let $\lambda=2V_\theta/(\sigma^2\alpha)$ and note that V_θ satisfies equation (32). □

EC.2. Derivation of Equation (25)

Define

$$\tilde{G}(x,p,q,r):=e^{-\frac{2}{\sigma^2}p}\left(\frac{\sigma^2}{2}\alpha r-\frac{\alpha}{2}q^2+k(x)-\lambda(x)\left(\frac{\sigma^2}{2}\alpha r-\frac{\alpha}{2}q^2+\eta\right)\right).$$

Note that we still have the Euler-Lagrange equation:

$$h\left(\frac{\partial\tilde{G}}{\partial\phi}-\frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial\tilde{G}}{\partial\phi'}+\frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial\tilde{G}}{\partial\phi''}\right)-g\left(\frac{\partial H}{\partial\phi}\right)=0.$$

Similar to equation (EC.10), using the definitions of $\tilde{G}$ and H (where the definition of H can be found in §EC.1), we can deduce that

$$\begin{aligned}
&\frac{\partial H}{\partial\phi}=-\frac{2}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}, \quad \frac{\partial\tilde{G}}{\partial\phi}=-\frac{2}{\sigma^2}\tilde{G},\\
&\frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial\tilde{G}}{\partial\phi'}=\frac{2\alpha}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}(\phi')^2-\alpha e^{-\frac{2}{\sigma^2}\phi}\phi''+\alpha e^{-\frac{2}{\sigma^2}\phi}(\lambda'(x)\phi'+\lambda(x)\phi'')-\frac{2\alpha}{\sigma^2}e^{-\frac{2}{\sigma^2}\phi}(\phi')^2\lambda(x),\\
&\text{and}\quad \frac{\mathrm{d}^2}{\mathrm{d}x^2}\frac{\partial\tilde{G}}{\partial\phi''}=\frac{\mathrm{d}}{\mathrm{d}x}\frac{\partial\tilde{G}}{\partial\phi'}+\alpha\lambda'(x)\phi'e^{-\frac{2}{\sigma^2}\phi}-\frac{\alpha\sigma^2}{2}\lambda''(x)e^{-\frac{2}{\sigma^2}\phi}.
\end{aligned} \tag{EC.13}$$

Plugging equation (EC.13) back and with some arrangement, we have that

$$h\left(\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+k(z)-\lambda(z)\left(\frac{\sigma^2}{2}\alpha\theta'(z)-\frac{\alpha}{2}\theta^2(z)+\eta\right)+\frac{\sigma^2}{2}\alpha\left(\frac{\sigma^2}{2}\lambda''(z)-\theta(z)\lambda'(z)\right)\right)-g=0.$$

Letting $\tilde{\kappa}:=g/h$ in the preceding and noting that $\theta=\phi'$, we reach equation (25).

EC.3. Justifying the Use of an RBM to Model Wait-Time Dynamics

The demonstration below is largely adapted from the paper by Huang and Gurvich (2018), where the authors show that a universally near-optimal control policy for a single-server queue can be obtained via an intuitive Brownian control problem.

Consider an M/G/1 queue where a single server processes Poisson arrivals with service requirements distributed according to a general distribution. Let $A(t)$ denote the number of arrivals by time t, and let $\{\xi_i;t\geq 1\}$ denote the service requirements of consecutive jobs that enter the system. The process

$A := \{A(t); t \geq 0\}$ is a Poisson process with some rate λ, and $\{\xi_i; t \geq 1\}$ are independent random variables that follow a common distribution G with a mean of $1/\mu$.

The workload, which corresponds to the virtual wait time in the M/G/1 queue, can be represented as a process that evolves according to the following equation:

$$\begin{aligned} W(t) &= W(0) + \sum_{i=1}^{A(t)} \xi_i - (t - I(t)) \\ &= W(0) + (\varrho - 1)t + I(t) + \left(\sum_{i=1}^{A(t)} \xi_i - \varrho t \right), \end{aligned}$$

where $\varrho := \lambda/\mu$, and $I(t)$ denotes the cumulative idle time of the server by time t. At each time t, the compound Poisson input of the M/G/1 queue can be approximated by a central limit theorem, given by:

$$\sum_{i=1}^{A(t)} \xi_i \approx \varrho t + \mathcal{N}(t),$$

where $\mathcal{N}(t)$ is a zero-mean normal random variable with variance $\lambda \mathbb{E}[\xi_1^2] t$. As a result, it is reasonable to replace the input process with $\varrho t + \sqrt{\lambda \mathbb{E}[\xi_1^2]} B(t)$, where $B := B(t), t \geq 0$ is a standard Brownian motion. This approximation leads to the Brownian queue:

$$Z(t) = Z(0) + (\varrho - 1)t + L(t) + \sqrt{\lambda \mathbb{E}[\xi_1^2]} B(t),$$

where W and I have been replaced by Z and L, respectively, to emphasize that they are approximations. In the case of exponentially distributed service requirements, we have that $\sqrt{\lambda \mathbb{E}[\xi_1^2]} = \sqrt{2\lambda/\mu^2}$. Assuming $\lambda = \mu$ and letting $\sigma := \sqrt{\lambda \mathbb{E}[\xi_1^2]}$, we obtain:

$$Z(t) = Z(0) + \sigma B(t) + L(t),$$

which is the uncontrolled version of equation (1).

EC.4. Proofs of Auxiliary Results

Proof of Lemma EC.1. The desired conclusion follows directly from the comparison principle for first-order differential equations. □

Proof of Lemma EC.2. Since $w_\rho'(0) = \rho \geq 0$, it suffices to show that w_ρ does not have strict local minima. From (EC.2), we see that $w_\rho''(z) = -k'(z) < 0$ whenever $w_\rho'(z) = 0$. Hence, w_ρ does not have strict local minima. For the second part of the lemma, it suffices to demonstrate that w_ρ cannot be bounded from below by a finite constant. If so, then there exists a sequence $\{z_l\}$ such that $\lim_{l\to\infty} w_\rho'(z_l) = 0$ and $\lim_{l\to\infty} w_\rho(z_l) = b$ for some $b \in \mathbb{R}$. But this leads to a contradiction in view of (EC.2) and the fact that $\lim_{l\to\infty} k(z_l) = \infty$. □

Proof of Lemma EC.3. By definition, we have that $0 \in \mathcal{L}$. To show that $\mathcal{U}$ is not empty, we first claim that the set $\bar{\mathcal{U}}$ defined as $\bar{\mathcal{U}} := \{\rho \geq 0 : w_\rho(1) \geq 1\}$ is not empty. Since

$$w_\rho(1) = \rho + \int_0^1 (w_\rho(x))^2 \mathrm{d}x - \int_0^1 k(x)\mathrm{d}x \geq \rho - \int_0^1 k(x)\mathrm{d}x,$$

we know that $\lim_{\rho\to\infty} w_\rho(1) = \infty$. It follows that $\bar{\mathcal{U}}$ is nonempty, establishing the claim. Also, if letting $\rho_1 := \inf \bar{\mathcal{U}}$, then $\rho_1 < \infty$. Suppose, for the sake of contradiction, that $\mathcal{U} = \emptyset$. Then for any $\rho > \rho_1$, there exists $z_{1,\rho} > 1$ such that $w_\rho(z_{1,\rho}) = 1$. For $z \in [1, z_{1,\rho}]$, we have that $(w_\rho(z))^2 \geq w_\rho(z)$, and therefore $w_\rho(z) \geq \underline{w}_\rho(z)$, where $\underline{w}_\rho$ solves

$$\underline{w}_\rho'(z) - \underline{w}_\rho(z) + k(z) = \rho$$

with $\underline{w}_\rho(1) = w_\rho(1)$. A straightforward calculation yields

$$\underline{w}_\rho(z) = e^{z-1} w(1) + e^z \int_1^z e^{-x}(\rho - k(x))\mathrm{d}x.$$

The second term is strictly positive for sufficiently large ρ. Thus, $\underline{w}_\rho(z) > 1$ for all $z \in [1, z_{1,\rho}]$ when ρ is large enough. But this contradicts $w_\rho(z_{1,\rho}) = 1$. Therefore, $\mathcal{U}$ is not empty. □

Proof of Lemma EC.4. By Lemma EC.3, we know that $\rho^\star < \infty$. We next show that $\rho^\star \in \mathcal{U}$. To reach a contradiction, suppose that $\rho^\star \in \mathcal{L}$. Then, by Lemma EC.2, we know that there exists a maximum point of $w_{\rho^\star}$, denoted as $z^\star$, and let $w^\star := w_{\rho^\star}(z^\star) < \infty$. By the same lemma, for some given $\delta > 0$, there exists $\epsilon > 0$ such that $w_{\rho^\star}(z^\star + \epsilon) = w^\star - \delta$. On the other hand, for all $\rho > \rho^\star$, we have that $\rho \in \mathcal{U}$, and therefore,

$$w_\rho(z^\star + \epsilon) > w_\rho(z^\star) > w_{\rho^\star}(z^\star) = w^\star \quad \text{if} \quad \rho \in \mathcal{U}.$$

This is a contradiction since $w_\rho(z)$ is jointly continuous in ρ and z. Therefore, $\rho^\star \in \mathcal{U}$ and $w'_{\rho^\star}(z) > 0$ for all $z \geq 0$. □

Proof of Lemma EC.5. First, let a be a positive constant such that

$$\sqrt{a}(k(z) + a) - k'(z) > 0 \quad \text{for all} \quad z \geq 0.$$

Note that the existence of such a positive constant is ensured by Assumption 1. Next, for every $\rho \in (0, \rho^\star)$, let $\hat{z}_\rho$ be such that $w_\rho(\hat{z}_\rho) = \sup_{z\geq 0} w_\rho(z)$. Since $w'_\rho(\hat{z}_\rho) = 0$, we have that

$$w_\rho(\hat{z}_\rho) = \sqrt{k(\hat{z}_\rho) - \rho} < \sqrt{2k(\hat{z}_\rho) + a} \quad \text{for} \quad \rho \in (0, \rho^\star).$$

Now, let $f_\rho(z) := (w_\rho(z))^2 - 2k(z) - a$. We intend to argue that $f_\rho(z) \leq 0$ for all $z \in [0, \hat{z}_\rho]$. For that purpose, suppose by way of contradiction that this is not true. Then there must exist $\underline{z}_\rho$ and $\bar{z}_\rho$ with $0 \leq \underline{z}_\rho < \bar{z}_\rho < \hat{z}_\rho$

such that $f_\rho(\underline{z}_\rho) = f_\rho(\bar{z}_\rho) = 0$, and $f_\rho(z) > 0$ for all $z \in (\underline{z}_\rho, \bar{z}_\rho)$. In particular, the foregoing implies that $w_\rho(z) > \sqrt{2k(z)+a}$, and so

$$w_\rho'(z) = (w_\rho(z))^2 - k(z) + \rho > k(z) + a$$

for every $z \in (\underline{z}_\rho, \bar{z}_\rho)$ and $\rho \in (0, \rho^\star)$. Thus,

$$f_\rho'(z) = 2w_\rho(z)w_\rho'(z) - 2k'(z) > 2\sqrt{a}(k(z)+a) - 2k'(z) > 0 \quad \text{for all} \quad z \in (\underline{z}_\rho, \bar{z}_\rho).$$

This, however, contradicts our hypothesis which holds that $f_\rho(\underline{z}_\rho) = f_\rho(\bar{z}_\rho) = 0$. Therefore, $f_\rho(z) \le 0$ for all $z \in [0, \hat{z}_\rho]$, from which the desired result follows. $\square$