Enhancing Make-to-Order Manufacturing Agility: When Flexible Capacity Meets Dynamic Pricing

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

Shiwei Chai, Anand Paul
Department of Information Systems and Operations Management
University of Florida, Gainesville, Florida 32608, USA, s.chai@ufl.edu, paulaa@ufl.edu

Lingjiong Zhu
Department of Mathematics
Florida State University, Tallahassee, Florida 32306, USA, zhu@math.fsu.edu

Abstract: The rise of online marketplaces has raised customer expectations regarding customization and lead time. It poses significant challenges to manufacturing firms and prompts a move from make-to-stock to a more flexible make-to-order system. Compared to make-to-stock settings, make-to-order systems cannot smooth fluctuations in demand using available stock. While viewing dynamic pricing as a useful strategy to balance supply with demand, many manufacturing firms can also create capacity flexibility. In that scenario, system costs could be cut by managing capacity and demand simultaneously. In this paper, we consider a make-to-order production environment with base and surge capacity as well as the ability to adjust product pricing. Our main focus is on operational decision-making, assuming that base and surge capacity are fixed, but activating the surge capacity incurs a setup cost. Initially, we propose a stochastic control model to reflect this complex decision problem. However, our initial model leads to an intractable dynamic programming problem. To overcome this, we convert the problem to a more tractable diffusion control problem. This approach helps to reveal the conditions under which utilizing flexible capacity is more advantageous than relying solely on fixed capacity. When flexible capacity is advantageous, we provide a solution to the diffusion control problem that can guide optimal capacity and price adjustments. We discover an interesting interplay between capacity adjustment and dynamic pricing. In particular, we find that the price, which aims at reducing congestion, may not monotonically increase with the congestion level when capacity adjustments incur a fixed cost.

Key words: make-to-order production; flexible capacity; dynamic pricing; queueing; diffusion models

History: Received: March 2023; Accepted March 2024 by Michael Pinedo, after 2 revisions.

---

1 Introduction

The rise of online marketplaces such as Amazon, eBay, and Walmart has brought significant challenges to manufacturing firms. For one thing, customers are now expecting highly customized items with even shorter lead times. This is also true in business-to-business sectors, where individualized, engineering-to-order approaches are becoming more common. Another trend is the ongoing digitization of manufacturing, which results in mass customization and a shorter product life cycle. For those reasons, manufacturers are urged to move from a traditional make-to-stock (MTS) manufacturing environment to a more flexible make-to-order (MTO) manufacturing environment because demand for highly customized products cannot be economically produced to stock (similar to customer services that cannot be physically stored). However, a major challenge in implementing the MTO approach is matching supply with demand. This is because fluctuations in demand cannot be smoothed by available stock, potentially causing long delays in order processing in periods of high demand and idle capacity in periods of low demand. While longer wait times lead to customer dissatisfaction and/or loss of customer lifetime value, idle capacity results in investment waste.

From a demand-control perspective, dynamic pricing has been recognized as a popular strategy for increasing supply chain efficiency in MTO production setups. For instance, Tesla adjusts the price of its electric vehicles based on supply and demand, in addition to customers' selected features and options. Similarly, Nike offers custom-designed sneakers on its website, where prices vary based on colors, material choices, and market conditions. To further balance demand and supply, many manufacturing firms have capacity flexibility at their disposal. A common approach to achieving capacity flexibility in manufacturing is to combine base capacity with surge capacity. Base capacity refers to the amount of production that can be sustained over a long period of time, while surge capacity allows for temporary increases in production to meet unexpected demand or short-term spikes in orders. In the automotive industry, car manufacturers can achieve capacity flexibility by opening and closing production lines. For instance, a factory with $n+1$ production lines can choose between running n or $n+1$ lines. In this case, n of the lines constitute the base capacity, while the remaining line can be turned on and off as needed, thereby serving as surge capacity (Wu and Chao 2014). Alternatively, surge capacity can be procured by commissioning an external capacity provider (ECP) who offers *manufacturing as a service*.[1] For example, in the shoe manufacturing industry, an ECP can be a 3D printing bureau that accepts outsourced orders from a shoe factory, enabling the shoe factory to expand its otherwise rigid internal production capacity as needed.

In this paper, we study a queueing model that serves as a parsimonious mathematical representation of an MTO manufacturing system offering multiple products. The system features both base and surge production capacity, and it adapts pricing to balance supply with demand. We model the base capacity as the primary production server and the surge capacity as the secondary server. The primary server is continuously operational, whereas the secondary server can be switched on and off as needed to meet fluctuating demand. However, activating the secondary server incurs a fixed setup cost of C, necessitating careful management of switching operations. As the system does not keep inventory, delays in order fulfillment pose a concern. In addition, the system can adjust product prices dynamically, influencing order arrival rates. The goal is to identify a joint capacity adjustment and dynamic pricing strategy that maximizes long-term average profit, calculated as profit from product sales minus waiting costs. Notably, capacity adjustments are path-dependent

[1] An ECP provides on-demand flexible production capabilities through its own production facilities, allowing for remote production if it is logistically feasible (Rauschecker et al. 2014).

decisions as they involve determining the timing of activating and deactivating the secondary server. To jointly optimize capacity adjustment and pricing decisions, we adopt a rigorous stochastic control framework that accounts for the stochastic uncertainties in both demand and production.

The problem of finding an optimal joint capacity adjustment and dynamic pricing strategy is generally complex and not amenable to exact analysis. Therefore, we apply standard approximation techniques commonly used in the literature. We assume that the manufacturing firm operates with a large demand volume and is functioning at a critical capacity, which we will define later. In this scenario, we can use the traditional heavy-traffic regime as an approximation to the original stochastic control problem. By analyzing the solution to the diffusion control problem (DCP) that arises in the heavy-traffic regime, we can readily convert it into a control strategy for the original MTO system. We use a diffusion approximation to analyze the queueing system in our model for two main reasons. First, the manufacturing firm we model operates at a high demand volume with in-house capacity nearly matching the demand, which is a suitable parameter regime to apply heavy-traffic approximation. Second, a diffusion approximation allows decisions to be optimized “under very weak distributional restrictions on model input” (Bradley and Glynn 2002).

We outline the contributions of this paper in three aspects.

Modeling: We utilize a formal stochastic control framework to investigate joint capacity adjustment and dynamic pricing. The two control levers, capacity adjustment and pricing, generate complex feedback loops in the system dynamics, resulting in a challenging decision problem. We overcome this challenge by deriving and solving an approximating DCP. We then translate the solution to the DCP into an easy-to-implement joint capacity adjustment and pricing strategy. In contrast to the prior research, our model incorporates both a fixed cost of expediting and stochastic expedited processing times. We further generalize the model to allow for on-hand inventory, enabling us to model a make-to-stock setting.

Methodology: This paper establishes the well-posedness of the solution to the Bellman equation that arises from the DCP. Distinct from the extant literature on optimal switching that needs to deal with a pair of linear differential equations (thereby admitting explicit solutions), our work requires dealing with a pair of nonlinear differential equations that do not satisfy the global Lipschitz continuity condition and cannot be explicitly solved. We overcome the technical challenge by analyzing the structural properties of the differential equations. This technical hurdle methodologically distinguishes our paper from prior works on optimal switching.

Managerial Insights: Our model and analysis provide several useful insights into the control of MTO systems that have access to both flexible capacity and real-time pricing capability. Taking a single-product system as an example, we find that when the fixed cost is not prohibitively high, meaning that surge capacity is economically valuable, the manufacturing firm should use a capacity adjustment strategy that activates surge capacity when the number of outstanding orders exceeds a certain level b and turns it off when it falls below another level a. Furthermore, we uncover a rich interplay between on-demand capacity adjustments and dynamic pricing decisions. Interestingly, we show that the price does not always increase with the congestion level. Instead, the manager may choose to raise prices first to curb customer demand as the congestion level rises to a critical point, after which lowering the price will best serve the manager's interests. To the best of our knowledge, this is the first paper showing a non-monotonic structure for congestion pricing when the congestion level is unobservable. We attribute this "abnormality" to the capacity expansion and shrinkage mechanisms and elaborate on it later in the paper.

The remainder of this paper is structured as follows: Section 2 summarizes the relevant literature. Section 3 introduces our model. Due to the intractability, we formulate and solve an approximating DCP in Section 4, along with the main theoretical results. We also briefly demonstrate how these results carry over to an MTS system in the same section. Section 5 describes our proposed policy and elaborates on the atypical pricing structure. We present extensive numerical studies in Section 6 and conclude in Section 7.

2 Literature Review

Pricing in queues. Our approach to modeling dynamic pricing and its impact on demand is based on a stream of research studying dynamic pricing for queues. Yoon and Lewis (2004) study a problem of dynamic pricing and admission control for a system where arrival and service rates are nonstationary and customers are sensitive to prices. They establish several structural properties of the optimal policy, including the monotonicity of the optimal prices in the state of the system under various cost structures. Ata and Shneorson (2006) consider the problem of dynamically controlling the arrival and service rates in a service facility to optimize long-run average system welfare. They propose a solution to determine the optimal dynamic prices and service rates a system manager should set when serving delay-sensitive, rational customers. Kim and Randhawa (2018) examine the value of dynamic pricing to maximize revenues in queueing systems. They demonstrate that when demand volume is on the order of n, static pricing causes a revenue loss of the order of $n^{1/2}$, whereas dynamic pricing can reduce the loss to the order of $n^{1/3}$, and a two-price policy can achieve most of the benefits of dynamic pricing. In an earlier work, Hall et al. (2009) observe a similar result: what they call the "constant price up to cutoff state" (a special two-price policy in Kim and Randhawa (2018)) results in revenue very close to that gained by a general dynamic pricing policy. More recently, Ata and Barjesteh (2023) consider the dynamic control of a multiclass make-to-stock queue, where multiple types of products are produced and stored in inventory to satisfy customer demand. The paper incorporates pricing and outsourcing decisions, extending the classic model studied in Wein (1992). Gao and Huang (2023) consider a model slightly more general than that in Ata and Barjesteh (2023) and establish the existence of a unique smooth solution to the associated Bellman equation. They also prove the asymptotic optimality of the proposed policy. Our study is most closely related to Çelik and Maglaras (2008) in that both their paper and ours consider a control problem faced by an MTO manufacturer who offers multiple products to customers to maximize profit through dynamic pricing, scheduling, and using surge capacity. In addition, both their paper and ours conduct analysis based on the approximating DCPs. There are, however, several notable distinctions between these two works. First, in their work, orders can be instantaneously expedited via surge capacity, whereas we assume the order processing time on the surge capacity follows some general distributions. This leads to very different DCPs: the controlled process in Çelik and Maglaras (2008) is a one-dimensional workload process, whereas the controlled process in our paper consists of both a one-dimensional workload process and a stochastic process taking binary values indicating the status of the surge capacity. Second, our model incorporates both a setup cost and a per-unit-of-time manufacturing cost related to the surge capacity, whereas in Çelik and Maglaras (2008) the system is assumed to incur a cost per expedited order. Our consideration of setup cost C leads to an atypical structure of pricing scheme, as shown in §5.2, a feature absent in Çelik and Maglaras (2008). Third, Çelik and Maglaras (2008) approximate the revenue function by a Taylor expansion, making their Bellman equation a Riccati equation, whereas we do not make approximations to our profit function. Although this difference adds another level of complexity to our analysis, it shrinks the approximation errors, as shown in §6.1.1.

Optimal switching. This paper falls into a problem category known as "optimal switching", which involves switching costs considered as fixed investments required to realize the operational advantages of an appropriate regime. Duckworth and Zervos (2001) study an optimal two-regime switching problem and derive and solve the Bellman equation. Zervos et al. (2013) consider the problem of an investor who seeks to maximize the expected discounted cash flow by sequentially buying and selling one share of an asset while incurring fixed transaction costs. They model the underlying asset price using a general one-dimensional Itô diffusion and solve the resulting stochastic control problem in a closed analytic form. The aforementioned papers consider the discounted cost criterion and therefore are not applicable to the average cost case, which is the focus of this paper. Wu and Chao (2014) investigate a control problem for a stochastic production/inventory system with two production modes and frame the problem as an optimal switching problem for Brownian motion. Adopting an average cost criteria, they establish that the optimal production policy is of the (s, S) type. To solve the Bellman equation, Wu and Chao (2014) use a similar approach to ours by considering a parametric family of function pairs and identifying a pair whose area of intersection is precisely the switching cost. However, unlike our approach, which involves solutions to nonlinear differential equations, the pair of functions in their analysis are linear differential equations that are explicitly solvable. We make methodological advances by providing an explicit construction of a solution to the Bellman equation, not relying on explicit formulas. Along the way, we establish useful properties of the solution to the Bellman equation. These properties, in turn, facilitate the identification of important structures of our control policy. Concurrently, Sun and Liu (2023) study an optimal switching problem that arises from optimizing short-term capacity adjustment decisions in the context of many-server queues; therein, they also consider customer scheduling and examine the effect of capacity adjustment decisions on scheduling.

Joint rate and impulse control. Our problem reduces to an instance of joint rate and impulse control if the speed of service at the secondary server is infinite. Amidst this problem category, Yao (2017) analyzes an infinite-horizon, continuous-review stochastic inventory system with price-dependent cumulative customer demand modeled as a Brownian motion and excess demand backlogged. The author shows that inventory control follows the $(s^{\star}, S^{\star})$ policy and characterizes the optimal state-dependent pricing strategy analytically. Cao and Yao (2018) investigate a stochastic inventory system where pricing and inventory controls are modeled as rate and impulse controls, respectively. Using stochastic control techniques, they demonstrate that inventory control is of the control-band type and that the optimal drift rate is first increasing and then decreasing in the relevant domain. Sun and Zhu (2024a) study a robust impulse control problem that arises from the need to devise sequencing and outsourcing strategies in the face of model misspecification. Their control problem involves a drift rate controlled by nature that serves as an imaginary adversary to promote robustness.

Controlling queues in heavy traffic. From a methodological standpoint, this paper is related to the stream of research that studies queueing controls in heavy traffic, which can be seen in various works such as Plambeck (2004), Ata et al. (2005), Ghosh and Weerasinghe (2007, 2010), Ghamami and Ward (2013), and Liu and Sun (2022). These works often utilize heavy traffic approximations that result in either drift-rate or singular control problems for diffusion models, which are more amenable to analysis.

Demand and/or capacity management in make-to-stock systems. Last but not least, this paper is related to the literature on demand and/or capacity management in make-to-stock systems. We remark that we allow on-hand inventory, which extends our model to that of a make-to-stock system. Bradley and Glynn (2002) consider a manufacturing system that produces a single product using a make-to-stock approach. In their model, inventory is managed through a base-stock policy, and capacity and inventory are treated as joint decision variables. Although capacity and inventory decisions are optimized jointly in Bradley and Glynn (2002), they occur on different time scales: capacity decisions are regarded as strategic decisions, whereas inventory decisions are operational. By contrast, in our model, capacity adjustments fall under operational decision-making. Some make-to-stock systems can prioritize orders based on profitability and other factors, which is yet another type of capacity management. For example, in Janakiraman et al. (2018), fixed capacity is allocated among a finite set of products, making the allocation of scarce capacity among a suite of products a key ingredient of their problem. Demand control, particularly in the form of pricing, has also been investigated in the context of a make-to-stock environment. For example, Federgruen and Heching (1999) examine the joint determination of pricing and inventory replenishment strategies in both finite and infinite horizon models. Their objective is to maximize the total expected discounted profit, or its time-average value.

3 Model

Notations. Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ be a filtered probability space supporting all random quantities of interest. We use $1_{\{\cdot\}}$ to denote an indicator function and $\Delta X(t) := X(t) - X(t^-)$ the jump of some cdlg process X at time t. We let $[x]^+ := \max(0,x)$ and $[x]^- := -\min(0,x)$ denote the positive and negative parts of x, respectively. For two continuous functions $f,g:\mathbb{R}\to\mathbb{R}$, we say $x\in\mathbb{R}$ is an "intersect point" if $f(x)=g(x)$, and x is a "cross point" if x is an "intersect point" and $f-g$ changes sign at x. We use bold-face letters to represent vectors and write the inner product of $\boldsymbol{a}$ and $\boldsymbol{b}$ as $\langle \boldsymbol{a}, \boldsymbol{b}\rangle$. We say $\boldsymbol{a}\geq \boldsymbol{b}$ if $\boldsymbol{a}$ and $\boldsymbol{b}$ have the same dimension and the inequality holds component-wise.

Consider an MTO production system producing multiple products, indexed by $k=1,\dots,K$, to a market of price-sensitive customers. Conforming to the common language in queueing theory, we will refer to a request for one unit of product k as a class k job. We model the arrival process of class k jobs as a non-homogeneous Poisson process. Specifically, the number of class k jobs that have arrived up to time t is

$$A_k(t) := N_k\left(\int_0^t \lambda_k(u)\mathrm{d}u\right),$$

where N_k is a unit-rate Poisson process and is independent of anything else. Jobs that cannot be processed immediately upon arrival are held in a dedicated queue with an infinite buffer, and jobs in the same dedicated queue are processed on a first-come, first-served basis. We call $\boldsymbol{\lambda}(t) := (\lambda_k(t))$ the instantaneous demand rate vector at time t, and use $\{\boldsymbol{\lambda}(t): t\geq 0\}$ to represent the K-dimensional demand rate process. The system manager can determine the demand rate vector at time t by quoting a price vector $\boldsymbol{p}(t) := (p_k(t))$, where $p_k(t)$ represents the price for product k at time t. We assume that a demand function Ξ exists to capture the price-sensitivity of demand by mapping each price vector to an instantaneous demand rate vector, i.e., $\boldsymbol{\lambda}(t) = \Xi(\boldsymbol{p}(t))$. Following the literary convention, we assume that Ξ has an inverse function so that we can define the profit rate π as a function of the demand rate vector; that is, $\pi(\boldsymbol{\lambda}) := \langle \boldsymbol{\lambda}, \Xi^{-1}(\boldsymbol{\lambda}) - \boldsymbol{q}\rangle$, where $\boldsymbol{q} := (q_k)$ is the vector of the product-specific per-unit production costs. We assume π admits a unique maximum point $\bar{\boldsymbol{\lambda}}$ that lies in the interior of the feasible region $\{\boldsymbol{\lambda}\geq \boldsymbol{0} : \pi(\boldsymbol{\lambda})\geq 0\}$. We refer to $\bar{\boldsymbol{\lambda}}$ as the *nominal demand*. Because we can use Ξ^{-1} to infer the price from the demand rate, we will interchangeably refer to $\boldsymbol{\lambda}$ and $\boldsymbol{p}$ as the pricing decisions.

The time to produce a class k product on the base capacity (primary server) is characterized by a general distribution G_k with a mean of $1/\mu_k$ and a squared coefficient of variation of ν_k^2. Accordingly, we can define a renewal process S_k to represent the number of class k jobs produced by the primary server up to time t if the server were continuously working, where $S_k(t) := \max\{n : \sum_{i=1}^n \xi_i \leq t\}$ and ξ_i are independently and identically distributed (i.i.d.) in G_k. Let $T_k(t)$ be the total amount of time devoted to product k on the primary server, such that the total number of jobs completed on the primary server up to time t is $S_k(T_k(t))$. Since the primary server is continuously operational, we have $\sum_{k=1}^K T_k(t)\leq t$ and

$$L_1(t) := t - \sum_{k=1}^K T_k(t) \tag{1}$$

represents the primary server's idle time.

Aside from the pricing decisions, the system manager can adjust short-term capacity by activating and deactivating the surge capacity, modeled as a secondary server. In the rest of this paper, we use secondary server and surge capacity interchangeably. Similarly, we characterize the production time of a class k job on the surge capacity (secondary server) by a general distribution H_k with a mean of $1/\gamma_k$, and the number of class k jobs produced on the secondary server up to time t by the corresponding renewal process $\Gamma_k(t)$ if the secondary server were continuously working. The short-term capacity adjustment results in two operating modes. The system is said to be in its "on" mode if the secondary server is operational and in its "off" mode if the secondary server is deactivated. The transition from one mode to the other is instantaneous, and these transitions form a sequence of decisions made by the system manager, which are modeled by an adapted, finite variation, and cdlg process $Y(t) \in \{0,1\}$ with 1 indicating the secondary server is in the on mode (activated) and 0 the off mode (deactivated). We also define $\boldsymbol{Y}(t) := (Y_k(t))$, where each component $Y_k(t)$ is an adapted, finite variation, and cdlg process taking values in $\{0,1\}$ with 1 indicating the secondary server is currently manufacturing product k and 0 otherwise. Hence, at any time t, $\sum_{k=1}^{K} Y_k(t) \leq Y(t)$, and we can define

$$L_2(t) := \int_0^t \left[Y(u) - \sum_{k=1}^{K} Y_k(u) \right] \mathrm{d}u \tag{2}$$

to denote the secondary server's cumulative idle time up to t when it is operational. Apparently, $\{L_2(t); t \geq 0\}$ forms a nondecreasing process. In addition, we assume that switching from the off to the on mode triggers a setup cost of C, whereas switching from the on to the off mode is free. For greater generality, we assume keeping the secondary server activated incurs an additional manufacturing cost of c per unit of time.

Let $X_k(t)$ denote the number of class k jobs in the system (in queue or in service) at time t. It follows from the conservation of flow that

$$X_k(t) = X_k(0) + A_k(t) - S_k(T_k(t)) - \Gamma_k \left(\int_0^t Y_k(u) du \right). \tag{3}$$

We say a controlled process $(\boldsymbol{\lambda}, \boldsymbol{T}, \boldsymbol{Y}, Y)$ is admissible, if it is non-anticipating, and keeps $X_k(t) \geq 0$ for all k and $t \geq 0$. In addition, for each class k, the system incurs a waiting cost $h_k(x)$ per unit of time whenever the number of class k jobs equals x. We impose the following assumption on $h_k(x)$ purely for technical reasons.

ASSUMPTION 1. *$h_k(0) = 0$ and $\lim_{x\to\infty} h_k(x) = \infty$ for any $k = 1, 2, \ldots, K$. Also, for any $a > 0$, there exists some $\tilde{a} > 0$ such that $h_k(x) > a h_k'(x) - \tilde{a}$ holds for any $x > 0$ and any $k = 1, 2, \ldots, K$.*

The above assumption is made because we need the cost rate function h_k to grow faster than its derivative at some point in our mathematical proof. To understand the implications of this assumption, we note that it is satisfied by all polynomial functions of the form $h_k(x) = a_{k1} x^{a_{k2}}$ for $a_{k1} > 0$ and $a_{k2} \geq 1$. Models in the classical inventory control literature often assume linear or increasing convex holding cost functions to derive optimal policies; for instance, the optimality of base stock policies in a multiple-period inventory model with backlogging requires this assumption (see Theorem 3, Chapter 10 by Karlin and Scarf (1958)).

The manager's objective is to find an admissible $(\boldsymbol{\lambda}, \boldsymbol{T}, \boldsymbol{Y}, Y)$ to maximize

$$\liminf_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \pi(\boldsymbol{\lambda}(u))\mathrm{d}u - \sum_{k=1}^K \int_0^t h_k(X_k(u))\mathrm{d}u - c\int_0^t Y(u)\mathrm{d}u - C\sum_{u\le t}[\Delta Y(u)]^+\right], \tag{4}$$

where the notation $\sum_{u\le t}[\Delta Y(u)]^+ := \sup_{P\in\mathcal{P}_t}\sum_{u\in P}[Y(u)-Y(u^-)]^+$ represents positive jumps of Y accumulated by time t, and $\mathcal{P}_t$ is defined as the class of all partitions over $[0,t]$:

$$\mathcal{P}_t := \left\{\{t_i\}_{i=1}^n : 0 = t_0 < t_1 < \ldots < t_n = t, \text{ for any } n \in \{1,2,\ldots\}\right\}.$$

In (4), the first term represents the profit; the second term describes the cost of holding orders; the third term is the surge capacity cost; and the last term captures expenses associated with switching operations.

Problem (4) is complex and does not lend itself to tractable solutions via exact analysis. As a result, it makes sense to consider approximation techniques. As a starting point, notice that if we temporarily ignore all randomness in the system and any potential capacity constraints, the system manager will simply use the "fluid-optimal" price $\bar{\boldsymbol{p}} := \Xi^{-1}(\bar{\boldsymbol{\lambda}})$, as this price maximizes profit rate. This motivates the consideration of an operating regime under which the base capacity "matches" nominal demand, resulting in moderate congestion caused by stochastic variability. Consequently, congestion effects are considered to be second-order. To refine the fluid-optimal pricing solution, we implement a second-order correction, resulting in a DCP, which we formally describe in (4). Formulating and solving the DCP that approximates the original problem will be the main focus of the next section.

REMARK 1. The fixed setup cost is a key feature in our model; this feature differentiates our model from many other models of make-to-stock production with a secondary capacity source. We recall that in the classical multiple-period inventory control literature, there is a clear dividing line between models that incorporate a fixed order cost and those that do not; the optimal policies in the former class of models are generally of the (s,S) type, whereas the optimal policies in the latter class are generally of the order-up-to type. Analogously, we find that our fixed cost model yields some insights that are absent from similar models that omit a fixed setup cost for the surge capacity, as we will discuss in the sequel.

4 DCP and the Equivalent Workload Problem

4.1 The DCP formulation

To better understand our approximation framework, it is instructive to view our original decision problem as a member of a series of problems indexed by Λ, with relevant model primitives scaled so that

$$\Xi^{-1}(\cdot) := \sqrt{\Lambda}\hat{\Xi}^{-1}(\cdot/\Lambda), \quad q_k := \sqrt{\Lambda}\hat{q}_k, \quad h_k(\cdot) := \sqrt{\Lambda}\hat{h}_k(\cdot/\sqrt{\Lambda}), \quad c := \sqrt{\Lambda}\hat{c} \quad \text{and} \quad C := \sqrt{\Lambda}\hat{C}, \tag{5}$$

where $\hat{\Xi}^{-1}, \hat{q}_k, \hat{h}_k, \hat{c}$, and $\hat{C}$ are corresponding baseline quantities or functions that do not scale with Λ. A simple algebraic manipulation shows that under the scaling condition (5), the nominal demand rate of the problem Λ is $\bar{\boldsymbol{\lambda}} = \Lambda\hat{\boldsymbol{\lambda}}$, where $\hat{\boldsymbol{\lambda}}$ solves $\max_{\boldsymbol{\lambda}} \langle \boldsymbol{\lambda}, \hat{\Xi}^{-1}(\boldsymbol{\lambda}) - \boldsymbol{q} \rangle$. This implies that $\sum_{k=1}^{K} \bar{\lambda}_k$ increases linearly in Λ, so it is also convenient to treat $\bar{\lambda} := \sum_{k=1}^{K} \bar{\lambda}_k$ as a scaling parameter reflecting the system scale. The main purpose of applying the scaling conditions mentioned above is to establish an appropriate objective function, with all terms in (4) being of equal order, as $\bar{\lambda} \to \infty$. Otherwise, they are not of equal importance for $\bar{\lambda}$ large enough; that is, some terms will be dominated by others and therefore become insignificant as $\bar{\lambda} \to \infty$. For further understanding, we provide in EC.3 a further discussion on the order of each term in the objective as well as the baseline quantities. Because we are mostly concerned with deriving the DCP rather than formally establishing a proper notion of asymptotic optimality, in what follows, we continue to use h_k, c, C, Ξ^{-1}, and q_k (hence π) with the understanding that they need to satisfy appropriate scaling conditions so that they are compatible with the scaling imposed on μ and γ that will be introduced in a moment.

The MTO system has high customer demand, a large base production capacity, and is critically loaded in the following sense:

$$\sum_{k=1}^{K} \rho_k = 1 - \psi \quad \text{for} \quad \rho_k := \bar{\lambda}_k / \mu_k, \tag{6}$$

where ψ is a quantity perceived to be of order $1/\sqrt{\bar{\lambda}}$. More formally, we can let $\psi := \hat{\psi}/\sqrt{\bar{\lambda}}$, where $\hat{\psi}$ is a constant that does not scale with $\bar{\lambda}$. To see that the scaling condition (6) can arise naturally in practice, consider a manufacturer who has a primary goal of maximizing profit and a secondary goal of making efficient use of the production capacity. Then it would make sense for the manufacturer to set a base capacity level that approximately matches demand, such that $\sum_{k=1}^{K} \rho_k \approx 1$. Indeed, as noted by Bradley and Glynn (2002), the high cost of capacity can force a firm to set capacity at a level that leads to high utilization at the production facility, forcing the facility into a "heavy traffic" regime where the system can be approximated by a Brownian motion-based model.

The particular choice of the scaling condition $\psi := \hat{\psi}/\sqrt{\bar{\lambda}}$ is standard in modeling the heavy-traffic operating regime (see p. 1136 in Çelik and Maglaras (2008)). A key insight is that when using the fluid-optimal price and capacity level (6), the system capacity will naturally have a high utilization with a mild level of congestion overall, and the resulting congestion-related costs will be moderate. This means that we will need to fine-tune the corresponding pricing decisions and develop an appropriate pricing scheme. Specifically, we consider pricing schemes of the following form:

$$\boldsymbol{\lambda}(t) = \bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta}(t) \qquad \text{for} \quad t \geq 0, \tag{7}$$

where $\boldsymbol{\vartheta}$ is a correction term seen to be an order of magnitude smaller than $\bar{\lambda}$. More formally, we let $\boldsymbol{\vartheta} := \sqrt{\bar{\lambda}}\boldsymbol{\theta}$, where θ is a control process independent of the scaling parameter $\bar{\lambda}$.

We assume the two vectors, $\boldsymbol{\mu} := (\mu_k)$ and $\boldsymbol{\gamma} := (\gamma_k)$, are proportional and satisfy the scaling condition:

$$\frac{\gamma_k}{\mu_k} = \bar{\gamma}, \tag{8}$$

where we have defined $\bar{\gamma} := \hat{\gamma}/\sqrt{\bar{\lambda}}$ for some $\hat{\gamma}$ that does not scale with $\bar{\lambda}$. This assumption simply means that if a product requires more production time from the primary server, it will also consume more machine time on the secondary server. This makes sense in practice because the length of time it takes to manufacture a product is typically determined by its inherited characteristics, such as manufacturing complexity. Equation (8) also means that the processing speed of the secondary server will be in the second order, which is a reasonable assumption because the base capacity already matches the nominal demand in the first order.

To find the diffusion process that approximates $\mathbf{X} := (X_k)$, we apply the strong approximation to $\mathbf{A} := (A_k)$ and $\mathbf{S} := (S_k)$ in a manner similar to that in Çelik and Maglaras (2008) to get

$$A_k(t) = \bar{\lambda}_k t - \int_0^t \vartheta_k(u)\mathrm{d}u + \sqrt{\bar{\lambda}_k}\hat{A}_k(t) + \varepsilon_{ak}(t) \quad \text{and} \quad S_k(T_k(t)) = \mu_k \rho_k t - \mu_k V_k(t) + \nu_k \sqrt{\bar{\lambda}_k}\hat{S}_k(t) + \varepsilon_{pk}(t), \tag{9}$$

where $\hat{A}_k$ and $\hat{S}_k$ are two independent, standard Brownian motions capturing stochastic fluctuations in order arrivals and completions; in Equation (9), $\varepsilon_{ak}(t)$ and $\varepsilon_{pk}(t)$ are error terms from the strong approximation, and they are an order of magnitude smaller than $\sqrt{\bar{\lambda}}$ (for every fixed t). In addition, we have defined

$$V_k(t) := \rho_k t - T_k(t) \tag{10}$$

as the deviation of the actual allocated time to process class k jobs from the "nominal" allocated time $\rho_k t$. Next, by applying the functional strong law of large numbers type approximation for renewal processes, we get

$$\Gamma_k\left(\int_0^t Y_k(u)\mathrm{d}u\right) = \gamma_k \int_0^t Y_k(u)\mathrm{d}u + \varepsilon_{rk}(t), \tag{11}$$

where $\varepsilon_{rk}(t)$ is an approximation error term that has an order of magnitude smaller than $\sqrt{\bar{\lambda}}$. Plugging (7)–(11) into (3) and ignoring all the error terms, we arrive at the desired diffusion approximation for X_k (which we recall is the process that tracks the number of class k jobs over time and is defined in (3)). In particular, the approximating diffusion process, which we denote by Z_k, is given as the solution to the following stochastic integral equation:

$$Z_k(t) = Z_k(0) - \int_0^t \vartheta_k(u)\mathrm{d}u + \mu_k V_k(t) - \gamma_k \int_0^t Y_k(u)\mathrm{d}u + \sigma_k B_k(t), \tag{12}$$

where $\sigma_k := \sqrt{\bar{\lambda}_k(1+\nu_k^2)}$, Y_k tracks the operating mode of the system, and (B_k) are K independent standard Brownian motions.

Define $\delta(\boldsymbol{\vartheta}) := \pi(\bar{\boldsymbol{\lambda}}) - \pi(\bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta})$, which can be regarded as the profit loss due to the deviation from the nominal demand rate. Then, using Z_k in (12) to approximate X_k, we arrive at the diffusion approximation for the objective function in (4) which is given by

$$\limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))\mathrm{d}u + \int_0^t \sum_{k=1}^K h_k(Z_k(u))\mathrm{d}u + c\int_0^t Y(u)\mathrm{d}u + C\sum_{u\le t}[\Delta Y(u)]^+\right]. \tag{13}$$

4.2 Equivalent Workload Formulation

The above DCP is a multidimensional stochastic control problem. To achieve dimensional reduction, we next develop a one-dimensional formulation that is equivalent to the DCP, and we refer to it as the workload problem. To that end, let $\boldsymbol{m} = (m_k) := (1/\mu_k)$ and define the one-dimensional workload process as

$$W(t) := \sum_{k=1}^{K} m_k Z_k(t) \quad \text{for} \quad t \geq 0. \tag{14}$$

Multiplying (12) by m_k and adding over $k = 1, \ldots, K$, plus using (1), (2), (6), (10), we obtain

$$W(t) = W(0) - \psi t - \sum_{k=1}^{K} \int_0^t m_k \vartheta_k(u)\mathrm{d}u - \bar{\gamma} \int_0^t Y(u)\mathrm{d}u + \sigma B(t) + L(t) \quad \text{for} \quad t \geq 0, \tag{15}$$

where $\sigma := \sqrt{\sum_{k=1}^{K} \bar{\lambda}_k (1 + \nu_k^2) m_k^2}$, $B(t)$ a standard Brownian motion, and $L(t) := L_1(t) + \bar{\gamma} L_2(t)$ is a non-decreasing process representing the cumulative weighted sum of idleness from the primary and secondary servers (during activation time).

To formally state the workload problem, let $\bar{h}(\cdot)$ be defined as

$$\bar{h}(w) := \min \left\{ \sum_{k=1}^{K} h_k(z_k) : \sum_{k=1}^{K} m_k z_k = w \right\} \quad \text{for} \quad w \geq 0, \tag{16}$$

so that $\bar{h}$ is the work-based waiting cost function. With these preparations, we can state the workload problem as follows: We seek the triplet $(\boldsymbol{\vartheta}^\star, L^\star, Y^\star)$ that solves the constrained optimization problem

$$\text{minimize} \quad \limsup_{t \to \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \delta(\boldsymbol{\vartheta}(u))\mathrm{d}u + \int_0^t \bar{h}(W(u))\mathrm{d}u + c \int_0^t Y(u)\mathrm{d}u + C \sum_{u \leq t} [\Delta Y(u)]^+ \right] \tag{17}$$

$$\text{subject to (15), } W(t) \geq 0, \text{ and } L(t) \text{ is non-decreasing with } I(0) = 0. \tag{18}$$

It is known from the Skorokhod map that for any controlled processes $\boldsymbol{\vartheta}$ and Y, there exists a regulator I, which is a nondecreasing process satisfying

$$\int_0^t 1_{\{W(u)>0\}} dI(u) = 0, \qquad \text{for any } t \geq 0, \tag{19}$$

yields the minimal idleness in order to keep W non-negative (see, e.g., Chapter 2.2 on pp. 20–21 (Harrison 2013) for the reference). Let $L^\star$ be the process satisfying (19). Then the workload problem is reduced to seeking $(\boldsymbol{\vartheta}^\star, Y^\star)$ to solve the problem (17).

4.3 The Solution

We apply Bellman's principle of optimality to characterize the solution to the workload problem. To that end, let v denote the (relative) value function associated with the workload problem and define the convex conjugate function $g(x) := \sup_{\boldsymbol{\vartheta}} \left\{ x \boldsymbol{\vartheta}^T \boldsymbol{m} - \delta(\boldsymbol{\vartheta}) \right\}$. Note that $g(x) \geq 0$ for any $x \in \mathbb{R}$, because $\vartheta = 0$ is a feasible solution to $\sup_{\vartheta} \left\{ x \boldsymbol{\vartheta}^T \boldsymbol{m} - \delta(\boldsymbol{\vartheta}) \right\}$ and $\delta(\mathbf{0}) = 0$. We will assume that g is locally Lipschitz continuous and exhibits super-linear growth (i.e., for any $l \in \mathbb{R}$, we have $g(x) \geq lx$ for all x large enough). With reference to the general control theory, we expect that the value function v, in conjunction with some constant $\eta^\star$, solves the following Bellman equation, which takes the form of the quasi-variational inequality

$$\min\left\{\frac{\sigma^2}{2}v_{ww}(y,w) - (\psi + \bar{\gamma}y)\, v_w(y,w) - g(v_w(y,w)) + \bar{h}(w) + cy - \eta,\right.$$
$$\left. v(1,w) + C - v(0,w),\ v(0,w) - v(1,w)\right\} = 0, \tag{20}$$

subject to the boundary condition $v_w(y,0) = 0$ and the requirement that $v_w(y,w)$ exhibits polynomial growth as $w \to \infty$. In average cost dynamic programming, $\eta^\star$ is interpreted as a guess for the optimal average cost. Intuitively speaking, the first term in the quasi-variational inequality (20) states the optimality condition if the status of the secondary server does not change at the state (y,w), whereas the second and third terms establish the optimality conditions for the system's transition from off to on and from on to off, respectively. If a solution v to Equation (20) exists, then we can extract an optimal $\boldsymbol{\vartheta}^\star$, given as

$$\boldsymbol{\vartheta}^\star(y,w) := \arg\max_{\boldsymbol{\vartheta}} \left\{ v_w(y,w)\boldsymbol{\vartheta}^T \boldsymbol{m} - \delta(\boldsymbol{\vartheta}) \right\}. \tag{21}$$

We next draw from our intuition to postulate the structure of the Bellman equation solution. Note that if costs associated with on-demand production capacity are not exorbitant, it would clearly be advantageous for the system manager to use surge capacity on a temporary basis. The intuition leads us to conjecture that the optimal control strategy is a sequential switching policy comprising the following actions: If the system is currently operating in its off mode, then it is optimal to remain in that mode if W is below a threshold, say $w_1^\star$, and switch to its on mode once W rises above $w_1^\star$. On the other hand, if the system is currently operating in its on mode, then it is optimal to remain in that mode if W is above a certain level, say $w_0^\star$, and switch to its off mode as soon as W drops below $w_0^\star$. Clearly, this strategy is well-defined if $w_0^\star < w_1^\star$. Moreover, if this strategy, henceforth denoted as $Y^\star$, is indeed optimal, we should be able to find $v_0(w) := v(0,w), v_1(w) := v(1,w)$ and $\eta^\star > 0$, such that

$$\frac{\sigma^2}{2}v_0''(w) - \psi v_0'(w) - g(v_0'(w)) + \bar{h}(w) = \eta^\star \quad \text{for} \quad w \in [0, w_1^\star), \tag{22}$$

$$\frac{\sigma^2}{2}v_1''(w) - (\psi + \bar{\gamma})v_1'(w) - g(v_1'(w)) + \bar{h}(w) + c = \eta^\star \quad \text{for} \quad w > w_0^\star, \tag{23}$$

$$v_0(w) = v_1(w) \quad \text{for} \quad w \in [0, w_0^\star], \quad \text{and} \quad v_0(w) = v_1(w) + C \quad \text{for} \quad w \geq w_1^\star, \tag{24}$$

with the boundary condition

$$v_0'(0) = 0 \quad \text{and the requirement that } v_1'(w) \text{ grows polynomially as } w \to \infty, \tag{25}$$

plus a set of optimality conditions derived from the "principle of smooth fit":

$$v_0'(w_0^\star) = v_1'(w_0^\star) \quad \text{and} \quad v_0'(w_1^\star) = v_1'(w_1^\star). \tag{26}$$

Because (22) and (23) do not involve the zero-order term (the unknown function itself), they can be reduced to a pair of first order differential equations for $f_0(w) := v_0'(w)$ and $f_1(w) := v_1'(w)$. The observation leads to the consideration of the class of functions $\{f_0(\cdot,\eta);\eta\in\mathbb{R}\}$ where $f_0(\cdot,\eta)$ solves

$$\frac{\sigma^2}{2}f_0'(w)-\psi f_0(w)-g(f_0(w))+\bar{h}(w)-\eta=0, \tag{27}$$

subject to the boundary condition $f_0(0)=0$ and the function class $\{f_1(\cdot,\eta);\eta\in\mathbb{R}\}$ where $f_1(\cdot,\eta)$ is the solution to the following differential equation

$$\frac{\sigma^2}{2}f_1'(w)-(\psi+\bar{\gamma})f_1(w)-g(f_1(w))+\bar{h}(w)+c-\eta=0, \tag{28}$$

subject to the requirement that $f_1(w)$ exhibits polynomial growth as $w\to\infty$. For those equations specified by (27), we also let $w_{\eta,\infty} := \inf\{w : \lim_{x\uparrow w} f_0(x)=\pm\infty\}$. Also, it is straightforward to verify that the requirement (24) leads to

$$\int_{w_0^\star}^{w_1^\star}[f_0(w,\eta^\star)-f_1(w,\eta^\star)]\,dw=C. \tag{29}$$

Also, by appealing to (26), we obtain

$$f_0(w_0^\star,\eta^\star)=f_1(w_0^\star,\eta^\star)\quad\text{and}\quad f_0(w_1^\star,\eta^\star)=f_1(w_1^\star,\eta^\star). \tag{30}$$

In sum, the mission of constructing a solution to the Bellman equation (20) boils down to finding variables $\eta^\star$, $w_0^\star$, and $w_1^\star$ such that (29) and (30) hold. Our next result is concerned with the properties of f_0 and f_1.

PROPOSITION 1. *(i) For all $w<w_{\eta,\infty}\wedge w_{\eta',\infty}$, we have $f_0(w,\eta)<f_0(w,\eta')$ if $\eta<\eta'$; in particular, there exists a unique $\eta_0>0$ such that $f_0(w,\eta_0)$ grows to infinity at a polynomial rate with $w_{\eta,\infty}=\infty$, and $\lim_{w\to w_{\eta,\infty}} f_0(w,\eta)=-\infty$ for any $\eta<\eta_0$. (ii) For all w, we have $f_1(w,\eta)>f_1(w,\eta')$ if $\eta<\eta'$; particularly, there exists a unique $\eta_1>0$ such that $f_1(0,\eta_1)=0$, and $f_1(0,\eta)>0$ for any $\eta<\eta_1$.*

We relegate the proof of Proposition 1 to the e-companion. The most intricate part of the proof is establishing the stated polynomial growth rate. For this purpose, we adapt and generalize some of the methodological tools developed in the recent paper by Sun and Zhu (2024b), where the authors establish a desired polynomial growth rate for a Bellman equation in the form of a Riccati equation.

To spell out conditions under which the triplet $(\eta^\star,w_0^\star,w_1^\star)$ exists, we first focus on two benchmark policies corresponding to $f_0(\cdot,\eta_0)$ and $f_1(\cdot,\eta_1)$ in Proposition 1. The first policy, denoted as S_0, requires $Y\equiv 0$ regardless of the system state; this corresponds to the case where the production system relies entirely on in-house capacity, thereby forgoing the option of using on-demand capacity. The second policy, denoted as S_1, sets $Y\equiv 1$; this describes the scenario where the system always retains the on-demand production capacity, thereby incurring costs associated with flexible capacity at all times. With respect to ergodic control of diffusion processes, we can interpret η_0 as the long-run average cost incurred by the production system under the static policy S_0 and η_1 as the long-run average cost under the static policy S_1. A key intuition is that the constant $\eta^\star$, if it exists, ought to satisfy $\eta^\star \le \bar{\eta} := \min(\eta_0, \eta_1)$, because both S_0 and S_1 are admissible policies. The following result is concerned with the number of crossings that $f_0(\cdot, \eta)$ and $f_1(\cdot, \eta)$ can have for each η that is less than $\bar{\eta}$.

LEMMA 1. *There exists $\underline{\eta} \le \bar{\eta}$ such that $f_0(\cdot, \eta)$ and $f_1(\cdot, \eta)$ do not intersect on $\mathbb{R}^+$ for $\eta < \underline{\eta}$, intersect on $\mathbb{R}^+$ but do not cross for $\eta = \underline{\eta}$, and cross at exactly two points on $\mathbb{R}^+$ for $\eta \in (\underline{\eta}, \bar{\eta})$.*

The proof of the lemma is provided in the e-companion. Based on Lemma 1, if $\underline{\eta} = \bar{\eta}$, then there is no such $\eta < \bar{\eta}$ that $f_0(\cdot, \eta)$ and $f_1(\cdot, \eta)$ will ever cross. We interpret this case to mean that no matter how small the fixed cost C is, sequential switching is never optimal. If, however, $\underline{\eta} < \bar{\eta}$, then by Lemma 1, we know that $f_0(\cdot, \eta)$ and $f_1(\cdot, \eta)$ will cross at exactly two points for all $\eta \in (\underline{\eta}, \bar{\eta})$. We interpret this case to mean that, for a sufficiently small fixed cost, it is optimal to switch between the different modes.

To avoid triviality, in the rest of the paper we will focus on the latter scenario by further assuming that $\underline{\eta} < \bar{\eta}$. It is noteworthy that in this case the two functions, $f_0(\cdot, \bar{\eta})$ and $f_1(\cdot, \bar{\eta})$ will not only coincide at $w = 0$ or $w = \infty$ but also cross at some finite point; we denote these two points of intersection by $\bar{w}_0$ and $\bar{w}_1$ with $\bar{w}_0 < \bar{w}_1$. The next result provides the condition under which an optimal sequential switching policy exists.

THEOREM 1. *Suppose $C < \bar{C} := \int_{\bar{w}_0}^{\bar{w}_1} [f_0(w, \bar{\eta}) - f_1(w, \bar{\eta})] dw$. Then the following statements are true. (i) There exists a triplet $(\eta^\star, w_0^\star, w_1^\star)$ satisfying (27)–(30). (ii) The "sequential switching policy" characterized by $(w_0^\star, w_1^\star)$ and the pricing scheme $\boldsymbol{\vartheta}^\star$ given by (21) are jointly optimal for the DCP (13). On the other hand, when $C \ge \bar{C}$, then (iii) $\eta^\star = \bar{\eta} = \min(\eta_0, \eta_1)$, and if $\eta_0 < \eta_1$ then the static policy S_0 is optimal and otherwise the static policy S_1 is optimal.*

The proof of Theorem 1 can be found in the e-companion. The theorem states the conditions under which the optimal policy for managing surge capacity is of sequential switching or static type. To be more precise, sequential switching is optimal if the switching cost is not too high, i.e., $C < \bar{C}$. When $C \ge \bar{C}$, an optimal policy for managing the surge capacity is to set the static policy S_0 by always sticking to the off mode if $\eta_0 < \eta_1$, and to set S_1 by always sticking to the on mode if $\eta_0 \ge \eta_1$.

4.4 Allowing For On-Hand Inventory

We can extend our model to an MTS system by allowing $X_k(t) < 0$ (recall that $X_k(t)$ is defined in (3)), where $[X_k(t)]^-$ represents the number of class k products stored as inventory. Assume in this setting the inventory storage has capacity M, and each unit of class k product takes space s_k, so that $\sum_{k=1}^K s_k [X_k(t)]^- \le M$. Accordingly, one can define the minimum workload that system can maintain by $-b := \min \left\{ w : \sum_{k=1}^K m_k z_k = w, \sum_{k=1}^K s_k [z_k]^- \le M \right\}$ for some $b > 0$. In addition, we modify the meaning of h_k to let it be a non-negative continuous function with $h_k(0) = 0$ that is strictly increasing in positive x

and strictly decreasing in negative x. When $x<0$, $h_k(x)$ represents the cost rate of holding inventory $-x$. Similarly, we can define the corresponding work-based inventory/waiting cost function

$$\bar{h}(w)=\min\left\{\sum_{k=1}^{K}h_k(z_k):\sum_{k=1}^{K}m_kz_k=w,\sum_{k=1}^{K}s_k[z_k]^-\le M\right\}\quad\text{for}\quad w\ge -b.$$

We can easily show that $\bar{h}(w)$ is also strictly increasing in positive w and strictly decreasing in negative w.

The corresponding workload problem can be deduced analogously. Specifically, the workload process for the MTS system is the solution to the following stochastic integral equation:

$$W(t)=W(0)-\psi t-\int_0^t\vartheta(u)\mathrm{d}u-\bar{\gamma}\int_0^t Y(u)\mathrm{d}u+\sigma B(t)+L_b(t), \tag{31}$$

where the last term in (31) is a "regulator" ensuring that W is always greater than or equal to the negative threshold, i.e.,

$$W(t)\ge -b\quad\text{for all}\quad t\ge 0. \tag{32}$$

In particular, L_b is a one-sided regulator that is non-decreasing and continuous with $L_b(0)=0$. It is known that for a fixed control strategy Y, a regulator L_b satisfying

$$\int_0^t 1_{\{W(u)>-b\}}dL_b(u)=0,\qquad\text{for any } t\ge 0, \tag{33}$$

yields the minimal idleness in order to keep $W\ge -b$. So under the make-to-stock extension, the DCP further simplifies to the one that seeks a pair $(\boldsymbol{\vartheta}^\star,Y^\star)$ to minimize (13) subject to (31), (32) and (33).

As in §4.3, the solution to the new problem can be characterized via the quasi-variational inequality (20) subject to the boundary condition $v_w(y,-b)=0$ and the requirement that $v_w(y,w)$ exhibits polynomial growth as $w\to\infty$. We can show that the main result (Theorem 1) still holds under the make-to-stock extension.

THEOREM 2. *The conclusion in Theorem 1 holds under the extension to the make-to-stock model.*

We relegate the proof of Theorem 2 to the e-companion. The main technical challenge therein is that under the make-to-stock extension, we no longer have the monotonicity of the waiting cost function $\bar{h}(x)$ in $x\ge -b$ because $W(t)$ can take negative values. Indeed, $\bar{h}(x)$ is strictly decreasing in negative x and strictly increasing in positive x. In addition, the control over the moments of the W process gets more complicated when $W(t)$ can also take negative values.

5 Policy Recommendations

In this section, we provide policy recommendations based on the solution to the workload problem. Because the workload problem is equivalent to the corresponding DCP, in the following, when we mention the DCP solution, we really mean the solution derived from the workload problem.

We interpret the DCP solution in the context of the corresponding original system by proposing a dynamic control policy, utilizing the results in Theorem 1. Note that we can numerically solve the equations (27) and (28) through the finite difference method. The DCP solution can then be interpreted in a way as some "near-optimal" admissible policy to the original control problem for the MTO system, i.e., problem (4). Intuitively, because the MTO system is facing relatively high demands, the optimal control policy for the original control problem should not be "far" from those of the DCP. This framework for interpreting the DCP solution in the context of the original queueing system was pioneered by Harrison (1988) and is widely used in the studies of control problems in queueing systems.

There are three control levers related to (4): dynamic pricing, capacity management rules, and scheduling. In the sequel, we first interpret the three control levers separately in the original MTO system in §5.1. Then, in §5.2, we expose the interesting structure of the pricing scheme under this system. Finally, we explain the managerial insights into the interplay between dynamic pricing and capacity management in §5.3.

5.1 Interpretation of the DCP Solution

Dynamic Pricing. Given the solution $v_w(y,w)$ to the quasi-variational inequality and the current system state (y,w), we derive the optimal demand rate adjustment term $\boldsymbol{\vartheta}^\star(y,w)$, which is the solution to the equation

$$\boldsymbol{m} v_w(y,w) - \nabla\delta(\boldsymbol{\vartheta}) = \boldsymbol{0}.$$

We can then derive a pricing policy:

$$\boldsymbol{p}^\star(y,w) = \Xi^{-1}\left(\bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta}^\star(y,w)\right).$$

To gain some intuition, it helps to consider a single-product MTO system. In this case, the proposed pricing policy reduces to $p^\star(y,w) = \Xi^{-1}\left(\bar{\lambda} - \vartheta^\star(y,w)\right)$, where $\vartheta^\star$ solves $m v_w(y,w) - \delta'(\vartheta) = 0$ (the subscripts indicating the classes of jobs are suppressed). Hence, the DCP solution roughly states that the proposed pricing policy should have a very similar structure as $v_w(y,\cdot)$. Based on our proof arguments of Theorem 1, we know that when $C < \bar{C}$, $v_w(0,\cdot)$ could be a non-monotonic function with a single maximum point, although $v_w(1,\cdot)$ is always increasing. This suggests it could be the case that as the congestion level increases but is not too high, the manager should raise prices to limit customer demand. When the system is too congested, however, it may be in the manager's best interest to lower the price. Indeed, as we will demonstrate in §5.2, under a logistic demand model, a pricing policy with a structure of first increasing then decreasing frequently appears to be optimal. This rather intriguing finding reveals a strong interaction between dynamic pricing and the setup cost of the surge capacity. Previous studies show that in the absence of the setup cost, the pricing scheme is a monotonically increasing function of the system congestion level (see, e.g., Çelik and Maglaras (2008), Ata and Barjesteh (2023)). The present paper seems to be the first to reveal a non-monotonic pricing scheme structure when queues are unobservable. We offer a detailed explanation of this atypical pricing scheme in §5.3.

Capacity Management Rule. As we have discussed in Theorem 1, depending on whether $C < \bar{C}$ or $C \geq \bar{C}$, a dynamic or static capacity management policy is optimal for the approximating problem. Hence, we propose the following capacity management rule $Y^\star(\cdot)$:

- If $C \geq \bar{C}$, we apply a static capacity management rule. Let $Y^\star(y,w) \equiv \arg\min_i \{\eta_i : i = 0,1\}$ (let $Y^\star(y,w) \equiv 1$ if $\eta_0 = \eta_1$), where η_0 and η_1 are defined in Proposition 1, such that the secondary server is always off/on.
- If $C < \bar{C}$, we apply a sequential switching rule defined by $w_0^\star$ and $w_1^\star$ as in Theorem 1: $Y^\star(y,w) = 1$ if $y = 0$ and $w > w_1^\star$; $Y^\star(y,w) = 0$ if $y = 1$ and $w < w_0^\star$.

Scheduling Rule. The workload problem defined in (16) indicates the optimal allocation of the total workload into different classes for the DCP. Therefore, given the workload w, the solution

$$\boldsymbol{x}^\star(w) = \arg\min \left\{ \sum_{k=1}^{K} h_k(x_k) : \sum_{k=1}^{K} m_k x_k = w \right\} \tag{34}$$

(with ties broken consistently and arbitrarily) suggests we try to keep the queue length of each class as close to $\boldsymbol{x}^\star(w)$ as possible. We propose the following scheduling rule:

- At each time, whenever the primary server is available or the secondary server is both on and available and there are jobs waiting in some queues, serve the head-of-the-line job from class

$$i \in \arg\max_k \{X_k(t) - x_k^\star(w)\}$$

with ties broken consistently and arbitrarily.

As in Ata and Tongarlak (2013) and a few related studies, different waiting cost functions result in different scheduling rules. For instance, when all waiting cost functions are linear, the proposed scheduling rule reduces to the celebrated $c\mu$ rule. As another example, when all waiting cost functions are quadratic, the schedule becomes a queue-ratio rule that strives to maintain various queue lengths at some target ratio.

5.2 The Observed Atypical Pricing Scheme

In this subsection, we provide concrete examples to illustrate our proposed policy. To make things simple, we consider a single-product MTO system, and we can therefore suppress all the subscripts indicating the job classes. Customers' willingness-to-pay follows a logistic distribution $G(x) = 1/(1+e^{-(x-r)/s})$, where r and s are two model parameters representing location and scale, respectively. The demand rate function is then $\Xi(p) = \Lambda G^c(p)$, where Λ signifies the potential demand arrival rate and $G^c(p) := 1 - G(p)$ denotes the tail probability. The inverse demand rate function is thus $\Xi^{-1}(\lambda) = r + s\log((\Lambda-\lambda)/\lambda)$. We choose the following model primitives: $\Lambda = 78.327$, $r = 500$, $s = 30$, $c = 200$, $C = 600$, and $q = 400$. The nominal demand rate is $\bar{\lambda} = \arg\max_\lambda \lambda \cdot (\Xi^{-1}(\lambda) - q) = 50$. Production times at both the primary and secondary servers are exponentially distributed with the rate parameters $\mu = 42.929$ and $\gamma = 14.142$, respectively.

Figures 1(a) and 1(b) illustrate the proposed pricing schemes as a function of the number of pending jobs in the system when the waiting cost rates are $h(x) = x$ and $h(x) = 0.1x^2$, respectively. To compare with the optimal control policy, we plot both the pricing policies derived from our DCP and a Markov Decision Process (MDP) formulation (please refer to EC.4 for details in the MDP formulation). The solid and dashed lines represent the pricing policies under MDP and DCP, with different colors differentiating whether the surge capacity is on or off. In particular, the legend "price (base capacity)" indicates the price curve when the surge capacity is off, whereas "price (surge capacity)" indicates the price curve when the surge capacity is on. Furthermore, the up-pointing and down-pointing triangles denote the thresholds of switching derived from the solutions of the MDP and DCP (notice that they are μw_0^* and μw_1^* from the definition of workload (14)), respectively. One can see that the pricing policy derived from the DCP is very close to their optimal counterparts and shares the same structure. The performance gap between the two policies will be studied in Section 6. We can also see that the optimal policies for managing surge capacity are of either the threshold or static type. This has been proven for the DCP (Theorems 1). However, our extensive numerical experiments reveal that a solution to an MDP exhibits the same structure, as illustrated in the two figures.

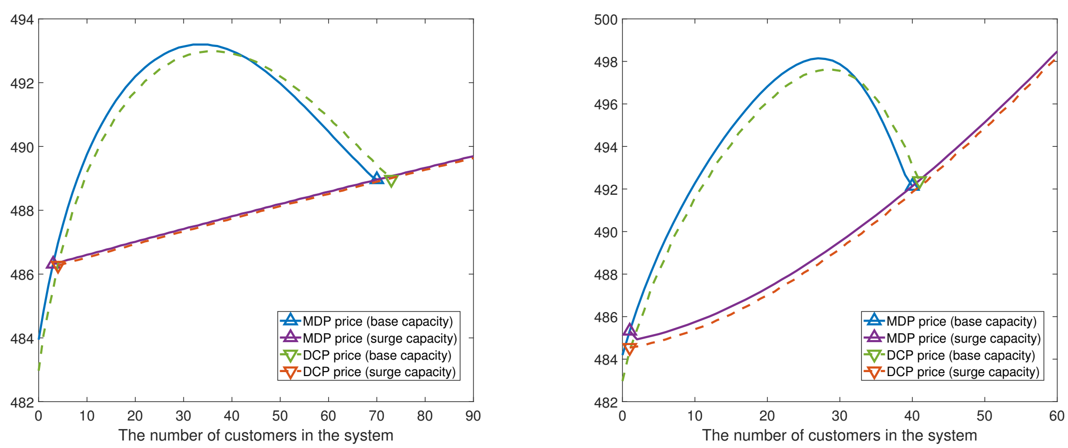

(a) $h(x) = x$ (b) $h(x) = 0.1x^2$

Figure 1 Dynamic Pricing Schemes and Sequential Switching Rules Derived From MDPs and DCPs

When we compare Figures 1(a) and 1(b), we can see that the quadratic waiting costs tend to result in a more aggressive pricing strategy: The pricing curves under both base and surge capacities in Figure 1(b) are steeper than those in Figure 1(a). Additionally, the thresholds for the sequential switching policy in Figure 1(b) are lower than those shown in Figure 1(a). These findings reinforce our intuition about reducing waiting costs through dynamic pricing and expanding capacity. When the waiting cost function is quadratic, the system waiting cost rate grows rapidly with the congestion level, requiring a higher price and lower thresholds for activating the surge capacity to effectively prevent system congestion levels from growing too high. The downward slope of pricing policy with base capacity in Figure 1(b) is much steeper than that shown in Figure 1(a), implying the system should quickly accumulate backlog and thus spend less time in states that are congested but not so congested as to justify activating the secondary server.

5.3 Interaction Between Dynamic Pricing and Capacity Management

To explain the atypical structure of the pricing rule, consider the manager's ideal scenario: rely on the primary server, set the price at $\bar{p}$, and maintain a low level of congestion. However, this perfect situation is not achievable because, with a base capacity that matches nominal demand, stochasticity can often overload the system, leading to significant waiting costs. To combat congestion, managers have two levers of control: they can raise the product price to lower the effective arrival rate or activate the secondary server to increase capacity. However, the existence of the setup cost discourages switching the surge capacity on and off too frequently, and the higher C is, the less frequently the surge capacity should be activated.

Zooming in on the details, when facing a nontrivial setup cost, the system manager should activate the secondary server only when the system is heavily congested, which further implies that pricing should be the main device to adjust demand when the congestion level is low or moderate. When only the base capacity is operational, the system manager is obliged to offer a higher price to curb demand as the system becomes more congested. However, since a higher price also implies a higher profit loss, reducing system congestion by raising prices may become more cost-ineffective as the system congestion level further increases. When the congestion level falls short but is near the activating threshold, it is a reasonable strategy for the system manager to lower the price to both reduce profit loss and accumulate enough backlog to make the best use of the activation of surge capacity. Although lowering the price makes the system even more congested in the short term, the long-term holding cost would not be too high, provided that the system could accumulate enough backlog and activate the surge capacity in a brief period of time.

6 Numerical Studies

In this section, we conduct numerical experiments to examine system performance under the proposed policy. In §6.1, we focus on a single-product system, and we study a two-product system in §6.2.

6.1 Single Product

The basic parameter settings of the single-product systems we study in this subsection are the same as those described in §5.2. For better illustration, we define $z_0^\star := \mu w_0^\star$ and $z_1^\star := \mu w_1^\star$ to indicate the threshold of switching the secondary server with respect to the number of pending jobs in the system. In §6.1.1, we numerically study the performance gap of policies derived from the DCP in this paper and a Taylor series-based DCP (T-DCP) in the literature, compared to the MDP solution. We elaborate on the effect of the setup cost C on the system in §6.1.2. We next study the effect of the load parameter ψ in §6.1.3. Finally, in §6.1.4 we show how values of per-unit surge capacity cost c affect the system performance.

6.1.1 Performance Gap of DCP and T-DCP Compared to MDP

One noteworthy fact in this paper is that we do not use any function approximation techniques to the profit loss function δ, in contrast to a stream of papers using second-order Taylor expansions as the approximation to the revenue or profit function (Çelik and Maglaras 2008, Kim and Randhawa 2018, Ata and Barjesteh 2023). More specifically, if $\|\boldsymbol{\lambda}-\bar{\boldsymbol{\lambda}}\|$ is not too large, then

$$\pi(\boldsymbol{\lambda}) \approx \tilde{\pi}(\boldsymbol{\lambda}) := \pi(\bar{\boldsymbol{\lambda}}) + \nabla\pi(\bar{\boldsymbol{\lambda}})(\boldsymbol{\lambda}-\bar{\boldsymbol{\lambda}}) + \frac{1}{2}(\boldsymbol{\lambda}-\bar{\boldsymbol{\lambda}})^T\nabla^2\pi(\bar{\boldsymbol{\lambda}})(\boldsymbol{\lambda}-\bar{\boldsymbol{\lambda}}).$$

So it is reasonable to substitute π by $\tilde{\pi}$ in the objective function. The upside of this approximation is that the resulting Bellman equation is a Riccati equation, which may be solved explicitly. However, it may cause a non-negligible error if the actual arrival rates deviate from $\bar{\boldsymbol{\lambda}}$ too much. We henceforth call the DCP derived from this second-order Taylor expansion the Taylor series-based DCP (T-DCP).

To study the performance gap of the policies generated by our DCP and T-DCP, we make a comparison of the system performance under the two policies with the optimal policy computed through an MDP formulation. The details of how to formulate our problem as an MDP are provided in EC.4. Tables 1 and 2 present the numerical results, corresponding to the holding cost function $h(x)=x$ and $h(x)=0.1x^2$, respectively. The optimal long-run average cost can be derived by solving the MDP, which is denoted as $\bar{\eta}^*$ in the tables. We use $\hat{\eta}_{\text{DCP}}$ and $\hat{\eta}_{\text{T-DCP}}$ to represent the estimated long-run average cost through conducting Monte Carlo simulation programs using the policies derived from solving the DCP and T-DCP, respectively. The values in the parentheses indicate 95% confidence levels. "Gap" is defined as $(\hat{\eta}_k-\bar{\eta}^*)/\hat{\eta}_k$ for $k=\text{DCP},\text{T-DCP}$ to measure the differences between the two policies and the optimal ones.

Table 1 Gaps of DCP and T-DCP Compared to MDP With $h(x)=x$

	MDP	DCP		T-DCP	
C	$\bar{\eta}^*$	$\hat{\eta}_{\text{DCP}}$	Gap (%)	$\hat{\eta}_{\text{T-DCP}}$	Gap (%)
0	100.294	100.159(0.423)	-0.14	100.298(0.414)	0.00
200	120.616	120.622(0.408)	0.00	121.075(0.378)	0.38
400	130.028	130.212(0.406)	0.14	130.489(0.389)	0.35
600	136.503	136.855(0.681)	0.26	137.224(0.711)	0.53
800	141.083	141.654(0.695)	0.40	142.652(0.671)	1.10
1,000	144.217	144.559(0.685)	0.24	146.388(0.713)	1.48
"Static Off"	148.319	148.238(0.675)	-0.05	149.34(0.624)	0.68
"Static On"	207.348	207.827(0.779)	0.23	207.409(0.808)	0.03

Tables 1 and 2 show that the policy derived from our DCP is indeed near-optimal since the gaps are all below about 0.5%, under various values of C. In contrast, the performance gap between the T-DCP and MDP policies enlarges as C increases, and the worst gap is approximately 3% when $C=1,000$ with the quadratic waiting cost function $h(x)=0.1x^2$. Our numerical results show that solving the DCP without approximating the profit function can keep the performance gaps extremely small under various scenarios. Notably, gaps

Table 2 Gaps of DCP and T-DCP Compared to MDP With $h(x) = 0.1x^2$

	MDP	DCP		T-DCP	
C	$\bar{\eta}^*$	$\hat{\eta}_{\text{DCP}}$	Gap (%)	$\hat{\eta}_{\text{T-DCP}}$	Gap (%)
0	108.001	107.621(0.417)	-0.35	108.141(0.395)	0.13
200	143.942	144.135(0.423)	0.13	144.845(0.396)	0.62
400	160.031	160.161(0.718)	0.08	161.617(0.668)	0.98
600	169.403	169.839(0.629)	0.26	172.341(0.632)	1.70
800	174.1	174.388(0.632)	0.17	178.714(0.705)	2.58
1,000	175.896	176.832(0.693)	0.53	181.394(0.636)	3.03
"Static Off"	176.495	176.509(0.679)	0.01	176.967(0.683)	0.26
"Static On"	209.386	208.669(0.811)	-0.34	209.289(0.831)	0.30

should always be non-negative, but due to simulation errors, we may sometimes obtain negative estimated gaps, and the negative gaps typically imply the performance of the proposed policy is very close to that under the optimal one.

6.1.2 The Effect of Setup Cost C

We next study the effect of the setup cost parameter C, which is an important feature in our model to study the surge capacity. The simulation results are reported in Tables 3 and 4 for waiting cost functions $h(x) = x$ and $h(x) = 0.1x^2$, respectively. We present how the thresholds $z_0^\star, z_1^\star$ and average costs $\hat{\eta}_{\text{DCP}}$ are affected by various values of C. In addition, the columns "Profit Loss", "Waiting Cost", "Surge Capacity Cost", and "Setup Cost" report the estimated values of individual cost rates of the first to the last terms in (13), with Z_k replaced by X_k, the state variables in pre-limit systems.

In practice, the setup cost could be related to some physical costs. For example, C can be used to model the fixed fee component of contractual arrangements with a contract manufacturer that take the form of a fixed fee plus a variable fee. The setup cost can also be regarded as a system design parameter that prevents the surge capacity from switching on and off too frequently. Indeed, our formulation is closely related to the following constrained problem:

$$\text{minimize} \quad \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))\mathrm{d}u + \sum_{k=1}^K \int_0^t h_k(Z_k(u))\mathrm{d}u + c\int_0^t Y(u)\mathrm{d}u\right], \tag{35}$$

subject to the constraint

$$\limsup_{t\to\infty} \frac{\mathbb{E}\left[\sum_{u\le t}[\Delta Y(u)]^+\right]}{t} \le \upsilon, \tag{36}$$

for some constant υ. The constraint (36) requires the long-run average switching frequency of the surge capacity to be bounded from above by some "budget" υ. This is effectively a constrained Markov decision process, and our formulation (13) can, in a loose sense, be regarded as a Lagrangian relaxation of the constrained problem, where C is the Lagrange multiplier. We do not rigorously justify the equivalence between the constrained and the Lagrangian problems. But we encourage interested readers to refer to the

Table 3 Numerical Results With $h(x) = x$

C	$(z_0^\star, z_1^\star)$	$\hat{\eta}_{\mathrm{DCP}}$	Profit Loss	Waiting Cost	Surge Capacity Cost	Setup Cost	Switching Frequency
0	(19.42, 19.42)	100.159(0.423)	19.876	16.392	63.891	0.000	1.665
200	(6.607, 53.153)	120.622(0.408)	31.755	25.803	51.547	11.517	0.058
400	(5.105, 64.565)	130.212(0.406)	41.815	29.621	43.923	14.852	0.037
600	(4.204, 72.973)	136.855(0.681)	52.604	31.643	36.597	16.012	0.027
800	(3.604, 79.279)	141.654(0.695)	64.682	32.672	29.113	15.188	0.019
1,000	(3.303, 84.985)	144.559(0.685)	77.761	32.629	21.212	12.956	0.013
"Static off"	-	148.238(0.675)	119.506	28.733	0	0	0
"Static on"	-	207.827(0.779)	0.687	7.14	200	0	0

Table 4 Numerical Results With $h(x) = 0.1x^2$

C	$(z_0^\star, z_1^\star)$	$\hat{\eta}_{\mathrm{DCP}}$	Profit Loss	Waiting Cost	Surge Capacity Cost	Setup Cost	Switching Frequency
0	(12.514, 12.514)	107.621(0.417)	20.399	17.295	69.927	0	1.935
200	(3.303, 31.231)	144.135(0.423)	38.167	29.364	55.732	20.872	0.104
400	(2.102, 36.937)	160.161(0.718)	59.314	35.102	41.208	24.537	0.061
600	(1.502, 41.141)	169.839(0.629)	85.412	37.672	26.581	20.174	0.034
800	(1.201, 45.345)	174.388(0.632)	111.33	36.952	13.514	12.592	0.016
1,000	(0.901, 49.249)	176.832(0.693)	130.749	35.114	5.447	5.521	0.006
"Static off"	-	176.509(0.679)	143.945	32.564	0	0	0
"Static on"	-	208.669(0.811)	0.429	8.24	200	0	0

textbook by Altman (1999); see also Chai et al. (2023) and the references therein for the relevant problems developed in the same vein.

We emphasize that if the system manager has a budget for the cost associated with the switching frequency of the surge capacity, it can be achieved by carefully fine-tuning the parameter C. A higher value of C implies a higher penalty cost per switch and hence discourages switching on and off the surge capacity too frequently, while a lower value of C would stimulate switches. We corroborate the above intuition through simulation programs summarized in the "Switching Frequency" columns in both Tables 3 and 4. One noteworthy result is that the switching frequency becomes prohibitively high when the setup cost is zero. For example, our results in Table 4 show that the optimal policy requires the system to switch on and off the surge capacity nearly twice a day when $C = 0$ and $h(x) = 0.1x^2$. This may not be a viable policy for some manufacturing systems, especially those that require extra efforts during each capacity expansion, such as staffing, training, etc. To make a comparison, under the linear waiting cost, if $C = 600$, then the switching frequency is only 0.027 and $1/0.027 \approx 37$, so that on average the system only activates the secondary server once per 37 days.

Tables 3 and 4 also provide insights on how the individual costs are affected by C. We notice that both the profit loss and waiting cost increase in C, but the surge capacity cost decreases in C. This result is intuitive because a higher penalty leads to a lower frequency of switching, which makes the system less flexible to reduce waiting costs through dynamic capacity adjustment and hence more reliable on pricing control. Less flexibility also encourages the manager to be less reliant on the surge capacity, so its usage also decreases.

Finally, our numerical results show the setup cost may not be a monotone function with respect to C, as it is a product of C and a term (switching frequency) decreasing in C.

6.1.3 The Effect of Load Parameter ψ

We next investigate how the load parameter ψ affects the system performance by fixing $C = 600$ and keeping other parameters unchanged. Table 5 summarizes the numerical results as we vary ψ changes from 0 to -0.39. We focus on the critically loaded and overloaded regimes because only in these regimes does surge capacity play an important role in reducing waiting costs.

Table 5 The Effect of Load Parameter ψ With $h(x) = x$

ψ	$(z_0^\star, z_1^\star)$	$\bar{\eta}$	$\hat{\eta}_{\text{DCP}}$	Gap	Profit Loss	Waiting Cost	Surge Capacity Cost	Setup Cost	Switching Frequency
0	"Static Off"	27.719	27.876(0.174)	0.56	9.213	18.663	0	0	0
-0.08	(12.913, 88.589)	69.05	68.98(0.367)	-0.10	40.9	25.276	1.808	0.996	0.002
-0.16	(4.204, 72.973)	136.503	136.855(0.681)	0.26	52.604	31.643	36.597	16.012	0.027
-0.27	(1.201, 70.871)	196.867	196.601(0.744)	-0.14	29.169	36.903	103.739	26.79	0.045
-0.39	"Static On"	228.24	227.871(0.759)	-0.16	9.153	18.717	200	0	0

We can see from Table 5 that the nominal traffic intensity has a great impact on the average cost, where as ψ decreases from 0 to -0.39, $\hat{\eta}_{\text{DCP}}$ has more than eightfold increase from about 28 to about 228. In addition, as the nominal traffic intensity increases, the capacity adjustment strategy changes from "Static Off" to dynamic adjustment and finally to "Static On". We also see that the long-run average profit loss, waiting cost, and setup cost are all non-monotonic functions in ψ, whereas the surge capacity cost grows rapidly as ψ decreases.

6.1.4 The Effect of Per-Unit Surge Capacity Cost c

The parameter c measures how expensive it is to continuously use the surge capacity. Intuitively speaking, for any fixed setup cost C, we expect that the system should have a "Static On" policy if c is too small and a "Static Off" policy if c is too large. We expect a dynamic capacity adjustment strategy to arise when c is in some middle range, an intuition which we verify through the numerical examination in Table 6, where we vary c from 50 to 400.

Table 6 The Effect of Per-Unit Surge Capacity Cost c With $h(x) = x$

c	$(z_0^\star, z_1^\star)$	$\bar{\eta}$	$\hat{\eta}_{\text{DCP}}$	Gap (%)	Profit Loss	Waiting Cost	Surge Capacity Cost	Setup Cost	Switching Frequency
50	"Static On"	57.348	57.494(0.388)	0.25	0.35	7.144	50	0	0
100	"Static On"	107.348	107.314(0.823)	-0.03	0.166	7.147	100	0	0
200	(4.204, 72.973)	136.503	137.091(0.756)	0.43	52.824	31.639	36.59	16.039	0.027
300	(10.511, 80.781)	146.784	147.022(0.665)	0.16	99.606	29.824	13.073	4.518	0.008
400	"Static Off"	148.28	148.077(0.654)	-0.14	119.36	28.717	0	0	0

By comparing the numerical results in the first two rows when $c = 50, 100$ and the rest of the data in Table 6, we see an interesting implication indicated by negligible profit loss and unchanged waiting cost: Pricing policy keeps at a nearly static optimal level when c is relatively small. In other words, the system behaves like an $M/M/2$ queue for a range of c being small. If, however, c increases beyond this range but is not too large, the manager should combine dynamic pricing and dynamic capacity adjustment to restrict waiting costs. Furthermore, when c becomes too large, the manager relies more on pricing as we see the profit loss increase and eventually solely uses dynamic pricing while keeping the secondary server off. This result again shows the intricate interaction between dynamic pricing and the setup cost of surge capacity.

6.2 Two Products

In this subsection, we study a two-product MTO system to illustrate the performance of our model in multi-product settings. §6.2.1 illustrates the structure of the pricing rule under a multinomial logit model in a multi-product setting. In §6.2.2, we study the effect of price sensitivity on system performance.

6.2.1 Pricing Structure

We now consider an MTO manufacturer offering two products, indexed by $k = 1, 2$, with the price vector $\boldsymbol{p}$. Following the literature, we assume the probability of purchasing product k follows a multinomial logit model, such that

$$\mathbb{P}\left(\text{purchase } k|\boldsymbol{p}\right) = \frac{e^{U_k}}{1 + \sum_{k=1}^{2} e^{U_k}},$$

where $U_k := b_k^0 - b_k^1 p_k$ for some model primitives b_k^0, b_k^1. Then we have the demand rate function $\Xi_k(\boldsymbol{p}) = \Lambda \mathbb{P}\left(\text{purchase } k|\boldsymbol{p}\right)$ where we recall that Λ denotes the arrival rate of potential customers. Through some algebraic manipulation, we can derive the inverse demand rate function

$$\Xi_k^{-1}(\boldsymbol{p}) = -\frac{1}{b_k^1} \log\left(\frac{\lambda_k}{\Lambda - \sum_{k=1}^{2} \lambda_k}\right) + \frac{b_k^0}{b_k^1}, \quad \text{for} \quad k = 1, 2.$$

We set the basic parameters as follows: $\Lambda = 75$, $q_1 = q_2 = 400$, $c = 200$, $C = 600$, $\hat{\psi} = -1$, $\hat{\gamma} = 2$, $b_1^0 = b_2^0 = 15$, and $b_1^1 = b_2^1 = 0.03$. The nominal demand rate can be calculated as $\bar{\boldsymbol{\lambda}} = (25, 25)$ through a gradient ascent method. Assuming the production rates of the two products are identical, then we can derive the service rates of the primary server $\boldsymbol{\mu} = (43.8, 43.8)$ by (6), and $\boldsymbol{\gamma} = (12.39, 12.39)$ by (8). The waiting cost functions are assumed to be $h_1(x) = x$ and $h_2(x) = 1.2x$. It is worth mentioning that based on our assumption on the waiting cost functions, by (34) we have $x_1^\star(W(t)) = \mu_1 W(t)$ and $x_2^\star(W(t)) = 0$ such that the scheduling rule is to always give priority to class 2 products.

Based on our model parameter settings, we can see that the profit rate π is a symmetric function with respect to demand rates λ_1 and λ_2, which implies the optimal demand rates and hence the optimal pricing rule from the DCP should be the same for products 1 and 2. Being aware of this, we only plot the pricing schemes of class 1 products in Figures 2 and 3, corresponding to the pricing decisions with surge capacity off and on, respectively. In addition, the panels 2(a) and 3(a) present the pricing curves, where the *x*- and *y*-axis represent the number of pending requests for products of class 1 and 2 in the system, respectively. For better illustration on pricing, panels 2(b) and 3(b) present two-dimensional color plots with brighter colors representing higher values.

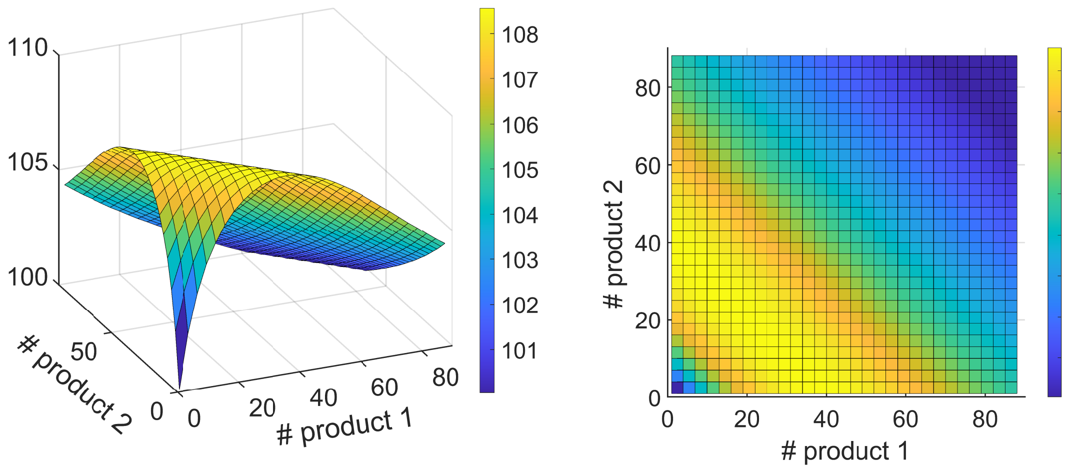

(a) Pricing curve (b) Pricing colormap

Figure 2 Pricing Schemes of Class 1 with Base Capacity

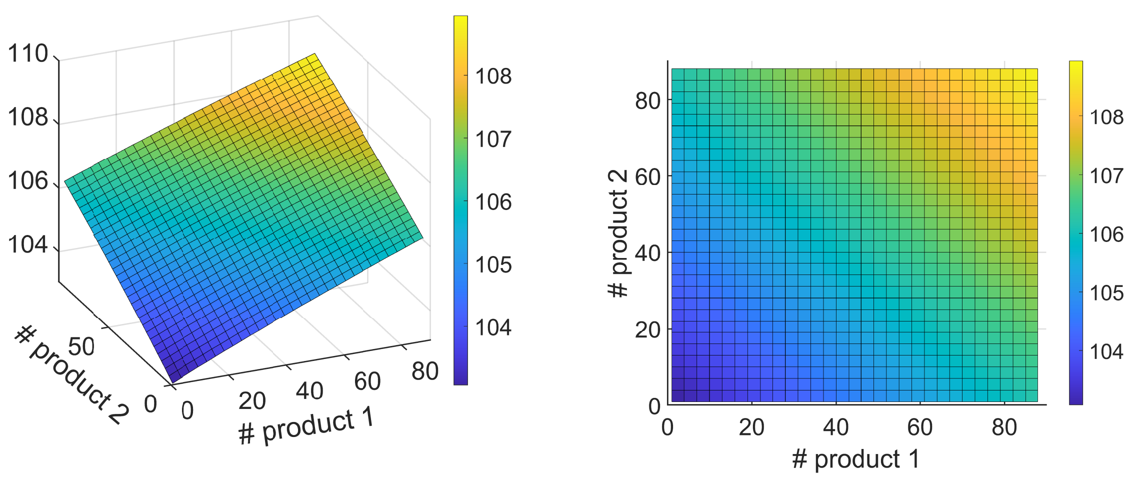

(a) Pricing curve (b) Pricing colormap

Figure 3 Pricing Schemes of Class 1 with Surge Capacity

We can see from Figures 2 and 3 that the pricing rule shares a similar structure as that depicted by Figure 1: When the surge capacity is off, the proposed pricing policy first increases and then decreases with respect to the congestion level; but when the surge capacity is on, the proposed pricing policy is monotonically increasing with respect to the congestion level. Furthermore, it is easy to see from the panels 2(b) and 3(b) that the product-specific pricing schemes are structurally the same along all lines parallel to a diagonal line. This is because we derive the pricing rule based on the system workload, which is a weighted sum of pending jobs from all the classes.

6.2.2 Effect of Price Sensitivity Between the Two Products

We now study how the price sensitivity between the two products affects the system performance. To this end, we keep all basic parameters fixed, but vary b_2^1 from 0.03 to 0.034. Notice that by changing the demand function, we need to recalculate $\bar{\boldsymbol{\lambda}}$, $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ accordingly. The numerical results are shown in Table 7.

Table 7 The Effect of Different Price Sensitivity Across Two Products

b_2^1	$(w_0^\star, w_1^\star)$	$\hat{\eta}_{\text{DCP}}$	$\hat{\eta}_{\text{T-DCP}}$	Profit Loss	Waiting Cost	Surge Capacity Cost	Setup Cost	Switching Frequency
0.03	(0.115, 1.672)	137.752(0.852)	138.123(0.827)	47.471	33.186	41.22	15.876	0.026
0.031	(0.13, 1.712)	131.712(0.756)	132.565(0.769)	52.576	31.927	33.856	13.354	0.022
0.032	(0.14, 1.747)	128.306(0.707)	128.904(0.759)	59.263	30.73	27.272	11.041	0.018
0.033	(0.155, 1.782)	123.905(0.709)	125.295(0.725)	61.735	29.918	22.779	9.473	0.016
0.034	(0.16, 1.812)	122.153(0.677)	122.128(0.724)	65.438	29.494	19.253	7.969	0.013

By comparing the columns "$\hat{\eta}_{\text{DCP}}$" and "$\hat{\eta}_{\text{T-DCP}}$" of Table 7, we first find that the policies derived from our DCP outperform the ones derived from the T-DCP in multi-product scenarios. Table 7 also shows that the long-run average cost $\hat{\eta}_{\text{DCP}}$ decreases as the customers become more price-sensitive to the product 2. This is because when customers are more price-sensitive, to achieve the same demand rates, the manager needs to set a lower price compared to the case of less price-sensitive customers, which leads to the higher profit rate loss. Moreover, as customers become more price-sensitive, the system manager should rely more on dynamic pricing and less on surge capacity, as suggested by the results in the columns "Profit Loss", "Surge Capacity Cost", and "Setup Cost". The less dependency on the surge capacity also implies the switching frequency should decrease in price sensitivity, as expected in the simulation results shown in the last column of Table 7.

7 Concluding Remarks

Maintaining capacity flexibility can be a cost-effective option for manufacturers who produce MTO products to meet the fluctuating demand of their customers. However, adjusting capacity levels while making pricing decisions is a challenging task. This paper introduces a formal stochastic control framework that jointly considers capacity adjustment with a setup cost and dynamic pricing, which, to the best of our knowledge, has not been done before. We present a nearly explicit joint capacity adjustment and pricing strategy that depends on the congestion level in the system. Capacity adjustments follow a switching strategy that determines when to use surge capacity and when to shut it down, based on a set of switching boundaries. We have also observed an unusual non-monotonic feature for the pricing rule, which highlights the complex interplay between the two control levers. In particular, if the setup cost associated with the secondary server is not too high, when the congestion level goes from low to moderate, the manager should increase the price. However, if the congestion keeps rising, it may be wise to *decrease* price and then consider activating the secondary server if the backlogged demand approaches a critically high level. Further, as a rule of thumb, the manager should use dynamic pricing rather than surge capacity as a lever when the customers are price-sensitive.

References

Altman, Eitan. 1999. *Constrained Markov Decision Processes*. Chapman and Hall/CRC press.

Ata, Bariş, Nasser Barjesteh. 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.

Ata, Bariş, JM Harrison, Lawrence A Shepp. 2005. Drift rate control of a Brownian processing system. *Annals of Applied Probability*, 15 (2), 1145-1160.

Ata, Bariş, Shiri Shneorson. 2006. Dynamic control of an M/M/1 service system with adjustable arrival and service rates. *Management Science*, 52 (11), 1778-1791.

Ata, Barış, Mustafa H Tongarlak. 2013. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems*, 74 (1), 65-104.

Bradley, James R, Peter W Glynn. 2002. Managing capacity and inventory jointly in manufacturing systems. *Management Science*, 48 (2), 273-288.

Cao, Ping, Dacheng Yao. 2018. Optimal drift rate control and impulse control for a stochastic inventory/production system. *SIAM Journal on Control and Optimization*, 56 (3), 1856-1883.

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

Chai, Shiwei, Xu Sun, Hossein Abouee-Mehrizi. 2023. Customer scheduling in large service systems under model uncertainty. *Operations Research*, In Press.

Duckworth, Kate, Mihail Zervos. 2001. A model for investment decisions with switching costs. *Annals of Applied Probability*, 11 (1), 239-260.

Federgruen, Awi, Aliza Heching. 1999. Combined pricing and inventory control under uncertainty. *Operations research*, 47 (3), 454-475.

Gao, Xuefeng, Junfei Huang. 2023. Asymptotically optimal control of make-to-stock systems. *Mathematics of Operations Research*, In Press.

Ghamami, Samim, Amy R Ward. 2013. Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: Asymptotic optimality of a two-threshold policy. *Mathematics of Operations Research*, 38 (4), 761-824.

Ghosh, Arka P, Ananda P Weerasinghe. 2007. Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Systems*, 55 (3), 147-159.

Ghosh, Arka P, Ananda P Weerasinghe. 2010. Optimal buffer size and dynamic rate control for a queueing system with impatient customers in heavy traffic. *Stochastic Processes and their Applications*, 120 (11), 2103-2141.

Hairer, Martin. 2021. Convergence of Markov Processes [Lecture Notes]. `https://www.hairer.org/notes/Convergence.pdf`.

Hall, Joseph M, Praveen K Kopalle, David F Pyke. 2009. Static and dynamic pricing of excess capacity in a make-to-order environment. *Production and Operations Management*, 18 (4), 411-425.

Harrison, J Michael. 1988. Brownian models of queueing networks with heterogeneous customer populations. W. Fleming, P.L. Lions, eds., *Stochastic Differential Systems, Stochastic Control Theory and Applications*. Springer, New York, 147-186.

Harrison, J Michael. 2013. *Brownian Models of Performance and Control*. Cambridge University Press.

Janakiraman, Ganesh, Mahesh Nagarajan, Senthil Veeraraghavan. 2018. Simple policies for managing flexible capacity. *Manufacturing & Service Operations Management*, 20 (2), 333-346.

Karatzas, Ioannis, Steven E Shreve. 1991. *Brownian Motion and Stochastic Calculus*. Springer New York.

Karlin, Samuel, Herbert Scarf. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, 155-178.

Kim, Jeunghyun, Ramandeep S Randhawa. 2018. The value of dynamic pricing in large queueing systems. *Operations Research*, 66 (2), 409-425.

Liu, Weiliang, Xu Sun. 2022. Energy-aware and delay-sensitive management of a drone delivery system. *Manufacturing & Service Operations Management*, 24 (3), 1294-1310.

Plambeck, Erica L. 2004. Optimal leadtime differentiation via diffusion approximations. *Operations Research*, 52 (2), 213-228.

Rauschecker, Ursula, Daniel Stock, Matthias Stöhr, Alexander Verl. 2014. Connecting factories and related IT environments to manufacturing clouds. *International Journal of Manufacturing Research*, 9 (4), 389-407.

Sun, Xu, Weiliang Liu. 2023. Expanding service capabilities through an on-demand workforce. *Operations Research*, In Press.

Sun, Xu, Xiaohan Zhu. 2024a. Dynamic control of a make-to-order system under model uncertainty. Working Paper.

Sun, Xu, Xiaohan Zhu. 2024b. Dynamic incentives in service contracting. Working Paper.

Varadhan, S.R.S. 2011. Advanced Topics in Probability [Notes on Lecture 16]. `https://www.math.nyu.edu/~varadhan/Spring11/topics16.pdf`.

Wein, Lawrence M. 1992. Dynamic scheduling of a multiclass make-to-stock queue. *Operations Research*, 40 (4), 724-735.

Wu, Jingchen, Xiuli Chao. 2014. Optimal control of a Brownian production/inventory system with average cost criterion. *Mathematics of Operations Research*, 39 (1), 163-189.

Yao, Dacheng. 2017. Joint pricing and inventory control for a stochastic inventory system with Brownian motion demand. *IISE Transactions*, 49 (12), 1101-1111.

Yoon, Seunghwan, Mark E Lewis. 2004. Optimal pricing and admission control in a queueing system with periodically varying parameters. *Queueing Systems*, 47 (3), 177-199.

Zervos, Mihail, Timothy C Johnson, Fares Alazemi. 2013. Buy-low and sell-high investment strategies. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 23 (3), 560-578.

E-Companion

EC.1 Proofs of Main Results

Proof of Proposition 1. To prove Proposition 1, we need the following auxiliary lemmas:

LEMMA EC.1. *If $\eta_1 < \eta_2$, then $f_0(w,\eta_1) < f_0(w,\eta_2)$ for all $w \in (0, w_{\eta_1,\infty} \wedge w_{\eta_2,\infty})$.*

LEMMA EC.2. *Let $\mathcal{L}_0 := \{\eta > 0 : \exists w \in (0, w_{\eta,\infty}),\ f_0'(w,\eta) < 0\}$. If $\eta \in \mathcal{L}_0$, then $f_0(w,\eta)$ is quasi-concave and $\lim_{w \to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. In addition, $\mathcal{L}_0$ is nonempty.*

LEMMA EC.3. *The set $\mathcal{U}_0 := \{\eta > 0 : \eta \notin \mathcal{L}_0\}$ is nonempty.*

LEMMA EC.4. *Let $\eta_0 := \sup \mathcal{L}_0$. Then $\eta_0 \in \mathcal{U}_0$, and $f_0'(w,\eta_0) \geq 0$ for all $w \geq 0$.*

LEMMA EC.5. *For each $\eta \in \mathcal{L}_0$, $f_0(w,\eta) < g^{-1}(2\bar{h}(w) + a)$ for some positive constant a.*

The proofs of the auxiliary lemmas are postponed to EC.2. Part (i) of Proposition 1 follows directly from Lemmas EC.1–EC.5.

Next, to prove part (ii) of Proposition 1, we first define a class of functions $\{\xi_\alpha(x,\eta); \eta \in \mathbb{R}\}$ which are solutions to (28), but with the left boundary conditions $\xi_\alpha(0,\eta) = \alpha$ and no restriction to the right boundary condition. Let $w^\alpha_{\eta,\infty} := \inf\{w : \lim_{x \uparrow w} \xi_\alpha(x,\eta) = \pm\infty\}$. We first need Proposition EC.1:

PROPOSITION EC.1. *For all $w < w^\alpha_{\eta,\infty} \wedge w^\alpha_{\eta',\infty}$, we have $\xi_\alpha(w,\eta) < \xi_\alpha(w,\eta')$ if $\eta < \eta'$; in particular, there exists a unique $\eta_\alpha > 0$ such that $\xi_\alpha(w,\eta_\alpha)$ grows to infinity at a polynomial rate with $w_{\eta,\infty} = \infty$, and $\lim_{w \to w_{\eta,\infty}} \xi_\alpha(w,\eta) = -\infty$ for any $\eta < \eta_\alpha$.*

Let $c_0 := c - (\psi + \bar{\gamma})\alpha - g(\alpha)$. Analogously, the proof of Proposition EC.1 relies on the following lemmas:

LEMMA EC.6. *If $\eta_1 < \eta_2$, then $\xi_\alpha(w,\eta_1) < \xi_\alpha(w,\eta_2)$ for all $w \in (0, w^\alpha_{\eta_1,\infty} \wedge w^\alpha_{\eta_2,\infty})$.*

LEMMA EC.7. *Let $\mathcal{L}_\alpha := \left\{\eta > c_0 : \exists w \in (0, w^\alpha_{\eta,\infty}),\ \xi_\alpha'(w,\eta) < 0\right\}$. If $\eta \in \mathcal{L}_\alpha$, $\xi_\alpha(w,\eta)$ is quasi-concave and $\lim_{w \to w^\alpha_{\eta,\infty}} \xi_\alpha(w,\eta) = -\infty$. In addition, $\mathcal{L}_\alpha$ is nonempty.*

LEMMA EC.8. *The set $\mathcal{U}_\alpha := \left\{\eta > c_0 : \eta \notin \mathcal{L}^\alpha_\eta\right\}$ is nonempty.*

LEMMA EC.9. *Let $\eta_\alpha := \sup \mathcal{L}_\alpha$. Then $\eta_\alpha \in \mathcal{U}_\alpha$, and $\xi_\alpha'(w,\eta_\alpha) \geq 0$ for all $w \geq 0$.*

LEMMA EC.10. *For each $\eta \in \mathcal{L}_\alpha$, $\xi_\alpha(w,\eta) < g^{-1}(2\bar{h}(w) + a)$ for some positive constant a.*

In the sequel, we denote η_α as $\eta(\alpha)$ (with a little abuse of notation) to stress that $\eta(\alpha)$ is a function of α. We next show that $\eta(\alpha)$ is a strictly decreasing and continuous function of α.

LEMMA EC.11. *If $\alpha_1 > \alpha_2$, then $\eta(\alpha_1) < \eta(\alpha_2)$ and $\xi_{\alpha_1}(w,\eta(\alpha_1)) > \xi_{\alpha_2}(w,\eta(\alpha_2))$. In addition, $\eta(\cdot)$ is a continuous mapping and $\eta(\alpha) \to -\infty$ as $\alpha \to \infty$.*

From Lemma EC.11, the inverse function $\eta^{-1}(\cdot)$ is well defined. Let $f_1(w,\eta) = \xi_{\eta^{-1}(\eta)}(w,\eta)$. We know from Proposition EC.1 that $f_1(w,\eta)$ corresponds to the desired class of functions in part (ii). Thus, the proof of the Proposition 1 is now complete. $\square$

Proof of Lemma 1. To begin with, we observe that the two functions $f_0(\cdot,\bar{\eta})$ and $f_1(\cdot,\bar{\eta})$ will touch at either $w=0$ or $w=\infty$. On the other hand, $f_1(w,0)>0$ for any $w\in[0,\infty]$ and $f_0(w,0)\le 0$ for any $w\in[0,w_{0,\infty})$, which implies that $f_1(w,0)-f_0(w,0)>0$ for any $w\in[0,w_{0,\infty})$. Moreover, by Proposition 1, for any w, $f_0(w,\eta)$ is increasing in η and $f_1(w,\eta)$ is decreasing in η. Therefore, there exists some $\underline{\eta}\in(0,\bar{\eta}]$ such that for any $\eta<\underline{\eta}$, $f_0(\cdot,\eta)$ and $f_1(\cdot,\eta)$ do not intersect, and for $\eta=\underline{\eta}$, $f_0(\cdot,\eta)$ and $f_1(\cdot,\eta)$ intersect but do not cross each other, and for any $\eta\in(\underline{\eta},\bar{\eta})$, $f_0(\cdot,\eta)$ and $f_1(\cdot,\eta)$ cross at least twice. In the remainder of the proof, we will write f_0 and f_1 in place of $f_0(\cdot,\eta)$ and $f_1(\cdot,\eta)$, respectively, and let $\tilde{f}:=f_0-f_1$. We intend to argue that f_0 and f_1 can cross at most twice for any $\eta\in(\underline{\eta},\bar{\eta})$.

Suppose, by way of contradiction, that the two functions cross more than twice. Then $\tilde{f}$ must have crossed the horizontal line $y=0$ at least four times, two times from below and two times from above. On the other hand, from (27) and (28) it follows that

$$\tilde{f}'(w)=\frac{2}{\sigma^2}\left[c-\bar{\gamma}f_1(w)\right] \quad \text{for all } w \text{ such that} \quad \tilde{f}(w)=0. \tag{EC.1}$$

However, by Proposition 1, we know that f_1 is strictly increasing on $[0,\infty)$ for all $\eta\le\eta_1$. This implies that $\tilde{f}(w)$ can cross the horizontal line $y=0$ at most twice thanks to (EC.1), leading to a contradiction. Therefore, f_0 and f_1 can cross at most twice for any $\eta\in(\underline{\eta},\bar{\eta})$. The proof is thus complete. $\square$

Now, we are ready to prove Theorem 1.

Proof of Theorem 1. We first prove (i). For any $\eta\in(\underline{\eta},\bar{\eta})$, by Lemma 1, the functions f_0 and f_1 cross at two points. Let $w_0(\eta)<w_1(\eta)$ denote two points where the functions f_0 and f_1 cross. When $\eta\to\underline{\eta}$, $w_1(\eta)-w_0(\eta)\to 0$ and $\int_{w_0(\eta)}^{w_1(\eta)}[f_0(w,\eta)-f_1(w,\eta)]dz\to 0$. When $\eta=\bar{\eta}$, $w_0(\eta)=\bar{w}_0$ and $w_1(\eta)=\bar{w}_1$. By Proposition 1, for any w, $f_1(w,\eta)$ is decreasing in η and $f_0(w,\eta)$ is increasing in η. Therefore, $w_0(\eta)$ is decreasing in η and $w_1(\eta)$ is increasing in η. As a result, $\int_{w_0(\eta)}^{w_1(\eta)}[f_0(w,\eta)-f_1(w,\eta)]dz$ is increasing in η. Hence, as η increases from $\underline{\eta}$ to $\bar{\eta}$, $\int_{w_0(\eta)}^{w_1(\eta)}[f_0(w,\eta)-f_1(w,\eta)]dz$ increases from 0 to $\bar{C}=\int_{\bar{w}_0}^{\bar{w}_1}[f_0(w,\bar{\eta})-f_1(w,\bar{\eta})]dz$, which is greater than C by our assumption. Hence, we conclude that there exists some $\eta^\star\in(\underline{\eta},\bar{\eta})$ such that $\int_{w_0(\eta^\star)}^{w_1(\eta^\star)}[f_0(w,\eta^\star)-f_1(w,\eta^\star)]dz=C$. Denoting $w_0^\star=w_0(\eta^\star)$ and $w_1^\star=w_1(\eta^\star)$, we complete the proof of (i).

Next, let us prove (ii). We use a verification argument consisting of two steps.

Verification Argument: Step 1.

Let v be such that (a) $v_w(0,w)=f_0(w,\eta^\star)$ on $[0,w_1^\star)$ and $v_w(0,w)=f_1(w,\eta^\star)$ on $[w_1^\star,\infty)$, (b) $v_w(1,w)=f_0(w,\eta^\star)$ on $[0,w_0^\star]$ and $v_w(0,w)=f_1(w,\eta^\star)$ on $(w_0^\star,\infty)$, and (c) $v(0,w_0^\star)=v(1,w_0^\star)$. Then, for any $y\in\{0,1\}$, $v(y,w)$ is a twice continuously differentiable function on $[0,\infty)$ except at $w_0^\star$ and $w_1^\star$. Define $\mathcal{K}=\mathbb{R}^+\setminus\{w_0^\star,w_1^\star\}$. It is easy to verify that this function $v(y,w)$ satisfies (22)–(26) and thus is a solution of Equation (20) when $w\in\mathcal{K}$. Then for any $w\in\mathcal{K}$ and $y\in\{0,1\}$,

$$\frac{\sigma^2}{2}v_{ww}(y,w)-(\psi+\bar{\gamma}y)\,v_w(y,w)-g(v_w(y,w))+\bar{h}(w)+cy\ge\eta^\star, \tag{EC.2}$$

and

$$0 \leq v(0,w) - v(1,w) \leq C, \tag{EC.3}$$

with $v_w(y,0) = 0$ and $v_w(y,w)$ grows polynomially in w as $w \to \infty$.

Note that based on our construction, $v_w(y,\cdot)$ is continuously differentiable on $\mathbb{R}_+ \setminus \{w_0^\star, w_1^\star\}$. Since the occupation time of Itô processes at any single point is zero, we can arbitrarily assign value to $v_{ww}(y,w)$ on $\{w_0^\star, w_1^\star\}$ for calculus computation and the result will not be changed (see e.g. § 3.6 of Karatzas and Shreve (1991)). Hence, in the rest of this section, we shall stipulate $v_{ww}(y,w)$ as its right limit, i.e., $v_{ww}(y,w) = \lim_{h\to 0^+}(v_w(y,w+h) - v_w(y,w))/h$. Then, by applying the Itô's formula and taking expectations, we get

$$\begin{aligned}
&\mathbb{E}[v(Y(t),W(t))] - v(Y(0),W(0)) \\
&= -\psi\mathbb{E}\left[\int_0^t v_w(Y(u),W(u))du\right] - \mathbb{E}\left[\int_0^t \boldsymbol{m}^T\boldsymbol{\vartheta}(u)v_w(Y(u),W(u))du\right] \\
&\quad - \bar{\gamma}\mathbb{E}\left[\int_0^t Y(u)v_w(Y(u),W(u))du\right] + \frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t v_{ww}(Y(u),W(u))du\right] \\
&\quad + \mathbb{E}\left[v_w(Y(t),0)L(t)\right] + \mathbb{E}\left[\sum\nolimits_{u\leq t}\Delta v(Y(u),W(u))\right],
\end{aligned} \tag{EC.4}$$

where we used

$$\sigma\mathbb{E}\left[\int_0^t v_w(Y(u),W(u))dB(u)\right] = 0,$$

which holds since $v_w(y,w)$ has at most polynomial growth in w and by Lemma EC.12 part (ii)

$$\mathbb{E}\left[\int_0^t (W(u))^k du\right] < \infty,$$

for any $k \in \mathbb{N}$. Since $v_w(y,0) = 0$, we have $\mathbb{E}[v_w(Y(t),0)L(t)] = 0$. Moreover, the definition $g(x) = \sup_{\boldsymbol{\vartheta}}\{x\boldsymbol{m}^T\boldsymbol{\vartheta} - \delta(\boldsymbol{\vartheta})\}$ and (EC.2) imply that

$$\begin{aligned}
&-\psi\mathbb{E}\left[\int_0^t v_w(Y(u),W(u))du\right] - \mathbb{E}\left[\int_0^t \boldsymbol{m}^T\boldsymbol{\vartheta}(u)v_w(Y(u),W(u))du\right] \\
&\quad - \bar{\gamma}\mathbb{E}\left[\int_0^t Y(u)v_w(Y(u),W(u))du\right] + \frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t v_{ww}(Y(u),W(u))du\right] \\
&\geq -\psi\mathbb{E}\left[\int_0^t v_w(Y(u),W(u))du\right] - \mathbb{E}\left[\int_0^t g\left(v_w(Y(u),W(u))\right)du\right] - \mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))du\right] \\
&\quad - \bar{\gamma}\mathbb{E}\left[\int_0^t Y(u)v_w(Y(u),W(u))du\right] + \frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t v_{ww}(Y(u),W(u))du\right] \\
&\geq \eta^\star t - \mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))du\right] - \mathbb{E}\left[\int_0^t \bar{h}(W(u))du\right] - c\mathbb{E}\left[\int_0^t Y(u)du\right],
\end{aligned}$$

and (EC.3) implies that

$$\mathbb{E}\left[\sum\nolimits_{u\leq t}\Delta v\left(Y(u),W(u)\right)\right] \geq -\mathbb{E}\left[C\sum\nolimits_{u\leq t}[\Delta Y(u)]^+\right].$$

Hence, we have

$$\begin{aligned}
&\mathbb{E}\left[v\left(Y(t),W(t)\right)\right] - v(Y(0),W(0)) \\
&\geq \eta^\star t - \mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))du\right] - \mathbb{E}\left[\int_0^t \bar{h}(W(u))du\right] - c\mathbb{E}\left[\int_0^t Y(u)du\right] - \mathbb{E}\left[C\sum\nolimits_{u\leq t}[\Delta Y(u)]^+\right].
\end{aligned} \tag{EC.5}$$

To proceed, we need the following Lemma.

LEMMA EC.12. *Regardless of the choice of Y, we have (i)* $\limsup_{t\to\infty}\mathbb{E}\left[(W(t))^k\right]<\infty$ *for any* $k>0$*; (ii) for any* $k,t>0$, $\mathbb{E}\left[\int_0^t (W(s))^k \mathrm{d}s\right]<\infty$.

To ease the flow of ideas, we defer the proof of this lemma to the end. Next, we divide both sides of (EC.5) by t, send $t\to\infty$, and then appeal to Lemma EC.12 part(i) plus the fact that $v_w(y,w)$ has at most polynomial growth, to get that for any admissible strategy Y,

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}(u))du+\int_0^t \bar{h}(W(u))du+c\int_0^t Y(u)du+C\sum_{u\le t}[\Delta Y(u)]^+\right]\ge\eta^\star.$$

Verification Argument: Step 2. Next, we will show that for $Y^\star$, $W^\star$ and optimal pricing $\boldsymbol{\vartheta}^\star$, we have

$$\limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}^\star(u))du+\int_0^t \bar{h}(W^\star(u))du+c\int_0^t Y^\star(u)du+C\sum_{u\le t}[\Delta Y^\star(u)]^+\right]=\eta^\star.$$

Following the same procedures as in **Step 1**, (i.e., applying the generalized Itô's formula for non-smooth function, making substitutions to the term with local time, and taking expectations), we get

$$\begin{aligned}
&\mathbb{E}[v(Y^\star(t),W^\star(t))]-v(Y(0),W(0))\\
&=-\psi\mathbb{E}\left[\int_0^t v_w(Y^\star(u),W^\star(u))du\right]-\mathbb{E}\left[\int_0^t \boldsymbol{m}^T\boldsymbol{\vartheta}^\star(u)v_w(Y^\star(u),W^\star(u))du\right]\\
&\quad-\bar{\gamma}\mathbb{E}\left[\int_0^t Y^\star(u)v_w(Y^\star(u),W^\star(u))du\right]+\frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t v_{ww}(Y^\star(u),W^\star(u))du\right]\\
&\qquad+\mathbb{E}\left[\sum\nolimits_{u\le t}\Delta v(Y^\star(u),W^\star(u))\right]\\
&=-\psi\mathbb{E}\left[\int_0^t v_w(Y^\star(u),W^\star(u))du\right]-\mathbb{E}\left[\int_0^t g\left(v_w(Y^\star(u),W^\star(u))\right)du\right]-\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}^\star(u))du\right]\\
&\quad-\bar{\gamma}\mathbb{E}\left[\int_0^t Y^\star(u)v_w(Y^\star(u),W^\star(u))du\right]+\frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t v_{ww}(Y^\star(u),W^\star(u))du\right]\\
&\qquad+\mathbb{E}\left[\sum\nolimits_{u\le t}\Delta v(Y^\star(u),W^\star(u))\right],
\end{aligned}$$

where we used the definition of the optimal pricing strategy $\boldsymbol{\vartheta}^\star(u)$ such that $g(v_w(Y^\star(u),W^\star(u)))=v_w(Y^\star(u),W^\star(u))\boldsymbol{m}^T\boldsymbol{\vartheta}^\star(u)-\delta(\boldsymbol{\vartheta}^\star(u))$ and

$$\sigma\mathbb{E}\left[\int_0^t v_w(Y^\star(u),W^\star(u))dB(u)\right]=0,$$

which holds since $v_w(y,w)$ has at most polynomial growth in w and by Lemma EC.12 part (ii)

$$\mathbb{E}\left[\int_0^t (W^\star(u))^k du\right]<\infty,$$

for any $k\in\mathbb{N}$, and $\mathbb{E}[v_w(Y^\star(t),0)L(t)]=0$ since $v_w(y,0)=0$. We can verify that

$$\begin{aligned}
&\mathbb{E}[v(Y^\star(t),W^\star(t))]-v(Y(0),W(0))\\
&=\eta^\star t-\mathbb{E}\left[\int_0^t \delta(\boldsymbol{\vartheta}^\star(u))du\right]-\mathbb{E}\left[\int_0^t \bar{h}(W^\star(u))du\right]-c\mathbb{E}\left[\int_0^t Y^\star(u)du\right]-\mathbb{E}\left[C\sum\nolimits_{u\le t}[\Delta Y^\star(u)]^+\right].
\end{aligned}$$

If we divide both hand sides of the above equation by t and let t go to infinity, we obtain the desired result by appealing to Lemma EC.12 part (i) and the fact that $v_w(y,w)$ has at most polynomial growth. This completes the proof of (ii).

Finally, let us prove (iii). Suppose that $C \geq \bar{C}$. Let $v_w(0,w) = f_0(w,\bar{\eta})$ on $[0,\bar{w}_1)$ and $v_w(0,w) = f_1(w,\bar{\eta})$ on $[\bar{w}_1,\infty)$. Also, let $v_w(1,w) = f_0(w,\bar{\eta})$ on $[0,\bar{w}_0]$, and $v_w(0,w) = f_1(w,\bar{\eta})$ on $(\bar{w}_0,\infty)$. For any $y \in \{0,1\}$, $v(y,w)$ is a twice continuously differentiable function on $[0,\infty)$ except at $\bar{w}_0$ and $\bar{w}_1$. By writing $v_y(w) = v(y,w)$, we have

$$\frac{\sigma^2}{2} v_0''(w) - \psi v_0'(w) - g(v_0'(w)) + \bar{h}(w) = \bar{\eta} \quad \text{for} \quad w \in [0,\bar{w}_1), \tag{EC.6}$$

$$\frac{\sigma^2}{2} v_1''(w) - (\psi + \bar{\gamma}) v_1'(w) - g(v_1'(w)) + \bar{h}(w) + c = \bar{\eta} \quad \text{for} \quad w > \bar{w}_0. \tag{EC.7}$$

In addition,

$$v_0(w) = v_1(w) \quad \text{for} \quad 0 \leq w \leq \bar{w}_0, \quad \text{and} \quad v_0(w) = v_1(w) + \bar{C} \quad \text{for} \quad w \geq \bar{w}_1, \tag{EC.8}$$

subject to the boundary conditions $v_0'(0) = 0$, and $v_1'(w)$ grows polynomially in w as $w \to \infty$, as well as a set of optimality conditions: $v_0'(\bar{w}_0) = v_1'(\bar{w}_0) \quad \text{and} \quad v_0'(\bar{w}_1) = v_1'(\bar{w}_1)$. Then for any $w \neq \bar{w}_0, \bar{w}_1$ and $y \in \{0,1\}$, we have

$$\frac{\sigma^2}{2} v_{ww}(y,w) - (\psi + \bar{\gamma} y)\, v_w(y,w) - g\left(v_w(y,w)\right) + \bar{h}(w) + cy \geq \bar{\eta}, \tag{EC.9}$$

and

$$0 \leq v(0,w) - v(1,w) = \bar{C} \leq C, \tag{EC.10}$$

with $v_w(y,0) = 0$ and $v_w(y,w)$ grows polynomially in w as $w \to \infty$. By adapting the same argument as in **Step 1** of the **Verification Argument**, we obtain

$$\limsup_{t\to\infty} \frac{1}{t} \mathbb{E}\left[\int_0^t \delta(\vartheta(u))du + \int_0^t \bar{h}(W(u))du + c\int_0^t Y(u)du + C\sum_{u\leq t}[\Delta Y(u)]^+\right] \geq \bar{\eta} = \min(\eta_0,\eta_1).$$

The above inequality becomes an equality by using the S_0 strategy ($Y \equiv 0$) or S_1 strategy ($Y \equiv 1$) depending on whether $\eta_0 \leq \eta_1$ or $\eta_0 > \eta_1$. This completes the proof of (iii). □

Proof of Theorem 2. The proof of Theorem 2 departs from that of Theorem 1 mainly because the lower threshold for the approximating diffusion process $W(t)$ is changed from 0 to $-b$. So the boundary conditions need to be modified accordingly. Despite this departure, the proofs of the two theorems are largely identical, except that we need to make two non-trivial modifications, one on Lemma EC.2 and the other on Lemma EC.12.

The following lemma extends Lemma EC.2 that is essential for proving Proposition 1.

LEMMA EC.13. *Let* $\mathcal{L}_0 := \{\eta > 0 : \exists w \in (-b, w_{\eta,\infty}),\ f_0'(w,\eta) < 0\}$. *If* $\eta \in \mathcal{L}_0$, *then* $f_0(w,\eta)$ *is quasi-concave and* $\lim_{w \to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. *In addition,* $\mathcal{L}_0$ *is nonempty.*

The proof of Lemma EC.13 is much more sophisticated than the proof of Lemma EC.2 due to the fact that the waiting cost function $\bar{h}(x)$ is no longer monotonically increasing everywhere. Indeed, $\bar{h}(x)$ is strictly decreasing in negative x. We will provide the proof of Lemma EC.13 in §EC.2.

Next, to show Theorem 2, since the approximating diffusion process $W(t)$ can take negative values, we need the following modification of Lemma EC.12:

LEMMA EC.14. *Regardless of the choice of* Y*, we have (i)* $\limsup_{t\to\infty} \mathbb{E}\left[|W(t)|^k\right] < \infty$ *for any* $k > 0$*; (ii) for any* $k, t > 0$, $\mathbb{E}\left[\int_0^t |W(s)|^k \mathrm{d}s\right] < \infty$.

The proof of Lemma EC.14 is more involved than the proof of Lemma EC.12 due to the fact that the Z process can take negative values. We will provide the proof of Lemma EC.14 in §EC.2.

The rest of the proof of Theorem 2 follows similarly from the proof of Theorem 1. □

EC.2 Proofs of Auxiliary Lemmas

Proof of Lemma EC.1. Since $g(\cdot)$ is a locally Lipschitz continuous function (by our hypothesis), there exists a continuous function $f_0(\cdot,\eta)$ satisfying (27) on $[0, w_{\eta,\infty})$ for each $\eta \in \mathbb{R}$. If $\eta_1 < \eta_2$, then $f_0'(0,\eta_1) < f_0'(0,\eta_2)$. Hence, there exists some $w > 0$ such that $f_0(w,\eta_1) < f_0(w,\eta_2)$. Then the lemma follows by applying the standard comparison theorem of ordinary differential equations. □

Proof of Lemma EC.2. Assume $\eta \in \mathcal{L}_0$ and $f_0(w,\eta)$ exists. To prove $f_0(w,\eta)$ is quasi-concave, we need to show there exists only one local maximum point w_0 such that $f_0(w_0,\eta) \geq f_0(w,\eta)$ for all $w \in [0, w_{\eta,\infty})$ and $\eta \in \mathcal{L}_0$.

Let $w_0 := \inf\{w > 0 : f_0'(w,\eta) < 0\}$. Then $f_0'(w_0,\eta) = 0$, and $f_0(w_0,\eta) \geq f_0(w,\eta)$ for all $w \in (0, w_0)$. In addition, there exists some $\varepsilon > 0$ such that $f_0'(w,\eta) < 0$ and $f_0(w,\eta) < f_0(w_0,\eta)$ on $(w_0, w_0 + \varepsilon)$. We claim for all $w \in (w_0, w_{\eta,\infty})$, $f_0'(w,\eta) \leq 0$ and $f_0(w,\eta) < f_0(w_0,\eta)$. Assume by contradiction that this is not true, then there exists some $w_1 > w_0$ such that $f_0'(w_1,\eta) > 0$ and $f_0(w_1,\eta) < f_0(w_0,\eta)$. We can deduce there exists $\bar{w}_0 \in (w_0, w_1)$ such that $f_0(\bar{w}_0,\eta) = f_0(w_1,\eta)$, $f_0'(\bar{w}_0,\eta) \leq 0$. We have

$$\frac{\sigma^2}{2}\left(f_0'(\bar{w}_0,\eta) - f_0'(w_1,\eta)\right) = -\left(\bar{h}(\bar{w}_0) - \bar{h}(w_1)\right),$$

which is a contradiction, because the left-hand side is strictly negative but the right-hand side is non-negative.

Next, we prove $\lim_{w \to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. If $w_{\eta,\infty} < \infty$, due to $f_0'(w,\eta) \leq 0$ for all $w \in (w_0, w_{\eta,\infty})$, $\lim_{w \to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. If, on the other hand, $w_{\eta,\infty} = \infty$, assume $\lim_{w\to\infty} f_0(w,\eta) = l > -\infty$. We have $\lim_{w\to\infty} f_0'(w,\eta) = 0$, so $\lim_{w\to\infty} \bar{h}(w) = \eta + \psi l - g(l)$, which is a contradiction since the left-hand side is ∞ but the right-hand side is finite.

Finally, to prove $\mathcal{L}_0 \neq \emptyset$, we first notice that for all $\eta < 0$, $f_0'(0,\eta) < 0$. Based on the above arguments, we have $f_0(w,\eta) < 0$, $f_0'(w,\eta) \leq 0$ for all $w \in (0, w_{\eta,\infty})$ and $\lim_{w\to\infty} f_0(w,\eta) = -\infty$. By the continuity of $f_0(w,\eta)$ with respect to η, we have $f_0(w,0) = \lim_{\eta\uparrow 0} f_0(w,\eta) \leq 0$ and $f_0'(w,0) = \lim_{\eta\uparrow 0} f_0'(w,\eta) \leq 0$. But since $f_0(w,0) \equiv 0$ is not possible, there must exist some $w_1 > 0$ such that $f_0(w_1,0) < 0$. Using the continuity of $f_0(w,\eta)$ with respect to η again, we know that there exists some $\eta_1 > 0$ such that $f_0(w_1,\eta_1) < 0$. Hence, there exists some $w \in (0,w_1)$ such that $f_0'(w,\eta_1) < 0$, so $\eta_1 \in \mathcal{L}_0$. $\square$

Proof of Lemma EC.3. Assume by contradiction that $\mathcal{U}_0 = \emptyset$. Then for all $\eta > 0$, by Lemma EC.2 $\lim_{w\to w_{\eta,\infty}} f_0(w,\eta) = -\infty$.

To proceed, we first claim the set defined as $\bar{\mathcal{U}}(w,y) := \{\eta : f_0(w,\eta) \geq y\}$ is nonempty for any $w \geq 1, y \geq 0$. Since $g(x) \geq 0$, we have

$$
\begin{aligned}
f_0'(w,\eta) &\geq \frac{2}{\sigma^2}\left(\eta + \psi f_0(w,\eta) - \bar{h}(w)\right), \\
f_0(w,\eta) &\geq \frac{2}{\sigma^2} e^{\frac{2\psi}{\sigma^2}} \int_0^w e^{-\frac{2\psi}{\sigma^2}} \left(\eta - \bar{h}(x)\right) \mathrm{d}x.
\end{aligned}
$$

So, given w, the right-hand side could be larger than any y if η is sufficiently large. Let $\bar{\eta}(w,y) := \inf \bar{\mathcal{U}}(w,y)$.

By the super-linear growth assumption of function $g(\cdot)$, we have $g(x) > lx$ for any $l \in \mathbb{R}$ when x is sufficiently large. In particular, if we let $l = -\psi + 1$, then there exists some $\bar{x}$ such that $g(x) > (-\psi+1)x$ for all $x \geq \bar{x}$. Let $\eta = \bar{\eta}(1,\bar{x})$, then we have $g(f_0(1,\eta)) + \psi f_0(1,\eta) > f_0(1,\eta)$.

Hence, we can deduce

$$
f_0'(1,\eta) > \frac{2}{\sigma^2}\left(f_0(1,\eta) - \bar{h}(1) + \eta\right),
$$

which suggests there exists $\bar{w}$($\bar{w}$ could be ∞) such that $f_0(w,\eta) \geq \bar{x}$ on $(1,\bar{w})$. We also have $f_0'(w,\eta) > \frac{2}{\sigma^2}\left(f_0(w,\eta) - \bar{h}(w) + \eta\right)$ on $(1,\bar{w})$. Next, for any $w \in (1,\bar{w})$, we have

$$
\begin{aligned}
f_0(w,\eta) &\geq f_0(1,\eta) + \frac{2}{\sigma^2} e^{\frac{2}{\sigma^2}w} \int_1^w e^{-\frac{2}{\sigma^2}x} \left(-\bar{h}(x) + \eta\right) \mathrm{d}x \\
&= f_0(1,\eta) + \eta\left(e^{\frac{2}{\sigma^2}(w-1)} - 1\right) - \frac{2}{\sigma^2} e^{\frac{2}{\sigma^2}w} \int_1^w e^{-\frac{2}{\sigma^2}x} \bar{h}(x) \mathrm{d}x.
\end{aligned}
$$

By the polynomial growth of $\bar{h}(\cdot)$, we have $\bar{h}(x) \leq Ax^k + B$ for some constants $A, B \in \mathbb{R}^+$ and some integer $k > 0$ for all $x > 0$. Then, we have

$$
\begin{aligned}
f_0(w,\eta) &> f_0(1,\eta) + \eta\left(e^{\frac{2}{\sigma^2}(w-1)} - 1\right) - \frac{2}{\sigma^2} e^{\frac{2}{\sigma^2}w} \int_1^w e^{-\frac{2}{\sigma^2}x} (Ax^k + B) \mathrm{d}x \\
&> f_0(1,\eta) + \eta\left(e^{\frac{2}{\sigma^2}(w-1)} - 1\right) - \frac{2}{\sigma^2} e^{\frac{2}{\sigma^2}w} \int_1^\infty e^{-\frac{2}{\sigma^2}x} (Ax^k + B) \mathrm{d}x \\
&= f_0(1,\eta) + \left(\frac{\eta}{e^{\frac{2}{\sigma^2}}} - M\right) e^{\frac{2}{\sigma^2}w} - \eta,
\end{aligned}
$$

where M is some constant. Then there exists some η such that $\frac{\eta}{e^{\frac{2}{\sigma^2}}} - M > 0$ and $\left(\frac{\eta}{e^{\frac{2}{\sigma^2}}} - M\right) e^{\frac{2}{\sigma^2}w} - \eta > 0$. Denoting an arbitrary such η satisfying the inequality as $\hat{\eta}$, and let $\eta = \max\{\hat{\eta}, \bar{\eta}(1,\bar{x})\}$, we have $f_0(w,\eta) >$

$f_0(1,\eta)$ on $(1,\bar{w})$. Also, $f_0(w,\eta)$ is greater than a strictly increasing function on $(1,\bar{w})$. By the continuity of f_0, this suggests $f_0(w,\eta) > f_0(1,\eta)$ for all $1 \le w \le w_{\eta,\infty}$. However, this contradicts to $\lim_{w\to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. Hence we conclude that $\eta \in \mathcal{U}_0$. $\square$

Proof of Lemma EC.4. By contradiction, if $\eta_0 \in \mathcal{L}_0$, then $\lim_{w\to w_{\eta_0,\infty}} f_0(w,\eta_0) = -\infty$. Hence, there exists some $w_0 > 0$ such that $f_0(w_0,\eta_0) < 0$. By the continuity of $f_0(w,\eta)$ with respect to η, there exists $\varepsilon > 0$ such that $f_0(w_0,\eta_0+\varepsilon) < 0$, which contradicts the definition of η_0. Therefore, $\eta_0 \in \mathcal{U}_0$. We can also deduce that $\eta_0 = \inf \mathcal{U}_0$.

Next, by definition of $\mathcal{U}_0$, $f_0'(w,\eta_0) \ge 0$ for all $w \in (0, w_{\eta_0,\infty})$. We prove $w_{\eta_0,\infty} = \infty$. Assume by contradiction $w_{\eta_0,\infty} < \infty$. Then for any number $a \in \mathbb{R}$, there exists some w such that $f_0(w,\eta_0) > a$. By the continuity of $f_0(w,\eta)$ with respect to η, there exists $\varepsilon > 0$ such that $f_0(w,\eta-\varepsilon) > a$. Next, by the similar arguments as in the proof of Lemma EC.3, we can deduce that $f_0(w,\eta-\varepsilon)$ is greater than some strictly increasing function, which suggests $\eta - \varepsilon \in \mathcal{U}_0$. However, this contradicts to the definition of $\mathcal{U}_0$, so we conclude $w_{\eta_0,\infty} = \infty$. $\square$

Proof of Lemma EC.5. For clarity, we will assume in this proof that $\psi = 0$. The case for $\psi \neq 0$ can be proven in a similar manner, albeit with cumbersome notation. To start, for some $\tilde{w} > 0$, let $a > 0$ be such that

$$(2/\sigma^2)g'(\tilde{w})(\bar{h}(w)+a) - 2\bar{h}'(w) > 0. \tag{EC.11}$$

Note that the existence of such a constant a is ensured by Assumption 1. Notice that $\bar{h}$ may not be differentiable at some w. In that case, we can integrate (EC.11) and start from the inequality in the integral form to reach the same results. Define $\varphi_\eta(w) := g(f_0(w,\eta)) - 2\bar{h}(w) - a$, where the value of a will be defined more clearly in (EC.12). Now, let $\hat{w}_\eta$ be such that $f_0(\hat{w}_\eta,\eta) = \sup_{w\ge 0} f_0(w,\eta)$. Since $f_0'(\hat{w}_\eta,\eta) = 0$, we know that $g(f_0(\hat{w}_\eta,\eta)) - \bar{h}(\hat{w}_\eta) = -\eta < 0$. Hence, $\varphi_\eta(\hat{w}_\eta) < -\bar{h}(\hat{w}_\eta) - a \le 0$. Next, we argue that $\varphi_\eta(w) \le 0$ for all $w \in [0,\hat{w}_\eta]$. Suppose this is not true. Then there must exist some $\underline{w}_\eta$ and $\bar{w}_\eta$ with $0 < \underline{w}_\eta < \bar{w}_\eta < \hat{w}_\eta$ such that $\varphi_\eta(\underline{w}_\eta) = \varphi_\eta(\bar{w}_\eta) = 0$, and $\varphi_\eta(w) > 0$ for all $w \in (\underline{w}_\eta, \bar{w}_\eta)$. In particular, the foregoing implies that $g(f_0(w,\eta)) > 2\bar{h}(w) + a$, and so

$$\frac{\sigma^2}{2} f_0'(w,\eta) = g(f_0(w,\eta)) - \bar{h}(w) + \eta = g(f_0(w,\eta)) - 2\bar{h}(w) + \bar{h}(w) + \eta > \bar{h}(w) + a$$

for every $w \in (\underline{w}_\eta, \bar{w}_\eta)$. We note that the above arguments hold for any $a > 0$.

On the other hand, by the definition of φ_η and (EC.11), we can find some a such that

$$\varphi_\eta'(w) = g'(f_0(w,\eta))f_0'(w,\eta) - 2h'(w) > (2/\sigma^2)g'(g^{-1}(2\bar{h}(\underline{w}_\eta)))(\bar{h}(w)+a) - 2\bar{h}'(w) > 0. \tag{EC.12}$$

This, however, is a contradiction due to our hypothesis which holds that $\varphi_\eta(\underline{w}_\eta) = \varphi_\eta(\bar{w}_\eta) = 0$. As a result, $\varphi_\eta(w) \le 0$ for all $w \in [0,\hat{w}_\eta]$, from which the desired result follows. $\square$

The proofs of Proposition EC.1 and the related Lemmas EC.6-EC.10 follow the similar arguments as the proofs of Lemmas EC.1-EC.5, so they are omitted.

Proof of Lemma EC.11. We first prove if $\alpha_1 > \alpha_2$, then $\xi_{\alpha_1}(w,\eta(\alpha_1)) > \xi_{\alpha_2}(w,\eta(\alpha_2))$. Suppose by way of contradiction that this is not true. Then there exists $\bar{w}$ such that $\xi_{\alpha_1}(\bar{w},\eta(\alpha_1)) = \xi_{\alpha_2}(\bar{w},\eta(\alpha_2))$ and $\xi'_{\alpha_1}(\bar{w},\eta(\alpha_1)) \le \xi'_{\alpha_2}(\bar{w},\eta(\alpha_2))$. Then we can deduce that

$$\frac{\sigma^2}{2}\left(\xi'_{\alpha_1}(\bar{w},\eta(\alpha_1)) - \xi'_{\alpha_2}(\bar{w},\eta(\alpha_2))\right) = \eta(\alpha_1) - \eta(\alpha_2) \le 0.$$

By the uniqueness of solutions to ODE, $\eta(\alpha_1) < \eta(\alpha_2)$, so that there exists some $w > \bar{w}$ such that $\xi_{\alpha_1}(w,\eta(\alpha_1)) < \xi_{\alpha_2}(w,\eta(\alpha_2))$. Next, using the comparison theorem for differential equations as in the proof of Lemma EC.1, we know that $\xi_{\alpha_1}(w,\eta(\alpha_1)) < \xi_{\alpha_2}(w,\eta(\alpha_2))$ must hold for all $w \in (\bar{w},\infty)$. However, by the continuity of $\xi_\alpha(w,\eta)$ with respect to η, there exists some $\varepsilon > 0$ such that $\xi_{\alpha_2}(w,\eta(\alpha_2)-\varepsilon) > \xi_{\alpha_1}(w,\eta(\alpha_1))$ for all $w > \bar{w}$, which suggests $\xi_{\alpha_2}(w,\eta(\alpha_2)-\varepsilon)$ grows to infinity as w grows large. But this contradicts to the definition $\eta(\alpha) = \inf \mathcal{U}_\alpha$. Hence, $\xi_{\alpha_1}(w,\eta(\alpha_1)) > \xi_{\alpha_2}(w,\eta(\alpha_2))$ on $(0,\infty)$.

Next, we prove $\eta(\alpha_1) < \eta(\alpha_2)$ holds for all $\alpha_1 > \alpha_2$. Assume by contradiction this is not true, then there exists some $\alpha_1 > \alpha_2$ such that $\eta(\alpha_1) \ge \eta(\alpha_2)$. Using the standard comparison theorem for differential equations and the continuity of $\xi_\alpha(w,\eta)$ with respect to η again, we know that there exists $\varepsilon > 0$ such that $\xi_{\alpha_1}(w,\eta(\alpha_1)-\varepsilon)$ grows to infinity, which leads to the contradiction. Hence, $\eta(\alpha)$ is a decreasing function.

Next, we show $\eta(\cdot)$ is a continuous mapping. Consider an increasing sequence $\{\alpha_n\}$ with α being the limit. For ease of notation, we write $(\xi_{\alpha_n}(w,\eta(\alpha_n)),\eta(\alpha_n))$ as $(\xi_n(w),\eta_n)$. We aim to show that $\eta_n \to \eta(\alpha)$ as $n \to \infty$. The aforementioned arguments show that $\{\xi_n(w)\}$ is an increasing sequence and $\xi_n(w) \le \xi_\alpha(w,\eta(\alpha))$ for each fixed w. Hence $\xi_\infty(w) := \lim_{n\to\infty}\xi_n(w)$ is well defined for each fixed w. Integrating (28), we have

$$\frac{\sigma^2}{2}(\xi_n(w) - \alpha_n) = \int_0^w \left[(\psi+\bar{\gamma})\xi_n(x) + g(\xi_n(x)) - \bar{h}(x) - c + \eta\right] \mathrm{d}x.$$

Sending $n \to \infty$, we have

$$\frac{\sigma^2}{2}(\xi_\infty(w) - \alpha) = \int_0^w \left[(\psi+\bar{\gamma})\xi_\infty(x) + g(\xi_\infty(x)) - \bar{h}(x) - c + \eta\right] \mathrm{d}x.$$

By the uniqueness of the solution, we have $\xi_\infty(w) = \xi_\alpha(w,\eta(\alpha))$. Similar arguments also apply to a decreasing sequence of $\{\alpha_n\}$ with α being the limit. Hence, we have proved $\eta(\alpha)$ is a continuous mapping. The last part of the lemma can be established by using a routine proof-by-contradiction argument, which is elementary and thus omitted. □

Proof of Lemma EC.12. We first prove part (i). The infinitesimal generator for $(Y(t),W(t))$ process can be written as

$$\mathcal{L}f(y,w) = \begin{cases} \frac{\sigma^2}{2}f_{ww}(y,w) - \psi f_w(y,w) - \bar{\gamma}y f_w(y,w) - \boldsymbol{m}^T\boldsymbol{\vartheta}^\star(y,w)f_w(y,w), & \text{if} \quad w > 0, \\ \iota f_w(y,0), & \text{if} \quad w = 0, \end{cases} \tag{EC.13}$$

where ι is some constant (the derivation of the generator can be found on page 5 in Varadhan (2011)), and $\boldsymbol{\vartheta}^\star(y,w) = \arg\max_{\boldsymbol{\vartheta}}\{v_w(y,w)\boldsymbol{m}^T\boldsymbol{\vartheta} - \delta(\boldsymbol{\vartheta})\}$.

Define the Lyapunov function as $f(y,w) = w^k + 1$. We have $f(y,0) = 1$. For any $k \geq 2$, we can compute that

$$\mathcal{L}(w^k+1) = \frac{\sigma^2}{2}k(k-1)w^{k-2} - \psi k w^{k-1} - \bar{\gamma} y k w^{k-1} - \boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w) k w^{k-1}, \tag{EC.14}$$

for any $w > 0$.

Let $F(w) := w^{k-1} - w^{k-2}$. Then $F'(w) = (k-1)w^{k-2} - (k-2)w^{k-3} \geq 0$ if and only if $w \geq \frac{k-2}{k-1}$. So, $w^{k-2} \leq w^{k-1} - C_1$ where $C_1 = F(\frac{k-2}{k-1})$. Then we have

$$\begin{aligned}\mathcal{L}(w^k+1) &\leq \left(\frac{\sigma^2}{2}(k-1) - \psi - \boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w)\right) k w^{k-1} - \frac{\sigma^2}{2}k(k-1)C_1 \\ &= \left(\frac{\sigma^2}{2}(k-1) - \psi - \boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w)\right) k (w^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1.\end{aligned}$$

If $\frac{\sigma^2}{2}(k-1) - \psi < 0$, $\mathcal{L}(w^k+1) \leq \left(\frac{\sigma^2}{2}(k-1) - \psi\right) k (w^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1$. Let $\phi(x) := x^{\frac{k-1}{k}}$, then ϕ is a strictly concave function with $\phi(0) = 0$ and $\phi(x)$ increases to infinity as $x \to \infty$. Thus, $\phi(w^k+1) - \phi(w^k) = \int_{w^k}^{w^k+1} \phi'(t)\mathrm{d}t$ strictly decreases such that $\phi(w^k+1) - \phi(w^k) \leq \phi(1) - \phi(0) = 1$. Therefore, we have

$$\mathcal{L}(w^k+1) \leq k\left(\frac{\sigma^2}{2}(k-1) - \psi\right)\left((w^k+1)^{\frac{k-1}{k}} - 1\right) - \frac{\sigma^2}{2}k(k-1)C_1.$$

By Theorem 4.1 on page 16 of Hairer (2021), we have $\limsup_{t\to\infty} \mathbb{E}[([W(t)]^k + 1)^{\frac{k-1}{k}}] < \infty$. This implies $\limsup_{t\to\infty} \mathbb{E}[W(t)^{k-1}] < \infty$. Since $k \geq 2$ can be any large number, we have the desired result.

Next, suppose $\frac{\sigma^2}{2}(k-1) - \psi \geq 0$. By our assumption on the super linear growth of function g, we know that $\boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w)$ grows to infinity as $w \to \infty$. Therefore, for any constant $\tau \in \mathbb{R}$, we can find a $\bar{w}$ such that $\boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w) \geq \min(\boldsymbol{m}^T \boldsymbol{\vartheta}^\star(0,w), \boldsymbol{m}^T \boldsymbol{\vartheta}^\star(1,w)) \geq \tau$ for all $w \geq \bar{w}$. Let $\tau = \frac{\sigma^2}{2}(k-1) - \psi + 1$. Hence,

$$\boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w) \geq \left(\frac{\sigma^2}{2}(k-1) - \psi + 1\right) 1_{w \geq \bar{w}}.$$

Thus, we have

$$\mathcal{L}(w^k+1) \leq \begin{cases} \left(\frac{\sigma^2}{2}(k-1) - \psi\right) k (w^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1, & \text{if } \quad w < \bar{w}, \\ -k(w^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1, & \text{if } \quad w \geq \bar{w}, \end{cases}$$

which implies

$$\mathcal{L}(w^k+1) \leq -k(w^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1 + \left(\frac{\sigma^2}{2}(k-1) - \psi + 1\right) k (\bar{w}^k)^{\frac{k-1}{k}}. \tag{EC.15}$$

Use the above arguments and Theorem 4.1 of Hairer (2021) again, we have the desired result. So the proof of part (i) is complete.

Next, to prove part (ii), we apply Theorem 16.2 and Remark 16.1 in the notes of Varadhan (2011). For $f(w) = w^k$ with $k > 1$, it has polynomial growth and we have $f'(0) = 0$. Therefore, $f(W(t)) - f(W(0)) -$

$\int_0^t \mathcal{L}W(s)\mathrm{d}s$ is a martingale so that $\mathbb{E}\int_0^t \mathcal{L}W(s)\mathrm{d}s = \mathbb{E}[f(W(t))] - \mathbb{E}[f(W(0))]$. We can see that $\mathcal{L}(w^k) = \mathcal{L}(w^k+1)$. From our previous calculations in Equation (EC.15) in part (i), we get $\mathcal{L}(w^k) \leq -aw^{k-1}+b$ for some $a,b>0$. This gives

$$\mathbb{E}[W(t)^k] - \mathbb{E}[W(0)^k] = \mathbb{E}\left[\int_0^t \mathcal{L}W(s)^k \mathrm{d}s\right] \leq -\mathbb{E}\left[\int_0^t aW(s)^{k-1}\mathrm{d}s\right] + bt,$$

so that

$$\mathbb{E}\left[\int_0^t aW(s)^{k-1}\mathrm{d}s\right] \leq bt + \mathbb{E}[W(0)^k] - \mathbb{E}[W(t)^k] \leq bt + \mathbb{E}[W(0)^k]. \tag{EC.16}$$

Since it holds for any $k>1$, the proof of part (ii) is done.

□

Proof of Lemma EC.13. Assume $\eta \in \mathcal{L}_0$ and $f_0(w,\eta)$ exists. To prove $f_0(w,\eta)$ is quasi-concave, we need to show there exists only one local maximum point w_0 such that $f_0(w_0,\eta) \geq f_0(w,\eta)$ for all $w \in [-b, w_{\eta,\infty})$ and $\eta \in \mathcal{L}_0$. Let $w_0 := \inf\{w > -b : f_0'(w,\eta) < 0\}$. Then $f_0'(w_0,\eta) = 0$, and $f_0(w_0,\eta) \geq f_0(w,\eta)$ for all $w \in (-b, w_0)$. Let w_2 be the smallest value of w such that $w > w_0$ and w_2 is a local minimum of $f_0(w,\eta)$; if there exists no w satisfying these conditions then it is immediate that $f_0(w,\eta)$ is quasi-concave.

Case 1: Suppose $w_2 > 0$. Since w_2 is a local minimum, it follows that there exist $\bar{w}_0, w_1$ such that $w_2 > \bar{w}_0 > 0$, $w_1 > w_2$, $f_0(\bar{w}_0,\eta) = f_0(w_1,\eta)$, $f_0'(\bar{w}_0,\eta) < 0$, and $f_0'(w_1,\eta) > 0$. We have

$$\frac{\sigma^2}{2}\left(f_0'(\bar{w}_0,\eta) - f_0'(w_1,\eta)\right) = -\left(\bar{h}(\bar{w}_0) - \bar{h}(w_1)\right),$$

which is a contradiction because the left-hand side is strictly negative but the right-hand side is strictly positive.

Case 2: Suppose $w_2 \leq 0$. The tangent line to $f_0(w,\eta)$ at the point $w = w_2$ meets $f_0(w,\eta)$ at some point $w = w_3$ where $-b \leq w_3 < w_0$. To verify that w_3 exists, since $f_0(-b,\eta) = 0$, we need only check that $f_0(w,\eta)$ does not have a root in the interval $(-b, w_2)$; for if $f_0(c,\eta) = 0$ for $c \in (-b, w_2)$ then using the facts that $\bar{h}(w)$ is strictly decreasing in w for all $w < 0$ and that the slope of $f_0(w,\eta)$ is positive at $w = -b$ and negative at $w = c$ it follows that (27) cannot hold at both $w = -b$ and $w = c$, a contradiction. It is possible that $f_0(w,\eta)$ has a root at $w = w_2$; this situation corresponds to $w_3 = -b$. We have thus verified that $w = w_3$ exists, where $-b \leq w_3 < w_0$.

Writing (27) for $w = w_3$ and $w = w_2$ and subtracting we get the equation below where the left hand side is positive and the right hand side is strictly negative (since $\bar{h}(w)$ is strictly decreasing in $(-b, 0)$), yielding a contradiction:

$$\frac{\sigma^2}{2}\left(f_0'(w_3,\eta) - f_0'(w_2,\eta)\right) = \bar{h}(w_2) - \bar{h}(w_3).$$

This completes the proof of quasi-concavity.

Next, we prove $\lim_{w\to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. If $w_{\eta,\infty} < \infty$, due to $f_0'(w,\eta) \leq 0$ for all $w \in (w_0, w_{\eta,\infty})$,

$\lim_{w\to w_{\eta,\infty}} f_0(w,\eta) = -\infty$. If, on the other hand, $w_{\eta,\infty} = \infty$, assume $\lim_{w\to\infty} f_0(w,\eta) = l > -\infty$. We have $\lim_{w\to\infty} f_0'(w,\eta) = 0$, so $\lim_{w\to\infty} \bar{h}(w) = \eta + \psi l - g(l)$, which is a contradiction since the left-hand side is ∞ but the right-hand side is finite.

Finally, to prove $\mathcal{L}_0 \neq \emptyset$, we first notice that for all $\eta < 0$, $f_0'(-b,\eta) < 0$. Based on the above arguments, we have $f_0(w,\eta) < 0$, $f_0'(w,\eta) \leq 0$ for all $w \in (-b, w_{\eta,\infty})$ and $\lim_{w\to\infty} f_0(w,\eta) = -\infty$. By the continuity of $f_0(w,\eta)$ with respect to η, we have $f_0(w,0) = \lim_{\eta\uparrow 0} f_0(w,\eta) \leq 0$ and $f_0'(w,0) = \lim_{\eta\uparrow 0} f_0'(w,\eta) \leq 0$. But since $f_0(w,0) \equiv 0$ is not possible, there must exist some $w_1 > 0$ such that $f_0(w_1,0) < 0$. Using the continuity of $f_0(w,\eta)$ with respect to η again, we know that there exists some $\eta_1 > 0$ such that $f_0(w_1,\eta_1) < 0$. Hence, there exists some $w \in (-b, w_1)$ such that $f_0'(w,\eta_1) < 0$, so $\eta_1 \in \mathcal{L}_0$. □

Proof of Lemma EC.14. We first prove part (i). The infinitesimal generator for $(Y(t), W(t))$ process can be written as

$$\mathcal{L}f(y,w) = \begin{cases} \frac{\sigma^2}{2} f_{ww}(y,w) - \psi f_w(y,w) - \bar{\gamma} y f_w(y,w) - \boldsymbol{m}^T \boldsymbol{\vartheta}^\star(y,w) f_w(y,w), & \text{if} \quad w > -b, \\ \iota f_w(y,-b), & \text{if} \quad w = -b, \end{cases} \tag{EC.17}$$

where ι is some constant (the derivation of the generator can be found on page 5 in Varadhan (2011)), and $\boldsymbol{\vartheta}^\star(y,w) = \arg\max_{\boldsymbol{\vartheta}} \{v_w(y,w)\boldsymbol{m}^T\boldsymbol{\vartheta} - \delta(\boldsymbol{\vartheta})\}$.

Define the Lyapunov function as $f(y,w) = (w+b)^k + 1$. We have $f(y,-b) = 1$. For any $k \geq 2$, we can compute that

$$\mathcal{L}((w+b)^k+1) = \frac{\sigma^2}{2} k(k-1)(w+b)^{k-2} - \psi k(w+b)^{k-1} - \bar{\gamma} y k(w+b)^{k-1} - \boldsymbol{\vartheta}^\star(y,w) k(w+b)^{k-1} \tag{EC.18}$$

for any $w > -b$.

Let $F(w) := (w+b)^{k-1} - (w+b)^{k-2}$. Then $F'(w) = (k-1)(w+b)^{k-2} - (k-2)(w+b)^{k-3} \geq 0$ if and only if $w \geq \frac{k-2}{k-1} - b$. Therefore, $(w+b)^{k-2} \leq (w+b)^{k-1} - C_1$ where $C_1 = F(\frac{k-2}{k-1} - b)$. Then we have

$$\begin{aligned} \mathcal{L}((w+b)^k+1) &\leq \left(\frac{\sigma^2}{2}(k-1) - \psi - \boldsymbol{m}^T\boldsymbol{\vartheta}^\star(y,w)\right) k(w+b)^{k-1} - \frac{\sigma^2}{2}k(k-1)C_1 \\ &= \left(\frac{\sigma^2}{2}(k-1) - \psi - \boldsymbol{m}^T\boldsymbol{\vartheta}^\star(y,w)\right) k((w+b)^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1. \end{aligned}$$

If $\frac{\sigma^2}{2}(k-1) - \psi < 0$, then

$$\mathcal{L}((w+b)^k+1) \leq \left(\frac{\sigma^2}{2}(k-1) - \psi\right) k((w+b)^k)^{\frac{k-1}{k}} - \frac{\sigma^2}{2}k(k-1)C_1.$$

Let $\phi(x) := x^{\frac{k-1}{k}}$, then ϕ is a strictly concave function in $x > 0$ with $\phi(0) = 0$ and $\phi(x)$ increases to infinity as $x \to \infty$. Thus, $\phi((w+b)^k+1) - \phi((w+b)^k) \leq \phi(1) - \phi(0) = 1$. Therefore, we have

$$\mathcal{L}((w+b)^k+1) \leq k\left(\frac{\sigma^2}{2}(k-1) - \psi\right)\left(((w+b)^k+1)^{\frac{k-1}{k}} - 1\right) - \frac{\sigma^2}{2}k(k-1)C_1.$$

By Theorem 4.1 on page 16 of Hairer (2021), we have $\limsup_{t\to\infty}\mathbb{E}[([W(t)+b]^k+1)^{\frac{k-1}{k}}]<\infty$. Since $\mathbb{E}[([W(t)+b]^k+1)^{\frac{k-1}{k}}1_{-b\le W(t)\le 0}]\in[0,(b^k+1)^{\frac{k-1}{k}}]$, we have

$$\limsup_{t\to\infty}\mathbb{E}[([W(t)+b]^k+1)^{\frac{k-1}{k}}1_{W(t)>0}]<\infty.$$

This implies

$$\begin{aligned}\limsup_{t\to\infty}\mathbb{E}[|W(t)|^{k-1}] &= \limsup_{t\to\infty}\left(\mathbb{E}[|W(t)|^{k-1}1_{W(t)>0}]+\mathbb{E}[|W(t)|^{k-1}1_{-b\le W(t)\le 0}]\right)\\ &\le \limsup_{t\to\infty}\left(\mathbb{E}[([W(t)+b]^k+1)^{\frac{k-1}{k}}1_{W(t)>0}]+b^{k-1}\right),\end{aligned}$$

where $k\ge 2$. Since $k\ge 2$ can be any large number, we have the desired result.

Next, suppose $\frac{\sigma^2}{2}(k-1)-\psi\ge 0$. By our assumption on the super linear growth of function g, we know that $\boldsymbol{m}^T\boldsymbol{\vartheta}^\star(y,w)$ grows to infinity as $w\to\infty$. Therefore, for any constant $\tau\in\mathbb{R}$, we can find a $\bar{w}$ such that $\boldsymbol{m}^T\boldsymbol{\vartheta}^\star(y,w)\ge\min(\boldsymbol{m}^T\boldsymbol{\vartheta}^\star(0,w),\boldsymbol{m}^T\boldsymbol{\vartheta}^\star(1,w))\ge\tau$ for all $w\ge\bar{w}$. Let $\tau=\frac{\sigma^2}{2}(k-1)-\psi+1$. Hence,

$$\boldsymbol{\vartheta}^\star(y,w)\ge\left(\frac{\sigma^2}{2}(k-1)-\psi+1\right)1_{w\ge\bar{w}}.$$

So, we have

$$\mathcal{L}((w+b)^k+1)\le\begin{cases}\left(\frac{\sigma^2}{2}(k-1)-\psi\right)k((w+b)^k)^{\frac{k-1}{k}}-\frac{\sigma^2}{2}k(k-1)C_1, & \text{if} \quad w<\bar{w},\\ -k((w+b)^k)^{\frac{k-1}{k}}-\frac{\sigma^2}{2}k(k-1)C_1, & \text{if} \quad w\ge\bar{w},\end{cases}$$

which implies

$$\mathcal{L}((w+b)^k+1)\le -k((w+b)^k)^{\frac{k-1}{k}}-\frac{\sigma^2}{2}k(k-1)C_1+\left(\frac{\sigma^2}{2}(k-1)-\psi+1\right)k(\bar{w}^k)^{\frac{k-1}{k}}. \tag{EC.19}$$

Use the above arguments and Theorem 4.1 of Hairer (2021) again, we have the desired result. Thus the proof of part (i) is complete.

Next, to prove part (ii), we apply Theorem 16.2 and Remark 16.1 in the notes of Varadhan (2011). For $f(w)=(w+b)^k$ with $k>1$, it has polynomial growth and we have $f'(-b)=0$. Therefore, $f(W(t))-f(W(0))-\int_0^t\mathcal{L}W(s)\mathrm{d}s$ is a martingale so that $\mathbb{E}\int_0^t\mathcal{L}W(s)\mathrm{d}s=\mathbb{E}[f(W(t))]-\mathbb{E}[f(W(0))]$. We can see that $\mathcal{L}((w+b)^k)=\mathcal{L}((w+b)^k+1)$. From our previous calculations in Equation (EC.19) in part (i), we get $\mathcal{L}((w+b)^k)\le-\bar{a}(w+b)^{k-1}+\bar{b}$ for some $\bar{a},\bar{b}>0$. This gives

$$\mathbb{E}[(W(t)+b)^k]-\mathbb{E}[(W(0)+b)^k]=\mathbb{E}\left[\int_0^t\mathcal{L}(W(s)+b)^k\mathrm{d}s\right]\le-\mathbb{E}\left[\int_0^t\bar{a}(W(s)+b)^{k-1}\mathrm{d}s\right]+\bar{b}t,$$

so that

$$\mathbb{E}\left[\int_0^t\bar{a}(W(s)+b)^{k-1}\mathrm{d}s\right]\le\bar{b}t+\mathbb{E}[(W(0)+b)^k]-\mathbb{E}[(W(t)+b)^k]\le\bar{b}t+\mathbb{E}[(W(0)+b)^k], \tag{EC.20}$$

and therefore

$$\begin{aligned}\mathbb{E}\left[\int_0^t\bar{a}|W(s)|^{k-1}\mathrm{d}s\right] &= \mathbb{E}\left[\int_0^t\left(\bar{a}|W(s)|^{k-1}1_{-b\le W(s)\le 0}+\bar{a}(W(s))^{k-1}1_{W(s)>0}\right)\mathrm{d}s\right]\\ &\le \bar{a}tb^{k-1}+\mathbb{E}\left[\int_0^t\bar{a}(W(s)+b)^{k-1}\mathrm{d}s\right]\le\bar{a}tb^{k-1}+\bar{b}t+\mathbb{E}[(W(0)+b)^k].\end{aligned}$$

Since it holds for any $k>1$, the proof of part (ii) is done. $\square$

EC.3 Comment on the Scaling Conditions and Baseline Quantities/Functions

We first argue that, under the scaling conditions specified in (5), all terms in the objective function given by (13) have an order of $\sqrt{\Lambda}$ as $\Lambda \to \infty$. The third and fourth terms are clearly of order $\sqrt{\Lambda}$ by the scaling conditions on c and C, namely, $c = \sqrt{\Lambda}\hat{c}, C = \sqrt{\Lambda}\hat{C}$, plus the fact that $Y(t)$ takes values in $\{0,1\}$. The queue length in heavy traffic is known to be of order $\sqrt{\bar{\lambda}}$, and hence $\sqrt{\Lambda}$. By the scaling condition for h_k, we know that the second term in the objective is also of order $\sqrt{\Lambda}$. To show that the first term is also of order $\sqrt{\Lambda}$, we first use Taylor's series expansion of $\hat{\Xi}^{-1}$ at $\bar{\boldsymbol{\lambda}}$:

$$\hat{\Xi}^{-1}(\boldsymbol{\lambda}) \approx \hat{\Xi}^{-1}(\bar{\boldsymbol{\lambda}}) + \nabla\hat{\Xi}^{-1}(\bar{\boldsymbol{\lambda}})(\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}) + o(\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|),$$

where $\nabla\hat{\Xi}^{-1}(\bar{\boldsymbol{\lambda}})$ represents the Jacobian matrix of $\hat{\Xi}^{-1}$ at $\bar{\boldsymbol{\lambda}}$ in this context. Hence, we have

$$\begin{aligned}
\delta(\boldsymbol{\vartheta}) &= \pi(\bar{\boldsymbol{\lambda}}) - \pi(\bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta}) \\
&= \sqrt{\Lambda}\left\langle \bar{\boldsymbol{\lambda}}, \hat{\Xi}^{-1}\left(\bar{\boldsymbol{\lambda}}/\Lambda\right) - \hat{\boldsymbol{q}} \right\rangle - \sqrt{\Lambda}\left\langle \bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta}, \hat{\Xi}^{-1}\left((\bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta})/\Lambda\right) - \hat{\boldsymbol{q}} \right\rangle \\
&\approx \sqrt{\Lambda}\left\langle \bar{\boldsymbol{\lambda}}, \hat{\Xi}^{-1}\left(\bar{\boldsymbol{\lambda}}/\Lambda\right) - \hat{\boldsymbol{q}} \right\rangle - \sqrt{\Lambda}\left\langle \bar{\boldsymbol{\lambda}} - \boldsymbol{\vartheta}, \hat{\Xi}^{-1}\left(\bar{\boldsymbol{\lambda}}/\Lambda\right) - \nabla\hat{\Xi}^{-1}\left(\bar{\boldsymbol{\lambda}}/\Lambda\right)(\boldsymbol{\vartheta}/\Lambda) - \hat{\boldsymbol{q}} + O(1/\Lambda) \right\rangle \\
&= \frac{-1}{\sqrt{\Lambda}}\left\langle \boldsymbol{\vartheta}, \nabla\hat{\Xi}^{-1}\left(\bar{\boldsymbol{\lambda}}/\Lambda\right)\boldsymbol{\vartheta} \right\rangle + O(\sqrt{\Lambda}),
\end{aligned}$$

where we have used the first order condition

$$\hat{\Xi}^{-1}\left(\frac{\bar{\boldsymbol{\lambda}}}{\Lambda}\right) - \hat{q} + (\nabla\hat{\Xi}^{-1})^T\left(\frac{\bar{\boldsymbol{\lambda}}}{\Lambda}\right)\frac{\bar{\boldsymbol{\lambda}}}{\Lambda} = \mathbf{0}.$$

Hence, the first term is also of order $\sqrt{\Lambda}$. In conclusion, all terms in the objective function are of the same order of $\sqrt{\Lambda}$.

EC.4 An Exact MDP Solution to Multi-Product Systems

EC.4.1 State Descriptions and Transition Rates

To apply the MDP framework, we need further restrictions on assumptions about the processing times of primary and secondary servers. Let the processing time of class k jobs on the primary server be exponentially distributed with rate μ_k and the processing time on the secondary server be exponentially distributed with rate γ_k.

We next define the state variables for the exact continuous-time MDP solution framework under a multi-product system. Let $(M_1(t), \mathbf{Q}(t), M_2(t), Y(t))$ be the state descriptor, where $\mathbf{Q}(t)$ is the K-dimensional vector representing the number of jobs waiting in buffers and $Y(t) = 0$ or 1 is a scalar denoting the secondary server is either on or off. The two new defined scalars $M_1(t)$ and $M_2(t)$ represent whether or not the primary and secondary servers are busy and, if so, which class of jobs they are working on. To be more specific,

$M_1(t) \in \{0,1,\dots,K\}$, where $M_1(t)=0$ indicates the primary server is idle while $M_1(t)=k>0$ represents the server is busy processing the class k jobs. $M_2(t)$ is defined in a similar way.

Next, we describe the state transitions under different actions. There are three control levers at the system manager's discretion: pricing, scheduling, and activating or deactivating the surge capacity. Pricing control can be represented by $\boldsymbol{\lambda}$, the arrival rates. Let $m_1(m_2) \in \{0,1,\dots,K\}$ be the priority rule of accepting the class $m_1(m_2)$ jobs to the primary server (secondary server), if there exists some class $m_1(m_2)$ jobs waiting in the queue for $m_1,m_2>0$. 0 indicates to make the server idle, and under the non-idling policy, this action only happens when $Q_k(t)=0$ for all k. Henceforth, we use the tuple $(\boldsymbol{\lambda},m_1,m_2)$ to represent the control action, without changing the status of surge capacity.

We assume the system is non-preemptive, such that both the primary and secondary servers have to complete processing the current in-service jobs before accepting new ones. In addition, we assume jobs are routed to the primary server if both primary and secondary servers are idle. We next write down the transition rate from state X to X' under the control $(\boldsymbol{\lambda},m_1,m_2)$ as $P_{(\boldsymbol{\lambda},m_1,m_2)}(X,X')$, where $X=(M_1,\mathbf{Q},M_2,Y)$ and $X'=(M_1',\mathbf{Q}',M_2',Y')$. Let $\mathbf{e}_k$ be the K-dimensional vector with the k-th component being 1 and all other components being 0. We divide the discussion into two scenarios: $Y=0$ or $Y=1$.

When $Y=0$, we have for $k=1,\dots,K$, the state transitions due to arrival events are

$$
\begin{aligned}
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),(M_1,\mathbf{Q}+\mathbf{e}_k,M_2,0))=\lambda_k, \quad \text{if} \quad M_1\neq 0,\\
&P_{(\boldsymbol{\lambda},m_1,m_2)}((0,\mathbf{Q},M_2,0),(k,\mathbf{Q},M_2,0))=\lambda_k, \quad \text{if} \quad M_1=0.
\end{aligned}
$$

The transition rates due to service completions when $Y=0$ are

$$
\begin{aligned}
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),(m_1,\mathbf{Q}-\mathbf{e}_{m_1}1_{\{m_1\neq 0\}},M_2,0))=\mu_{M_1}1_{\{M_1\neq 0\}},\\
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),(M_1,\mathbf{Q},0,0))=\gamma_{M_2}1_{\{M_2\neq 0\}}.
\end{aligned}
$$

When $Y=1$, we have for $k=1,\dots,K$, the state transitions due to arrival events are

$$
\begin{aligned}
&P_{(\boldsymbol{\lambda},m_1,m_2)}((0,\mathbf{0},M_2,1),(k,\mathbf{0},M_2,1))=\lambda_k, \quad \text{if} \quad M_1=0,\\
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{0},0,1),(M_1,\mathbf{0},k,1))=\lambda_k, \quad \text{if} \quad M_1\neq 0 \quad \text{but} \quad M_2=0,\\
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,1),(M_1,\mathbf{Q}+\mathbf{e}_k,M_2,1))=\lambda_k, \quad \text{if} \quad M_1\neq 0 \quad \text{and} \quad M_2\neq 0.
\end{aligned}
$$

For the service completion, when $Y=1$, we have

$$
\begin{aligned}
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,1),(m_1,\mathbf{Q}-\mathbf{e}_{m_1}1_{\{m_1\neq 0\}},M_2,1))=\mu_{M_1}1_{\{M_1\neq 0\}},\\
&P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,1),(M_1,\mathbf{Q}-\mathbf{e}_{m_2}1_{\{m_2\neq 0\}},m_2,1))=\gamma_{M_2}1_{\{M_2\neq 0\}}.
\end{aligned}
$$

In addition to the transition rates discussed above in the case that surge capacity status is kept unchanged, we also need to specify the immediate system change when the surge capacity is activated or deactivated. This immediate system change can be characterized by a map T_m. Specifically, we have

$$\begin{aligned}
&T_m((M_1,\mathbf{Q},M_2,0)) = (M_1,\mathbf{Q},M_2,1), \quad \text{if} \quad M_2 \neq 0,\\
&T_m((M_1,\mathbf{Q},M_2,0)) = (M_1,\mathbf{Q}-\mathbf{e}_m 1_{\{m\neq 0\}},m,1), \quad \text{if} \quad M_2 = 0,\\
&T_m((M_1,\mathbf{Q},M_2,1)) = (M_1,\mathbf{Q}+\mathbf{e}_{M_2} 1_{\{M_2\neq 0\}},0,0).
\end{aligned}$$

EC.4.2 Uniformization and Optimality Conditions

Directly applying the continuous-time MDP framework is not amenable for computation. Hence, we apply the uniformization technique to convert the continuous-time MDP to its embedded discrete-time MDP. To this end, let us make a further assumption on the arrival rates such that $0 \leq \sum_{k=1}^{K} \lambda_k(t) \leq \Lambda$ for all $t \geq 0$. This assumption is not restrictive since Λ can be regarded as the external job arrival rates. To ease the notation, we define a constant

$$\bar{\alpha} := \Lambda + \max_k \mu_k + \max_k \gamma_k.$$

We are now ready to write down the Bellman equation of the embedded discrete-time MDP, based on the defined state descriptors and transition rates. Let V be the relative value function and η be the long-run average cost. Also, let $P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),\cdot)$ denote the row of transition rate matrix. We have

$$\begin{aligned}
&V((M_1,\mathbf{Q},M_2,0)) + \frac{\eta}{\bar{\alpha}}\\
&= \max\Bigg\{ \max_{\boldsymbol{\lambda},m_1,m_2} \Bigg[\frac{1}{\bar{\alpha}}\Bigg(\langle P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),\cdot),V\rangle + \pi(\boldsymbol{\lambda}) - \sum_k h_k(Q_k + 1_{\{M_1=k\}} + 1_{\{M_2=k\}})\Bigg)\\
&\qquad + \Bigg(1 - \frac{\|P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),\cdot)\|_1}{\bar{\alpha}}\Bigg)V((M_1,\mathbf{Q},M_2,0))\Bigg],\\
&\quad -C + \max_{\boldsymbol{\lambda}',m_1',m_2',m}\Bigg[\frac{1}{\bar{\alpha}}\Bigg(\langle P_{(\boldsymbol{\lambda}',m_1',m_2')}(T_m(M_1,\mathbf{Q},M_2,0),\cdot),V\rangle + \pi(\boldsymbol{\lambda}') - c\\
&\qquad - \sum_k h_k(T_m(M_1,\mathbf{Q},M_2,0)(2,k) + 1_{\{M_1=k\}} + 1_{\{M_2=k\}})\Bigg)\\
&\qquad + \Bigg(1 - \frac{\|P_{(\boldsymbol{\lambda}',m_1',m_2')}(T_m(M_1,\mathbf{Q},M_2,0),\cdot)\|_1}{\bar{\alpha}}\Bigg)V(T_m(M_1,\mathbf{Q},M_2,0))\Bigg]\Bigg\},
\end{aligned}$$

and

$$\begin{aligned}
&V((M_1,\mathbf{Q},M_2,1)) + \frac{\eta}{\bar{\alpha}}\\
&= \max\Bigg\{ \max_{\boldsymbol{\lambda},m_1,m_2} \Bigg[\frac{1}{\bar{\alpha}}\Bigg(\langle P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,1),\cdot),V\rangle + \pi(\boldsymbol{\lambda}) - \sum_k h_k(Q_k + 1_{\{M_1=k\}} + 1_{\{M_2=k\}}) - c\Bigg)\\
&\qquad + \Bigg(1 - \frac{\|P_{(\boldsymbol{\lambda},m_1,m_2)}((M_1,\mathbf{Q},M_2,0),\cdot)\|_1}{\bar{\alpha}}\Bigg)V((M_1,\mathbf{Q},M_2,1))\Bigg],
\end{aligned}$$

$$\max_{\boldsymbol{\lambda}',m_1',m_2',m} \left[\frac{1}{\bar{\alpha}}\Bigg(\langle P_{(\boldsymbol{\lambda}',m_1',m_2')}\left(T_m(M_1,\mathbf{Q},M_2,1),\cdot\right),V\rangle + \pi(\boldsymbol{\lambda}') \right.$$
$$- \sum_k h_k(T_m(M_1,\mathbf{Q},M_2,1)(2,k) + 1_{\{M_1=k\}} + 1_{\{M_2=k\}})\Bigg)$$
$$\left. + \left(1 - \frac{\|P_{(\boldsymbol{\lambda}',m_1',m_2')}\left(T_m(M_1,\mathbf{Q},M_2,1),\cdot\right)\|_1}{\bar{\alpha}}\right) V\left(T_m(M_1,\mathbf{Q},M_2,1)\right)\right]\Bigg\},$$

where $\|\cdot\|_1$ denotes the L^1-norm. The optimal policy can then be derived by solving the Bellman equations.